University of California
Santa Barbara

Noise-State Calculus for Dynamic Games with Strategic Information

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Statistics and Applied Probability

by

Samuel Babichenko

Committee in charge:

Professor Tomoyuki Ichiba, Chair
Professor Jean-Pierre Fouque
Professor Michael Ludkovski

September 2026

The Dissertation of Samuel Babichenko is approved.

Professor Jean-Pierre Fouque

Professor Michael Ludkovski

Professor Tomoyuki Ichiba, Committee Chair

June 2026

Noise-State Calculus for Dynamic Games with Strategic Information

Copyright © 2026
by
Samuel Babichenko

To Viki

Acknowledgements

My first thanks go to my grandmother, Lyudmila, who left her job and her friends to move across the world and raise me through the first years of my life. Without her, I would not have been in a position to write this dissertation.

I am grateful to my advisor, Tomoyuki Ichiba, for his guidance and support throughout my doctoral studies, and to Javier Birchenall for valuable discussions on the economic applications of this work. I thank my committee members Jean-Pierre Fouque and Michael Ludkovski for their time and feedback.

I thank my girlfriend, Viki, for her love and for taking on my physical and mental burden throughout writing this. Without her, I would not have been able to finish this dissertation.

I am grateful to the rest of my family: my sister Anna, my cousin Max, my mother Lena, my father Leonard, my aunt Ira, my uncle Andrew, and my late grandfather Genrikh. Finally, I thank my friends Thiha Aung, Olivier Mulkin, Yan Lashchev, and Daniel Naylor for their encouragement and companionship throughout my time as a doctoral student.

Vita

Samuel Babichenko

Education

Ph.D. in Statistics and Applied Probability, emphasis in Financial Mathematics, UC Santa Barbara	2026
M.A. in Statistics, UC Santa Barbara	2026
B.S. in Mathematics, UC San Diego	2021
A.S. in Physics, Chemistry, and Mathematics, Orange Coast College	2019

Professional Employment

Teaching Assistant, Department of Statistics and Applied Probability, UC Santa Barbara	2021–2026

Presentations

"Forecasting and Manipulating the Forecasts of Others." Southern California Quantitative Finance Forum (SCQF), UC Santa Barbara, April 14, 2026.

Publications

Babichenko, Sam. "Forecasting and Manipulating the Forecasts of Others."
arXiv preprint arXiv:2603.12140, 2026.

Fields of Study

Major Field: Statistics and Applied Probability.
Studies in dynamic games with private information, stochastic filtering, and market microstructure with Professor Tomoyuki Ichiba.

Abstract

Noise-State Calculus for Dynamic Games with Strategic Information
by
Samuel Babichenko

Dynamic games with dispersed private information look hard, even when linearized, because each player needs forecasts of others' forecasts, without end. The regress comes from a loop. Actions move the state, the state generates observations, observations update beliefs, and beliefs determine actions. Current approaches become tractable by cutting the loop. This dissertation builds a noise-state calculus that keeps the loop intact.

Each player keeps estimates of the primitive shocks instead of belief hierarchies about the endogenous state. Under perfect information, these estimates are the shocks themselves, and the calculus reduces to ordinary impulse responses. Beliefs, prices, and policy rules are deterministic impulse responses in continuous-time linear-quadratic-Gaussian games. Equilibrium is a fixed point in those functions and is solved numerically. A finite-deviation identity verifies a computed Nash equilibrium against arbitrary admissible deviations.

The information wedge is the shadow price of changing an opponent's beliefs. It breaks the separation principle and vanishes when the loop is cut. In a two-player benchmark, a planner can starve an inefficient player by reallocating signal precision.

Signals and actions may arrive late without changing the calculus. Stationary systems describe kernels in terms of ages and measure distance from finite-dimensionality. In Kyle–Back markets, traders differ in signal precision within and across markets and face trader-specific price impact. On symmetric networks of local markets, shocks average out but strategic restraint does not, lowering mean orders while nearly invisible in responses to aggregate shocks. When some can identify deviations and others cannot, naive players attribute effects to ordinary shocks while privy players respond to labeled deviations. In computed markets, strategic market makers put more weight on inventory in their quotes when order flow is public and bear lower real inventory cost, with price discovery unchanged.

Contents

Notation

Times, indices, and conventions

$t,\ s,\ \tau$	calendar times
$u,\ r$	primitive source times
$a,\ b$	ages; ℓ the two-sided lag
$X_t(u)$ vs. $x(a)$	finite-horizon objects carry calendar time as a subscript and source time as an argument; stationary objects are functions of ages or lags
$X_t,\ X_t(u)$	a process and its kernel share a letter: the state, and its response to the shock born at u
$D_t^i,\ D_t^i(u)$	the control and its policy kernel
$B_X^X(t),\ G^{XX,i}(t),\ \sigma(t)$	exogenous coefficients carry time in parentheses
$X_t,\ D_t^i(u),\ H_t^{k,i}(u)$	processes, kernels, and equilibrium gains carry it as a subscript
$i,\ j,\ k \in \mathcal{I}$	player indices, as superscripts
$x(a)$ vs. $X_t(u)$	lowercase letters are the stationary current-value counterparts of uppercase finite-horizon objects
$D_{W,t}^i(u),\ d_W^i(a)$	a subscript W marks a kernel written in primitive-shock coordinates
$\mathcal{F},\ \mathcal{B},\ \ldots$	calligraphic letters are sets, σ-fields, and operators; the wedge $\mathcal{V}_t^i$ is the one calligraphic process

Accents

$\widehat{W},\ \widehat{X}$	hats are conditional expectations
$\widetilde{X},\ \widetilde{C}$	tildes are unresolved parts
$\bar{D},\ \bar{h}$	bars are deterministic means
$\check{f}$	checks are transforms
C; $\mathsf{C}_t^i(u,r)$; $\widetilde{C}_t^{i,Y}(u)$	an observation kernel; a conditional covariance; the unresolved observation row

Letters whose meaning is local to a chapter

z	a control spike (Chapter 4's raw order spike is the same object) or a transform variable, never a time
$v = B_i^X(t)z$; $v_i(s)$	the state impulse a spike induces; a deviation seed (Chapter 6)
$\delta,\ \delta_i$; $\delta w_t^k(u)$	a first variation, one with respect to an origin-i seed; the noise-state perturbation in density form
$\Delta_i,\ \Delta_r$; Δ	observation delays; in the appendix to Chapter 1, a finite difference
ζ	the noise-trader variance asymmetry, $\Sigma_Z = \mathrm{diag}(1+\zeta, 1-\zeta)$ (Chapter 4)
$v,\ w$; r	vertices and a group displacement (Chapter 5)
q	the number of assets (Chapter 4); the delayed source date $q_k = s - \Delta_k$ (Chapter 2)
(p, q)	Laplace variables (Chapter 3); naive and privy indices (Chapter 6)
$j \trianglerighteq i$; $\mathcal{P}_i,\ \mathcal{N}_i$	player j observes deviations whose origin is player i; the players privy and naive to an origin-i deviation (Chapter 6)

The following table collects the objects carried across chapters and marks where each is developed. Chapter-specific objects, proof-local variables, and one-use abbreviations are defined in place.

Core finite-horizon objects

Symbol	Object	Scope
$W = (W^0, W^1, \ldots, W^n)$	primitive Brownian shock path, with common and player-specific coordinates	General
$\mathcal{I}$; i, j, k	player set and generic player indices	General
$\mathcal{F}_t,\ \mathcal{F}_t^i$	primitive and player-i information filtrations	General
E^i	block selector for player i's noise channel, $W^i = E^i W$	General
$X_t,\ X_t(u)$	physical state and its primitive response to the shock born at u	Ch. 1
$B_X^X(t),\ B_i^X(t),\ \sigma(t)$	drift, control gains, and noise loading of the state, $dX_t = (B_X^X(t)X_t + \sum_i B_i^X(t)D_t^i)dt + \sigma(t)\, dW_t^0$	Ch. 1
$Y_t^i,\ B_X^{Y,i},\ P^i(t)$	player i's observation process, observation map, and precision $B_X^{Y,i\top} B_X^{Y,i}$	Ch. 1
dI_t^i	player i's innovation, $dY_t^i - B_X^{Y,i}\widehat{X}_t^i\, dt$	Ch. 1

Symbol	Object	Scope
$\widehat{W}_t{}^i(u)$	noise-state: player i's conditional estimate of W_u at time t	Ch. 1
$\Pi^i,\ F_t^i(u,s)$	direct projection and filter kernel; together, the blueprint of player i's beliefs	Ch. 1
$\widetilde{X}_t^i(u)$	unresolved response: the part of $X_t(u)$ that player i's information does not resolve	Ch. 1
$\delta w_t^k(u)$	noise-state perturbation in density form, $d_u\,\delta\widehat{W}_t^k(u) = \delta w_t^k(u)\,du$	Chs. 1, 2, 4, 6
$D_t^i,\ D_t^i(u),\ D_{W,t}^i(u)$	control, its policy kernel, and its control kernel in primitive-shock coordinates	Ch. 1
$H_t^{X,i}$	state price: player i's adjoint to the physical state	Ch. 1
$H_t^{k,i}(u)$	belief price: player i's adjoint to player k's noise-state coordinate at source time u	Ch. 1
$\mathcal{V}_t^i$	information wedge: the price of moving opponents' noise-states, entering the dynamics of the state price	Chs. 1, 2, 4, 5, 6
$\mathfrak{B}^k,\ \mathfrak{B}^{k\top}$	filter contraction: integrates a belief price against player k's own filter row; and its transpose	Chs. 1, 2, 4, 6
ρ	exponential discount rate	Chs. 3, 4

Finite-horizon and stationary forms

The stationary chapter changes coordinates rather than economic objects. Calendar dates become ages or lags, and finite-horizon uppercase objects become lowercase current-value objects; processes keep their letters:

Finite-horizon form	Stationary form	Interpretation
$X_t(s)$	$x(t-s) = x(a)$	state-response kernel
$\widetilde{X}_t^i(s)$	$\tilde{x}^i(t-s) = \tilde{x}^i(a)$	unresolved response
$D_t^i(u)$	$d^i(t-u) = d^i(a)$	noise-state policy kernel
$D_{W,t}^i(s)$	$d_W^i(t-s) = d_W^i(a)$	control kernel in primitive-shock coordinates
$F_t^i(u,s)$	$f^i(t-u, t-s) = f^i(b,a)$	filter kernel
$\widehat{W}_t^i(u)$	$\widehat{W}_t^i(t-a)$	noise-state (a process; stays uppercase)
$H_t^{X,i}$	$\bar{h}^{X,i}$ and $h^{X,i}(\ell)$	mean and two-sided state price
$H_t^{k,i}(u)$	$\bar{h}^{k,i}(b)$ and $h^{k,i}(b,\ell)$	belief price
$\mathcal{V}_t^i$	$\bar{v}^i$ and $v^i(\ell)$	information wedge

Glossary

Term	Meaning	Where
noise-state	a player's conditional estimate of the primitive shock path, $\widehat{W}^i_t(u) = \mathbb{E}[W_u \mid \mathcal{F}^i_t]$; the coordinates in which strategies, filters, and prices are written	Introduction; Thm. 1.5
blueprint	the pair (Π^i, F^i) of direct projection and filter kernel, the deterministic map from primitive shocks to a player's estimated shocks	Thm. 1.5
unresolved response	the part $\widetilde{X}^i_t(u)$ of the state's impulse response that a player's information has not resolved; the filtering gain	Sec. 1.3.2
loop, chord	the state–observation–belief–action feedback cycle, and the direct edge by which an action enters what others observe without passing through the state	Introduction, Fig. 1
spike	a marginal control impulse $z\,\delta_t$; the deviation whose first variation gives the first-order conditions (Chapter 4's raw order spike is the same object)	Sec. 1.2.3; Sec. 4.2
belief price	$H^{k,i}_t(u)$, the marginal cost to player i of moving player k's estimate of the shock born at u	Thm. 1.8
information wedge	$\mathcal{V}^i_t$, the adjoint price of moving opponents' noise-states; the only term through which belief manipulation enters the state-price equation	Thm. 1.8
birth	the entry of a new source coordinate into a filter or adjoint at the current date (diagonal birth); a delayed public signal adds a second birth when its news arrives, at $u = t - \Delta_i$	Thm. 1.5; Chapter 2
disclosure window	the interval between a private action and the public arrival of the news that explains it	Chapter 2, Fig. 2.1
early peak	the peak in the stationary state price of a signal-noise shock a short lag after it arrives; a noise shock harms only through the responses it triggers, and no player has responded yet when it arrives	Sec. 3.2
seed	the displacement a spike leaves in observers' filters, an old-history density plus a birth atom; in Chapter 6, the labeled state impulse a monitored deviation creates	Sec. 4.3; Def. 6.1
execution, position	the two parts of trading profit: execution is earned at the fill, the gap between the new quote and the price paid; position is earned afterwards on shares already held; their sum, the price-profit rate, is zero-sum across the market	Sec. 4.8
book gap	$\mathbb{E}[Q^i_t(V_t - P_t)]$, the average undervaluation of the shares held; a level, not a flow	Sec. 4.8
monitoring relation	$j \trianglerighteq i$: player j observes player i's deviations and knows they are player i's; reflexive and transitive	Def. 6.1
privy, naive	a player who sees a deviation, knows who made it, and responds to it; one who does not and treats its effect as more primitive shock	Sec. 6.4
blip of irrationality	a single forced departure after which the deviating player resumes equilibrium play; pins down what it does after the blip, which is what privy players respond to	Sec. 6.4.2
transparent, opaque market	the market with and without published order flow, in which the trader is privy or naive to quote deviations	Sec. 6.8

Introduction

This dissertation develops a calculus for the dynamics of strategic information flow.

No one knows much. In life, we only see a corner of the world, and we see beyond it by watching what others say or do. With each of us knowing so little, it is surprising that an economy can function at all.

No one person knows everything needed to make a pencil, yet pencils manage to exist, costing almost nothing. Leonard Read's essay "I, Pencil" [98] traced the paths of cedar, graphite, lacquer, and rubber across the world to loggers, miners, and chemists who know nothing of each other's work. Why would any of these strangers bother? How would any of them know what to make, and how much?

Adam Smith answered the first question with decentralized exchange: people pursuing their own interests could coordinate without anyone directing the whole [107], with competition keeping anyone from steering it. In competitive markets where many people can offer the same thing, no one has much power over the price, and excess profits are competed away. Prices emerged from the exchange itself, set by no one.

Friedrich Hayek, writing "The Use of Knowledge in Society" at the end of the socialist calculation debate, answered the second [53]. The debate asked whether a planning board could replicate market outcomes by solving the economy's equations, and Hayek switched perspectives: the relevant knowledge, of the particular circumstances of time and place, exists only in dispersed fragments in countless minds, and no board could collect it. Prices, for Hayek, were a telecommunications system. The logger never sees beyond their own corner; the prices that reach them carry what the rest of the world needs them to know.

Prices may be useful for communication, but they must come from somewhere. Consider buying a car from a used-car salesman. The seller knows more about the car than you do [2]. But why should their knowing more cost you anything? You pay what you believe the car is worth, but the seller shapes what you learn. The same problem appears whenever actions both do things and signal, and formalizing that is hard.

The first obstacle was defining equilibrium under private information at all, which Harsanyi solved through priors over player types [52]. Muth proposed players hold rational expectations consistent with the beliefs the model implies [87]. In macroeconomics, the reasons to do so became visible in the Phillips curve. Inflation and unemployment had been stably related [96] for long enough that the relationship was treated as a policy lever [102], and that relationship changed once it was used. Lucas explained the change using Phelps's economy of islands [77, 95]. Isolated on each island, producers could not tell a general price rise due to inflation apart from a real movement in their own market and had to respond the same way to both, producing the correlation. Once the government began using inflation to manage unemployment, however, a price rise said less about conditions on the producer's island and more about government policy. Producers therefore no longer responded to price increases in the same way, and at the macro scale that showed up as a shift in the Phillips curve.

The idea rebuilt the field. Following Lucas's critique [79], relationships in economic data were no longer treated as potential levers to pull or equations to solve, but as features a model had to derive. Unfortunately, at the time, the lesson was applied mainly to the government, the one actor considered large enough for the effect to matter. In Lucas's models, private agents remained price takers whose actions move nothing observed by others.

A few years later, Grossman and Stiglitz carried Lucas's signal-extraction logic into asset markets and asked if prices could reveal everything traders know [47]. They cannot. If the price revealed everything, then gathering information would earn nothing, so nobody would pay to gather it. The invisible hand of the market could never reveal everything traders know.

Townsend later studied firms watching prices that reflect other firms' actions [109]. Actions came from beliefs, so each firm had to predict what other firms believed. Those firms did the same, so to determine their actions, those beliefs in turn had to include beliefs about others' beliefs, and so on. To get a solvable model, Townsend assumed that all private information is revealed after a fixed number of periods, capping the chain at a fixed length.

Keynes mentioned the same recursion decades earlier in *The General Theory* through his "beauty contest" analogy [66]. Readers of a newspaper entered a contest to choose the six prettiest faces out of a hundred, but prettiest was decided by the entries of the readers. Winning meant anticipating not which faces were prettiest, but which faces everyone else expected everyone else to choose, with nothing grounding those choices in the quality of the faces themselves. Beauty contests today describe stock markets ignoring fundamentals.

The closest the literature came to a full formal model of strategic information flow was the Kyle line of models, which began in finance in the mid-1980s. Kyle modeled an insider who learned something about a stock's future value and a market maker who sets the price after seeing only order flow muddied by uninformed participants [71]. Every trade the insider makes moves the price and shows the market maker part of what they know. Like the islanders, the market maker cannot tell the two sources apart, so the insider stays under the radar by trading slowly, and the information only enters the price over time. Kyle kept the core of the problem, but in a special case where it could all be solved, with one insider, informed once at the start of trading, and a competitive market maker whose belief is public through the price. When the finance community saw the whole problem in a solvable model, it pounced, and a generation of extensions followed [11, 13, 38, 56]. Extensions with several informed traders stayed solvable only under heavy symmetry, with every trader drawing a signal from the same distribution [23]. The tractable territory was treated as exhausted.

Over time, interest spilled into nearby approaches: mean-field limits, global games

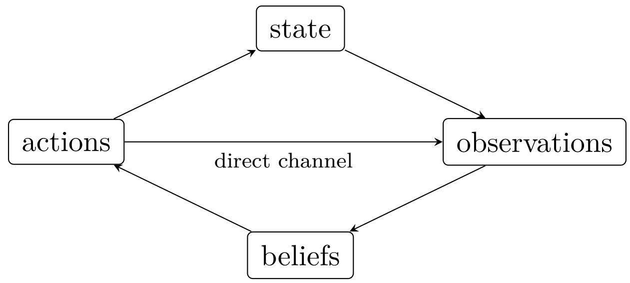

Figure 1: The feedback loop of dispersed information.

and beauty contests, information design. After a generation, students learn dispersed information through these ideas. Global games, for example, introduce private noise to select among equilibria; the same device has come to serve as a model of what dispersed private information does [5].

In the problems studied here, actions move the state of the system, the state generates observations, observations update beliefs. Beliefs complete the cycle by determining the actions that started it. A loop (Figure 1). There is also a chord, where an action can set what others observe directly. Every tractable approach beyond the Kyle line severs it somewhere, so a modeler's options can be organized by the cut each makes. The cuts fall on the privacy of beliefs, on what actions can teach, on the weight of one player, on the flow from observations to beliefs, or on time itself.

One option is to make private information public before the loop can keep running in private. Townsend reveals it after a finite number of periods, capping the hierarchy. Common-information methods do this with more finesse, pooling part of each player's history into a shared record and often restricting beliefs to that shared information, so the loop never completes a lap in private [89, 92]. Hambly et al. [49], for instance, make actions completely observable and restrict to linear strategies, so that each action reveals the beliefs behind it and past beliefs become common knowledge after a single step.

One option is to fix the information structure independently of play, as in Bayes correlated equilibrium [18], where actions respond to signals, but actions cannot change what other players learn. Another option is to take a mean-field limit [58, 73], so an individual player is too small to move the others' states. Another option is to keep the network but remove the private-learning problem, cutting the flow from observations to beliefs [46].

Signaling and cheap talk are mostly studied in static settings, so the loop runs once [33, 108]. In cheap talk the message changes what others observe, not the state, keeping the chord and nothing else. In reputation models actions are public, so beliefs about the long-run player are common [69]. In sequential social learning each agent acts once, so nothing others learn ever flows back [15, 21]. Beauty-contest models keep Keynes's recursion, but usually for one round and with the information structure fixed, so the loop runs once [5, 86]. Even information design, the subfield whose subject is what agents are allowed to observe, is studied mostly in static settings, so the loop runs once.

The two nearest neighbors sever the loop at its subtlest points. The frequency-domain solutions of the forecasting-the-forecasts literature [60, 64, 101] carry the belief hierarchy exactly, but only in stationary equilibria among negligible agents, so no individual action moves any signal and the edges out of actions are cut at the individual level. The insider-trading literature built on enlargement of filtrations [24, 29] keeps the chord and extends the Kyle line, but still for a single insider facing a market maker whose belief is public through the price, so the beliefs node carries no private tower. Expanding a filtration with an entire process, a dynamic endowment, is already hard on the purely probabilistic side [65].

Some tools must be taken out of their domain. Mean-field game theory, for instance, was built for a large anonymous population. However, the same limit has become a dominant way of handling dispersed information [4, 5, 76]. There a method designed for agents whose individual actions move nothing others observe gets applied to questions about what agents learn from one another.

Each approach removes a defining feature of dispersed information, restricting not just the questions that can be answered but the phenomena the model can produce. Approaches searching for tractability are often tempted by an elegant property of linear-quadratic-Gaussian (LQG) control known as the separation principle. A single controller can treat learning about the world and choosing its actions separately. When others are listening in on their behavior, actions become signals, and Townsend's hierarchy applies. Separation fails unless something else is imposed.

Ask any model that claims to study the interaction between information flow and strategic incentives whether its results deliver a separation principle: whether, in the end, information and incentives can be treated separately. If so, the interaction it set out to study is absent, akin to a method that lets people study correlations between variables only when those correlations are zero. Cutting the loop often makes separation close to an assumption. If separation arises because information has been prevented from responding to incentives, then the model has achieved tractability by removing the mechanism of interest.

A journal cannot keep publishing papers that stop at saying the problem is hard. Individual papers are usually explicit about the assumptions they make, but what compounds are the results existing tools could reach. An outside modeler sees decades of work on private information, learning, signaling, prices, and strategic interaction, and naturally assumes that the underlying decentralized problem has already been studied, that what remains consists of technical refinements for specialists. They mistake the field's principal results for the limits of what decentralized information can explain, but what looks like an economically exhausted subject might instead be a subject whose most expressive models have remained out of reach.

This dissertation is my attempt to allow the field to build the models it set out to build. I hope that the stories told through Hayek's "The Use of Knowledge in Society" and Leonard Read's "I, Pencil" can be put to paper more fully, and that I can someday enjoy reading the papers of a field that has kept its heart.

At a linear equilibrium, states, signals, and actions are linear functions of Gaussian shocks. Estimating an unknown then becomes a form of linear regression. The coefficients and the residual uncertainty depend on which sources a player watches, never on what those sources happened to say, so precision is common knowledge while estimates are not. Linearity also lets each player subtract the effect of their own actions from what they observe, so a player cannot manipulate themselves, only others. These two properties are both the sources of noise-state tractability and the limitations to what it can model. Economists already linearize to get tractable approximations to realistic models, but this had not been enough to handle decentralized information.

Art Moore, my linear algebra professor in community college who inspired me to do this PhD, once explained the difference between ordinary and partial differential equations by saying that PDEs are so much harder to analyze because they can describe so much more. A language capable of describing complex phenomena necessarily carries some of that complexity in its own structure. My ambition for the dissertation's central device, the noise-state, is to give economists a tractable way to write down that complexity. The resulting mathematics may be complex, but it is the simplest machinery I know that keeps the loop intact.

The noise-state

The noise-state is a recursive state for dynamic games with dispersed information. A player's noise-state represents what that player believes about the underlying sources of uncertainty. To reason about other players, one also tracks how their beliefs about those same sources evolve; by necessity, those beliefs are themselves functions of the underlying sources of uncertainty. In the linear-quadratic-Gaussian case, the relevant updates become deterministic kernels, so equilibrium can be characterized as a fixed point and solved numerically.

Let $W = (W^0, W^1, \ldots, W^n)$ be the primitive Brownian motion driving the state and the observation noises. Each player observes only their own signal, specified in each chapter; write $\mathcal{F}_t^i$ for everything player i knows when choosing at time t. Define

$$\widehat{W}_t^i(u) = \mathbb{E}[W_u \mid \mathcal{F}_t^i], \qquad 0 \leq u \leq t.$$

This is the player's estimate, at time t, of the primitive shock path that has already occurred.

Any process can be made Markov by declaring its whole history to be the state. What is needed is a Markov form that is stable under conditioning and best response. Say a process is linear if at each time it is a deterministic mean plus a deterministic impulse response kernel integrated against the primitive shocks, $L_t = \bar{L}(t) + \int_0^t L(t,s)\, dW_s$. The shocks' path is a sufficient statistic for every linear process at once, however path-dependent. The primitive shocks, not the endogenous state, are the natural coordinates. Under perfect information every player reads these coordinates directly, and the kernels are the model's ordinary impulse responses. Since each player observes those shocks through their own signals, the relevant statistic is the player's belief about them: the noise-state. It changes what a belief is a belief about. Unlike higher-order beliefs about endogenous states that are hard to close and unnatural to approximate, beliefs about primitive shocks are neither. By linearity everything can be written in terms of primitive shocks, including beliefs about them.

Common-information methods, introduced above, remove the tower the other way, by anchoring beliefs in information all players share. The loop can still survive, but the method buys tractability when what remains private is limited, usually alongside cuts made elsewhere. When most of the problem can be described through a shared state, the crux of the problem is not really in decentralization. Here signals are mostly private, so the shared state carries little. Each player instead builds coordinates inside their own filtration over the same primitive shocks.

Filtering reduces to a single deterministic object: the map from the primitive white noise dW to the player's estimated white noise $d\widehat{W}^i$. Explicitly,

$$d_u \widehat{W}_t^i(u) = \Pi^i\, dW_u + \left(\int_0^t F_t^i(u,s)\, dW_s \right) du,$$

a singular part $\Pi^i\, dW_u$ plus a continuous density part with deterministic kernel F^i, forming a blueprint for all the player's beliefs. Here d_u is an increment in the path index

u, rather than in the time of beliefs t. Conditional on the player's information, the shocks are Gaussian, with the noise-state providing the conditional mean and the blueprint determining the covariance.

The conditional expectation of a linear process is easy to represent. Write $\widehat{L}^i_t$ for $\mathbb{E}[L_t \mid \mathcal{F}^i_t]$. Conditioning changes the integrator and nothing else, $\widehat{L}^i_t = \bar{L}(t) + \int_0^t L(t,u)\, d_u \widehat{W}_t{}^i(u)$, the same mean and the same kernel now run against the noise-state. In primitive coordinates the kernel becomes

$$\widehat{L}^i(t,s) = L(t,s)\Pi^i + \int_0^t L(t,u) F^i_t(u,s)\, du,$$

mixing the primitive shocks, and discretizing turns that composition into matrix multiplication. Taking an expectation of an expectation is no harder than an expectation of any other linear process, so the tower closes.

The unrestricted best response to a noise-state linear strategy is noise-state linear. A noise-state linear strategy is a linear process in the player's noise-state rather than in the primitive shocks,

$$D^i_t = \bar{D}^i_t + \int_0^t D^i_t(u)\, d\widehat{W}_t{}^i(u).$$

The equilibrium problem becomes a deterministic fixed point in kernels.

The physical state and the players' noise-states can be read as one enlarged state of the system, $(X, \widehat{W}^1, \ldots, \widehat{W}^n)$. Standard optimal control already prices movements of the physical state through a running shadow price, while the noise-state formulation adds a shadow price for moving another player's noise-state, which I call the information wedge. It prices information externalities in both noncooperative and cooperative settings, such as in a team where a teammate acts to show the rest what they know. Every manipulation channel passes through it. It breaks separation, and vanishes whenever the loop is cut.

Stochastic games are typically posed as forward-backward systems whose backward equation includes a martingale-representation term. The conditional expectations and the shadow prices separate here because the forward equations, including those for the noise-state, are stochastic while the backward equations for the adjoints are deterministic ordinary differential equations.

Rational-expectations models describe today's behavior and ask what happens to the propagation of shocks as policy rules change. Policy changes still require recomputing the equilibrium, in the Lucas sense, but the kernels being recomputed are already indexed by the primitive shocks. A researcher can take an existing model with perfect information, specify who sees what, and solve again in the same coordinates to see how impulse responses and means change.

The accompanying Python package, `pypi.org/project/noisestate`, implements this workflow for supported finite-horizon and stationary LQG games.

Chapter roadmap

Baseline. Chapter 1, which closely follows the preprint [10], builds the calculus in the finite-horizon LQG game. It develops the noise-state linear class, proves that filtering and best responses stay inside the class, and derives the adjoint system in which the information wedge prices what an action teaches the other players. A finite-deviation identity then verifies a computed equilibrium against every admissible deviation, including those outside the linear class. The deliberately small two-player example separates the statistical value of information from its strategic value and shows how a planner can deliberately starve an inefficient player by how it allocates signal precision. An identical-interest variant closes the chapter. Interactive versions of the dissertation's computations are collected at sbabichenko.com/noisestate.

Delay. Chapter 2 lets news arrive late. Players may hear the same news after different delays, and what a player does can show up in a signal directly as well as through the state. Until the news arrives, a player can move what others believe; once it arrives, it explains the movement away, as when an insider's trade is disclosed weeks after it moved the market. Delays usually make filtering messy, but here the delay adds terms to the same equations rather than requiring a different method.

Stationarity. Chapter 3 rewrites the baseline equations for a game that has run long enough to become stationary, so objects of the calculus become functions of ages rather than dates. The stationary formulation is also the easiest to compute. It lets us measure how closely a finite-dimensional system can approximate the equilibrium. It is a natural initial condition for studying a change of regime, and the setting of Chapter 4.

Trading. Chapter 4 applies the stationary equations to Kyle–Back trading on an infinite horizon. The chord ran through the earlier chapters without being leaned on; here it is the object of study, because order flow is both an action and a signal, so a marginal order moves the current price and changes future filtering by the market maker and the other informed traders. Each trader knows its own order and interprets the remaining flow differently, so the rivals' responses make dynamic price impact trader-specific. In the computations, they can reverse a price move the market maker alone would leave permanent. This chapter keeps Kyle's market, with several traders of heterogeneous signal quality and continuous signals. A specialist may trade first in a correlated market with more noise-trader flow.

Aggregation. Chapter 5 turns to economies made of many small local markets, each with only a few players: a regional credit market with a handful of lenders. Each player is negligible nationally but consequential locally, so the local game keeps the wedge while the economy stays large. Graph symmetry then collapses the player labels: policy, filtering, and belief-price kernels depend only on relative position, and equilibrium is solved once per position rather than once per player. The idiosyncratic shocks average out across markets but the incentives do not. In the cycle economy every firm holds back its order because its supplier reads demand off it, and the aggregate mean order sits below its perfect-information level. The results are for finite networks, with the infinite-network limit stated as a conjecture.

Monitoring. Chapter 6 makes the chord player-dependent. A deviation may be monitored by some players and not others. A privy player sees the deviation labeled with who made it and does not filter it, while a naive player keeps filtering it as primitive shock. The feedback splits into a naive channel, which is the wedge, and a privy channel, in which the privy player responds to the deviation directly, and the adjoint becomes a rectangular gain between responder–origin pairs. The monitoring relation runs from the setting of Chapter 1, where no deviation is ever observed, to the classical closed-loop games in which every deviation is; both endpoints are recovered exactly (Proposition 6.13). Its worked example computes the same market under both relations: once the order flow is published the market maker puts more weight on inventory in its quote, because a trader who sees the flow can be steered, and most of transparency's gain is real inventory efficiency.

Chapter 1

Baseline Linear-Quadratic-Gaussian Games

One player's action changes what the others learn, their response feeds back into the state, and the first player must price it. The noise-state represents each player's beliefs about the shocks that generate histories, so higher-order beliefs are generated by composition rather than tracked as separate state variables. In continuous-time LQG, beliefs, prices, and policy kernels are deterministic impulse-response kernels, and equilibrium is a deterministic fixed point in those functions. The adjoint system contains a shadow price of changing opponents' posteriors, the information wedge; in a two-player benchmark it separates the statistical and strategic value of information and turns signal allocation into an instrument for controlling inefficient effort.

1.1 Introduction

Dynamic games with dispersed private information pervade economics. Firms infer competitors' pricing from market outcomes, traders learn from the order flow their own trades generate, central banks know their announcements reshape private agents' signals. In each setting, actions feed back into the information environment that agents optimize against, generating an infinite hierarchy of beliefs. Townsend [109] identified that hierarchy, and Sargent [104] computed equilibria in linear-Gaussian environments by positing low-dimensional ARMA laws rather than deriving them.

Beliefs here are beliefs about the underlying sources of randomness rather than hierarchies of beliefs about the endogenous state. A researcher who specifies payoffs, a state equation, and a signal structure gets back equilibrium impulse-response functions; the representation handles the belief hierarchy internally.

Call this representation the noise-state.

In a game of poker, the common approach describes each player's beliefs about opponents' hands, their beliefs about opponents' beliefs, and so on. The noise-state approach instead tracks each player's belief about the underlying source of randomness (here, the order the cards were shuffled) and how that belief depends on the truth; higher-order beliefs follow by composing these dependencies. The two representations are equivalent, but approximating a composition of deterministic dependencies is natural, whereas truncating a belief hierarchy is not.

A deviation enters both the physical state and opponents' noise-state coordinates. The Markov state consists of the physical state, the players' noise-states, and, when needed, the blueprints that translate histories into beliefs.

The construction extends beyond the LQG setting studied here, but the LQG benchmark is the clean case in which it becomes explicitly solvable, since the kernel form relies on conditional Gaussian filtering. It is also the base case for a broader noise-state program in conditional-Gaussian linear economies, including strategic trading, CARA risk aversion, and replicated finite symmetric networks in which local strategic feedback can affect aggregate response coefficients.

Characterizing equilibrium produces a shadow price of manipulation, the information wedge $\mathcal{V}_t^i$. Every effect of belief manipulation enters the dynamics of the state price

(Theorem 1.8) through that single term. In the two-player benchmark, the wedge explains why most gains from pooling information are strategic rather than purely statistical, and why optimal precision allocation can generate information starvation by concentrating signals on the player who moves the state more cheaply, leaving the other player with none. Changing signal precision changes equilibrium policy rules, not only posterior variances, so separation fails. As a diagnostic, Corollary 1.9 shows the wedge vanishes when observation laws are fixed, recovering the standard fixed-information LQG optimality system. Section 1.2 works these economic consequences out in a two-player game, and Section 1.4 builds the general theory. Proposition 1.12 gives a unique unilateral best response for every finite horizon. On short horizons, Theorem A.7 gives a unique equilibrium on a strategy ball, with geometric Picard convergence. For an arbitrary horizon, existence and uniqueness on a strategy ball hold under a weighted monotonicity condition (Theorem A.9) that Corollary A.11 can certify numerically.

1.1.1 Related literature

The noise-state formulation contributes to the rational-expectations literature on higher-order beliefs and dispersed information. Townsend's "forecasting the forecasts of others" problem and Sargent's signal-extraction examples raise the possibility that private information destroys finite-dimensional recursion even in linear environments [104, 109], though in Townsend's own model the regress collapses [94]. Existing approaches recover tractability by making signals exogenous to the best-response problem, making individual actions negligible, using public histories, or truncating the hierarchy. Here the private filters of finitely many players are endogenous, and beliefs about primitive shocks represent the hierarchy.

The analysis also relates to the social value of information and imperfect common knowledge. In Morris–Shin, Angeletos–Pavan, Woodford, and related work, signal structures affect behavior but are fixed when agents optimize [5, 6, 19, 76, 86, 118]. Here, endogenous belief prices add a strategic component to the social value of information absent under fixed signal laws.

Frequency-domain and finite-state methods exploit signal laws fixed in the best-response problem [51, 60, 64, 101]. Here, a deviation changes the state process opponents observe, so filters are part of the equilibrium response. In the LQG benchmark, the noise-state turns this joint filtering/control problem into a deterministic fixed point in impulse-response maps. Finite-order truncations approximate the hierarchy directly; noise-states generate higher-order beliefs by composing primitive-shock projections [90].

The noise-state formulation connects to decentralized stochastic control, LQG team theory, and common-information methods. Classical team results give linearity under fixed information structures, while common-information methods condition on a shared information base [82, 89, 97]. The noise-state approach instead keeps private filters player-specific and endogenous. Their strategic externality appears as the information wedge.

Mean-field LQG games buy tractability by making individual actions informationally negligible [26, 27, 58, 73]. This is appropriate for aggregate modes and removes local informational channels; the wedge is the finite-player object that disappears.

The precision-allocation exercise relates to information design and rational inattention [18, 63, 106, 114], but it does not solve for optimal attention or disclosure. It gives the equilibrium map from a fixed signal technology to behavior and welfare when precision changes policy rules through the wedge.

A recent control-theoretic literature studies belief manipulation directly. Vasal and Anastasopoulos [112] construct signaling equilibria for dynamic LQG games with asymmetric information through a sequential-decomposition backward recursion. The present chapter differs in carrying the manipulation price as an explicit adjoint kernel rather than encoding it in a fixed-point over prescription maps. An adversarial line [67, 119] models active deception against a fixed sequential hypothesis test, so the beliefs in play are prescribed mechanisms rather than equilibrium objects. In this chapter, each player's filter is the Bayes-consistent estimate given equilibrium play, manipulation operates only through the signal channel, and its price is the wedge.

1.2 Bilateral belief manipulation in a two-player game

A two-player tracking game isolates the local feedback channel (Figure 1.1). Every figure in this section, and the solver behind it, runs in the browser at sbabichenko.com/noisestate with adjustable parameters.

1.2.1 Primitives and objectives

Fix a horizon $[0, T]$ and a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supporting a standard three-dimensional Brownian motion $W = (W^0, W^1, W^2)$, where $\mathbb{F} = (\mathcal{F}_t)_{t\in[0,T]}$ is the completed, right-continuous natural filtration of W. The state $X = (X_t)_{t\in[0,T]}$ is the real-valued, $\{\mathcal{F}_t\}$-adapted process satisfying

$$dX_t = \left(D_t^1 + D_t^2\right) dt + \sigma\, dW_t^0, \qquad X_0 = x_0, \tag{1.2.1}$$

where $\sigma > 0$ is the fundamental volatility. Player $i \in \{1, 2\}$ does not observe X directly; instead they observe the private diffusion channel

$$dY_t^1 = \sqrt{p_1}\, X_t\, dt + dW_t^1, \qquad dY_t^2 = \sqrt{p_2}\, X_t\, dt + dW_t^2. \tag{1.2.2}$$

The numbers p_i are signal precisions. Player i chooses D^i using only their private observation history. Formally, $\mathbb{F}^i = (\mathcal{F}_t^i)_{t\in[0,T]}$ is the completed, right-continuous augmentation of $\sigma(Y_s^i : 0 \le s \le t)$, and an admissible control D^i is $\mathbb{F}^i$-progressively measurable and square-integrable, $\mathbb{E} \int_0^T (D_t^i)^2\, dt < \infty$. The two controls are adapted to different filtrations, D^1 to $\{\mathcal{F}_t^1\}$ and D^2 to $\{\mathcal{F}_t^2\}$, and in general neither filtration contains the other.

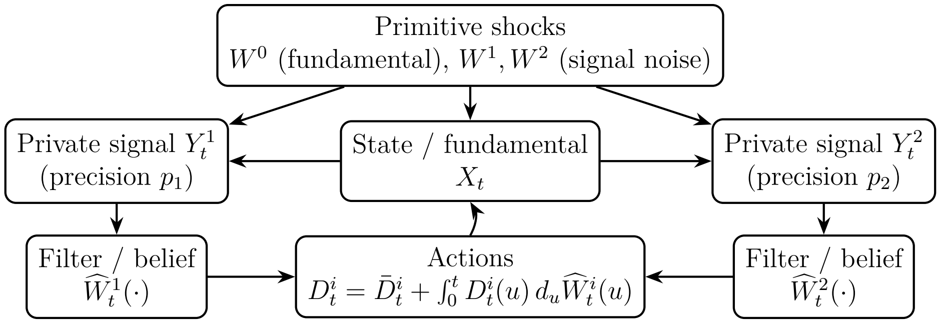

Figure 1.1: Information-to-action feedback in the worked example.

Running costs are symmetric tracking losses with quadratic effort penalty $r > 0$:

$$J^1(D^1; D^2) = \mathbb{E}\int_0^T \Big((X_t-1)^2+r\,(D_t^1)^2\Big)\, dt, \quad J^2(D^2; D^1) = \mathbb{E}\int_0^T \Big((X_t+1)^2+r\,(D_t^2)^2\Big)\, dt. \tag{1.2.3}$$

Player 1 wants the state near $+1$; player 2 wants it near -1.

Each player minimizes its own cost, taking the opponent's strategy as given. A pair $(D^{1\star}, D^{2\star})$ of admissible strategies is a *Nash equilibrium* if $J^1(D^{1\star}; D^{2\star}) \leq J^1(D^1; D^{2\star})$ for every admissible D^1, and symmetrically for player 2. Section 1.3.3 restates the definition for the n-player game.

1.2.2 Impulse responses and wedges

Equilibrium controls belong to the noise-state linear class (Definition 1.4) and can be written in two equivalent impulse-response coordinates, with the noise-state defined on the middle line:

$$\begin{aligned} D_t^i &= \bar{D}_t^i + \int_0^t D_t^i(u)\, d_u \widehat{W}_t^i(u), \\ \widehat{W}_t^i(u) &:= \mathbb{E}[W_u \mid \mathcal{F}_t^i], \\ D_t^i &= \bar{D}_t^i + \int_0^t D_{W,t}^i(u)\, dW_u. \end{aligned} \tag{1.2.4}$$

The first line is the strategy as the player implements it: a deterministic mean plus a reading of the noise-state. The third line is the same control composed with the player's filter and written directly over the primitive shocks. $D_t^i(u)$, the policy kernel, weights what the player has estimated and is measurable to the player, while $D_{W,t}^i(u)$ describes how the choice loads on the truth and involves every player's filter.

Proposition 1.1 (Best responses and equilibrium in the two-player game)**.** *For every finite horizon and every fixed linear strategy of the opponent, each player has a unique best response over the full admissible class, and that response is noise-state linear. There is $T^* > 0$ such that for $T \le T^*$ the game admits a noise-state linear equilibrium, unique in the strategy ball $\mathbb{B}_T^{M_0}$ of Theorem A.7, Picard iteration converges to it geometrically, and the equilibrium is Nash over the full admissible class. For an arbitrary horizon, any strategy ball satisfying the weighted monotonicity and residual-radius conditions of Theorem A.9 contains a unique fixed point, and that fixed point is a Nash equilibrium over the full admissible class.*

Proof. The coefficients are constant and bounded and the effort weights satisfy $r > 0$, so Assumption A.1 (the boundedness and ellipticity conditions of Appendix A.2) holds. Apply Proposition 1.12, Theorem A.7, and Theorem A.9 with $n = 2$ and $d = 1$. □

1.2.3 The life of a spike deviation

A spike by one player propagates through the loop of Figure 1.1. Player 1 adds an impulse $z\,\delta_t$ to its control. The state shifts by z and stays shifted. Player 2's innovation carries the part of the shift player 2 has not yet learned, $\sqrt{p_2}\,(\delta X_\tau - \delta \widehat{X}^2_\tau)\,d\tau$, and each coordinate of its noise-state moves through the unresolved response, $\widetilde{X}^2$, the share of the shock's effect player 2 has not yet resolved,

$$\partial_\tau\, \delta w^2_\tau(u) = p_2\, \widetilde{X}^2_\tau(u)\left(\delta X_\tau - \delta \widehat{X}^2_\tau\right), \qquad \delta \widehat{X}^2_\tau = \int_0^\tau X_\tau(s)\, \delta w^2_\tau(s)\, ds,$$

with the coordinate born at τ equal to $\sqrt{p_2}\,(\delta X_\tau - \delta \widehat{X}^2_\tau)$. Player 2's control moves with its coordinates and the state with the control. Player 2's target is -1, so the response opposes the shift.

Player 1 pays for the shift and for the response. The shift is priced at the running cost. The response is priced through player 2's noise-state, one coordinate at a time: $H_t^{2,1}(u)$ of (1.4.8), the belief price, is the marginal cost to player 1 of having moved player 2's estimate of the shock at u. Setting the variation to zero for every $\mathcal{F}_t^1$-measurable z,

$$2r\,D_t^1 = -\mathbb{E}\Big[\int_t^T \Big(2(X_\tau - 1) + \mathcal{V}_\tau^1\Big)\,d\tau\,\Big|\,\mathcal{F}_t^1\Big], \quad \mathcal{V}_\tau^1 = \sqrt{p_2}\,E^2 H_\tau^{2,1}(\tau) + p_2\int_0^\tau \widetilde{X}_\tau^2(r)\,H_\tau^{2,1}(r)\,dr, \tag{1.2.5}$$

which is (1.4.7) integrated, with the wedge (1.4.9) at the example's coefficients. Without player 2 only the first term remains. The second is the change in player 2's beliefs times their price. Theorem 1.8 is the same calculation for n players, a vector state, and a general observation map.

Player 2 sees the state only through its signal, so it cannot tell a unit spike by player 1 from a fundamental shock dW^0 of size $1/\sigma$. At $\sigma = 1$, player 2's curves in Figure 1.2 therefore show how player 2 responds to a unit spike by player 1. The left panel traces a single state shock through the actual state and through each player's posterior estimate of it: the Markov state of the game, the physical state together with every player's noise-state (Section 1.3). The shock moves the state. Each player's estimate catches up at a speed set by its signal precision. The right panel traces the same shock through the players' actions. Each player responds to the shock through what it has learned, the better-informed player earlier and more strongly. Both control responses die out at the horizon, when nothing remains to be gained from moving the state. These plots can be redrawn at any parameters in the browser version.

1.2.4 Computation

The unknowns of the fixed point are the kernels $D_t^i(u)$, $X_t(u)$, and $F_t^i(u,s)$; the adjoint equations (1.4.7)–(1.4.8) are solved for these kernels rather than for a value function on X_t, because each player's response depends on every opponent's filter. Equilibrium is

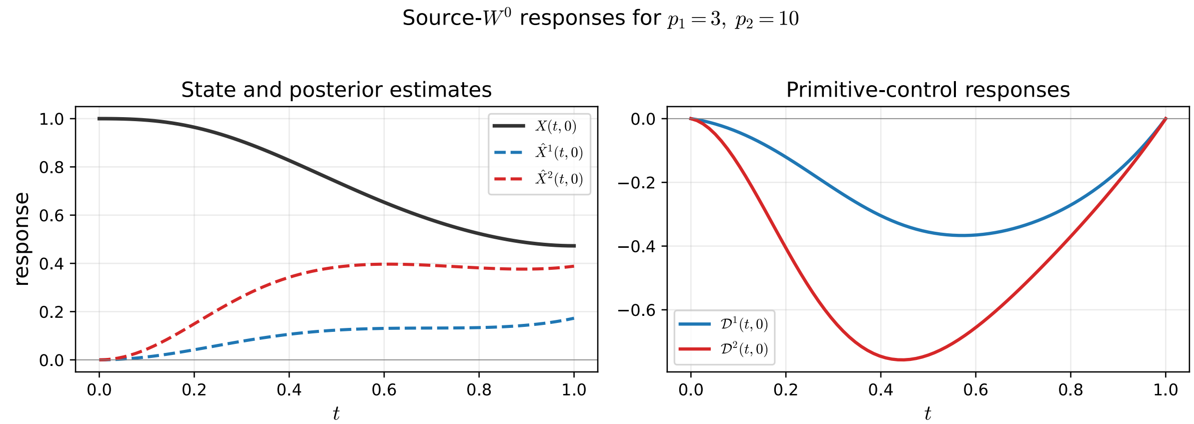

Figure 1.2: Impulse responses to the first state shock dW_0^0 with $p_1 = 3$, $p_2 = 10$, and $r_1 = r_2 = 0.1$. Left: state and posterior-estimate responses. Right: primitive-control responses $D^1_{W,t}(0)$ and $D^2_{W,t}(0)$.

computed as a fixed point of the simultaneous best-response map on the deterministic causal policy kernels, discretized on a uniform grid of $N = 40$ points; no Monte Carlo simulation is used. Starting from zero policy kernels, each outer iteration first solves the forward state–filter system generated by the current profile, then solves both players' backward physical and belief-price equations, and finally applies the first-order conditions to obtain a new pair of policy kernels. Both new kernels come from the same old profile, so the update is simultaneous rather than alternating.

Write $G(D)$ for this discretized best-response map and $D = (D^1, D^2)$. Convergence is measured by the normalized Frobenius residual over every time pair on the discrete causal triangle and every primitive-shock coordinate,

$$\operatorname{resid}(D) = \frac{\|G(D) - D\|_F}{\max\{1, \|G(D)\|_F\}}.$$

The first update is a relaxed Picard step with relaxation parameter 0.15; subsequent updates use depth-five Anderson acceleration with mixing parameter 0.6. The algorithm stops when the residual falls below 10^{-5}. Within each outer iteration, the coupled forward state–filter calculation uses two relaxed updates of the state impulse response and six filter sub-iterations. After the policy kernels converge, a separate relaxed fixed-point iteration, initialized at zero and solved to tolerance 10^{-10}, gives the deterministic mean controls and mean state. Parameter sweeps may be warm-started from the equilibrium at a neighboring parameter value.

Unless otherwise stated, $T = 1$, $x_0 = 0$, $r = 0.1$, $\sigma = 1$, and there is no terminal cost. In an optimized native C++ build with OpenMP, the symmetric benchmark $(p_1, p_2) = (3, 3)$ at $N = 40$ converges in 42 outer iterations and takes about 0.03 seconds; at $N = 160$ the same benchmark takes about 1.4 seconds.

The scheme is first order in the grid spacing. On the symmetric benchmark with targets $(1, -1)$, player 1's cost is 4.219, 4.160, and 4.131 at $N = 40$, 80, and 160, an observed order of 1.06; the policy kernels converge at the same rate. A single $N = 40$ run is therefore accurate to about 3 percent on costs and mean actions, and considerably worse on quantities small early in the horizon, such as the posterior estimates just after a shock and the information wedges. All figures and reported values below are Richardson extrapolations from two nested grids, $2f_{157} - f_{79}$, which agree with the extrapolation from $N = 40$ and $N = 79$ to within 0.3 percent except at the first two grid points of the wedge curves.

The solvers and figures of this section are the work of Claude Fable 5, a large language model built by Anthropic [7], working under the author's direction. The interactive version at sbabichenko.com/noisestate runs the same solver.

1.2.5 Mean actions, separation failure, and the welfare cost of belief manipulation

Since the stochastic integral in the first line of (1.2.4) has zero unconditional mean, $\mathbb{E}[D_t^i] = \bar{D}_t^i$ isolates the deterministic tug-of-war over the target ± 1, while the response term captures impulse responses to posterior shocks.

Separation failure. Under perfect information, the mean policy is independent of signal precision. Decentralized information breaks this. Changing signal precision changes how much of each shock remains unresolved for the opponent, which changes the belief price, so the policy rule moves independently of the state estimate (Remark 1.13). The mechanism is the information wedge.

Figure 1.3 plots the resulting failure. The equilibrium mean control path $\bar{D}_t^1$ shifts with the signal precision p, and the partial-information paths lie between two limits independent of p. With no information the players cannot see each other and the mean paths are the open-loop Nash equilibrium of the deterministic game, $\bar{D}_t^1 = (T-t)/r$. With perfect information they are its closed-loop Nash equilibrium, in which each player's feedback on the state is priced by the other. As p rises the mean control moves monotonically from the first to the second, and at $p = 10^3$ it is still ten percent above the closed-loop path, and at $p = 10^4$ seven percent. Under a separation principle the colored paths would sit on the closed-loop line at every precision. With $p_1 = p_2$, opposite targets, and $x_0 = 0$, the two mean controls cancel and $\bar{X}_t \equiv 0$ at every precision. The movement in Figure 1.3 is effort spent holding a path that does not change. It vanishes in models where information and incentives can be treated separately. Because the shift is in the mean $\bar{D}_t^1$ rather than in the zero-mean response term, it does not cancel when many independent copies of the market are averaged; Chapter 5 works this out.

Information pooling as a corrective intervention. Figure 1.4 compares private signals with a pooled signal of precision $p_1 + p_2$ under opposing targets (± 1) and (± 5) and a common target (0). At (± 1) and (0) pooling is Pareto-improving at every p_2, and the gain is an order of magnitude larger under competition. Since the target change affects only the deterministic mean system, the common-target case isolates the pure estimation benefit. Each player still pays its own effort cost, so this case is a Nash game rather than the team problem of Section 1.4.1. The excess gain under competition is the welfare cost of bilateral belief manipulation.

The private-signal curves also show who wants the precision. Under the common

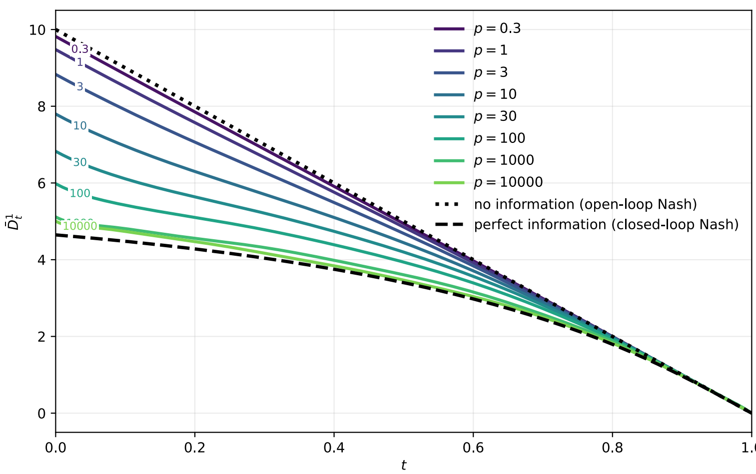

Figure 1.3: Equilibrium mean control $\bar{D}_t^1$ as signal precision p varies (spectral solution, both players at precision p). Dotted: the open-loop Nash path of the deterministic game, the no-information limit. Dashed: the closed-loop Nash path of the perfect-information game.

target the better-informed player has the higher cost. At $p_2 = 10$ player 2 pays 0.366 against player 1's 0.341. The better-informed player does more of the work, so each player would rather be the less precise one. Under opposing targets the ordering reverses, 3.70 against 3.83, and precision is an advantage over the opponent. In absolute terms, both players gain when either one is better informed, including player 1, whose own precision is fixed and whose opponent improves. The competitive case does not change that even though it raises every cost by an order of magnitude. At $p_2 = 1000$ the pooled cost and player 2's private cost have nearly met, since a player who observes the state almost exactly has little left to learn from a pooled signal. With stronger conflict they cross. The policy kernels do not depend on the targets, and the mean paths are linear in them, so at targets $(k, -k)$ the cost is $J^i = J^i_{\text{var}} + k^2\, \bar{J}^i$, with J^i_{var} the cost at the common target and $\bar{J}^i$ the mean part at $(1, -1)$. Pooling still lowers player 2's variance cost, but it raises its mean cost, because player 2 gives up the mean-path advantage of being the better informed. Above $p_2 \approx 250$ the second effect wins. Pooling then costs player 2 (0.9 at $p_2 = 1000$) while player 1 gains 18. Under strong conflict the better-informed player prefers to keep its edge. Under the common target player 1's private cost meets them too

(0.185 against 0.183), so an almost perfectly informed teammate is nearly as good as a shared signal. Under opposing targets it stays well above (3.17 against 2.44). An almost perfectly informed opponent is not.

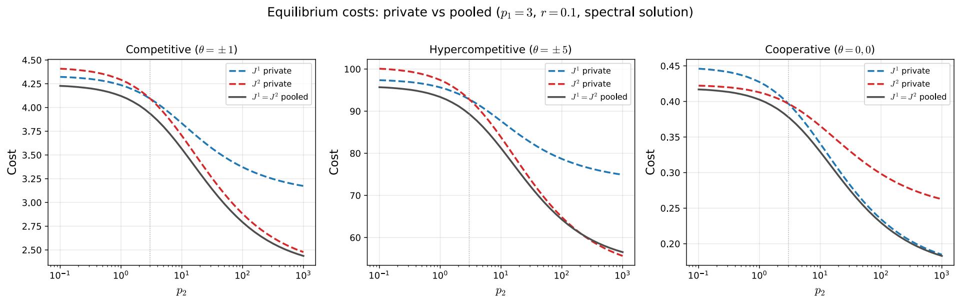

Figure 1.4: Equilibrium costs under private signals (dashed) and a pooled signal (solid), with $p_1 = 3$ fixed and p_2 varying from 0.1 to 1000 on a logarithmic axis (spectral solution). Left: opposing targets (± 1). Center: opposing targets (± 5). Right: common target (0). The vertical scales differ: costs are about 4, 80, and 0.4 in the three panels, since the mean part of the cost scales with the square of the target. The dotted line marks $p_2 = p_1$.

1.2.6 Optimal precision allocation

Information pooling is blunt. A planner with a fixed precision budget must still choose how to split it between the two players.

Setup. Retain the state dynamics (1.2.1) and observation channels (1.2.2) with opposing targets (± 1), but allow asymmetric precisions (p_1, p_2) subject to $p_1 + p_2 = \bar{P}$ and asymmetric effort costs (r_1, r_2) with $r_1 < r_2$, so player 1 is the more efficient mover of the state. A planner chooses the split to minimize aggregate equilibrium cost; for each candidate allocation Picard iteration solves the full equilibrium. Parameters: $\bar{P} = 20$, $T = 1$, with the fundamental volatility in (1.2.1) reduced to $\sigma = 0.5$. Because the players have opposing targets, their mean controls push in opposite directions rather than jointly stabilizing the state. Total destructive effort measures the intensity of this arms race,

$$\mathcal{E}^{\mathrm{dest}} = \int_0^T \left(r_1 \left(\mathbb{E}[D_t^1] \right)^2 + r_2 \left(\mathbb{E}[D_t^2] \right)^2 \right) dt.$$

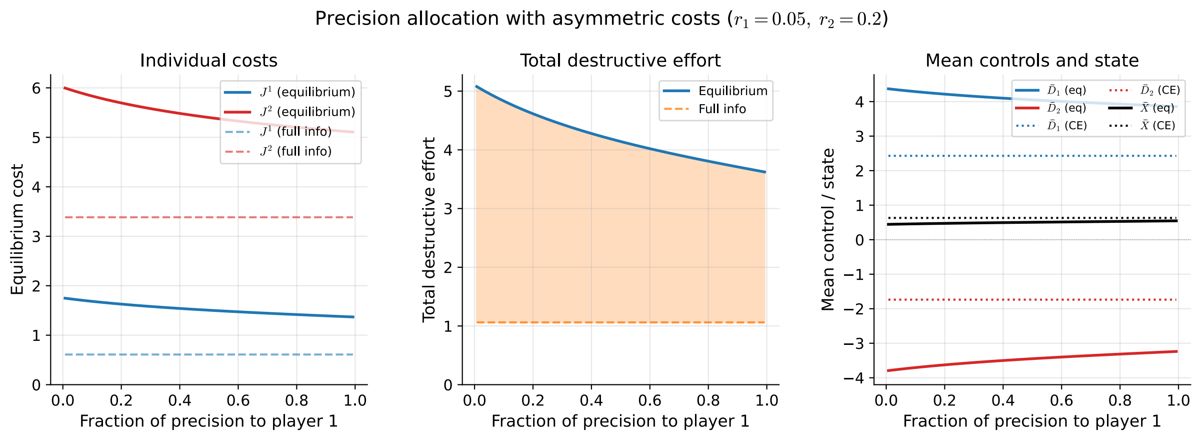

Figure 1.5: Equilibrium quantities across the precision split, $p_1 + p_2 = \bar{P} = 20$, with opposing targets, $\sigma = 0.5$, and asymmetric effort costs $(r_1, r_2) = (0.05, 0.2)$. Left: individual costs. Center: total destructive effort (the aggregate quadratic cost of the mean controls). Right: mean controls and mean state.

Information starvation. Under cooperation, a planner should assign precision to the efficient mover. Under competition, the natural guess is instead that a planner should balance precision to limit manipulation. Figure 1.5 shows the opposite. With asymmetric effort costs $(r_1, r_2) = (0.05, 0.2)$, concentrating precision on the efficient player shifts the mean state toward that player's target and sharply reduces destructive effort. On net, both players' costs fall as destructive effort moves toward the perfect-information benchmark.

Denying the inefficient player signal access weakens the bilateral arms race, a commitment logic in the sense of Schelling [105]. Giving all precision to the inefficient player produces the worst outcome, because the efficient player pushes hard against a well-informed opponent. The planner's optimum is not balanced precision but an information monopoly for the efficient player. The right panel shows separation failure again. Certainty-equivalent controls are invariant to the split, while equilibrium controls move with the belief prices.

1.3 The decentralized LQG game

Informal description of the game. The games here have a common evolving state, private noisy observation channels, and controls that feed back into the state and so into others' signals. These features generate the Townsend hierarchy of forecasting the forecasts of others [104, 109]; quadratic objectives and linear-Gaussian primitives make the problem tractable. The section generalizes Section 1.2: the same noise-state coordinates, filters, and prices, now with n players, vector states, and general linear dynamics.

1.3.1 State dynamics and objectives

Consider n players interacting over a finite horizon $[0, T]$ on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$.

State dynamics. The state $X_t \in \mathbb{R}^d$ evolves as

$$dX_t = \Big(B_X^X(t)X_t + \sum_{i=1}^{n} B_i^X(t)D_t^i\Big)dt + \sigma(t)\,dW_t^0, \qquad X_0 = x_0, \tag{1.3.1}$$

where $B_X^X(t) \in \mathbb{R}^{d\times d}$ is deterministic and bounded, W^0 is a d-dimensional Brownian motion, $\sigma(t) \in \mathbb{R}^{d\times d}$ is deterministic and bounded, and $D_t^i \in \mathbb{R}^d$ is player i's control, entering through the deterministic, bounded gain $B_i^X(t) \in \mathbb{R}^{d\times d}$ (the identity in the examples).

Objectives. Let $G_t^{X,i}$ denote the linear state-cost coefficient. In the baseline case it is deterministic and bounded.[1] Each player i incurs the cost

$$\begin{aligned}\mathcal{C}^i(D^i;D^{-i}):=&\frac{1}{2}\int_0^T\Big(X_t^\top G^{XX,i}(t)X_t+2(G_t^{X,i})^\top X_t+2(D_t^i)^\top G^{DX,i}(t)X_t+(D_t^i)^\top G^{DD,i}(t)D_t^i\Big)\,dt\\&+\frac{1}{2}\Big(X_T^\top G^{XX,i}(T)X_T+2(G_T^{X,i})^\top X_T\Big),\end{aligned}\tag{1.3.2}$$

where the joint running Hessian $\begin{pmatrix}G^{XX,i}(t) & G^{XD,i}(t)\\ G^{DX,i}(t) & G^{DD,i}(t)\end{pmatrix}$, with $G^{XD,i} = (G^{DX,i})^\top$, is symmetric positive semidefinite, $G^{DD,i}(t)$ and $G^{XX,i}(T)$ are positive semidefinite, and $D^{-i} := (D^j)_{j\neq i}$ denotes opponents' controls. Control coercivity needs the Schur complement

$$S^i(t) := G^{DD,i}(t) - G^{DX,i}(t)\Big(G^{XX,i}(t)\Big)^+ G^{XD,i}(t) \succeq \lambda I, \qquad \lambda > 0, \tag{1.3.3}$$

rather than $G^{DD,i} \succeq \lambda I$, which permits the form $(x-d)^2$. Semidefiniteness gives $\operatorname{range} G^{XD,i} \subseteq \operatorname{range} G^{XX,i}$, so the pseudo-inverse term is exact. The cross term $G^{DX,i}$ lets the price of the control depend on the state (an input bill paid at a market price, a trading profit $D_t^i(V_t - P_t)$); $G^{DX,i} \equiv 0$ is the tracking case, where $S^i = G^{DD,i}$. Each player i seeks to minimize

$$J^i(D^i; D^{-i}) := \mathbb{E}\Big[\mathcal{C}^i(D^i; D^{-i})\Big]. \tag{1.3.4}$$

Opponents' controls enter player i's cost only through X_t, which they move through (1.3.1) and which player i observes through (1.3.5).

[1] The formulas below also allow primitive-shock linear coefficients $G_t^{X,i} = \bar{G}_t^{X,i} + \int_0^t G_t^{X,i}(r)\,dW_r$ with deterministic mean and kernel; deterministic coefficients set $G_t^{X,i}(r) \equiv 0$. The examples use deterministic coefficients.

1.3.2 Information structure

Observations. Player i does not observe X_t directly. Instead, their private information comes through the diffusion signal

$$dY_t^i = B_X^{Y,i}(t)\, X_t\, dt + E^i\, dW_t, \quad Y_0^i = 0, \tag{1.3.5}$$

where $B_X^{Y,i}(t) \in \mathbb{R}^{d\times d}$ is a deterministic observation gain and $E^i \in \mathbb{R}^{d\times(n+1)d}$ is the block selector extracting player i's noise channel ($W_t^i := E^i W_t$). All Brownian motions $\{W^0, W^1, \ldots, W^n\}$ are mutually independent. The precision matrix $P^i(t) := B_X^{Y,i}(t)^\top B_X^{Y,i}(t)$ measures the Fisher information rate about the state per unit time. The block-selector form is without loss, since replacing E^i by a general invertible noise loading $\sigma^i(t)$ substitutes $(B_X^{Y,i})^\top(\sigma^i\sigma^{i,\top})^{-1}$ for $(B_X^{Y,i})^\top$ throughout. Throughout the main characterization, the gain $B_X^{Y,i}(\cdot)$ is fixed; the two-player benchmark studies comparative statics and planner-chosen precision.

Notation for primitive shocks. Collect all primitive shocks into a single vector $W_t := (W_t^0, \ldots, W_t^n)^\top \in \mathbb{R}^{(n+1)d}$. The direct projection $\Pi^i := E^{i,\top}E^i$ is the identity on the i-th block and zero elsewhere. The block-selector structure ensures each player's observation noise is an independent component of W. Player i (almost) directly observes $\Pi^i W$ and infers $(I - \Pi^i)W$ from drift-based learning (Theorem 1.5).

Player i's information at time t is their observation history, $\mathcal{F}_t^i := \sigma(Y_s^i : s \le t)$. An admissible control is $\mathcal{F}_t^i$-adapted with $\mathbb{E}[\int_0^T \|D_t^i\|^2\, dt] < \infty$, where $\|\cdot\|$ denotes the Euclidean norm on $\mathbb{R}^d$.

Unresolved response. For each player i, define

$$\widetilde{X}_t^i(u) := X_t(u) - \widehat{X}_t^i(u),$$

the component of the state's impulse response to the u-th shock that player i has not yet resolved at time t. It is large when the shock is recent or the observation channel weak, and it shrinks as observations accumulate. Formally, $\widetilde{X}_t^i(u) = \operatorname{Cov}(X_t,\, dW_u \mid \mathcal{F}_t^i)/du$; it is the filtering gain in the noise-state update (Theorem 1.5), with its explicit formula in (1.A.4).

1.3.3 Nash equilibrium

Each player's drift depends on opponents' estimates $\widehat{X}_t^k := \mathbb{E}[X_t \mid \mathcal{F}_t^k]$. The following equilibrium concept fixes opponents' strategy maps rather than their realized actions, which Section 1.4 needs to show that a best response to noise-state linear opponents is itself noise-state linear.

Definition 1.2 (Nash equilibrium in private-signal strategies)**.** Throughout, a player's choice is a strategy map: for each $t \in [0, T]$, the action D_t^i is a (progressively measurable) functional of the private observation history $Y^i_{\cdot\wedge t}$. Write $D_t^i = D_t^i[Y^i]$, where the bracket notation $D_t^i[\cdot]$ denotes the map from the observation path $(Y_s^i)_{s\le t}$ to the action at time t. Admissible strategy maps must satisfy $\mathbb{E}\int_0^T \|D_t^i\|^2\, dt < \infty$ and must also be causally well posed: against any fixed profile of the opponents' complete causal maps, the coupled state–observation–strategy equations admit a unique strong causal solution, adapted to the control-free observation filtration. Strategies use no independent randomization; any private randomization, if admitted, is included in the player's filtration.

A profile $D^* = (D^{*,1}, \ldots, D^{*,n})$ is a Nash equilibrium if for every player i and every admissible alternative strategy map $\widetilde{D}^i[\cdot]$,

$$J^i(D^{*,i}; D^{*,-i}) \le J^i(\widetilde{D}^i; D^{*,-i}).$$

Holding opponents' strategy maps fixed is substantive because signals are endogenous. A unilateral deviation changes the drifts of opponents' observation paths, and with them the likelihoods they assign to histories. Under admissible square-integrable controls this only changes drifts, so the induced observation laws are mutually absolutely continuous on every finite horizon. A perfect-Bayesian formulation is simpler here than in discrete signal games, because Bayes' rule applies on a common support and the fixed strategy maps are evaluated on whatever observation paths occur.

Noise-state. On a given profile, the *noise-state* is the conditional mean path of the aggregated primitive shock:

$$\widehat{W}_t^i(u) := \mathbb{E}[W_u \mid \mathcal{F}_t^i], \qquad 0 \le u \le t. \tag{1.3.6}$$

Against fixed opponent maps, player i's realized control is a known input. Proposition 1.6 shows that its full downstream effect on the state, on opponents' responses, and so on Y^i is a deterministic causal functional of the realized control path. Subtracting that effect gives a control-free Gaussian observation with the same filtration. Off the candidate profile, $\widehat{W}^i$ then denotes the conditional projection of W in this equivalent reference coordinate. The deviating player does not choose a filter. The actual conditional state estimate is the reference estimate plus the known state response, while the innovation and its covariance are unchanged. On the candidate profile this agrees with (1.3.6).

Definition 1.3 ((Primitive-shock) Linear Process)**.** An L^2 stochastic process $L = (L_t)_{t\in[0,T]}$ is a (primitive-shock) linear process if there exist deterministic functions

$$\bar{L} : [0,T] \to \mathbb{R}^m, \qquad L : [0,T]^2 \to \mathbb{R}^{m\times(n+1)d},$$

such that, for every $t \in [0,T]$,

$$L_t \;=\; \bar{L}(t) + \int_0^t L(t,s)\,dW_s, \qquad \int_0^t \|L(t,s)\|^2\,ds < \infty.$$

When convenient, extend $L(t,s)$ by 0 for $s > t$.

Definition 1.4 (Noise-state Linear control / strategy)**.** An admissible strategy map is noise-state linear if, for the fixed filter described above, there exist deterministic functions

$$\bar{D}^i : [0,T] \to \mathbb{R}^d, \qquad D^i : [0,T]^2 \to \mathbb{R}^{d\times(n+1)d},$$

such that, for every private history Y^i and every $t \in [0,T]$,

$$D^i_t[Y^i] = \bar{D}^i_t + \int_0^t D^i_t(u)\, d_u \widehat{W}^i_t[Y^i](u), \qquad \int_0^t \|D^i_t(u)\|^2\, du < \infty.$$

Extend $D^i_t(u)$ by 0 for $u > t$ and drop the argument $[Y^i]$ on the equilibrium path.

1.4 General theory

The primitive-shock space plays the role of a dynamic Harsanyi type space [52]: conditional beliefs are projections of $W = (W^0, \ldots, W^n)$ onto private signal histories. Higher-order beliefs are compositions of these projections. Player i's estimate of player k's estimate of the state is $\mathbb{E}[\widehat{X}^k_t \mid \mathcal{F}^i_t]$; since $\widehat{X}^k_t$ is a linear functional of $d\widehat{W}^k_t$ with the deterministic kernel of (1.4.1), its projection onto $\mathcal{F}^i_t$ is the same functional applied to $\mathbb{E}[d_s \widehat{W}^k_t(s) \mid \mathcal{F}^i_t]$, which is (1.4.1) for player k evaluated on player i's noise-state. Every order of belief is therefore a finite composition of the kernels (Π^k, F^k), and no new state variable is needed. Actions reshape the signals from which beliefs are formed, so those maps must be solved as a fixed point (Figure 1.6).

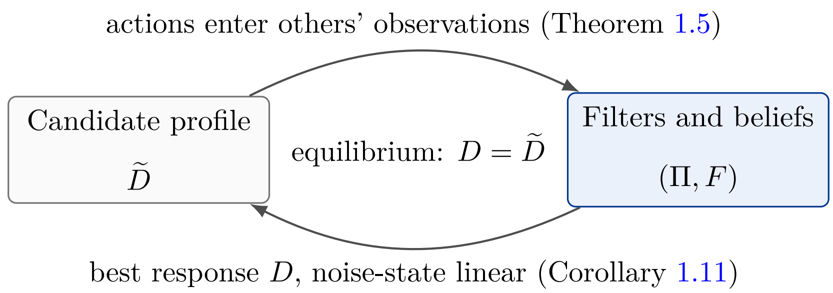

Figure 1.6: Noise-state fixed-point loop. Each map determines the other, so the profile must be solved as a fixed point.

Beliefs as deterministic kernels. The noise-state $\widehat{W}^i_t(\cdot)$ of (1.3.6) has a deterministic-kernel representation for its estimated-noise increments on the causal triangle $\Delta_T = \{(t,s) : 0 \le s \le t \le T\}$, extended by zero off Δ_T. Write d_u for increments in the path index (the source time u) and d_t for updates in time. The state has form $X_t = \bar{X}_t + \int_0^t X_t(s)\, dW_s$, with $X_t(s)$ the impulse-response kernel. Appendix 1.A gives the general linear-filtering result; here $C(t,s) = B^{Y,i}_X(t) X_t(s)$.

Theorem 1.5 (Beliefs are deterministic kernels: Filtering Closure)**.** *Fix a player i. Assume the state admits the primitive-shock linear form*

$$X_t = \bar{X}_t + \int_0^t X_t(s)\, dW_s,$$

and player i observes

$$dY^i_t = B^{Y,i}_X(t)\, X_t\, dt + E^i\, dW_t,$$

with deterministic $P^i(t) \succeq 0$ and block selector E^i.

Let $\Pi^i := E^{i,\top} E^i$ and $\widehat{X}^i_t := \mathbb{E}[X_t \mid \mathcal{F}^i_t]$, and define the innovation $dI^i_t := dY^i_t - B^{Y,i}_X(t)\, \widehat{X}^i_t\, dt$.

(i) The estimated-noise path increments admit the decomposition

$$d_u \widehat{W}^i_t(u) = \underbrace{\Pi^i\, dW_u}_{\textit{directly observed}} + \underbrace{\left(\int_0^t F^i_t(u,s)\, dW_s \right) du}_{\textit{indirect inference}}, \qquad u < t, \tag{1.4.1}$$

where the deterministic kernel F^i is given by

$$F^i_t(u,s) = \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top E^i \;+\; E^{i,\top} B^{Y,i}_X(u)\,\widetilde{X}^i_u(s) \;+\; \int_{\max(u,s)}^t \widetilde{X}^i_r(u)^\top P^i(r)\,\widetilde{X}^i_r(s)\,dr. \tag{1.4.2}$$

The pair (Π^i, F^i) is the blueprint of player i's beliefs.

(ii) The mixed (t,u)-update satisfies

$$d_t\Big(d_u \widehat{W}^i_t(u)\Big) = \widetilde{X}^i_t(u)^\top B^{Y,i}_X(t)^\top\, dI^i_t\, du.$$

Each old coordinate moves with the innovation, and $\widetilde{X}^i_t(u)$ is the gain. The singular term $\Pi^i\, dW_t$ is the coordinate's birth, its entry into the filter at the current date.

The unresolved response $\widetilde{X}^i$ is determined algebraically from F^i by (1.A.4); the filter kernel F^i is the unique solution of the forward evolution system (1.A.5)–(1.A.6) (Definition 1.16).

Proposition 1.6 (Against frozen opponents, the deviating player's own control is a known input to its filter)**.** *Fix player i and the opponents' complete noise-state linear strategy maps, including their filters. Let D^i and $\widetilde{D}^i$ be any two admissible controls generated by causally well-posed strategy maps, and write*

$$\Delta D^i := \widetilde{D}^i - D^i, \qquad \Delta X := X^{\widetilde{D}^i} - X^{D^i}.$$

The corresponding differences in the state, opponents' noise-states, and opponents' controls satisfy the exact deterministic causal system stated in (1.B.1)–(1.B.3). That system is linear in the realized input process ΔD^i even when the feedback rule generating $\widetilde{D}^i$ is nonlinear. No small-deviation approximation is used.

The own-observation difference is

$$dY^{i,\widetilde{D}^i}_t - dY^{i,D^i}_t = B^{Y,i}_X(t)\Delta X_t\, dt, \tag{1.4.3}$$

and ΔX_t is a deterministic causal functional of $\Delta D^i_{[0,t]}$. Taking $D^i \equiv 0$, define the control-free observation by subtracting this known response from the actual observation. Causal well-posedness gives an invertible causal correspondence between the actual and control-free histories, so they generate the same completed filtration. If $X^{i,0}$ and $Y^{i,0}$ denote the control-free state and observation, then

$$dI^i_t := dY^{i,0}_t - B^{Y,i}_X(t)\widehat{X}^{i,0}_t\,dt = dY^{i,\widetilde{D}^i}_t - B^{Y,i}_X(t)\Big(\widehat{X}^{i,0}_t + \Delta X_t[\widetilde{D}^i]\Big)dt \tag{1.4.4}$$

is the same standard Brownian innovation under every admissible own strategy. The deviating player's filter is then the conditional filter implied by the frozen environment and the known input.

Lemma 1.7 (Noise-state integrals as innovation integrals)**.** *Fix t and let $K_t : [0,t] \to \mathbb{R}^{m\times(n+1)d}$ be deterministic and square-integrable. In the reference innovation coordinate of Proposition 1.6,*

$$\int_0^t K_t(u)\,d_u\widehat{W}^i_t(u) = \int_0^t \left[K_t(s)E^{i,\top} + \int_0^s K_t(u)\,\widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top\,du\right]dI^i_s. \tag{1.4.5}$$

The closed linear span of the coordinates $\{\widehat{W}^i_t(u) : 0 \le u \le t\}$ contains every $\mathcal{F}^i_t$-measurable affine Gaussian random variable, so each has a deterministic noise-state representation. Equation (1.4.5) gives the corresponding innovation coefficient.

Proof. Integrating the mixed update in Theorem 1.5 from the birth time u to t gives, as a measure in u,

$$d_u\widehat{W}^i_t(u) = E^{i,\top}dI^i_u + \left(\int_u^t \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top\,dI^i_s\right)du.$$

Substitution followed by stochastic Fubini gives (1.4.5). For the span claim, conditional expectation is the orthogonal projection of linear functionals of the primitive shock onto the $\mathcal{F}^i_t$-measurable ones. The projected primitive increments $d_u\widehat{W}^i_t(u) = \mathbb{E}[dW_u \mid \mathcal{F}^i_t]$ then generate precisely that subspace. $\square$

Best-response verification and information wedges. The maximum-principle calculation identifies deterministic adjoint kernels and the information wedge $\mathcal{V}^i(t)$. The exact finite-deviation identity below then verifies the resulting candidate against the full admissible L^2 class.

Theorem 1.8 (Belief prices and information wedges)**.** *Fix a player i and suppose each opponent $k \neq i$ uses a noise-state linear strategy with deterministic impulse-response maps, which pins down a deterministic forward environment (state impulse response $X(\cdot,\cdot)$, filtering objects $(P^k, \widetilde{X}^k, F^k)$, and policy kernels D^k). There exist linear processes $H_t^{X,i}$ and $H_t^{k,i}(u)$, $k \neq i$, such that the following system holds for $t \in [0,T]$, with (1.4.8) holding for old-history coordinates $0 \leq u < t$. For each opponent $k \neq i$, define*

$$\mathcal{V}_t^{k,i} := B_X^{Y,k}(t)^\top E^k H_t^{k,i}(t) + P^k(t) \int_0^t \widetilde{X}_t^k(r)\, H_t^{k,i}(r)\, dr, \qquad \mathcal{V}_t^i := \sum_{k \neq i} \mathcal{V}_t^{k,i}. \tag{1.4.6}$$

The state price and belief prices satisfy

$$\frac{d}{dt} H_t^{X,i} = \underbrace{\left[-\left(G^{XX,i}(t)X_t + G_t^{X,i} + G^{XD,i}(t)D_t^i\right) - B_X^X(t)^\top H_t^{X,i}\right]}_{\text{perfect-information costate block}} - \underbrace{\mathcal{V}_t^i}_{\text{information wedge}}, \tag{1.4.7}$$

$$\frac{d}{dt} H_t^{k,i}(u) = -\left(B_k^X(t) D_t^k(u)\right)^\top H_t^{X,i} + X_t(u)^\top \mathcal{V}_t^{k,i}, \qquad k \neq i, \tag{1.4.8}$$

with terminal conditions $H_T^{X,i} = G^{XX,i}(T)X_T + G_T^{X,i}$ and $H_T^{k,i}(\cdot) = 0$. All coefficients, $P^k(t)$, $\widetilde{X}_t^k(r)$, $D_t^k(u)$, and $X_t(u)$, are determined by opponents' frozen strategy maps and do not change when player i deviates.

Proof. Appendix 1.B.2: the spike variation, the opponent filter variation, the coupled state–filter system, and the backward adjoints. The state price $H_t^{X,i}$ is the marginal cost of a spike that moves the state at t, and $H_t^{k,i}(u)$ the marginal cost of one that moves opponent k's estimate of the shock at u. Every channel through which player i's action moves what opponents believe, and every consequence of those moved beliefs, enters the dynamics of the state price only through the term $\mathcal{V}_t^i$ in (1.4.7). To write $\mathcal{V}_t^{k,i}$ explicitly, define the *filter contraction* of a kernel $M(r)$, $0 \le r \le t$, and its transpose action on a row-valued kernel $N(r)$ by

$$\mathfrak{B}^k M := E^k M(t) + \int_0^t \widetilde{C}_t^{k,Y}(r)\, M(r)\, dr, \qquad N\,\mathfrak{B}^{k\top} := N(t)E^{k\top} + \int_0^t N(r)\, \widetilde{C}_t^{k,Y}(r)^\top\, dr, \tag{1.4.9}$$

with $\widetilde{C}_t^{k,Y}(r) = B_X^{Y,k}(t)\widetilde{X}_t^k(r)$. Then $\mathcal{V}_t^{k,i} = B_X^{Y,k}(t)^\top\, \mathfrak{B}^k H_t^{k,i}$. The contraction $\mathfrak{B}^k H$ of (1.4.9) is opponent k's filter applied to a belief price: the birth term $E^k H(t)$ for the shock arriving now, plus the integral of H against the innovation kernel $\widetilde{C}_t^{k,Y}$ for the shocks already in the history. Chapters 2, 4, and 6 reuse it with $(E^k, \widetilde{C}^{k,Y})$ replaced by that chapter's observer's pair. In (1.4.8), $-D_t^k(u)^\top H_t^{X,i}$ is what player i pays, at the state price, for the action opponent k takes on a moved coordinate.

For a fixed profile, the continuation value is a functional of the physical state and the players' noise-states, with adjoints $H_t^{X,i} = \partial_X V_t^i$ and $H_t^{k,i}(u) = \delta V_t^i/\delta \widehat{W}_t^k(u)$. All coefficients in (1.4.7)–(1.4.8) are deterministic, so the mean and response-map coefficients decouple.

Corollary 1.9 (No observational externality)**.** *Suppose the state components that opponents observe are unaffected by player i's control. Then a unilateral deviation in D^i leaves every opponent's observation law, and hence the fixed filter for each $\widehat{W}^k$, $k \neq i$, unchanged; $\mathcal{V}_t^i \equiv 0$, the belief-price equations* (1.4.8) *decouple from the state price, and the optimality system reduces to the standard LQG game with fixed information structures.*

Conversely, a player can be made to ignore the effect of its own actions on the other players' beliefs by setting its information wedge to zero, which is how the exogenous-signal benchmarks of the later chapters are constructed.

Lemma 1.10 (Strict convexity on the fixed reference filtration)**.** *Fix opponents' complete causal maps, use the reference filtration of Proposition 1.6, and suppose the joint running*

Hessian is positive semidefinite and (1.3.3) holds. Then player i's objective is strictly convex and coercive on the admissible L^2 control space. For any two controls the state difference is the exact linear causal response to their control difference, and the quadratic part of the cost difference is half the joint form in $(\Delta X, \Delta D^i)$, which is nonnegative and dominates

$$\tfrac{1}{2}\lambda\,\mathbb{E}\int_0^T \|\Delta D_t^i\|^2\,dt$$

after completing the square in ΔX.

Proof. Proposition 1.6 makes the state and all induced opponent responses affine functions of the realized own-control input on one fixed adapted Hilbert space. Substituting them, the running integrand is half the joint form in $(\Delta X, \Delta D^i)$; completing the square in ΔX leaves $(\Delta D^i)^\top S^i \Delta D^i$ plus a nonnegative $G^{XX,i}$ term, so (1.3.3) gives the claim. □

A candidate that solves the stationarity system is a best response against every admissible deviation, including those outside the linear class.

Corollary 1.11 (Exact finite-deviation identity and best response over the full admissible class: Best Response Closure)**.** *Under the conditions of Theorem 1.8 and the Schur condition (1.3.3), let $D^{i,*}$ be a noise-state linear candidate whose deterministic kernels satisfy the stationarity system (1.B.13), and let $H^{X,i,*}$ be the corresponding state price. For every admissible strategy map $\widetilde{D}^i$ of any form, set*

$$\Delta D^i := \widetilde{D}^i - D^{i,*}, \qquad \Delta X := X^{\widetilde{D}^i} - X^{D^{i,*}}.$$

Then the following identity is exact:

$$\begin{aligned}
&J^i(\widetilde{D}^i;D^{-i}) - J^i(D^{i,*};D^{-i})\\
&\quad = \mathbb{E}\int_0^T (\Delta D_t^i)^\top \Big[G^{DD,i}(t)D_t^{i,*} + \mathbb{E}\Big[G^{DX,i}(t)X_t^* + B_i^X(t)^\top H_t^{X,i,*} \,|\, \mathcal{F}_t^i\Big]\Big]\,dt\\
&\qquad + \frac{1}{2}\mathbb{E}\int_0^T \Big[(\Delta X_t)^\top G^{XX,i}(t)\Delta X_t + 2(\Delta D_t^i)^\top G^{DX,i}(t)\Delta X_t + (\Delta D_t^i)^\top G^{DD,i}(t)\Delta D_t^i\Big]\,dt
\end{aligned}$$

$$+\frac{1}{2}\mathbb{E}\Big[(\Delta X_T)^\top G^{XX,i}(T)\Delta X_T\Big]. \tag{1.4.10}$$

The kernel stationarity equations make the first line on the right vanish for every adapted finite deviation. $D^{i,*}$ is then the unique best response over the full admissible L^2 class.

Proof. Appendix 1.B derives (1.4.10) by applying the physical and belief prices to the exact finite-difference system. Stationarity gives $G^{DD,i}D^{i,*}+\mathbb{E}[G^{DX,i}X^*+(B_i^X)^\top H^{X,i,*} \mid \mathcal{F}^i]=0$. The remaining terms are the joint quadratic form, nonnegative by Lemma 1.10, and bounded below by $\frac{1}{2}\lambda\mathbb{E}\int_0^T \|\Delta D_t^i\|^2 dt$ after completing the square. □

Corollary 1.11 takes a solution of the stationarity equations as given; the next proposition shows the best response exists.

Proposition 1.12 (Global solvability of unilateral best responses). *Fix a finite horizon T and complete linear causal strategy maps for all opponents, including their fixed filters from observations to controls. Under Assumption A.1, player i's optimization problem has a unique best response over the full admissible L^2 control class. By Corollary 1.11, this best response is noise-state linear.*

The proof is in Appendix A.2.

Remark 1.13 (Separation and the information wedge). Standard LQG separation holds the feedback gain fixed as signal precision changes. Here the FOC, $D_t^i = -(G^{DD,i}(t))^{-1}\mathbb{E}[G^{DX,i}(t)X_t + B_i^X(t)^\top H_t^{X,i} \mid \mathcal{F}_t^i]$, depends on the player's conditional expectation of the state price. When actions affect opponents' signals, that state price contains the wedge, which depends on opponents' filters and belief prices. Changing information then generally changes the policy rule itself, not only the posterior variance.

1.4.1 Identical-interest LQG games with control externalities

The baseline decentralized LQG game lets opponents' controls affect player i's payoff only through the state X_t. In the identical-interest variant, every player evaluates the same quadratic cost of the full control profile. The state and observation equations (1.3.1) and (1.3.5) are unchanged. Let

$$D_t := \Big((D_t^1)^\top, \dots, (D_t^n)^\top\Big)^\top \in \mathbb{R}^{nd}.$$

Fix deterministic weights

$$Q(t) = Q(t)^\top \succeq 0, \qquad R(t) = R(t)^\top \succ 0, \qquad Q_T = Q_T^\top \succeq 0,$$

together with deterministic linear terms q_t and q_T. The common cost is

$$\begin{aligned} C^{\mathrm{coop}}(D) := \frac{1}{2}\int_0^T \Big[X_t^\top Q(t) X_t + 2q_t^\top X_t + c(t) + D_t^\top R(t) D_t\Big]\, dt \\ + \frac{1}{2}\Big(X_T^\top Q_T X_T + 2q_T^\top X_T\Big). \end{aligned} \tag{1.4.11}$$

Every player has the same objective,

$$J^i(D^i; D^{-i}) := \mathbb{E}\left[C^{\mathrm{coop}}(D^i, D^{-i})\right], \qquad i = 1, \dots, n.$$

A profile D^* is a cooperative Nash equilibrium if, for every player i and every admissible private-signal strategy $\widetilde{D}^i$,

$$\mathbb{E}\left[C^{\mathrm{coop}}(D^*)\right] \le \mathbb{E}\left[C^{\mathrm{coop}}(\widetilde{D}^i, D^{*,-i})\right].$$

Any global solution of the decentralized team problem is a cooperative Nash equilibrium.

Identical-interest games with dispersed information are the team problem of Radner [97] and Marschak and Radner [82]; the common-information approach of Nayyar et al. [89] solves the case in which every player also observes a common signal history. Here the information structure remains decentralized: player i still computes the best response in the filtration $\mathcal{F}_t^i$, and a unilateral deviation still propagates through the other players' private filters. In the identical-interest case the wedge no longer prices an adversarial manipulation of beliefs.

The filtering equations are unchanged, since the observation system is unchanged. Only the source terms in the control first-order system change. Let

$$L_X(t) := Q(t)X_t + q_t, \qquad L_D(t) := R(t)D_t,$$
$$L_D^k(t) := \text{the } k\text{-th } d\text{-dimensional block of } L_D(t).$$

For player i, the state price satisfies

$$\frac{d}{dt}H_t^{X,i} = -L_X(t) - B_X^X(t)^\top H_t^{X,i} - \mathcal{V}_t^i, \qquad H_T^{X,i} = Q_T X_T + q_T. \tag{1.4.12}$$

For each opponent $k \neq i$, the belief price becomes

$$\frac{d}{dt}H_t^{k,i}(u) = -\Big(B_k^X(t)D_t^k(u)\Big)^\top H_t^{X,i} - \Big(D_t^k(u)\Big)^\top L_D^k(t) + X_t(u)^\top \mathcal{V}_t^{k,i}, \qquad H_T^{k,i}(\cdot) = 0. \tag{1.4.13}$$

The information wedge is the same filtering chain-rule object as before, (1.4.6). The new source term

$$-\Big(D_t^k(u)\Big)^\top L_D^k(t)$$

is the effort cost of the control opponent k takes, through its future posterior, in response to player i's deviation. Under the common objective, player i pays that cost too.

The best-response first-order condition is

$$\mathbb{E}\left[L_D^i(t) + B_i^X(t)^\top H_t^{X,i} \,\middle|\, \mathcal{F}_t^i\right] = 0. \tag{1.4.14}$$

Equivalently, writing the block decomposition

$$L_D^i(t) = R_{ii}(t)D_t^i + \sum_{j\neq i} R_{ij}(t)D_t^j,$$

and assuming $R_{ii}(t) \succ 0$, the cooperative Nash policy satisfies

$$D_t^i = -R_{ii}(t)^{-1}\mathbb{E}\left[B_i^X(t)^\top H_t^{X,i} + \sum_{j\neq i} R_{ij}(t)D_t^j \,\middle|\, \mathcal{F}_t^i\right]. \tag{1.4.15}$$

In the separable-effort case,

$$D_t^\top R(t) D_t = \sum_{j=1}^n (D_t^j)^\top R_j(t) D_t^j,$$

this reduces to

$$D_t^i = -R_i(t)^{-1} B_i^X(t)^\top \mathbb{E}\left[H_t^{X,i} \mid \mathcal{F}_t^i\right],$$

but the opponent-effort terms still enter the belief-price equations through $-\left(D_t^k(u)\right)^\top R_k(t) D_t^k$.

The objective remains quadratic and the state, filters, and frozen opponent maps remain linear in the unilateral best-response problem, so the best response is again noise-state linear. A fixed point of the resulting cooperative best-response system is then a Nash equilibrium in private-signal strategies with a common objective.

1.5 Conclusion

In the continuous-time LQG benchmark, beliefs, prices, and policy kernels are deterministic impulse-response kernels. Corollary 1.11 verifies that a fixed point in the noise-state linear class is a Nash equilibrium against arbitrary admissible L^2 deviations. The two-player examples show that the wedge changes welfare comparisons, information design, and the response to signal precision. Even with identical interests, teammates' effort enters the belief prices.

The following chapters keep the same objects and change the environment. Chapter 2 adds delayed observation rows, Chapter 3 rewrites the kernel system in stationary lag form, and Chapter 4 applies the adjoint logic to order-flow trading. Chapter 5 uses symmetry to solve finite local strategic-information games in relative-position coordinates, then aggregates their equilibrium outcomes across markets. Chapter 6 makes the observability of a deviation player-dependent.

Appendices to Chapter 1

The two appendices following Chapter 1 contain the canonical derivations used throughout the dissertation. Appendix 1.A derives the filtering closure and the primitive-noise representation of the noise state. Appendix 1.B derives the control first-order conditions, belief adjoints, and information wedge. Later chapters modify these objects rather than rederive them from scratch.

Proof-local notation. The symbols this appendix introduces for test functions, filter errors, and characteristic-function arguments are local to the proof in which they appear.

1.A Filtering

Theorem 1.5 follows by specializing a generic linear-filtering result to the physical observation model and then converting from innovations to primitive-shock coordinates.

1.A.1 Dynamics of the noise-state

Fix a player i. Recall the noise-state $\widehat{W}^i_t(u) := \mathbb{E}[W_u \mid \mathcal{F}^i_t]$, $0 \le u \le t$, and the observation channel

$$dY^i_t = \left(\bar{C}(t) + \int_0^t C_t(s)\, dW_s\right) dt + E^i\, dW_t, \qquad Y^i_0 = 0,$$

where $C_t(s)$ is deterministic. Since $\bar{C}$ is deterministic, set $\bar{C} \equiv 0$ without loss by replacing Y^i_t with $Y^i_t - \int_0^t \bar{C}(r)\, dr$.

Standing notation. Write $\mathbf{i} := \sqrt{-1}$ in bold (to distinguish it from the player index i) and $\mathcal{R} := L^\infty([0,t];\mathbb{R}^{d_Y})$; take all L^2 spaces to be complex; normalize the observation noise so that $E^i E^{i\top} = I_{d_Y}$, so $d\langle Y^i\rangle_t = I_{d_Y}\,dt$; and assume $C(t,\cdot) \in L^2([0,t];\mathbb{R}^{d_Y \times d_W})$ for each t, so that $\widehat{W}^i_t(u) = \mathbb{E}[W_u \mid \mathcal{F}^i_t] \in L^2$.

Lemma 1.14 (Test functionals). *For $\vec{\lambda} \in \mathcal{R}$ let φ^λ solve*

$$d\varphi^\lambda_s = \mathbf{i}\,\varphi^\lambda_s\,(dY^i_s)^\top \vec{\lambda}(s), \qquad \varphi^\lambda_0 = 1.$$

Each φ^λ_t is bounded, and $\{\varphi^\lambda_t : \vec{\lambda} \in \mathcal{R}\}$ is total in $L^2(\mathcal{F}^i_t)$. Consequently an $\mathcal{F}^i_t$-adapted L^2 process equals $\mathbb{E}[W_u \mid \mathcal{F}^i_t]$ as soon as it matches the test integrals

$$\mathbb{E}[\,\cdot\,\varphi^\lambda_t] = \mathbb{E}[W_u\,\varphi^\lambda_t] \qquad \textit{for every } \vec{\lambda} \in \mathcal{R}.$$

Proof. Each φ^λ_t is $\mathcal{F}^i_t$-measurable, and Itô's formula with $d\langle Y^i\rangle_s = I_{d_Y}\,ds$ gives $d|\varphi^\lambda_s|^2 = |\varphi^\lambda_s|^2|\vec{\lambda}(s)|^2\,ds$; so $|\varphi^\lambda_t| = \exp(\frac{1}{2}\int_0^t |\vec{\lambda}|^2\,ds)$ is deterministic and φ^λ_t is bounded. For a Brownian motion the family is total by [14, Lemma B.39]. The drift $\int_0^\cdot C(\cdot,s)\,dW_s$ of Y^i is square-integrable, so $\mathrm{Law}(Y^i) \ll \mathrm{Law}(\beta)$ for the Brownian motion β with the same diffusion coefficient (Girsanov), with density Λ. Totality transfers. If $\mathbb{E}[Z\,\varphi^\lambda_t] = 0$ for all $\vec{\lambda}$, then $Z\Lambda$ is orthogonal to every φ^λ_t under $\mathrm{Law}(\beta)$, so $Z\Lambda = 0$ by B.39 and then $Z = 0$. $\square$

Proposition 1.15 (Dynamics of $\widehat{W}^i_t(u)$). *The observation follows*

$$dY^i_t = \Big(\int_0^t C_t(s)\,dW_s\Big)dt + E^i\,dW_t, \qquad Y^i_0 = 0,$$

with $C(t,\cdot) \in L^2([0,t];\mathbb{R}^{d_Y \times d_W})$ for each t. Then there exist deterministic matrix-valued functions $B^{W,i}_Y(u,t)$ and $B^{W,i}_W(u;t,s)$ such that, for each fixed u,

$$d_t\widehat{W}^i_t(u) = \Big(\int_0^t B^{W,i}_W(u;t,s)\,d_s\widehat{W}^i_t(s)\Big)\,dt \;+\; B^{W,i}_Y(u,t)\,dY^i_t, \tag{1.A.1}$$

with

$$B_Y^{W,i}(u,t) = \int_0^t \partial_s \mathsf{C}_t^i(u,s)\, C_t(s)^\top\, ds + E^{i\top}\mathbf{1}_{\{u\geq t\}}, \qquad B_W^{W,i}(u;t,s) = -B_Y^{W,i}(u,t)\, C_t(s),$$

where $\mathsf{C}_t^i(u,s) := \mathrm{Cov}(W_u, W_s \mid \mathcal{F}_t^i)$ *is the deterministic conditional error-covariance kernel.*

Proof. **Step 1 (Test family).** Let φ^λ ($\vec{\lambda} \in \mathcal{R}$) be the test functionals of Lemma 1.14; each is bounded, solves $d\varphi_s^\lambda = \mathbf{i}\,\varphi_s^\lambda (dY_s^i)^\top \vec{\lambda}(s)$, and $W_u\varphi_t^\lambda, \widehat{W}_t^i(u)\varphi_t^\lambda \in L^2$, justifying the interchanges below.

Step 2 (Characterizing identity). As φ_t^λ is $\mathcal{F}_t^i$-measurable, the tower property gives $\mathbb{E}[\widehat{W}_t^i(u)\varphi_t^\lambda] = \mathbb{E}[W_u\varphi_t^\lambda]$ for all $\vec{\lambda}$ ($\star$). By Lemma 1.14 it suffices to substitute the ansatz (1.A.1), differentiate ($\star$) in t, and match.

Step 3 (Left side). Split $dY_t^i = E^i\, dW_t + (\int_0^t C(t,s)\, dW_s)\, dt$; only the martingale part has quadratic variation, so $d\langle \widehat{W}^i(u), \varphi^\lambda\rangle_t = \mathbf{i}\,\varphi_t^\lambda B_Y^{W,i}(u,t)\vec{\lambda}(t)\, dt$ (using $E^iE^{i\top} = I$; the drift affects the dt-comparison, not the bracket). Itô's product rule gives

$$\begin{aligned}\frac{d}{dt}\mathbb{E}[\widehat{W}_t^i(u)\varphi_t^\lambda] = \mathbf{i}\,\mathbb{E}\Big[\varphi_t^\lambda\Big(\widehat{W}_t^i(u)\int_0^t C(t,s)\, d_s\widehat{W}_t^i(s)^\top + B_Y^{W,i}(u,t)\Big)\Big]\vec{\lambda}(t)\\ + \mathbb{E}\Big[\varphi_t^\lambda\int_0^t \Big(B_W^{W,i}(u;t,s) + B_Y^{W,i}(u,t)C_t(s)\Big)\, d_s\widehat{W}_t^i(s)\Big].\end{aligned}$$

Step 4 (Right side). With $M_s^i(u,r) = \mathsf{C}_s^i(u,r) + \widehat{W}_s^i(u)\widehat{W}_s^i(r)^\top$ (Gaussian conditional law) and the conditional Itô isometry,

$$\begin{aligned}\frac{d}{dt}\mathbb{E}[W_u\varphi_t^\lambda] = \mathbf{i}\,\mathbb{E}\Big[\varphi_t^\lambda\Big(\int_0^t \partial_s\mathsf{C}_t^i(u,s)C_t(s)^\top ds\\ + \widehat{W}_t^i(u)\int_0^t C(t,s)\, d_s\widehat{W}_t^i(s)^\top\\ + E^{i\top}\mathbf{1}_{\{u\geq t\}}\Big)\Big]\vec{\lambda}(t).\end{aligned}$$

Step 5 (Matching). Equate the derivatives for all $\vec{\lambda}, t$; the common

$\widehat{W}^i_t(u)\int_0^t C(t,s)\,d_s\widehat{W}^i_t(s)^\top$ terms cancel. Varying the free endpoint $\vec{\lambda}(t)$ (which leaves φ^λ_t fixed) splits the identity into its $\vec{\lambda}(t)$-linear and $\vec{\lambda}(t)$-independent parts: *(i)* the $\vec{\lambda}(t)$-coefficient, deterministic after cancellation, gives B^Y; *(ii)* the remainder $\mathbb{E}[\varphi^\lambda_t\int_0^t(B+B^YC)\,d_s\widehat{W}^i_t(s)]=0$ for all $\vec{\lambda}$, so separation (Lemma 1.14) and Itô isometry force $B=-B^YC$. □

1.A.2 Specialization to the physical observation model

Specialize to the observation equation of Section 1.3,

$$dY^i_t = B^{Y,i}_X(t)\,X_t\,dt + E^i\,dW_t,$$

which corresponds to setting $C_t(s) = B^{Y,i}_X(t)\,X_t(s)$ in the general setting of Proposition 1.15.

Innovation process. Define the innovation

$$dI^i_t := dY^i_t - B^{Y,i}_X(t)\,\widehat{X}^i_t\,dt,$$

where $\widehat{X}^i_t := \mathbb{E}[X_t \mid \mathcal{F}^i_t]$. By the innovations theorem for Gaussian linear filtering (Liptser–Shiryaev [75], Theorem 7.12), I^i is a standard $(\mathcal{F}^i,\mathbb{P})$-Brownian motion.

Innovation form of the dynamics. Substituting $C_t(s) = B^{Y,i}_X(t)\,X_t(s)$ into Proposition 1.15 and using $dI^i_t = dY^i_t - B^{Y,i}_X(t)\,\widehat{X}^i_t\,dt$, the dynamics (1.A.1) become

$$d_t\widehat{W}^i_t(u) = \left(\int_0^t \partial_s \mathsf{C}^i_t(u,s)\,X_t(s)^\top\,ds\right)B^{Y,i}_X(t)^\top\,dI^i_t \;+\; E^{i,\top}\mathbf{1}_{\{u\ge t\}}\,dI^i_t. \tag{1.A.2}$$

The noise-state moves only with the innovation, through a covariance gain and an indicator term.

Identifying $\widetilde{X}^i$ as the filtering gain. For $u < t$, the indicator term in (1.A.2) is locally constant in u, so the mixed (t,u)-update takes the form

$$d_t\Big(d_u \widehat{W}^i_t(u)\Big) = \widetilde{X}^i_t(u)^\top B^{Y,i}_X(t)^\top \, dI^i_t \, du, \qquad 0 \le u < t \le T. \tag{1.A.3}$$

The gain in (1.A.3) is the conditional cross-covariance density, divided by dt, of dW_u and $B^{Y,i}_X(t)(X_t - \widehat{X}^i_t)\, dt$ given $\mathcal{F}^i_t$. Since $X_t - \widehat{X}^i_t = \int_0^t \widetilde{X}^i_t(s)\, dW_s$, Itô isometry gives $\widetilde{X}^i_t(u)^\top B^{Y,i}_X(t)^\top$.

Algebraic identity for $\widetilde{X}^i$ in terms of F^i. Define the induced estimated-state kernel $\widehat{X}^i_t(u) := X_t(u)\Pi^i + \int_0^t X_t(v)\, F^i_t(v,u)\, dv$. Since $\widetilde{X}^i_t(u) := X_t(u) - \widehat{X}^i_t(u)$,

$$\widetilde{X}^i_t(u) = X_t(u)(I - \Pi^i) - \int_0^t X_t(v)\, F^i_t(v,u)\, dv. \tag{1.A.4}$$

1.A.3 The blueprint: direct projection Π^i and filter kernel F^i

Define F^i by its evolution equation, with $\widetilde{X}^i$ given by (1.A.4).

Definition 1.16 (Blueprint system)**.** Fix player i. A deterministic causal kernel F^i on Δ_T solves the *blueprint system* if, with $\widetilde{X}^i$ defined by (1.A.4), it satisfies $F^i_0 = 0$ and

$$\partial_t F^i_t(u,s) = \widetilde{X}^i_t(u)^\top P^i(t)\, \widetilde{X}^i_t(s), \qquad u, s < t, \tag{1.A.5}$$

$$F^i_t(u,t) = \widetilde{X}^i_t(u)^\top B^{Y,i}_X(t)^\top E^i, \qquad F^i_t(t,u) = E^{i,\top} B^{Y,i}_X(t)\, \widetilde{X}^i_t(u). \tag{1.A.6}$$

Since (1.A.4) expresses $\widetilde{X}^i$ as a deterministic linear functional of F^i, the blueprint system (1.A.5)–(1.A.6) is a closed forward ODE for F^i, quadratic through the substitution; its right side is locally Lipschitz in F^i on bounded sets, so the solution is unique. It fixes the map from the primitive shocks to the player's estimated shocks.

Convert the innovations-based characterization of Subsections 1.A.1–1.A.2 into primitive-shock coordinates. What is needed is a deterministic kernel F^i such that for every $t \in [0,T]$ and $u < t$,

$$d_u \widehat{W}^i_t(u) = \Pi^i \, dW_u + \left(\int_0^t F^i_t(u,s) \, dW_s \right) du.$$

Π^i carries the contribution at $s = u$, and F^i carries the density across the other source times.

Derivation

Return to the innovation form of the estimated-noise dynamics (equation (1.A.2)): for $u \le t$,

$$d_u \widehat{W}^i_t(u) = E^{i,\top} \, dI^i_u + \left(\int_u^t \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top \, dI^i_s \right) du,$$

where the first term is the direct observation at time u and the second is the accumulated drift-based revision over $(u,t]$. To pass from innovations to primitive shock, expand each dI^i_s in terms of dW.

Expanding the innovation. Under the physical measure the innovation decomposes as

$$dI^i_t = E^i \, dW_t \; + \; B^{Y,i}_X(t) \, (X_t - \widehat{X}^i_t) \, dt.$$

Substituting the linear forms and using $\widetilde{X}^i_t(r) := X_t(r) - \widehat{X}^i_t(r)$:

$$dI^i_t = E^i \, dW_t + B^{Y,i}_X(t) \int_0^t \widetilde{X}^i_t(r) \, dW_r \, dt. \tag{1.A.7}$$

The drift is the state error $X_t - \widehat{X}^i_t$ written against the primitive shocks.

Expanding the direct term. Substituting (1.A.7) at $t = u$ into $E^{i,\top}\, dI^i_u$:

$$E^{i,\top}\, dI^i_u = \underbrace{E^{i,\top} E^i\, dW_u}_{= \Pi^i\, dW_u} + E^{i,\top} B^{Y,i}_X(u) \int_0^u \widetilde{X}^i_u(r)\, dW_r\, du. \tag{1.A.8}$$

The second term is absolutely continuous in u with primitive-shock density $E^{i,\top} B^{Y,i}_X(u)\, \widetilde{X}^i_u(s)$ for $s \le u$. In kernel coordinates over the primitive shock, the direct term is the singular diagonal part, a Dirac delta at $s = u$, while F^i collects everything with a density. Carrying the delta as the matrix coefficient Π^i keeps F^i a function.

Expanding the integral term. Substituting (1.A.7) into $\int_u^t \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top\, dI^i_s$:

$$\begin{aligned}\int_u^t \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top\, dI^i_s = &\underbrace{\int_u^t \widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top\, E^i\, dW_s}_{\text{(a) martingale}} \\ &+ \underbrace{\int_u^t \widetilde{X}^i_s(u)^\top P^i(s) \int_0^s \widetilde{X}^i_s(r)\, dW_r\, ds}_{\text{(b) abs. continuous}}.\end{aligned}$$

Term (a) contributes $\widetilde{X}^i_s(u)^\top B^{Y,i}_X(s)^\top E^i$ for $s \in (u, t]$. Term (b) has a *deterministic* integrand, so the stochastic Fubini theorem [113, Thm. 2.2] licenses exchanging the outer time integral $\int_u^t (\cdots)\, ds$ with the inner Itô integral $\int_0^s (\cdots)\, dW_r$; its hypothesis here reduces to the deterministic square-integrability

$$\int_u^t\!\int_0^s \left|\widetilde{X}^i_s(u)^\top P^i(s)\, \widetilde{X}^i_s(r)\right|^2 dr\, ds < \infty, \tag{1.A.9}$$

which holds because the impulse-response kernels $\widetilde{X}^i$ and the precision $P^i := (B^{Y,i}_X)^\top B^{Y,i}_X$ are bounded on the finite horizon $[0, T]$. Re-expressing the integration region $\{u \le s \le$

$t,\ 0 \le r \le s\}$ as $\{0 \le r \le t,\ \max(u,r) \le s \le t\}$ then gives

$$(\mathrm{b}) = \int_0^t \left(\int_{\max(u,r)}^t \widetilde{X}_s^i(u)^\top P^i(s)\, \widetilde{X}_s^i(r)\, ds \right) dW_r.$$

Term (b) contributes, at each source time r, the inference accumulated after $\max(u,r)$.

Assembling the density. Collecting the coefficient of $dW_s\, du$ from all three contributions, with causality ($\widetilde{X}_s^i(u) = 0$ for $u > s$) zeroing the irrelevant indicators:

$$F_t^i(u,s) = \widetilde{X}_s^i(u)^\top B_X^{Y,i}(s)^\top E^i \ + \ E^{i,\top} B_X^{Y,i}(u)\, \widetilde{X}_u^i(s) \ + \ \int_{\max(u,s)}^t \widetilde{X}_r^i(u)^\top P^i(r)\, \widetilde{X}_r^i(s)\, dr. \tag{1.A.10}$$

Kernel representation. Fix $t \in [0,T]$; the primitive Brownian motion W has dimension $(n+1)d$. Under the boundedness assumptions on $\widetilde{X}^i$, $B_X^{Y,i}$, E^i, and P^i, the expression in (1.A.10) defines a deterministic kernel

$$F_t^i \in L^2\Big([0,t]^2; \mathbb{R}^{(n+1)d \times (n+1)d}\Big).$$

Applying stochastic Fubini to the preceding identity gives, for every deterministic $h \in L^2([0,t]; \mathbb{R}^{(n+1)d})$,

$$\int_0^t h(u)^\top \Big(d_u \widehat{W}_t^i(u) - \Pi^i\, dW_u \Big) = \int_0^t \left(\int_0^t h(u)^\top F_t^i(u,s)\, du \right) dW_s. \tag{1.A.11}$$

The notation

$$d_u \widehat{W}_t^i(u) = \Pi^i\, dW_u + \left(\int_0^t F_t^i(u,s)\, dW_s \right) du$$

is understood in the integrated sense of (1.A.11).

The kernel is unique up to Lebesgue-null sets. Suppose F_t^i and G_t^i both satisfy (1.A.11), and write $K_t^i = F_t^i - G_t^i$. Then, for every $h \in L^2([0,t];\mathbb{R}^{(n+1)d})$,

$$\int_0^t \left(\int_0^t h(u)^\top K_t^i(u,s)\,du\right) dW_s = 0 \qquad \text{a.s.}$$

The Itô isometry implies

$$\int_0^t \left|\int_0^t h(u)^\top K_t^i(u,s)\,du\right|^2 ds = 0.$$

The Hilbert–Schmidt operator generated by K_t^i then vanishes on $L^2([0,t];\mathbb{R}^{(n+1)d})$, so

$$\int_0^t\int_0^t \|K_t^i(u,s)\|_{\mathrm{F}}^2\,du\,ds = 0.$$

Accordingly, $F_t^i = G_t^i$ almost everywhere on $[0,t]^2$.

1.B Control stationarity and finite-deviation verification

This appendix derives the stationarity condition (1.B.13), the closed backward system for the adjoint kernels, and the exact finite-deviation identity used in Corollary 1.11. Nothing here assumes that the best response is noise-state linear. Opponents' complete causal maps stay frozen, while the deviating player may use any causally well-posed square-integrable observation-feedback strategy.

1.B.1 Frozen opponents, exact finite differences, and the reference innovation

Fix player i and the opponents' complete noise-state linear maps, including their filters. Let D^i and $\widetilde{D}^i$ be any two admissible own controls and define

$$\Delta D^i_t := \widetilde{D}^i_t - D^i_t, \qquad \Delta X_t := X_t^{\widetilde{D}^i} - X_t^{D^i}.$$

For $k \neq i$, the direct primitive-shock part of the two fixed filters cancels, so write

$$d_u \widehat{W}_t^{k,\widetilde{D}^i}(u) - d_u \widehat{W}_t^{k,D^i}(u) = \Delta w^k_t(u)\, du.$$

The differences satisfy the exact system

$$\partial_t \Delta w^k_t(u) = \widetilde{X}^k_t(u)^\top P^k(t) \left(\Delta X_t - \int_0^t X_t(r) \Delta w^k_t(r)\, dr \right), \qquad 0 \le u < t, \tag{1.B.1}$$

$$\Delta w^k_t(t) = E^{k,\top} B_X^{Y,k}(t) \left(\Delta X_t - \int_0^t X_t(r) \Delta w^k_t(r)\, dr \right), \tag{1.B.2}$$

$$\frac{d}{dt} \Delta X_t = B_X^X(t) \Delta X_t + B_i^X(t) \Delta D^i_t + \sum_{k \neq i} B_k^X(t) \int_0^t D^k_t(u) \Delta w^k_t(u)\, du, \qquad \Delta X_0 = 0. \tag{1.B.3}$$

The opponent-control difference is $\Delta D^k_t = \int_0^t D^k_t(u) \Delta w^k_t(u)\, du$. Subtracting the two closed-loop systems gives these identities. They are exact and linear in the realized input ΔD^i; the rule generating that input may be nonlinear in the observation history.

The own-observation difference is

$$dY_t^{i,\widetilde{D}^i} - dY_t^{i,D^i} = B_X^{Y,i}(t) \Delta X_t\, dt.$$

Since (1.B.1)–(1.B.3) have deterministic causal coefficients, ΔX_t is known from the realized own-control difference up to time t. Taking $D^i \equiv 0$ defines the control-free pair

$(X^{i,0}, Y^{i,0})$. The actual observation comes from $Y^{i,0}$ by adding the known response of the realized control, and causal well-posedness gives the reverse reconstruction. The completed filtrations coincide. With

$$\widehat{X}_t^{i,0} := \mathbb{E}[X_t^{i,0} \mid \mathcal{F}_t^{i,0}],$$

the common reference innovation is

$$dI_t^i := dY_t^{i,0} - B_X^{Y,i}(t)\widehat{X}_t^{i,0}\,dt = dY_t^{i,\widetilde{D}^i} - B_X^{Y,i}(t)\Big(\widehat{X}_t^{i,0} + \Delta X_t[\widetilde{D}^i]\Big)dt. \tag{1.B.4}$$

It is a standard Brownian motion whose law does not depend on the own strategy, so the deviating player filters against the frozen environment and this known input.

1.B.2 Spike variation and the first-variation system

Fix $t \in [0,T)$ and a control spike $z \in L^2(\Omega, \mathcal{F}_t^i; \mathbb{R}^d)$. The spike perturbation $D_s^{i,\rho,\varepsilon} := D_s^i + \rho\, z\, \mathbf{1}_{[t,t+\varepsilon]}(s)$ induces the normalized first variation $\delta X_s := \lim_{\varepsilon\downarrow 0} \frac{1}{\varepsilon}\frac{\partial}{\partial\rho} X_s^{\rho,\varepsilon}\Big|_{\rho=0}$ with $\delta X_t = v := B_i^X(t)z$, the induced state impulse.

Opponents apply fixed linear maps, so $v \mapsto (\delta X_s, \{\delta w_s^k(\cdot)\}_{k\neq i})$ is linear with deterministic coefficients. The spike-induced state response is known to the deviating player and enters the reference conditional mean, so it creates no new own-innovation increment. The reference innovation (1.B.4) is unchanged by construction.

Opponents' filter response. For each $k \neq i$, $\delta(d_u \widehat{W}_s^k(u)) = \delta w_s^k(u)\,du$ with $\delta w_t^k(\cdot) \equiv 0$. Existing coordinates satisfy

$$\partial_s \delta w_s^k(u) = \widetilde{X}_s^k(u)^\top P^k(s)\Big(\delta X_s - \int_0^s X_s(r)\,\delta w_s^k(r)\,dr\Big), \qquad 0 \le u < s, \tag{1.B.5}$$

and the coordinate born at $u = s$ satisfies

$$\delta w_s^k(s) = E^{k,\top} B_X^{Y,k}(s)\Big(\delta X_s - \int_0^s X_s(r)\, \delta w_s^k(r)\, dr\Big). \tag{1.B.6}$$

The opponent control variation is $\delta D_s^k = \int_0^s D_s^k(u)\, \delta w_s^k(u)\, du$, giving the coupled state system

$$\frac{d}{ds}\delta X_s = B_X^X(s)\, \delta X_s + \sum_{k \neq i} B_k^X(s) \int_0^s D_s^k(u)\, \delta w_s^k(u)\, du, \qquad \delta X_t = v. \tag{1.B.7}$$

After the spike, the state moves through its own dynamics and through the opponents' control variations.

1.B.3 Transition kernels

The coupled system (1.B.5)–(1.B.7), with birth condition (1.B.6), has bounded deterministic coefficients and admits a unique deterministic two-parameter evolution family.

Definition 1.17 (Block transition kernels)**.** For $s \geq t$, define $\Phi^{XX}(s,t) \in \mathbb{R}^{d\times d}$ and $\Phi^{Xk}(s,t,u)$ ($k \neq i$, $u \in [0,t]$) by

$$\delta X_s = \Phi^{XX}(s,t)\, v + \sum_{k \neq i} \int_0^t \Phi^{Xk}(s,t,u)\, \delta w_t^k(u)\, du,$$

for initial data $\delta X_t = v$ and the given $\delta w_t^k(\cdot)$. In the spike deviation $\delta w_t^k \equiv 0$, so $\delta X_s = \Phi^{XX}(s,t)\, v$. The notation $\Phi^{Xk}(s,t,t)$ denotes the birth value for the coordinate born at time t.

These blocks satisfy

$$\begin{aligned}\frac{\partial}{\partial t}\Phi^{XX}(s,t) = -\Phi^{XX}(s,t)B_X^X(t) - \sum_{k\neq i} \Phi^{Xk}(s,t,t)E^{k,\top}B_X^{Y,k}(t)\\ - \sum_{k\neq i}\int_0^t \Phi^{Xk}(s,t,r)\, \widetilde{X}_t^k(r)^\top P^k(t)\, dr,\end{aligned} \tag{1.B.8}$$

$$\frac{\partial}{\partial t}\Phi^{Xk}(s,t,u) = -\Phi^{XX}(s,t)B^X_k(t)D^k_t(u) + \Phi^{Xk}(s,t,t)E^{k,\top}B^{Y,k}_X(t)X_t(u) \\ + \int_0^t \Phi^{Xk}(s,t,r)\,\widetilde{X}^k_t(r)^\top P^k(t)\,X_t(u)\,dr, \tag{1.B.9}$$

with $\Phi^{XX}(t,t) = I$ and $\Phi^{Xk}(t,t,u) = 0$. The terms involving $\Phi^{Xk}(s,t,t)E^{k,\top}B^{Y,k}_X(t)$ come from the birth condition (1.B.6).

1.B.4 First variation of costs and the adjoint kernels

Fix player i. Using $\delta X_s = \Phi^{XX}(s,t)v$ and the linear form of X_s, the normalized cost variation is

$$\delta J^i_t(z) = \mathbb{E}\Big[z^\top\Big(G^{DD,i}(t)D^i_t + G^{DX,i}(t)\Big[\bar{X}_t + \int_0^T X_t(u)\,d\widehat{W}^i_t(u)\Big] \\ + B^X_i(t)^\top\Big[\bar{H}^{X,i}_t + \int_0^T H^{X,i}_t(u)\,d\widehat{W}^i_t(u)\Big]\Big)\;\Big|\;\mathcal{F}^i_t\Big], \tag{1.B.10}$$

where the running marginal cost of the state now includes the cross term with the player's own control, $G^{XD,i}(s)D^i_s = G^{XD,i}(s)\Big[\bar{D}^i_s + \int_0^s D^i_{W,s}(u)\,dW_u\Big]$, and

$$\bar{H}^{X,i}_t := \int_t^T \Phi^{XX}(s,t)^\top\Big[G^{XX,i}(s)\bar{X}_s + \bar{G}^{X,i}_s + G^{XD,i}(s)\bar{D}^i_s\Big]\,ds \\ + \Phi^{XX}(T,t)^\top\Big[G^{XX,i}(T)\bar{X}_T + \bar{G}^{X,i}_T\Big], \tag{1.B.11}$$

$$H^{X,i}_t(u) := \int_t^T \Phi^{XX}(s,t)^\top G^{XX,i}(s)X_s(u)\,ds \\ + \int_t^T \Phi^{XX}(s,t)^\top\Big[G^{X,i}_s(u) + G^{XD,i}(s)D^i_{W,s}(u)\Big]\,ds \\ + \Phi^{XX}(T,t)^\top\Big[G^{XX,i}(T)X_T(u) + G^{X,i}_T(u)\Big]. \tag{1.B.12}$$

The upper limit in (1.B.10) may be written as T since $d_u\widehat{W}^i_t(u) = 0$ for $u > t$.

Stationarity condition (globally necessary and sufficient). Setting $\delta J_t^i(z) = 0$ for all $\mathcal{F}_t^i$-measurable z gives

$$
\begin{aligned}
&G^{DD,i}(t) D_t^i + G^{DX,i}(t)\Big[\bar{X}_t + \int_0^T X_t(u)\, d\widehat{W}_t^i(u)\Big] \\
&\quad + B_i^X(t)^\top \Big[\bar{H}_t^{X,i} + \int_0^T H_t^{X,i}(u)\, d\widehat{W}_t^i(u)\Big] = 0 \qquad \text{a.s.,}
\end{aligned}
\tag{1.B.13}
$$

with equilibrium coefficients

$$
\begin{aligned}
\bar{D}_t^i &= -G^{DD,i}(t)^{-1}\Big[G^{DX,i}(t)\bar{X}_t + B_i^X(t)^\top \bar{H}_t^{X,i}\Big], \\
D_t^i(u) &= -G^{DD,i}(t)^{-1}\Big[G^{DX,i}(t) X_t(u) + B_i^X(t)^\top H_t^{X,i}(u)\Big].
\end{aligned}
$$

Equivalently, $D_t^i = -(G^{DD,i}(t))^{-1}\mathbb{E}[G^{DX,i}(t)X_t + B_i^X(t)^\top(\bar{H}_t^{X,i} + \int_0^T H_t^{X,i}(u)\, dW_u) \mid \mathcal{F}_t^i]$. The player cannot act on shocks it has not observed, so the control is the conditional expectation of the state price given its own history. This is not a separation principle (Remark 1.13), because the state price itself is an equilibrium object shaped by opponents' filters and belief prices.

1.B.5 Closed backward system for the adjoints

Define the belief-price coefficients, the prices of moving opponent k's noise-state coordinates,

$$
\begin{aligned}
\bar{H}_t^{k,i}(u) &:= \int_t^T \Phi^{Xk}(s,t,u)^\top G^{XX,i}(s)\bar{X}_s\, ds \\
&\quad + \int_t^T \Phi^{Xk}(s,t,u)^\top \Big[\bar{G}_s^{X,i} + G^{XD,i}(s)\bar{D}_s^i\Big]\, ds \\
&\quad + \Phi^{Xk}(T,t,u)^\top \Big[G^{XX,i}(T)\bar{X}_T + \bar{G}_T^{X,i}\Big],
\end{aligned}
\tag{1.B.14}
$$

$$
\begin{aligned}
H_t^{k,i}(u,r) &:= \int_t^T \Phi^{Xk}(s,t,u)^\top G^{XX,i}(s) X_s(r)\, ds \\
&\quad + \int_t^T \Phi^{Xk}(s,t,u)^\top \Big[G_s^{X,i}(r) + G^{XD,i}(s) D_{W,s}^i(r)\Big]\, ds
\end{aligned}
$$

$$+ \Phi^{Xk}(T,t,u)^\top G^{XX,i}(T) X_T(r)$$
$$+ \Phi^{Xk}(T,t,u)^\top G^{X,i}_T(r). \tag{1.B.15}$$

Equivalently, introduce the noise-linear adjoint processes

$$H^{X,i}_t := \bar{H}^{X,i}_t + \int_0^T H^{X,i}_t(r)\, dW_r,$$
$$H^{k,i}_t(u) := \bar{H}^{k,i}_t(u) + \int_0^T H^{k,i}_t(u,r)\, dW_r, \qquad k \neq i.$$

Here and below, $H^{X,i}_t(u)$ denotes the kernel of the process $H^{X,i}_t$, while $H^{k,i}_t(u)$ denotes a process whose kernel carries the second argument $H^{k,i}_t(u,r)$. The argument count differs between the two families because $H^{k,i}$ is itself indexed by a coordinate. For each opponent $k \neq i$, set

$$\mathcal{V}^{k,i}_t := B^{Y,k}_X(t)^\top E^k H^{k,i}_t(t) + P^k(t) \int_0^t \widetilde{X}^k_t(r) H^{k,i}_t(r)\, dr, \qquad \mathcal{V}^i_t := \sum_{k \neq i} \mathcal{V}^{k,i}_t. \tag{1.B.16}$$

The first term carries the coordinate born at $u = t$ in (1.B.6), the second the accumulated drift inference from (1.B.5). Together they are the information wedge, the price player i pays for what player k learns from its action. In coefficients,

$$\bar{\mathcal{V}}^{k,i}_t = B^{Y,k}_X(t)^\top E^k \bar{H}^{k,i}_t(t) + P^k(t) \int_0^t \widetilde{X}^k_t(r) \bar{H}^{k,i}_t(r)\, dr, \tag{1.B.17}$$
$$\mathcal{V}^{k,i}_t(r) = B^{Y,k}_X(t)^\top E^k H^{k,i}_t(t,r) + P^k(t) \int_0^t \widetilde{X}^k_t(s) H^{k,i}_t(s,r)\, ds. \tag{1.B.18}$$

The adjoints satisfy the two backward equations

$$\frac{d}{dt} H^{X,i}_t = -\Big[G^{XX,i}(t) X_t + G^{X,i}_t + G^{XD,i}(t) D^i_t\Big] - B^X_X(t)^\top H^{X,i}_t - \mathcal{V}^i_t, \tag{1.B.19}$$
$$\frac{d}{dt} H^{k,i}_t(u) = -(B^X_k(t) D^k_t(u))^\top H^{X,i}_t + X_t(u)^\top \mathcal{V}^{k,i}_t, \qquad k \neq i. \tag{1.B.20}$$

The terminal conditions are

$$H_T^X = G^{XX,i}(T)X_T + G_T^{X,i}, \qquad H_T^k(\cdot) = 0, \qquad k \neq i. \tag{1.B.21}$$

1.B.6 Exact finite-deviation verification

Let $D^{i,*}$ be the noise-state linear candidate satisfying (1.B.13), and let $\widetilde{D}^i$ be any admissible finite deviation. Put $\Delta D^i := \widetilde{D}^i - D^{i,*}$ and let $(\Delta X, \{\Delta w^k\}_{k\neq i})$ solve (1.B.1)–(1.B.3). Expanding the quadratic objective gives

$$\begin{aligned}
&J^i(\widetilde{D}^i; D^{-i}) - J^i(D^{i,*}; D^{-i}) \\
&= \mathbb{E}\int_0^T \Big[\Big(G^{XX,i}(t)X_t^* + G_t^{X,i} + G^{XD,i}(t)D_t^{i,*}\Big)^\top \Delta X_t \\
&\qquad\quad + (D_t^{i,*})^\top G^{DD,i}(t)\Delta D_t^i + (\Delta D_t^i)^\top G^{DX,i}(t)X_t^*\Big]dt \\
&\quad + \mathbb{E}\Big[\Big(G^{XX,i}(T)X_T^* + G_T^{X,i}\Big)^\top \Delta X_T\Big] \\
&\quad + \frac{1}{2}\mathbb{E}\int_0^T \Big[(\Delta X_t)^\top G^{XX,i}(t)\Delta X_t + 2(\Delta D_t^i)^\top G^{DX,i}(t)\Delta X_t + (\Delta D_t^i)^\top G^{DD,i}(t)\Delta D_t^i\Big]\,dt \\
&\quad + \frac{1}{2}\mathbb{E}\Big[(\Delta X_T)^\top G^{XX,i}(T)\Delta X_T\Big].
\end{aligned} \tag{1.B.22}$$

Apply integration by parts to the paired forward–backward quantity

$$\Psi_t := H_t^{X,i,*\top}\Delta X_t + \sum_{k\neq i}\int_0^t H_t^{k,i,*}(u)^\top \Delta w_t^k(u)\,du.$$

Set

$$e_t^k := \Delta X_t - \int_0^t X_t(r)\Delta w_t^k(r)\,dr, \qquad k \neq i.$$

The product rule also picks up a birth contribution at the newly born coordinate $u = t$. Explicitly,

$$\frac{d}{dt}\Psi_t = \left(\frac{d}{dt}H_t^{X,i,*}\right)^\top \Delta X_t + H_t^{X,i,*\top}\frac{d}{dt}\Delta X_t$$

$$
\begin{aligned}
&+\sum_{k\neq i}\int_0^t\left[\left(\frac{d}{dt}H_t^{k,i,*}(u)\right)^{\top}\Delta w_t^k(u)+H_t^{k,i,*}(u)^{\top}\partial_t\Delta w_t^k(u)\right]du\\
&+\sum_{k\neq i}H_t^{k,i,*}(t)^{\top}\Delta w_t^k(t).
\end{aligned}
\tag{1.B.23}
$$

Substitute the exact finite-difference equations (1.B.1)– (1.B.3) and the adjoint equations (1.B.19)–(1.B.20). The physical-drift pairings and the opponent-control pairings cancel, leaving

$$
\begin{aligned}
\frac{d}{dt}\Psi_t = &-\Big(G^{XX,i}(t)X_t^*+G_t^{X,i}+G^{XD,i}(t)D_t^{i,*}\Big)^{\top}\Delta X_t+H_t^{X,i,*\top}B_i^X(t)\Delta D_t^i\\
&+\sum_{k\neq i}\Bigg\{-\mathcal{V}_t^{k,i,*\top}e_t^k+\left[P^k(t)\int_0^t\widetilde{X}_t^k(u)H_t^{k,i,*}(u)\,du\right]^{\top}e_t^k\\
&\qquad+\Big[B_X^{Y,k}(t)^{\top}E^kH_t^{k,i,*}(t)\Big]^{\top}e_t^k\Bigg\}.
\end{aligned}
\tag{1.B.24}
$$

The last two terms in braces are, respectively, the old-history and newly born components of $\mathcal{V}_t^{k,i,*\top}e_t^k$ by (1.B.16). The braces then vanish opponent by opponent, and

$$
\frac{d}{dt}\Psi_t=-\Big(G^{XX,i}(t)X_t^*+G_t^{X,i}+G^{XD,i}(t)D_t^{i,*}\Big)^{\top}\Delta X_t+H_t^{X,i,*\top}B_i^X(t)\Delta D_t^i. \tag{1.B.25}
$$

Integrating from 0 to T, the initial term vanishes because $\Delta X_0 = 0$ and the belief-coordinate integral is empty. At the other endpoint, $H_T^{k,i,*} = 0$, while $H_T^{X,i,*}=G^{XX,i}(T)X_T^*+G_T^{X,i}$. The terminal physical pairing is exactly the terminal linear term in (1.B.22), and the running pairing $(G^{XD,i}D^{i,*})^{\top}\Delta X$ cancels the cross linear term of (1.B.22). What remains from all linear terms is

$$
\mathbb{E}\int_0^T(\Delta D_t^i)^{\top}\Big[G^{DX,i}(t)X_t^*+B_i^X(t)^{\top}H_t^{X,i,*}\Big]\,dt.
$$

Since ΔD_t^i is adapted to the common reference filtration,

$$
\mathbb{E}\Big[(\Delta D_t^i)^{\top}\Big(G^{DX,i}(t)X_t^*+B_i^X(t)^{\top}H_t^{X,i,*}\Big)\Big]=\mathbb{E}\Big[(\Delta D_t^i)^{\top}\mathbb{E}\Big[G^{DX,i}(t)X_t^*+B_i^X(t)^{\top}H_t^{X,i,*}\,|\,\mathcal{F}_t^i\Big]\Big].
$$

Substituting gives the exact identity (1.4.10). The stationarity condition annihilates its linear term for every finite adapted deviation, while the remaining quadratic terms are nonnegative. Under (1.3.3), equality is possible only when $\Delta D^i = 0$ in L^2.

1.B.7 CARA / risk-sensitive extension: exponential utility FOC

Theorem 1.18 (Risk-sensitive extension)**.** *Take $G^{DX,i} \equiv 0$, the tracking case. Replace (1.3.4) with the entropic risk measure*

$$J^i_{\theta_i} := \theta_i^{-1} \log \mathbb{E}\Big[e^{\,\theta_i \mathcal{C}^i}\Big], \qquad \theta_i > 0, \tag{1.B.26}$$

where $\mathcal{C}^i$ is the random cost (1.3.2). A Taylor expansion gives $J^i_{\theta_i} = \mathbb{E}[\mathcal{C}^i] + \frac{\theta_i}{2} \operatorname{Var}(\mathcal{C}^i) + O(\theta_i^2)$, so θ_i governs risk aversion (the risk-neutral case $\theta_i \to 0$ recovers (1.3.4)).

Under a linear profile, the stationarity condition is identical to (1.B.13) with $\widehat{W}^i_t$ replaced by the risk-adjusted noise-state

$$\widehat{W}^{i,\theta}_t := \Big(I - \theta_i C_t K_t\Big)^{-1} \Big(\widehat{W}^i_t + \theta_i C_t k_t\Big),$$

where $C_t := \mathsf{C}^i_t(u,r) = \operatorname{Cov}(W_u, W_r \mid \mathcal{F}^i_t)$ is the deterministic conditional covariance kernel of W given $\mathcal{F}^i_t$ (deterministic by Gaussianity of the conditional law under the linear profile), K_t is the symmetric quadratic kernel of $\mathcal{C}^i$ in the representation (1.B.28), and k_t is the corresponding linear kernel (1.B.27). The formula is valid when $I - \theta_i C_t^{1/2} K_t C_t^{1/2} \succ 0$, which is precisely finiteness of $\mathbb{E}[e^{\theta_i \mathcal{C}^i}]$. The filtering structure is unchanged, but certainty equivalence fails; the system is the multi-agent, endogenous-information analogue of linear-exponential-quadratic-Gaussian control [62] and risk-sensitive certainty equivalence [115].

The spike variation of Section 1.B.2 carries over; differentiating the objective (1.B.26) in the spike amplitude gives

$$\delta J^i_{\theta_i,t}(z) = \mathbb{E}\Big[Z^i\,\delta\mathcal{C}^i_t(h)\Big], \qquad Z^i := \frac{e^{\theta_i \mathcal{C}^i}}{\mathbb{E}[e^{\theta_i \mathcal{C}^i}]},$$

with $\delta\mathcal{C}^i_t(h)$ identical to (1.B.10).

Quadratic structure of the cost. Expanding $\mathcal{C}^i$ through the linear forms $X_s = \bar{X}_s + \int_0^s X_s(u)\,dW_u$ and $D^i_s = \bar{D}^i_s + \int_0^s D^i_{W,s}(u)\,dW_u$ and applying stochastic Fubini gives the representation

$$\mathcal{C}^i = c_0 + \int_0^T \ell(u)^\top dW_u + \tfrac{1}{2}\int_0^T\!\!\int_0^T dW_u^\top K(u,r)\,dW_r,$$

where c_0 absorbs deterministic and trace terms, and the linear and symmetric quadratic kernels are

$$\begin{aligned}
\ell(u) := \int_u^T \Big[& X_t(u)^\top\Big(G^{XX,i}(t)\bar{X}_t + \bar{G}^{X,i}_t\Big) + G^{X,i}_t(u)^\top \bar{X}_t \\
& + D^i_{W,t}(u)^\top G^{DD,i}(t)\,\bar{D}^i_t\Big]\,dt \\
& + \; X_T(u)^\top\Big(G^{XX,i}(T)\bar{X}_T + \bar{G}^{X,i}_T\Big) + G^{X,i}_T(u)^\top \bar{X}_T, \qquad (1.B.27) \\
K(u,r) := \int_{\max(u,r)}^T \Big[& X_t(u)^\top G^{XX,i}(t)\,X_t(r) + X_t(u)^\top G^{X,i}_t(r) \\
& + G^{X,i}_t(u)^\top X_t(r) + D^i_{W,t}(u)^\top G^{DD,i}(t)\,D^i_{W,t}(r)\Big]\,dt \\
& + \; X_T(u)^\top G^{XX,i}(T)\,X_T(r) + X_T(u)^\top G^{X,i}_T(r) + G^{X,i}_T(u)^\top X_T(r). \qquad (1.B.28)
\end{aligned}$$

Tilted conditional FOC. Since $\mathcal{C}^i$ is quadratic in W, the weight Z^i tilts a Gaussian measure by a quadratic form, preserving Gaussianity. Setting $\delta J^i_{\theta_i,t}(z) = 0$ for all $\mathcal{F}^i_t$-measurable z and applying the tower property gives

$$G^{DD,i}(t)D^i_t + B^X_i(t)^\top \mathbb{E}^{Q^i}\left[\bar{H}^{X,i}_t + \int_0^T H^{X,i}_t(u)\,dW_u \,\middle|\, \mathcal{F}^i_t\right] = 0,$$

where $\mathbb{E}^{Q^i}[\,\cdot \mid \mathcal{F}_t^i] := \mathbb{E}[Z^i(\cdot) \mid \mathcal{F}_t^i]/\mathbb{E}[Z^i \mid \mathcal{F}_t^i]$.

Risk-adjusted noise-state. Since $\bar{H}_t^{X,i} + \int_0^T H_t^{X,i}(u)\, dW_u$ is linear in W with deterministic kernels, the FOC reduces to computing the tilted conditional mean $\widehat{W}_t^{i,\theta} := \mathbb{E}^{Q^i}[W_u \mid \mathcal{F}_t^i]$. Under the deterministic-kernel profile, $\mathcal{L}(W \mid \mathcal{F}_t^i)$ is Gaussian with mean $m_t := \widehat{W}_t^i$ and deterministic covariance $C_t := \mathsf{C}_t^i$. The quadratic tilt shifts this to

$$C_{t,\theta}^{-1} = C_t^{-1} - \theta_i K_t, \qquad m_{t,\theta} = C_{t,\theta}(C_t^{-1} m_t + \theta_i k_t),$$

provided $I - \theta_i C_t^{1/2} K_t C_t^{1/2} \succ 0$. Here C_t and K_t act as integral operators on $L^2([0,T];\mathbb{R}^{(n+1)d})$. The kernel K of (1.B.28) is supported on $[0,T]^2$ and couples past and future coordinates. C_t is the conditional covariance of the full path on $[0,T]^2$, so restricting to $[0,t]$ would drop the past–future coupling that the tilt introduces. On this domain, $(I - \theta_i C_t K_t)^{-1}$ is the resolvent (Neumann series) $(I - \theta_i C_t K_t)^{-1} f = f + \theta_i C_t K_t f + \theta_i^2 (C_t K_t)^2 f + \cdots$, convergent precisely when $\|\theta_i C_t^{1/2} K_t C_t^{1/2}\|_{\mathrm{op}} < 1$. Rearranging gives

$$\widehat{W}_t^{i,\theta}(u) := \Big[(I - \theta_i C_t K_t)^{-1}\Big(\widehat{W}_t^i + \theta_i\, C_t k_t\Big)\Big](u), \tag{1.B.29}$$

and substituting into the tilted FOC gives

$$\boxed{G^{DD,i}(t) D_t^i \;+\; B_i^X(t)^\top \Big[\bar{H}_t^{X,i} \;+\; \int_0^T H_t^{X,i}(u)\, d\widehat{W}_t^{i,\theta}(u)\Big] \;=\; 0.} \tag{1.B.30}$$

The pathwise cost $\mathcal{C}^i$ is convex in the control path, $x \mapsto e^{\theta_i x}$ is convex and increasing, and $\theta_i^{-1} \log \mathbb{E}$ of a log-convex family is convex, so $J_{\theta_i}^i$ is convex in the control. Stationarity is also sufficient, and (1.B.30) characterizes the best response. The linear kernel k_t (1.B.27) shifts the conditional mean toward lower expected cost; the quadratic kernel K_t (1.B.28) reweights conditional precision toward the directions of W along which cost variance is large. As $\theta_i \to 0$, $\widehat{W}_t^{i,\theta} \to \widehat{W}_t^i$ and (1.B.30) recovers (1.B.13).

Chapter 2

Delayed Public Signals and Asynchronous Learning

Public information often arrives after private learning has already occurred. Earnings releases, official statistics, and market-wide disclosures are public signals, but different agents may receive them with different delays after acting on partial private information. Delay creates a disclosure window during which a private action can move what others infer; at its known end the manipulation loses its effect.

Extending the noise-state filtering and adjoint system to this timing requires centering the arriving delayed signal: at time t, player i observes $dZ_{t-\Delta_i}$ with the part it could already predict subtracted, so the delayed signal arrives as an innovation. The resulting filter kernel has two births, the private-signal birth at $u = t$ and the public-signal birth at $u = t - \Delta_i$.

Linear-quadratic stochastic differential games with delay have been studied with the delay in the controls [28]; delayed observation of shared information is the delayed-sharing pattern of decentralized control [88, 111, 117]. Here the delay is in the observation of a public signal, so it enters the filtering and adjoint kernels rather than the state dynamics. The same equations cover the opposite timing, an action that is seen at once but takes effect after a lag τ_i, as with time to build [70] or a purchase that is resold later. A delayed action changes the state dynamics and the first-order condition; the filter is untouched.

Section 2.1 sets up the delayed observation, Sections 2.2 and 2.3 derive the filter and verify it, and Section 2.4 prices the arrival in the adjoints and the first-order condition.

2.1 Delayed observation structure

Fix a finite player set $\mathcal{I} = \{1, \ldots, n\}$ and choose an arbitrary player i. Throughout this chapter, i denotes the player whose filter or control is described, while $k \in \mathcal{I} \setminus \{i\}$ denotes an opponent or another observer.

Player i observes a private signal

$$dY_t^i = \left(\bar{C}^{i,Y}(t) + \int_0^t C^{i,Y}(t,s)\, dW_s\right) dt + E^i\, dW_t. \tag{2.1.1}$$

There is also a public signal

$$dZ_t = \left(\bar{C}^Z(t) + \int_0^t C^Z(t,s)\, dW_s\right) dt + E^Z\, dW_t. \tag{2.1.2}$$

Each player may observe the public signal with a different delay. Player i's delay is $\Delta_i \geq 0$, and

$$\mathcal{F}_t^i = \sigma(Y_r^i : r \leq t) \vee \sigma(Z_r : r \leq t - \Delta_i). \tag{2.1.3}$$

The common-delay case is $\Delta_1 = \cdots = \Delta_n = \Delta$. In applications the public drift kernel C^Z may contain state, predictable-control, and uncontrolled components,

$$C^Z(t,s) = B_X^Z(t) X_t(s) + \sum_m B_m^Z(t) D_{W,t}^m(s) + U_t^Z(s), \tag{2.1.4}$$

where control-dependent drifts use predictable controls D_{t-}^m.

The state may also carry delayed control rows,

$$dX_t = \Big(B_X^X(t) X_t + \sum_m B_m^X(t) D_t^m + \sum_m \tilde{B}_m^X(t) D_{t-\tau_m}^m \Big) dt + \sigma(t)\, dW_t, \quad D_s^m := 0 \ (s<0), \quad \tau_m \ge 0, \tag{2.1.5}$$

with $\tilde{B}_m^X(s) := 0$ for $s > T$. A control chosen at $t - \tau_m$ depends only on noise born by then, so in the kernel state equation the delayed row is $\tilde{B}_m^X(t) \mathbf{1}_{\{s \le t-\tau_m\}} D_{W,t-\tau_m}^m(s)$, and the mean equation gains $\tilde{B}_m^X(t) \bar{D}_{t-\tau_m}^m$. The filtering results below do not change. $D_{t-\tau_m}^m$ adds no martingale term to dX; the delayed rows therefore enter the drift kernels $C^{i,Y}$ and C^Z only through $X_t(s)$. With $\tilde{B}_m^X \equiv 0$ everything reduces to Chapter 1.

Assumption 2.1 (orthogonal delayed public news)**.** The private and public martingale rows are normalized and orthogonal:

$$E^i E^{i,\top} = I, \qquad E^Z E^{Z,\top} = I, \qquad E^i E^{Z,\top} = 0. \tag{2.1.6}$$

Private rows do not directly measure $E^Z dW$. Players may infer part of $E^Z dW_{t-\Delta_i}$ indirectly through earlier drift effects, but this inference is predictable at $t-$, so it changes only the drift of the arriving Z-innovation, not its instantaneous covariance.

Let $s_t^i := t - \Delta_i$. At time t, the arriving public increment for player i is $dZ_{s_t^i}$ when $t \ge \Delta_i$. Define the pre-arrival projection of a primitive kernel K by

$$(\mathcal{P}_{t-}^i K)(u) := K(u)\Pi^i + \mathbf{1}_{\{u < s_t^i\}} K(u) \Pi^Z + \int_0^t K(r) F_{t-}^i(r,u)\, dr, \tag{2.1.7}$$

where $\Pi^i = E^{i,\top} E^i$ and $\Pi^Z = E^{Z,\top} E^Z$. The strict inequality excludes the just-arriving coordinate $u = s_t^i$. The complement $(I - \mathcal{P}_{t-}^i)K$ is the part of K that player i has not resolved before the arrival.

The private unresolved row is

$$\widetilde{C}_t^{i,Y}(u) := (I - \mathcal{P}_{t-}^i)\Big[C^{i,Y}(t,\cdot) \Big](u). \tag{2.1.8}$$

Define the centered delayed-public innovation row

$$\boxed{\widetilde{C}_t^{i,Z}(u) := \mathbf{1}_{\{t\geq\Delta_i\}}\left((I-\mathcal{P}_{t-}^i)\Big[C^Z(s_t^i,\cdot)\Big](u) - E^Z F_{t-}^i(s_t^i,u)\right).} \tag{2.1.9}$$

Since $E^Z F_{t-}^i(t-\Delta_i,\cdot)$ need not vanish, player i also subtracts its predictable estimate of the arriving Z-noise increment. The delayed public innovation is

$$\begin{aligned} dI_t^{i,Z} &:= \mathbf{1}_{\{t\geq\Delta_i\}}\left(dZ_{s_t^i} - \mathbb{E}[dZ_{s_t^i} \mid \mathcal{F}_{t-}^i]\right) \\ &= \mathbf{1}_{\{t\geq\Delta_i\}} E^Z dW_{s_t^i} + \int_0^t \widetilde{C}_t^{i,Z}(u)\, dW_u\, dt. \end{aligned} \tag{2.1.10}$$

Similarly,

$$dI_t^{i,Y} := dY_t^i - \mathbb{E}[dY_t^i \mid \mathcal{F}_{t-}^i] = E^i dW_t + \int_0^t \widetilde{C}_t^{i,Y}(u)\, dW_u\, dt. \tag{2.1.11}$$

By (2.1.6), $dI_t^{i,Y}$ and $dI_t^{i,Z}$ have block identity instantaneous covariance and no mixed quadratic variation. Figure 2.1 shows the geometry.

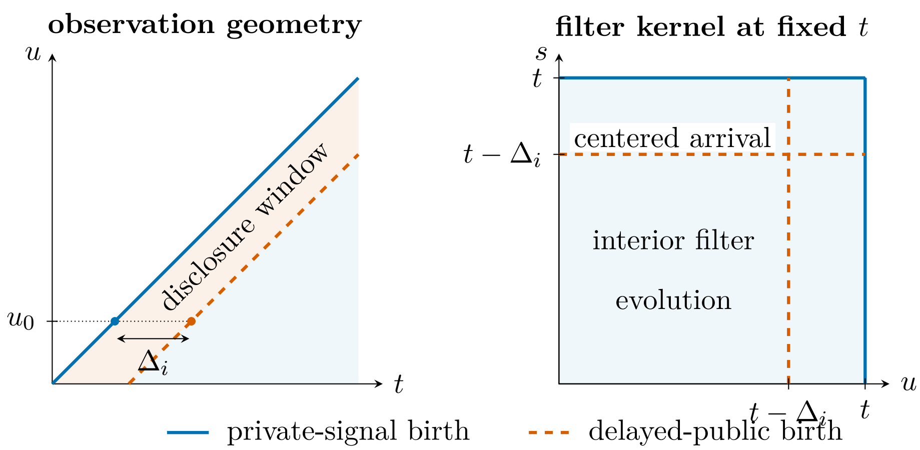

Figure 2.1: Observation and filter geometry under delay. Left: a source coordinate enters the private innovation at $t = u$ and reaches player i publicly at $t = u + \Delta_i$; the strip between is the disclosure window. Right: the delayed public signal creates interior birth lines at $t - \Delta_i$.

2.2 Filtering with delayed public signals

The theorem verifies a candidate kernel, taking the representation as given and deriving the equations it must satisfy. Existence of a kernel with the representation is part of the equilibrium fixed point and is not claimed here.

Theorem 2.2 (Filtering with a delayed public innovation row)**.** *For each $i \in \mathcal{I}$, let Δ_i be a known constant delay, maintain Assumption 2.1, and suppose a deterministic kernel F^i is such that the candidate noise-state has the representation*

$$d_u \widehat{W}_t^i(u) = \Pi^i \, dW_u + \mathbf{1}_{\{u<t-\Delta_i\}} \Pi^Z \, dW_u + \left(\int_0^t F_t^i(u,s) \, dW_s \right) du. \tag{2.2.1}$$

Define $\widetilde{C}^{i,Y}$ and $\widetilde{C}^{i,Z}$ by (2.1.8) and (2.1.9). Away from the two birth lines $u = t$ and $u = t - \Delta_i$, F^i satisfies

$$\partial_t F_t^i(u,s) = \widetilde{C}_t^{i,Y}(u)^\top \widetilde{C}_t^{i,Y}(s) + \widetilde{C}_t^{i,Z}(u)^\top \widetilde{C}_t^{i,Z}(s). \tag{2.2.2}$$

The private signal creates the usual birth at $u = t$,

$$F_t^i(u,t) = \widetilde{C}_t^{i,Y}(u)^\top E^i, \qquad F_t^i(t,u) = E^{i,\top} \widetilde{C}_t^{i,Y}(u), \tag{2.2.3}$$

and the delayed public signal creates a second birth at $s_t^i = t - \Delta_i$, the coordinate whose public news arrives at t,

$$F_t^i(u, s_t^i)\Big|_Z = \widetilde{C}_t^{i,Z}(u)^\top E^Z, \qquad F_t^i(s_t^i, u)\Big|_Z = E^{Z,\top} \widetilde{C}_t^{i,Z}(u). \tag{2.2.4}$$

Equivalently, with noncausal rows and empty integrals equal to zero,

$$\begin{aligned} F_t^i(u,s) = {} & \widetilde{C}_s^{i,Y}(u)^\top E^i + E^{i,\top} \widetilde{C}_u^{i,Y}(s) \\ & + \mathbf{1}_{\{s+\Delta_i \le t\}} \widetilde{C}_{s+\Delta_i}^{i,Z}(u)^\top E^Z + \mathbf{1}_{\{u+\Delta_i \le t\}} E^{Z,\top} \widetilde{C}_{u+\Delta_i}^{i,Z}(s) \end{aligned}$$

$$+\int_{\max\{u,s\}}^{t} \widetilde{C}_r^{i,Y}(u)^\top \widetilde{C}_r^{i,Y}(s)\,dr + \int_{\max\{u,s\}}^{t} \widetilde{C}_r^{i,Z}(u)^\top \widetilde{C}_r^{i,Z}(s)\,dr. \tag{2.2.5}$$

Formula (2.2.5) verifies that (2.2.1) is the conditional expectation $\mathbb{E}[W_u \mid \mathcal{F}_t^i]$.

2.3 Filtering derivation

The argument follows the proof of the baseline filtering theorem, Theorem 1.5. The only new ingredient is that the test martingale has two orthogonal measurement rows and the delayed public news is the centered innovation (2.1.10).

Exponential tests

Fix deterministic test functions r_Y, r_Z and define

$$d\phi_t = \mathbf{i}\phi_t\left[(dI_t^{i,Y})^\top r_Y(t) + (dI_t^{i,Z})^\top r_Z(t)\right], \qquad \phi_0 = 1. \tag{2.3.1}$$

The candidate $\widehat{W}_t^i(u)$ is the conditional expectation of W_u if and only if

$$\mathbb{E}[\widehat{W}_t^i(u)\phi_t] = \mathbb{E}[W_u\phi_t] \quad \text{for all } r_Y, r_Z. \tag{2.3.2}$$

Left-hand side

Write the candidate innovation-form dynamics for a fixed coordinate u as

$$d_t\widehat{W}_t^i(u) = B_Y^{W,i}(u,t)dI_t^{i,Y} + B_Z^{W,i}(u,t)dI_t^{i,Z}. \tag{2.3.3}$$

Since the two innovation rows have no mixed quadratic variation, Itô's product rule gives

$$\begin{aligned} d(\widehat{W}_t^i(u)\phi_t) &= \mathbf{i}\widehat{W}_t^i(u)\phi_t(dI_t^{i,Y})^\top r_Y(t) + \mathbf{i}\widehat{W}_t^i(u)\phi_t(dI_t^{i,Z})^\top r_Z(t) \\ &\quad + \phi_t B_Y^{W,i}(u,t)dI_t^{i,Y} + \phi_t B_Z^{W,i}(u,t)dI_t^{i,Z} \end{aligned}$$

$$+ \mathbf{i}\phi_t B_Y^{W,i}(u,t) r_Y(t)dt + \mathbf{i}\phi_t B_Z^{W,i}(u,t) r_Z(t)dt. \tag{2.3.4}$$

Using (2.1.11)–(2.1.10) to collect the dt terms,

$$\begin{aligned}\frac{d}{dt}\mathbb{E}[\widehat{W}_t^i(u)\phi_t] &= \mathbf{i}\mathbb{E}\Big[\phi_t\left(\widehat{W}_t^i(u)\int_0^t \widetilde{C}_t^{i,Y}(r)dW_r^\top + B_Y^{W,i}(u,t)\right)\Big] r_Y(t) \\ &\quad + \mathbf{i}\mathbb{E}\Big[\phi_t\left(\widehat{W}_t^i(u)\int_0^t \widetilde{C}_t^{i,Z}(r)dW_r^\top + B_Z^{W,i}(u,t)\right)\Big] r_Z(t). \end{aligned}\tag{2.3.5}$$

No additional predictable residual appears because $I^{i,Y}$ and $I^{i,Z}$ are already centered relative to $\mathcal{F}_{t-}^i$.

Right-hand side

For the private row,

$$\begin{aligned}\frac{1}{dt}\mathbb{E}[W_u(dI_t^{i,Y})^\top \mid \mathcal{F}_{t-}^i] &= \mathbb{E}\left[W_u\int_0^t \widetilde{C}_t^{i,Y}(r)dW_r^\top \mid \mathcal{F}_{t-}^i\right] + E^{i,\top}\mathbf{1}_{\{u\geq t\}} \\ &= \widehat{W}_{t-}^i(u)\int_0^t \widetilde{C}_t^{i,Y}(r)d\widehat{W}_{t-}^i(r)^\top + \int_0^t \partial_r C_{t-}^i(u,r)\widetilde{C}_t^{i,Y}(r)^\top dr \\ &\quad + E^{i,\top}\mathbf{1}_{\{u\geq t\}}, \end{aligned}\tag{2.3.6}$$

where $\mathsf{C}_t^i(u,r) = \mathrm{Cov}(W_u, W_r \mid \mathcal{F}_t^i)$. For the delayed public signal, the same calculation applies to the centered innovation (2.1.10). The martingale part arrives at s_t^i and the drift row is $\widetilde{C}^{i,Z}$, so

$$\begin{aligned}\frac{1}{dt}\mathbb{E}[W_u(dI_t^{i,Z})^\top \mid \mathcal{F}_{t-}^i] &= \widehat{W}_{t-}^i(u)\int_0^t \widetilde{C}_t^{i,Z}(r)d\widehat{W}_{t-}^i(r)^\top + \int_0^t \partial_r C_{t-}^i(u,r)\widetilde{C}_t^{i,Z}(r)^\top dr \\ &\quad + E^{Z,\top}\mathbf{1}_{\{t\geq\Delta_i\}}\mathbf{1}_{\{u\geq s_t^i\}}. \end{aligned}\tag{2.3.7}$$

Combining,

$$\frac{d}{dt}\mathbb{E}[W_u\phi_t]=\mathbf{i}\mathbb{E}\Big[\phi_t\Big(\widehat{W}_t^i(u)\int_0^t\widetilde{C}_t^{i,Y}(r)d\widehat{W}_t^i(r)^\top+\int_0^t\partial_r\mathsf{C}_t^i(u,r)\widetilde{C}_t^{i,Y}(r)^\top dr+E^{i,\top}\mathbf{1}_{\{u\geq t\}}\Big)\Big]r_Y(t)$$
$$+\mathbf{i}\mathbb{E}\Big[\phi_t\Big(\widehat{W}_t^i(u)\int_0^t\widetilde{C}_t^{i,Z}(r)d\widehat{W}_t^i(r)^\top+\int_0^t\partial_r\mathsf{C}_t^i(u,r)\widetilde{C}_t^{i,Z}(r)^\top dr+E^{Z,\top}\mathbf{1}_{\{t\geq\Delta_i\}}\mathbf{1}_{\{u\geq s_t^i\}}\Big)\Big]r_Z(t).$$
$$(2.3.8)$$

Coefficient matching

Matching (2.3.5) and (2.3.8) for arbitrary r_Y, r_Z gives

$$B_Y^{W,i}(u,t)=\int_0^t\partial_r\mathsf{C}_t^i(u,r)\widetilde{C}_t^{i,Y}(r)^\top dr+E^{i,\top}\mathbf{1}_{\{u\geq t\}}, \tag{2.3.9}$$
$$B_Z^{W,i}(u,t)=\int_0^t\partial_r\mathsf{C}_t^i(u,r)\widetilde{C}_t^{i,Z}(r)^\top dr+E^{Z,\top}\mathbf{1}_{\{t\geq\Delta_i\}}\mathbf{1}_{\{u\geq s_t^i\}}. \tag{2.3.10}$$

Differentiating these expressions in the source time u gives, away from jumps,

$$d_t(d_u\widehat{W}_t^i(u))=\widetilde{C}_t^{i,Y}(u)^\top dI_t^{i,Y}\,du+\widetilde{C}_t^{i,Z}(u)^\top dI_t^{i,Z}\,du. \tag{2.3.11}$$

Equation (2.3.9) jumps at $u=t$, giving the private birth, and (2.3.10) jumps at $u=s_t^i$, giving the delayed public birth.

Primitive-shock expansion and the two birth terms

Integrating (2.3.11) over observation times gives

$$d_u\widehat{W}_t^i(u)=E^{i,\top}dI_u^{i,Y}+\mathbf{1}_{\{u+\Delta_i\leq t\}}E^{Z,\top}dI_{u+\Delta_i}^{i,Z}$$
$$+\Big(\int_u^t\widetilde{C}_r^{i,Y}(u)^\top dI_r^{i,Y}\Big)du+\Big(\int_u^t\widetilde{C}_r^{i,Z}(u)^\top dI_r^{i,Z}\Big)du. \tag{2.3.12}$$

Substitute (2.1.11) and (2.1.10). The direct private term is

$$E^{i,\top}dI_u^{i,Y}=\Pi^i dW_u+E^{i,\top}\int_0^u\widetilde{C}_u^{i,Y}(s)dW_s\,du. \tag{2.3.13}$$

The direct delayed public term is

$$\mathbf{1}_{\{u+\Delta_i \le t\}} E^{Z,\top} dI^{i,Z}_{u+\Delta_i} = \mathbf{1}_{\{u+\Delta_i \le t\}} \Pi^Z dW_u + \mathbf{1}_{\{u+\Delta_i \le t\}} E^{Z,\top} \int_0^{u+\Delta_i} \widetilde{C}^{i,Z}_{u+\Delta_i}(s) dW_s \, du. \tag{2.3.14}$$

The accumulated martingale updates contribute

$$\int_u^t \widetilde{C}^{i,Y}_r(u)^\top E^i dW_r \, du, \qquad \int_{\max\{0,u-\Delta_i\}}^{t-\Delta_i} \widetilde{C}^{i,Z}_{s+\Delta_i}(u)^\top E^Z dW_s \, du. \tag{2.3.15}$$

The accumulated drift updates contribute, by stochastic Fubini,

$$\begin{aligned} &\int_0^t \left(\int_{\max\{u,s\}}^t \widetilde{C}^{i,Y}_r(u)^\top \widetilde{C}^{i,Y}_r(s) dr \right) dW_s \, du, \\ &\int_0^t \left(\int_{\max\{u,s\}}^t \widetilde{C}^{i,Z}_r(u)^\top \widetilde{C}^{i,Z}_r(s) dr \right) dW_s \, du. \end{aligned} \tag{2.3.16}$$

Collecting the coefficient of $dW_s \, du$ gives (2.2.5). The singular direct terms give the two direct pieces in (2.2.1). The candidate representation satisfies the test identities (2.3.2); uniqueness follows from Itô isometry as in the proof of Theorem 1.5.

2.4 First variations and delayed-public adjoints

A deviation misleads private learners today, but the manipulated component resurfaces when the public news arrives at delay Δ_k. Player i pays for moving the drift of the delayed public signal when that signal reaches opponent k, and its belief price jumps at that date; together these are the marginal cost of that second effect. Each opponent now carries two observation rows, the private row and the centered delayed public innovation.

Fix the deviating player i and freeze opponents' strategy maps and filters. For $k \neq i$, write $\delta w^k_s(u)$ for the density variation of opponent k's noise-state and δX_s for the physical-state variation. The induced opponent-control variation is

$$\delta D_s^k = \int_0^s D_s^k(u)\, \delta w_s^k(u)\, du. \tag{2.4.1}$$

Opponent k observes at time s the public signal generated at source time $q_k = s_s^k = s - \Delta_k$. Write $\delta b_{q_k}^Z$ for the variation of the predictable drift of that public signal. In the affine example $b_{q_k}^Z = B_X^Z(q_k)X_{q_k} + \sum_m B_m^Z(q_k)D_{q_k-}^m + U_{q_k}^Z$, this means

$$\delta b_{q_k}^Z = B_X^Z(q_k)\delta X_{q_k} + \sum_{m \neq i} B_m^Z(q_k) \int_0^{q_k} D_{q_k}^m(r)\delta w_{q_k}^m(r)\, dr, \tag{2.4.2}$$

with the right-hand side zero if q_k lies before the initial perturbation date. The deviating player's own instantaneous effect through $B_i^Z(q_k)D_{q_k}^i$ stays out of the forward variation system and enters the first-order condition below directly.

Forward first-variation system

Opponent k's private residual is

$$\eta_s^{k,Y} := B_X^{Y,k}(s)\left(\delta X_s - \int_0^s X_s(r)\delta w_s^k(r)\, dr\right). \tag{2.4.3}$$

The subtraction leaves the part of the state variation opponent k has not yet inferred. The centered delayed-public residual is

$$\eta_s^{k,Z} := \mathbf{1}_{\{s \geq \Delta_k\}}\left(\delta b_{q_k}^Z - \int_0^s C^Z(q_k, r)\delta w_{s-}^k(r)\, dr - E^Z \delta w_{s-}^k(q_k)\right), \qquad q_k = s - \Delta_k. \tag{2.4.4}$$

The last term is the first-variation counterpart of the centering correction $-E^Z F_{s-}^k(q_k, \cdot)$ in the delayed public innovation.

Existing noise-state coordinates satisfy

$$\partial_s \delta w_s^k(u) = \widetilde{C}_s^{k,Y}(u)^\top \eta_s^{k,Y} + \widetilde{C}_s^{k,Z}(u)^\top \eta_s^{k,Z}, \qquad u < s, \tag{2.4.5}$$

with two birth conditions, one at the current date and one at the arrival date,

$$\delta w_s^k(s)\Big|_Y = E^{k,\top}\eta_s^{k,Y}, \qquad \delta w_s^k(s-\Delta_k)\Big|_Z = \mathbf{1}_{\{s\geq\Delta_k\}} E^{Z,\top}\eta_s^{k,Z}. \tag{2.4.6}$$

The physical variation satisfies

$$\frac{d}{ds}\delta X_s = B_X^X(s)\delta X_s + \sum_{k\neq i} B_k^X(s)\int_0^s D_s^k(u)\delta w_s^k(u)\,du + \sum_{k\neq i}\tilde{B}_k^X(s)\int_0^{s-\tau_k} D_{s-\tau_k}^k(u)\delta w_{s-\tau_k}^k(u)\,du, \tag{2.4.7}$$

where $\delta w_r^k \equiv 0$ before the perturbation date. The last sum in (2.4.7) is what opponent k decided τ_k ago, on the beliefs it held then. The deviating player's spike now produces two impulses, $\delta X_t = v := B_i^X(t)z$ when it acts and a jump $\tilde{v} := \tilde{B}_i^X(t+\tau_i)z$ at the effect date. Equations (2.4.3)–(2.4.7) are the delayed-public analogue of the baseline forward variation system.

Forward transfer functions

For $s \geq t$, define the forward transfer functions by the response to initial data $(\delta X_t, \delta w_t^m(\cdot))$, $m \neq i$:

$$\delta X_s = \Phi^{XX}(s,t)v + \sum_{m\neq i}\int_0^t \Phi^{Xm}(s,t,u)\delta w_t^m(u)\,du, \tag{2.4.8}$$

$$\delta w_s^k(r) = \Phi^{kX}(s,t,r)v + \sum_{m\neq i}\int_0^t \Phi^{km}(s,t,r,u)\delta w_t^m(u)\,du. \tag{2.4.9}$$

The initial conditions are

$$\Phi^{XX}(t,t) = I, \quad \Phi^{kX}(t,t,r) = 0, \quad \Phi^{Xm}(t,t,u) = 0, \quad \Phi^{km}(t,t,r,u) = \mathbf{1}_{\{k=m\}}\delta(r-u), \tag{2.4.10}$$

where the last identity holds weakly.

Define the residual transfer coefficients obtained by substituting (2.4.8)–(2.4.9) into (2.4.3)–(2.4.4). For a physical initial perturbation v, set

$$\mathcal{Y}^{kX}(s,t) := B_X^{Y,k}(s)\left(\Phi^{XX}(s,t) - \int_0^s X_s(r)\Phi^{kX}(s,t,r)\,dr\right), \tag{2.4.11}$$

and

$$\begin{aligned} \mathcal{Z}^{kX}(s,t) := \mathbf{1}_{\{s\geq\Delta_k\}}\Bigg[\delta b^{Z,X}_{q_k;t} &- \int_0^s C^Z(q_k,r)\Phi^{kX}(s-,t,r)\,dr \\ &- E^Z\Phi^{kX}(s-,t,q_k)\Bigg], \qquad q_k = s-\Delta_k, \end{aligned} \tag{2.4.12}$$

where

$$\begin{aligned} \delta b^{Z,X}_{q_k;t} &:= B_X^Z(q_k)\Phi^{XX}(q_k,t) \\ &\quad + \sum_{\ell\neq i} B_\ell^Z(q_k)\int_0^{q_k} D^\ell_{q_k}(r)\Phi^{\ell X}(q_k,t,r)\,dr, \end{aligned} \tag{2.4.13}$$

with all transfer functions at $q_k < t$ read as zero in the first two terms. Similarly, for an initial belief perturbation in block m,

$$\mathcal{Y}^{km}(s,t,u) := B_X^{Y,k}(s)\left(\Phi^{Xm}(s,t,u) - \int_0^s X_s(r)\Phi^{km}(s,t,r,u)\,dr\right), \tag{2.4.14}$$

and

$$\begin{aligned} \mathcal{Z}^{km}(s,t,u) := \mathbf{1}_{\{s\geq\Delta_k\}}\Bigg[\delta b^{Z,km}_{q_k;t}(u) &- \int_0^s C^Z(q_k,r)\Phi^{km}(s-,t,r,u)\,dr \\ &- E^Z\Phi^{km}(s-,t,q_k,u)\Bigg]. \end{aligned} \tag{2.4.15}$$

where

$$\delta b^{Z,km}_{q_k;t}(u) := B^Z_X(q_k)\Phi^{Xm}(q_k,t,u) + \sum_{\ell\neq i} B^Z_\ell(q_k)\int_0^{q_k} D^\ell_{q_k}(r)\Phi^{\ell m}(q_k,t,r,u)\,dr. \tag{2.4.16}$$

The transfer functions solve the forward equations

$$\begin{aligned}
\partial_s\Phi^{XX}(s,t) &= B^X_X(s)\Phi^{XX}(s,t) + \sum_{k\neq i} B^X_k(s)\int_0^s D^k_s(r)\Phi^{kX}(s,t,r)\,dr \\
&\quad + \sum_{k\neq i}\tilde{B}^X_k(s)\int_0^{s-\tau_k} D^k_{s-\tau_k}(r)\Phi^{kX}(s-\tau_k,t,r)\,dr, \qquad (2.4.17)\\
\partial_s\Phi^{Xm}(s,t,u) &= B^X_X(s)\Phi^{Xm}(s,t,u) + \sum_{k\neq i} B^X_k(s)\int_0^s D^k_s(r)\Phi^{km}(s,t,r,u)\,dr \\
&\quad + \sum_{k\neq i}\tilde{B}^X_k(s)\int_0^{s-\tau_k} D^k_{s-\tau_k}(r)\Phi^{km}(s-\tau_k,t,r,u)\,dr, \qquad (2.4.18)\\
\partial_s\Phi^{kX}(s,t,r) &= \widetilde{C}^{k,Y}_s(r)^\top\mathcal{Y}^{kX}(s,t) + \widetilde{C}^{k,Z}_s(r)^\top\mathcal{Z}^{kX}(s,t), \qquad (2.4.19)\\
\partial_s\Phi^{km}(s,t,r,u) &= \widetilde{C}^{k,Y}_s(r)^\top\mathcal{Y}^{km}(s,t,u) + \widetilde{C}^{k,Z}_s(r)^\top\mathcal{Z}^{km}(s,t,u), \qquad (2.4.20)
\end{aligned}$$

for old-history coordinates, those already born and away from the two birth lines, together with the two boundary updates

$$\begin{aligned}
\Phi^{kX}(s,t,s)\Big|_Y &= E^{k,\top}\mathcal{Y}^{kX}(s,t), \\
\Phi^{km}(s,t,s,u)\Big|_Y &= E^{k,\top}\mathcal{Y}^{km}(s,t,u), \qquad (2.4.21)\\
\Phi^{kX}(s,t,s-\Delta_k)\Big|_Z &= \mathbf{1}_{\{s\geq\Delta_k\}}E^{Z,\top}\mathcal{Z}^{kX}(s,t), \\
\Phi^{km}(s,t,s-\Delta_k,u)\Big|_Z &= \mathbf{1}_{\{s\geq\Delta_k\}}E^{Z,\top}\mathcal{Z}^{km}(s,t,u). \qquad (2.4.22)
\end{aligned}$$

Transfer functions with a time argument below t are zero, so the delayed rows are inactive until $s = t+\tau_k$. The deviating player's delayed impulse stays outside the forward variation system; by linearity it contributes $\Phi^{XX}(s,t+\tau_i)\tilde{v}$.

2.4.1 State price, belief prices, and cost variation

Let

$$L_s^i := G^{XX,i}(s)X_s + G_s^{X,i} + G^{XD,i}(s)D_s^i, \qquad L_T^i := G^{XX,i}(T)X_T + G_T^{X,i}. \tag{2.4.23}$$

L_s^i and L_T^i are the running and terminal marginal costs of the physical state. Pairing the forward state response with future marginal costs gives the state price and the belief prices:

$$H_t^{X,i} := \int_t^T \Phi^{XX}(s,t)^\top L_s^i\, ds + \Phi^{XX}(T,t)^\top L_T^i, \tag{2.4.24}$$

$$H_t^{k,i}(u) := \int_t^T \Phi^{Xk}(s,t,u)^\top L_s^i\, ds + \Phi^{Xk}(T,t,u)^\top L_T^i, \qquad k \neq i. \tag{2.4.25}$$

Here Φ^{Xk} is the state response to an initial perturbation of opponent k's noise-state coordinate.

For a control spike $z \in L^2(\Omega, \mathcal{F}_t^i; \mathbb{R}^d)$, the normalized first variation of player i's cost is

$$\begin{aligned}\delta J_t^i(z) = &\mathbb{E}\Big[z^\top \Big(G^{DD,i}(t)D_t^i + G^{DX,i}(t)X_t + B_i^X(t)^\top H_t^{X,i} + \mathbf{1}_{\{t+\tau_i \le T\}} \tilde{B}_i^X(t+\tau_i)^\top H_{t+\tau_i}^{X,i}\Big) \,|\, \mathcal{F}_t^i\Big] \\ &+ \mathbb{E}\left[z^\top B_i^Z(t)^\top \sum_{k\neq i} \mathbf{1}_{\{t+\Delta_k \le T\}} \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t+\Delta_k \leftarrow t} \,|\, \mathcal{F}_t^i\right]\end{aligned} \tag{2.4.26}$$

when the deviating player's predictable action enters the public-signal drift through $B_i^Z(t)D_{t-}^i$. If $B_i^Z \equiv 0$, the second line drops out. The new term in the first line is the state price at the effect date.

Backward equations

The two birth conditions of the forward system, at the private date and at the delayed public date, enter the backward system through two forms of the filter contraction $\mathfrak{B}$

of Chapter 1, equation (1.4.9), which contracts a belief price against the observer's own filter. The filter contraction of the belief price is

$$\mathfrak{B}^k H_t^{k,i} = E^k H_t^{k,i}(t) + \int_0^t \widetilde{C}_t^{k,Y}(r) H_t^{k,i}(r)\, dr. \tag{2.4.27}$$

This is the price, for player i, of moving opponent k's private signal-source drift at time t. The delayed-public analogue, generated at source time s and observed at $s+\Delta_k$, defines the delayed contraction $\mathfrak{B}^{k,Z}$, built from the centered public innovation $(E^Z, \widetilde{C}^{k,Z})$; player i pays it at the source time s:

$$\left(\mathfrak{B}^{k,Z} H^{k,i}\right)_{s+\Delta_k \leftarrow s} := \mathbf{1}_{\{s+\Delta_k \le T\}} \left(E^Z H_{s+\Delta_k}^{k,i}(s) + \int_0^{s+\Delta_k} \widetilde{C}_{s+\Delta_k}^{k,Z}(r) H_{s+\Delta_k}^{k,i}(r)\, dr \right). \tag{2.4.28}$$

This is the price of moving the drift of the delayed public signal row at its source time s. The first term in each display is the transpose of the birth condition (2.4.6); the integral term is the transpose of the drift update (2.4.5).

Differentiating the transfer representations (2.4.24)–(2.4.25), using the forward equations (2.4.17)–(2.4.22), and collecting the coefficients of δX and δw^k gives the equations below. With delayed rows the forward system is a linear delay system, whose adjoint carries advanced arguments; the coefficient of $\delta w_t^k(u)$ then collects the state price at t and at $t+\tau_k$.

$$\frac{d}{dt} H_t^{X,i} = -L_t^i - B_X^X(t)^\top H_t^{X,i} - \sum_{k \ne i} \mathcal{V}_t^{k,i}, \tag{2.4.29}$$

where

$$\mathcal{V}_t^{k,i} = B_X^{Y,k}(t)^\top \mathfrak{B}^k H_t^{k,i} + \mathbf{1}_{\{t+\Delta_k \le T\}} B_X^Z(t)^\top \left(\mathfrak{B}^{k,Z} H^{k,i}\right)_{t+\Delta_k \leftarrow t}. \tag{2.4.30}$$

The belief price satisfies, for old-history coordinates other than the one whose public news arrives at t,

$$
\begin{aligned}
\frac{d}{dt} H_t^{k,i}(u) = & -\Big(B_k^X(t) D_t^k(u)\Big)^\top H_t^{X,i} - \mathbf{1}_{\{t+\tau_k \le T\}} \Big(\tilde{B}_k^X(t+\tau_k) D_t^k(u)\Big)^\top H_{t+\tau_k}^{X,i} \\
& + X_t(u)^\top B_X^{Y,k}(t)^\top \mathfrak{B}^k H_t^{k,i} \\
& + \mathbf{1}_{\{t \ge \Delta_k\}} C^Z(t-\Delta_k, u)^\top \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t \leftarrow t-\Delta_k} \\
& - \sum_{k' \neq i} \mathbf{1}_{\{t+\Delta_{k'} \le T\}} \Big(B_k^Z(t) D_t^k(u)\Big)^\top \Big(\mathfrak{B}^{k',Z} H^{k',i}\Big)_{t+\Delta_{k'} \leftarrow t}.
\end{aligned}
\tag{2.4.31}
$$

When the news about the coordinate $u = t$ arrives, at time $t + \Delta_k$, the centered delayed-public innovation adds a jump, dual to the centering term $-E^Z \delta w_{s-}^k(q_k)$ in (2.4.4):

$$
H_{t+\Delta_k -}^{k,i}(t) = H_{t+\Delta_k +}^{k,i}(t) - E^{Z,\top} \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t+\Delta_k \leftarrow t}. \tag{2.4.32}
$$

Equations (2.4.29)–(2.4.32) reduce to the baseline adjoint system when the delayed public signal is absent and $\tilde{B}^X \equiv 0$. The second term in (2.4.31) is the delayed-action term, the state price at $t + \tau_k$, which a backward sweep has already formed. Equation (2.4.29) is unchanged because a delayed control is noise-state linear, so its variation in (2.4.7) is driven by $\delta w_{s-\tau_m}^m$ and not by δX_s.

First-order condition

Since the best-response problem is strictly convex in the deviating player's L^2 control, stationarity of (2.4.26) for every $z \in L^2(\Omega, \mathcal{F}_t^i; \mathbb{R}^d)$ gives

$$
\boxed{
\begin{aligned}
G^{DD,i}(t) D_t^i + \mathbb{E}\Big[& G^{DX,i}(t) X_t + B_i^X(t)^\top H_t^{X,i} + \mathbf{1}_{\{t+\tau_i \le T\}} \tilde{B}_i^X(t+\tau_i)^\top H_{t+\tau_i}^{X,i} \\
& + B_i^Z(t)^\top \sum_{k \neq i} \mathbf{1}_{\{t+\Delta_k \le T\}} \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t+\Delta_k \leftarrow t} \,\Big|\, \mathcal{F}_t^i \Big] = 0.
\end{aligned}
}
\tag{2.4.33}
$$

Write $H^{X,i}_{t+\tau_i} = \bar{H}^{X,i}_{t+\tau_i} + \int_0^T H^{X,i}_{t+\tau_i}(u)\,dW_u$; then $\mathbb{E}[H^{X,i}_{t+\tau_i} \mid \mathcal{F}^i_t] = \bar{H}^{X,i}_{t+\tau_i} + \int_0^t H^{X,i}_{t+\tau_i}(u)\,d\widehat{W}^i_t(u)$, and apart from the public-drift term the equilibrium kernels are

$$D^i_t(u) = -G^{DD,i}(t)^{-1}\Big[G^{DX,i}(t)X_t(u) + B^X_i(t)^\top H^{X,i}_t(u) + \mathbf{1}_{\{t+\tau_i \le T\}}\tilde{B}^X_i(t+\tau_i)^\top H^{X,i}_{t+\tau_i}(u)\Big], \quad u \le t.$$

The coordinates $u \in (t, t+\tau_i]$, noise that arrives between the action and its effect, carry no information at t and drop out.

This is the mirror image of the delayed public signal. A delayed action moves opponents' beliefs when player i takes it and the state when it lands, so the player pays the wedge at t and the state price at $t+\tau_i$. An action that opponents see only through the delayed public signal moves the state at t and their beliefs when the news lands, so the player pays the state price at t and the wedge, through $\mathfrak{B}^{k,Z}$, at $t+\Delta_k$. Each effect is priced at the date it arrives and projected on the information of the action date. The finite-deviation identity of Corollary 1.11 carries over with the delayed public row and the delayed control rows appended, since both are deterministic causal functionals of the realized control.

2.5 Conclusion

Delayed observations add a second birth line, $u = t - \Delta_i$ in (2.2.4), to the filter kernel F^i; the noise-state coordinates $\widehat{W}^i_t(u)$ are unchanged. Theorem 2.2, the main filtering result, verifies the delayed noise-state filter representation. The two birth lines reappear in the adjoint as the public birth term $E^Z H^{k,i}_{s+\Delta_k}(s)$ in (2.4.28) and the jump at the arrival date (2.4.32), and in the first-order condition (2.4.33) through $\mathfrak{B}^{k,Z}$.

Chapter 3

Stationary Infinite-Horizon LQG

In the finite-horizon chapter every kernel depends on two calendar dates, the source date of a shock and the date of its evaluation. Here the calendar drops out and only ages remain, the age of the shock and the age of the belief about it. The causal triangle of Chapter 1 becomes a lag quadrant, the quarter plane of two nonnegative ages.

First, every impulse response function becomes time invariant, and a simple recipe turns the identities of Chapter 1 into stationary ones. The stationary system is therefore easier to compute and plot. Most work building on these methods would study stationary cases. An economic system that has settled into stationary equilibrium under one policy, like the economy behind an estimated Phillips curve, begins any new regime from that stationary state. The response to the shift is a nonstationary problem, so studying a regime change means solving the stationary problem first. Second, much of the literature studies the stationary case through the frequency domain [60, 64, 101], so working in the stationary case makes comparisons easier. Third, the stationary case makes it easier to see how close the system is to Markovian. A stationary linear Markov process has impulse responses that are combinations of exponentials, so how well such combinations fit the kernels here measures how far the equilibrium is from finite-dimensional. Finally, this chapter is the setting of the Kyle–Back analysis. The stationary deviation calculus of Chapter 4 runs in the lag coordinates introduced here, and its equilibrium condition plays the role of Proposition 3.1.

3.1 Stationary profiles, lag coordinates, and forward filtering

The time-homogeneous primitives are

$$dX_t = \left(B_X^X X_t + \sum_{m=1}^{n} B_m^X D_t^m + \sum_{m=1}^{n} \tilde{B}_m^X D_{t-\tau_m}^m \right) dt + \sigma \, dW_t^0, \tag{3.1.1}$$

$$dY_t^i = \left(B_X^{Y,i} X_t + \sum_{m=1}^{n} B_m^{Y,i} D_{t-}^m \right) dt + E^i \, dW_t, \qquad i = 1, \dots, n, \tag{3.1.2}$$

$$dZ_t = \left(B_X^Z X_t + \sum_{m=1}^{n} B_m^Z D_{t-}^m \right) dt + E^Z \, dW_t. \tag{3.1.3}$$

Two delays are allowed. A control may act on the state after a lag $\tau_m \geq 0$ (the $\tilde{B}_m^X$ rows; $D_s^m := 0$ for $s < 0$). The public signal Z is observed by player i with delay $\Delta_i \geq 0$, so that $\mathcal{F}_t^i = \sigma(Y_r^i : r \leq t) \vee \sigma(Z_r : r \leq t - \Delta_i)$, with the rows normalized and orthogonal as in Chapter 2: $E^i E^{i\top} = I$, $E^Z E^{Z\top} = I$, $E^i E^{Z\top} = 0$. Private rows may also carry other players' predictable controls (observed actions, $B_m^{Y,i}$). Setting $\tilde{B}_m^X = 0$, $B_m^{Y,i} = 0$ and dropping Z recovers the baseline. Here $W = (W^0, W^1, \dots, W^n)$,

$$\Pi^i = E^{i\top} E^i, \qquad P^i = (B_X^{Y,i})^\top B_X^{Y,i},$$

Player i's discounted objective is

$$J^i = \frac{1}{2} \mathbb{E} \int_0^\infty e^{-\rho t} \Big[X_t^\top G^{XX,i} X_t + 2(G^{X,i})^\top X_t + 2(D_t^i)^\top G^{DX,i} X_t + (D_t^i)^\top G^{DD,i} D_t^i \Big] \, dt, \quad \rho > 0, \tag{3.1.4}$$

with the joint running Hessian positive semidefinite.

All lags are calendar time minus source time:

$$a = t - s \geq 0, \qquad b = t - u \geq 0, \qquad \ell = t - s \in \mathbb{R}. \tag{3.1.5}$$

Delays are fixed ages. The variables a, b are ages of past state and noise-state coordinates. The adjoint lag ℓ is two-sided; future source times have $\ell < 0$. Brownian motions are extended to two-sided time for adjoint formulas.

The front-matter Notation reference collects the finite-horizon-to-stationary translation. Every lowercase object below is the corresponding Chapter 1 object after calendar time drops out. For adjoints the lowercase kernel is the current-value form, $h^{X,i}(t-s) := e^{\rho t} H_t^{X,i}(s)$ and $h^{k,i}(t-u, t-s) := e^{\rho t} H_t^{k,i}(u, s)$, which is why ρh appears in (3.2.9) in place of the finite-horizon time derivative.

A stationary noise-state linear profile has

$$D_t^i = \bar{d}^i + \int_0^\infty d^i(a)\, d\widehat{W}_t^i(a), \qquad d\widehat{W}_t^i(a) := d\mathbb{E}[W_{t-a} \mid \mathcal{F}_t^i]. \tag{3.1.6}$$

The same weight $d^i(a)$ applies at every date to the player's belief about a shock of age a. With the delayed public row the noise-state has the representation of Chapter 2 in age form, $d\widehat{W}_t^i(a) = \Pi^i\, dW_{t-a} + \mathbf{1}_{\{a>\Delta_i\}} \Pi^Z\, dW_{t-a} + \left(\int_0^\infty f^i(a,b)\, dW_{t-b}\, db \right) da$, with $\Pi^Z = E^{Z\top} E^Z$ and $f^i(a,b)$ the kernel from the coordinate of age a to the source of age b. Shocks older than the delay are also known through their public increment. The profile is admissible if the displayed integrals are square-integrable, the induced filtering operators are bounded, and the closed-loop response kernels below decay fast enough for the discounted cost and all integration-by-parts identities to be finite. The discounted boundary terms generated by integration by parts vanish at long horizons, which is the transversality used in Proposition 3.1. A stationary equilibrium in this class is a stationary profile for which no player can reduce (3.1.4) by a unilateral stationary noise-state linear deviation.

The finite-horizon state-response equation

$$\partial_t X_t(s) = B^X_X X_t(s) + \sum_{m=1}^{n} B^X_m D^m_{W,t}(s), \qquad X_s(s) = (\sigma, 0, \dots, 0) \tag{3.1.7}$$

becomes, with $X_t(s) = x(t-s)$ and $D^m_{W,t}(s) = d^m_W(t-s)$,

$$x'(a) = B^X_X x(a) + \sum_{m=1}^{n} B^X_m d^m_W(a) + \sum_{m=1}^{n} \tilde{B}^X_m \mathbf{1}_{\{a\geq\tau_m\}} d^m_W(a-\tau_m), \quad x(0) = (\sigma,0,\dots,0), \quad x(a)\to 0. \tag{3.1.8}$$

The stationary mean satisfies

$$0 = B^X_X \bar{X} + \sum_{m=1}^{n} B^X_m \bar{d}^m + \sum_{m=1}^{n} \tilde{B}^X_m \bar{d}^m. \tag{3.1.9}$$

For filter kernels $F^i_t(u,s) = f^i(b,a)$, so $\partial_t = \partial_b + \partial_a$. For adjoint kernels $H_t(s) = h(\ell)$, so $\partial_t = \partial_\ell$; with a noise-state coordinate $b = t-u$, the transport derivative is $\partial_b + \partial_\ell$.

Each instant a new shock is born unknown and every old shock is known a little better, so each shock's uncertainty falls over its lifetime while the cross-section of uncertainty over ages stays fixed.

The control kernel in primitive-shock coordinates is

$$d^i_W(a) = d^i(a)\Big(\Pi^i + \mathbf{1}_{\{a>\Delta_i\}}\Pi^Z\Big) + \int_0^\infty d^i(b) f^i(b,a)\,db. \tag{3.1.10}$$

This carries the policy kernel d^i into primitive coordinates: how player i's control answers a shock of age a. The unresolved state response is

$$\tilde{x}^i(a) = x(a)\Big(I - \Pi^i - \mathbf{1}_{\{a>\Delta_i\}}\Pi^Z\Big) - \int_0^\infty x(b) f^i(b,a)\,db. \tag{3.1.11}$$

This is the state's response to a shock less the part player i has already resolved, privately or through the public row. The private row's drift kernel is $c^{i,Y}(a) = B^{Y,i}_X x(a) + \sum_m B^{Y,i}_m d^m_W(a)$ and its unresolved part $\tilde{c}^{i,Y}(a) = c^{i,Y}(a)\Big(I - \Pi^i - \mathbf{1}_{\{a>\Delta_i\}}\Pi^Z\Big) - \int_0^\infty c^{i,Y}(b) f^i(b,a)\,db$; without observed actions $\tilde{c}^{i,Y} = B_X^{Y,i}\tilde{x}^i$. The delayed public row generated Δ_i ago has drift kernel $c^Z(a) = B_X^Z x(a) + \sum_m B_m^Z d_W^m(a)$ at its source time; its centered innovation kernel at the arrival date is (Chapter 2, in ages)

$$\tilde{c}^{i,Z}(a) = c^Z(a-\Delta_i)\Big(I - \Pi^i - \mathbf{1}_{\{a>\Delta_i\}}\Pi^Z\Big) - \int_0^\infty c^Z(b-\Delta_i) f^i(b,a)\, db - E^Z f^i(\Delta_i^-, a), \tag{3.1.12}$$

with $c^Z(b-\Delta_i) := 0$ for $b < \Delta_i$ (the row generated Δ_i ago loads only on shocks that existed then). In the last term player i subtracts what it had already inferred about the arriving public increment. The stationary filter is

$$(\partial_b + \partial_a) f^i(b,a) = \tilde{c}^{i,Y}(b)^\top \tilde{c}^{i,Y}(a) + \tilde{c}^{i,Z}(b)^\top \tilde{c}^{i,Z}(a), \qquad b, a > 0,\ b, a \neq \Delta_i, \tag{3.1.13}$$

$$f^i(b,0) = \tilde{c}^{i,Y}(b)^\top E^i, \qquad b > 0, \tag{3.1.14}$$

$$f^i(0,a) = E^{i\top}\tilde{c}^{i,Y}(a), \qquad a > 0, \tag{3.1.15}$$

$$f^i(b,\Delta_i)\Big|_Z = \tilde{c}^{i,Z}(b)^\top E^Z, \qquad f^i(\Delta_i,a)\Big|_Z = E^{Z\top}\tilde{c}^{i,Z}(a), \tag{3.1.16}$$

the last line being the second birth, at age Δ_i, when the public news about a shock arrives, with $\cdot|_Z$ the E^Z block. With $B_m^{Y,i} = 0$ the private source is $\tilde{x}^i(b)^\top P^i \tilde{x}^i(a)$ and the births are $\tilde{x}^i(b)^\top (B_X^{Y,i})^\top E^i$, $E^{i\top} B_X^{Y,i}\tilde{x}^i(a)$. Equivalently, for $b \neq a$ and, when $\Delta_i > 0$, away from the arrival age $b = \Delta_i$ or $a = \Delta_i$ (where the public birth adds the analogous term),

$$f^i(b,a) = \tilde{x}^i(b-a)^\top (B_X^{Y,i})^\top E^i + \int_0^a \tilde{x}^i(b-c)^\top P^i \tilde{x}^i(a-c)\, dc, \qquad b > a, \tag{3.1.17}$$

$$f^i(b,a) = E^{i\top} B_X^{Y,i}\tilde{x}^i(a-b) + \int_0^b \tilde{x}^i(b-c)^\top P^i \tilde{x}^i(a-c)\, dc, \qquad a > b. \tag{3.1.18}$$

The diagonal is a separate pointwise convention,

$$f^i(a,a) := f^i_{\text{diag}}(a), \tag{3.1.19}$$

and does not affect the integral equations.

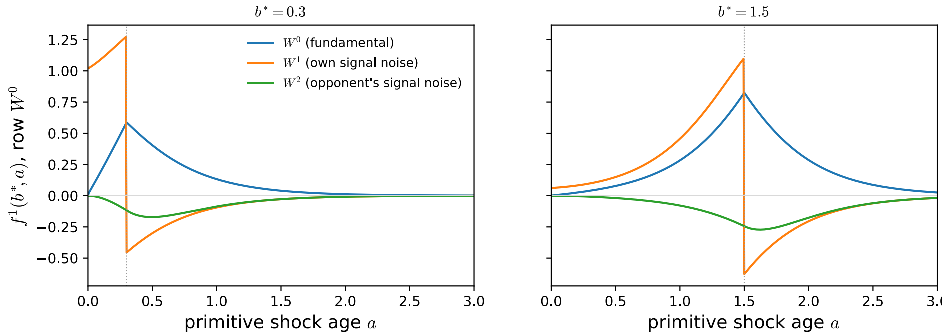

Figure 3.1: Channel decomposition of the W^0 channel of the stationary filter kernel $f^1(b^*, a)$ at two coordinate ages b^*. The loading on the opponent's signal noise is nonzero only because that noise moves the opponent's actions and through them the state.

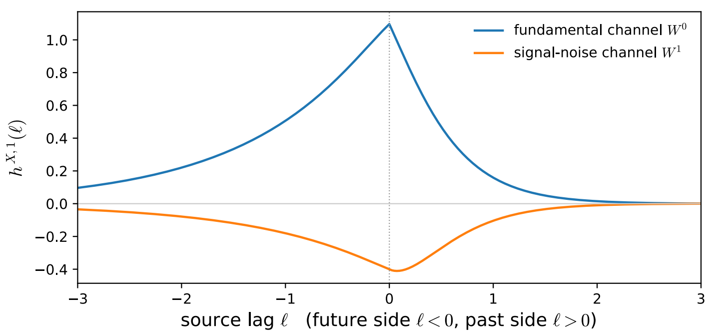

Figure 3.2: The stationary two-sided state price $h^{X,1}(\ell)$ in the symmetric two-player benchmark ($p_1 = p_2 = 3$, $r_1 = r_2 = 0.1$, average cost), from the spectral solution with 32 Chebyshev points per side.

The numerical benchmark is the stationary form of the tracking example of Chapter 1. The state is scalar with $dX_t = (D_t^1 + D_t^2)\, dt + dW_t^0$, each player observes $dY_t^i = \sqrt{p_i}\, X_t\, dt + dW_t^i$ with $p_1 = p_2 = 3$ and $r_1 = r_2 = 0.1$, and the objective is average cost. Section 3.2 also varies each precision from 1 to 20 with the other held at 3. Proposition 3.1 verifies the discounted system, so these computations solve the formal average-cost stationary equations, the $\rho = 0$ form of that system, rather than a verified optimum. Figure 3.1 shows the resulting kernel, with one curve per primitive shock.

3.2 Current-value adjoints

The unprojected state price is two-sided:

$$H_t^{X,i} = \bar{h}^{X,i} + \int_{-\infty}^{\infty} h^{X,i}(\ell)\, dW_{t-\ell}. \tag{3.2.1}$$

The player cannot act on future shocks, so the control uses the conditional expectation of this adjoint given $\mathcal{F}_t^i$.

$$\begin{aligned} D_t^i = -(G^{DD,i})^{-1}\mathbb{E}\Big[&G^{DX,i}X_t + (B_i^X)^\top H_t^{X,i} + e^{-\rho\tau_i}(\tilde{B}_i^X)^\top H_{t+\tau_i}^{X,i} \\ &+ \sum_{k\neq i}(B_i^{Y,k})^\top \mathfrak{B}^k H_t^{k,i} + \sum_{k\neq i} e^{-\rho\Delta_k}(B_i^Z)^\top \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t+\Delta_k\leftarrow t} \,\Big|\, \mathcal{F}_t^i\Big]. \end{aligned} \tag{3.2.2}$$

In the last two sums player i prices what its action does to opponent k's beliefs: $H^{k,i}$ is its belief price on opponent k's noise-state, and $\mathfrak{B}^k$ and $\mathfrak{B}^{k,Z}$ contract that price against k's private and delayed-public filter rows, both defined immediately after (3.2.4). The subscript ${}_{t+\Delta_k\leftarrow t}$ means a row generated at t and evaluated on arrival at $t+\Delta_k$. In (3.2.2) and in the proof of Proposition 3.1, $H_t^{X,i}$ and $H_t^{k,i}$ are the current-value adjoint processes, $e^{\rho t}$ times the costates of the discounted objective. Player i prices a shock's consequences through lags on both sides of the present. The negative lags show the benefit of acting on future shocks, and that benefit decays slowly with distance. The fundamental channel of $h^{X,1}$, its loading on the state noise, still holds a fifth of its peak two lags into the future, while two lags into the past it is nearly gone (Figure 3.2). The signal-noise channels, the loadings on the players' signal noises, behave differently at the boundary. A noise shock harms only through the responses it triggers, and at arrival no player has responded to it yet, so the channels' price peaks a short interval after the shock arrives. That near-boundary peak is the early peak. The peak lag shrinks as player 1's own precision grows, since faster filtering corrects the belief error sooner, from lag 0.11 at $p_1 = 1$ to 0.02 at $p_1 = 20$ (Figure 3.3). The depth is not monotone in p_1; it is largest near 4. The opponent's precision works the other way.

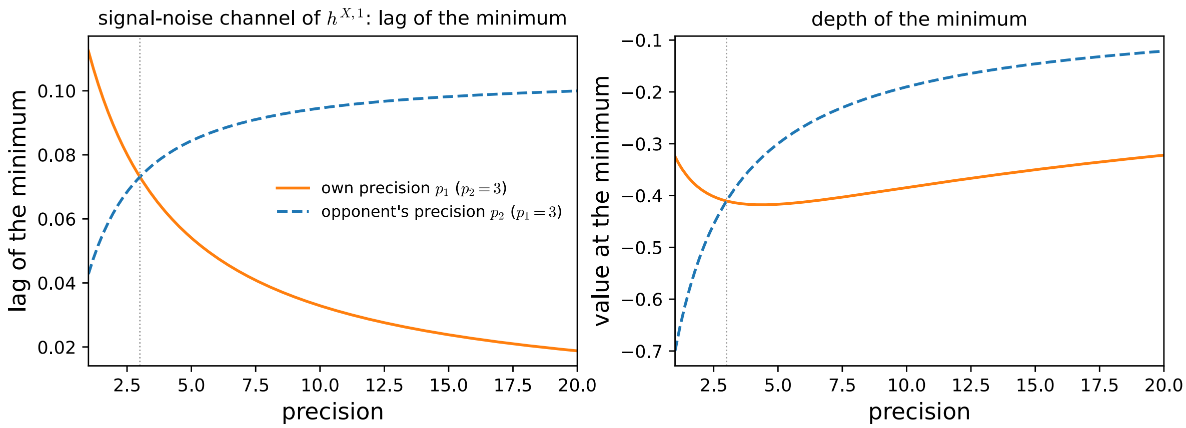

Figure 3.3: The early peak of $h^{X,1}$ against precision: lag and depth of the minimum of the signal-noise channel as the observer's own precision p_1 (solid) and the opponent's p_2 (dashed) vary, the other held at 3. Two hundred equilibria per curve from the spectral solver.

A more precise opponent responds to the shock sooner and more accurately, so the price of the noise shock falls, from -0.70 at $p_2 = 1$ to -0.12 at $p_2 = 20$, and its peak moves later.

Only lags $\ell \geq 0$ are estimable at time t, so

$$\begin{aligned}
\bar{d}^i &= -(G^{DD,i})^{-1}\Big[G^{DX,i}\bar{X} + (B_i^X + e^{-\rho\tau_i}\tilde{B}_i^X)^\top \bar{h}^{X,i} + \sum_{k\neq i}(B_i^{Y,k})^\top \mathfrak{B}^k \bar{h}^{k,i} \\
&\qquad + \sum_{k\neq i} e^{-\rho\Delta_k}(B_i^Z)^\top \Big(\mathfrak{B}^{k,Z}\bar{h}^{k,i}\Big)\Big], \\
d^i(a) &= -(G^{DD,i})^{-1}\Big[G^{DX,i}x(a) + (B_i^X)^\top h^{X,i}(a) + e^{-\rho\tau_i}(\tilde{B}_i^X)^\top h^{X,i}(a+\tau_i) \\
&\qquad + \sum_{k\neq i}(B_i^{Y,k})^\top \mathfrak{B}^k h^{k,i}(a) + \sum_{k\neq i} e^{-\rho\Delta_k}(B_i^Z)^\top \Big(\mathfrak{B}^{k,Z}h^{k,i}\Big)(a)\Big], \quad a > 0,
\end{aligned} \tag{3.2.3}$$

with the right trace $h^{X,i}(0^+)$ used at $a = 0$. Each term of (3.2.3) is an adjoint evaluated at the lag at which its effect occurs, and all are projected on $\mathcal{F}_t^i$. When player i's action takes effect τ_i later, the shock it reacts to now is τ_i older, so the player reads the state price τ_i further along the past-lag side; coordinates with age below τ_i, noise born between action and effect, drop out. An action the opponent sees directly is priced now, by contracting the belief price against the opponent's private filter row; an action that reaches the opponent only through the public signal is priced Δ_k later, when that signal arrives.

For opponent $k \neq i$, write

$$H_t^{k,i}(b) = \bar{h}^{k,i}(b) + \int_{-\infty}^{\infty} h^{k,i}(b,\ell)\, dW_{t-\ell}, \qquad b \geq 0. \tag{3.2.4}$$

The values at $b = 0$ are right traces. Write the private and delayed-public contractions of a belief price against opponent k's filter as

$$\begin{aligned} \mathfrak{B}^k h(\ell) &:= E^k h(0^+, \ell) + \int_0^\infty \tilde{c}^{k,Y}(b)\, h(b,\ell)\, db, \\ \Big(\mathfrak{B}^{k,Z} h\Big)(\ell) &:= E^Z h(\Delta_k^+, \ell + \Delta_k) + \int_0^\infty \tilde{c}^{k,Z}(b)\, h(b, \ell + \Delta_k)\, db, \end{aligned}$$

Define

$$\bar{v}^{k,i} = (B_X^{Y,k})^\top \mathfrak{B}^k \bar{h}^{k,i} + e^{-\rho\Delta_k} (B_X^Z)^\top \Big(\mathfrak{B}^{k,Z} \bar{h}^{k,i}\Big), \qquad \bar{v}^i = \sum_{k \neq i} \bar{v}^{k,i}, \tag{3.2.5}$$

$$v^{k,i}(\ell) = (B_X^{Y,k})^\top \mathfrak{B}^k h^{k,i}(\ell) + e^{-\rho\Delta_k} (B_X^Z)^\top \Big(\mathfrak{B}^{k,Z} h^{k,i}\Big)(\ell), \qquad v^i(\ell) = \sum_{k \neq i} v^{k,i}(\ell). \tag{3.2.6}$$

Without the public row or observed actions the first terms reduce to the baseline: birth term $(B_X^{Y,k})^\top E^k h^{k,i}(0^+, \ell)$ plus accumulated drift inference $P^k \int \tilde{x}^k(b) h^{k,i}(b,\ell)\, db$. When player i's own control enters opponent k's private row ($B_i^{Y,k} \neq 0$) the deviation moves that row directly. Its price $(B_i^{Y,k})^\top \mathfrak{B}^k h^{k,i}$ enters player i's first-order condition above rather than the state price, as the public-row term B_i^Z does. The sums $\bar{v}^i$ and v^i are the information wedge, the price of moving another player's beliefs, now including what the delayed public news will reveal.

The mean adjoint equations are

$$0 = \rho \bar{h}^{X,i} - (G^{XX,i} \bar{X} + G^{X,i} + G^{XD,i} \bar{d}^i) - (B_X^X)^\top \bar{h}^{X,i} - \bar{v}^i, \tag{3.2.7}$$

$$\frac{d}{db} \bar{h}^{k,i}(b) = \rho \bar{h}^{k,i}(b) - \Big((B_k^X + e^{-\rho\tau_k} \tilde{B}_k^X) d^k(b)\Big)^\top \bar{h}^{X,i} + r^{k,Y}(b)^\top \mathfrak{B}^k \bar{h}^{k,i}$$

$$+\, c^Z(b-\Delta_k)^\top \mathfrak{B}^{k,Z}\bar{h}^{k,i} - \sum_{k'\neq i} e^{-\rho\Delta_{k'}}\Big(B_k^Z d^k(b)\Big)^\top \mathfrak{B}^{k',Z}\bar{h}^{k',i}, \qquad \bar{h}^{k,i}(b)\to 0, \tag{3.2.8}$$

Set $x_+(\ell) = \mathbf{1}_{\ell\geq 0}x(\ell)$. Away from $\ell = 0$,

$$\begin{aligned}
\partial_\ell h^{X,i}(\ell) &= \rho h^{X,i}(\ell) - G^{XX,i}x_+(\ell) - G^{XD,i}d^i_{W,+}(\ell) - (B_X^X)^\top h^{X,i}(\ell) - v^i(\ell), \\
& h^{X,i}(\ell) \to 0 \ (|\ell|\to\infty), \qquad (3.2.9) \\
(\partial_b + \partial_\ell)h^{k,i}(b,\ell) &= \rho h^{k,i}(b,\ell) - (B_k^X d^k(b))^\top h^{X,i}(\ell) - e^{-\rho\tau_k}\Big(\tilde{B}_k^X d^k(b)\Big)^\top h^{X,i}(\ell+\tau_k) \\
&\quad + r^{k,Y}(b)^\top \mathfrak{B}^k h^{k,i}(\ell) + c^Z(b-\Delta_k)^\top\Big(\mathfrak{B}^{k,Z}h^{k,i}\Big)(\ell-\Delta_k) \\
&\quad - \sum_{k'\neq i} e^{-\rho\Delta_{k'}}\Big(B_k^Z d^k(b)\Big)^\top\Big(\mathfrak{B}^{k',Z}h^{k',i}\Big)(\ell), \qquad (3.2.10)
\end{aligned}$$

where $r^{k,Y}(b) = c^{k,Y}(b) - B_k^{Y,k}d^k(b)$ is the old-history coefficient of opponent k's private row (the opponent's own control is removed because its response to its beliefs is priced through d^k). In the baseline $r^{k,Y}(b)^\top\mathfrak{B}^k h^{k,i} = x(b)^\top v^{k,i}$ and $d^i_{W,+}(\ell) = \mathbf{1}_{\ell\geq 0}d^i_W(\ell)$. Opponent k's delayed control moves the state τ_k later, so player i prices that response at $\ell+\tau_k$. The delayed public signal enters twice. Opponent k now reads the public increment generated Δ_k ago, which loads on the coordinate through $c^Z(b-\Delta_k)$; player i prices it with the contraction at arrival, $(\mathfrak{B}^{k,Z}h^{k,i})(\ell-\Delta_k)$. The row generated now carries opponent k's response $B_k^Z d^k(b)$ and is observed by every opponent k' at $t+\Delta_{k'}$, priced by $e^{-\rho\Delta_{k'}}(\mathfrak{B}^{k',Z}h^{k',i})(\ell)$. At the arrival age $b=\Delta_k$ the belief price jumps by the centering trace,

$$h^{k,i}(\Delta_k^-,\ell) - h^{k,i}(\Delta_k^+,\ell) = -E^{Z\top}\Big[E^Z h^{k,i}(\Delta_k^+,\ell) + \int_0^\infty \tilde{c}^{k,Z}(b)h^{k,i}(b,\ell)\,db\Big], \tag{3.2.11}$$

the stationary form of the jump condition (2.4.32) at the arrival date. The belief price $h^{k,i}(b,\ell)$ decays as $b\to\infty$ and $|\ell|\to\infty$.

Proposition 3.1 (Stationary verification)**.** *Fix a stationary admissible profile. Suppose* *its kernels satisfy the forward equations* (3.1.8)–(3.1.19)*, the adjoint equations* (3.2.5)–(3.2.10)*, the jump condition* (3.2.11)*, the first-order condition* (3.2.2)*, and the stated decay/transversality conditions. Then each player is a stationary best response in the noise-state linear class. If the Schur complement* $S^i = G^{DD,i} - G^{DX,i}\left(G^{XX,i}\right)^{+} G^{XD,i} \succeq \lambda I$ *(*$G^{DX,i} = 0$ *is the tracking case, where* $S^i = G^{DD,i}$*), the stationary best response is unique.*

Proof. Apply the discounted spike variation of Appendix 1.B.2 on $[0, T]$ with opponents' stationary maps fixed, and let $T \to \infty$. The limit is taken as the definition of the stationary problem, with the two-sided adjoints (3.2.1) the objects verified against; convergence of the finite-horizon adjoints to them is not proved here. Write the kernels of (1.B.10) in lag coordinates and in current value. The first variation of J^i in a direction z is then

$$\delta J^i = \mathbb{E}\int_0^T e^{-\rho t}\, z_t^\top \Big(G^{DD,i} D_t^i + G^{DX,i} X_t + (B_i^X)^\top H_t^{X,i} + e^{-\rho\tau_i} (\tilde{B}_i^X)^\top H_{t+\tau_i}^{X,i} \\ + \sum_{k\neq i} (B_i^{Y,k})^\top \mathfrak{B}^k H_t^{k,i} + \sum_{k\neq i} e^{-\rho\Delta_k} (B_i^Z)^\top \Big(\mathfrak{B}^{k,Z} H^{k,i}\Big)_{t+\Delta_k \leftarrow t} \Big)\, dt$$

plus boundary terms. The finite-horizon adjoints $\bar{H}_t^X$, $H_t^X(\cdot)$ become the two-sided current-value adjoint (3.2.1), and (3.2.5)–(3.2.6) are the stationary form of the information wedge of Theorem 1.8. Integration by parts gives the boundary terms at 0, at T, and at lag zero, $\ell = 0$. The boundary terms at age zero are the axis birth terms (3.1.14)–(3.1.15); the terms at $\ell = 0$ cancel because bounded sources keep the kernels continuous there; the terms at the arrival age $b = \Delta_k$ are the centering trace, which the jump condition (3.2.11) absorbs; the terminal term vanishes as $T \to \infty$ by the transversality condition in the admissibility definition. The remaining variation is

$$\mathbb{E}\int_0^T e^{-\rho t}\, z_t^\top\, \mathbb{E}\Big[\, \cdot \,\Big|\, \mathcal{F}_t^i \Big]\, dt,$$

with the bracket of the display above inside the conditional expectation, and this is zero by (3.2.2). The cost is convex, so the profile is a best response (Corollary 1.11); under the Schur condition it is the unique one. $\square$

3.3 Transform identities

Delays enter the transforms as exponential factors. A delayed control row contributes $\tilde{B}^X_m e^{-p\tau_m}\check{d}^m_W(p)$ to (3.3.1), the delayed reads $h(\ell+\tau_k)$, $h(\ell+\Delta_k)$ become $e^{\mathbf{i}\omega\tau_k}\check{h}(\omega)$, $e^{\mathbf{i}\omega\Delta_k}\check{h}(\omega)$, and the births at the arrival age Δ_i add $e^{-p\Delta_i}$- and $e^{-q\Delta_i}$-weighted boundary transforms to (3.3.2). The formulas below are written for the baseline.

The transform formulas are not needed for verification, but they help when solving the stationary kernels. The objects transformed here include the belief prices and the wedge, which the competitive setting of the frequency-domain literature [60, 64, 101] forces to zero. For one-sided lag kernels set

$$\check{g}(p)=\int_0^\infty e^{-pa}g(a)\,da,\qquad \operatorname{Re}p>0,$$

and

$$\check{f}^i(p,q)=\int_0^\infty\int_0^\infty e^{-pb-qa}f^i(b,a)\,db\,da.$$

The state equation gives

$$(pI-B^X_X)\check{x}(p)=x(0)+\sum_{m=1}^{n}B^X_m\check{d}^m_W(p). \tag{3.3.1}$$

For the filter boundaries,

$$\check{f}^i_{a=0}(p):=\int_0^\infty e^{-pb}f^i(b,0)\,db=\check{\check{x}}^i(p)^\top(B^{Y,i}_X)^\top E^i,$$
$$\check{f}^i_{b=0}(q):=\int_0^\infty e^{-qa}f^i(0,a)\,da=E^{i\top}B^{Y,i}_X\check{\check{x}}^i(q).$$

Together these give

$$(p+q)\check{f}^i(p,q) = \check{f}^i_{a=0}(p) + \check{f}^i_{b=0}(q) + \check{\tilde{x}}^i(p)^\top P^i \check{\tilde{x}}^i(q). \tag{3.3.2}$$

For adjoints use the two-sided transform

$$\check{h}(\omega) = \int_{-\infty}^{\infty} e^{-\mathbf{i}\omega\ell} h(\ell)\, d\ell.$$

The kernel $h^{X,i}$ is continuous at $\ell = 0$, with a kink there rather than a jump (Section 3.A), so transforming (3.2.9) on the two half-lines produces no boundary term:

$$\Big((B^X_X)^\top - (\rho - \mathbf{i}\omega)I\Big)\check{h}^{X,i}(\omega) = -G^{XX,i}\check{x}_+(\omega) - G^{XD,i}\check{d}^i_{W,+}(\omega) - \check{v}^i(\omega). \tag{3.3.3}$$

For belief prices define

$$\begin{aligned}
\check{h}^{k,i}(p,\omega) &= \int_0^\infty \int_{-\infty}^\infty e^{-pb-\mathbf{i}\omega\ell} h^{k,i}(b,\ell)\, d\ell\, db, \\
\check{h}^{k,i}_0(\omega) &= \int_{-\infty}^\infty e^{-\mathbf{i}\omega\ell} h^{k,i}(0^+,\ell)\, d\ell,
\end{aligned}$$

with $h^{k,i}$ continuous across $\ell = 0$, since its sources are bounded. Then

$$(p + \mathbf{i}\omega - \rho)\check{h}^{k,i}(p,\omega) = \check{h}^{k,i}_0(\omega) - (B^X_k \check{d}^k(p))^\top \check{h}^{X,i}(\omega) + \check{x}(p)^\top \check{v}^{k,i}(\omega). \tag{3.3.4}$$

3.4 Rationality and finite-dimensional state

A one-sided kernel has a rational transform exactly when it is the impulse response of a finite-dimensional Markovian system. If the equilibrium kernels were rational, the noise-state would reduce to finitely many sufficient statistics and the stationary system would collapse to matrix Riccati equations. The primitives are rational, and filtering against a fixed observation structure preserves rationality. The equilibrium is different

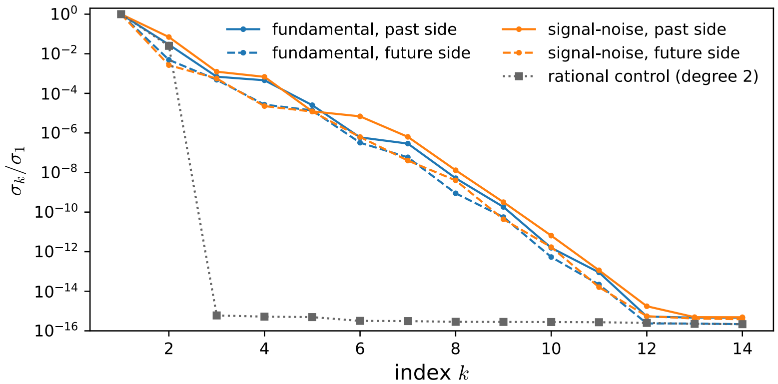

Figure 3.4: Singular values σ_k of the four one-sided pieces of $h^{X,1}$, with an exactly rational kernel of degree two as control.

because each player filters observations that contain the other players' filters. Each added level of beliefs about beliefs [109] raises the degree of the transform, and the equilibrium, sitting at the limit of the hierarchy, need not have a rational transform; whether the hierarchy collapses is model-specific [64, 94].

The computed kernels show this. Each one-sided kernel has a sequence of singular values $\sigma_1 \geq \sigma_2 \geq \cdots$ in which σ_{k+1} is the Hankel-norm error of the best k-dimensional approximation, so a rational kernel of degree k has $\sigma_{k+1} = 0$ [1, 44]. Figure 3.4 plots these values for the four one-sided pieces of $h^{X,1}$ against a rational control of degree two. Its third singular value is zero to machine precision. The equilibrium kernels' singular values fall by a factor of ten to thirty at each index and reach the accuracy of the computation, about 10^{-15}, by the twelfth without hitting zero. No low-order Markov state reproduces them exactly. Past the twelfth index the computation cannot distinguish a high degree from an infinite one.

The geometric decay is also a quantitative guarantee. A five-dimensional Markov state reproduces the kernel to better than 10^{-5}, and each further dimension gains another factor of ten. No k-dimensional model can do better than σ_{k+1}, and a truncated belief hierarchy is in general not the optimal k-dimensional approximation, so its error can be much larger than σ_{k+1}. Recent computational work on asymmetric-information LQG games restricts strategies to a finite-dimensional state from the outset [48, 92, 112]; the singular values lower-bound what any such restriction discards.

3.5 Conclusion

The stationary system is the finite-horizon noise-state calculus with calendar time removed, and its lag equations give the stationary notation the Kyle–Back chapter uses. Proposition 3.1 verifies that an admissible stationary profile solving these equations is a best response in the noise-state linear class, and the transform identities in Section 3.3 give an alternative representation for solving them. The uniqueness clause needs the Schur complement condition (1.3.3), which fails at zero trading cost, so the Kyle–Back computations run at a small positive cost and Chapter 4 verifies through second-order checks instead. Numerically, the computed kernels cannot be generated by any low-dimensional Markov state, though low-dimensional approximations are accurate (Section 3.4). Whether the infinite-horizon fixed point exists remains open.

3.A Computing the stationary equilibrium

Cutting off lags beyond L and placing them on a grid of fixed spacing turns the filter and adjoint kernels into vectors indexed by age. Delays are shifts of the age index. The kink of $h^{X,i}$ at $\ell = 0$ reappears at $\ell = -\tau_k$ and $\ell = -\Delta_k$, where the delayed kernels are evaluated, and at age Δ_i in the filter, so those points are panel boundaries in the spectral scheme.

The natural algorithm alternates between the filtering equations solved forward and the adjoint equations solved backward with damping, and it converges on a coarse grid but fails as the grid is refined. In classical filtering the update is nonexpansive because the observation structure is fixed in advance. Here the observation drift contains the equilibrium kernels themselves, so the projection changes from one iteration to the next. The combined step stops being a contraction as the kink at lag zero becomes steeper on finer grids.

Given the policy kernels, the filtering equations say that the belief loadings are the least-squares projection of each primitive shock onto the observation history. The solver computes that projection directly from the covariance matrix of the observations. That

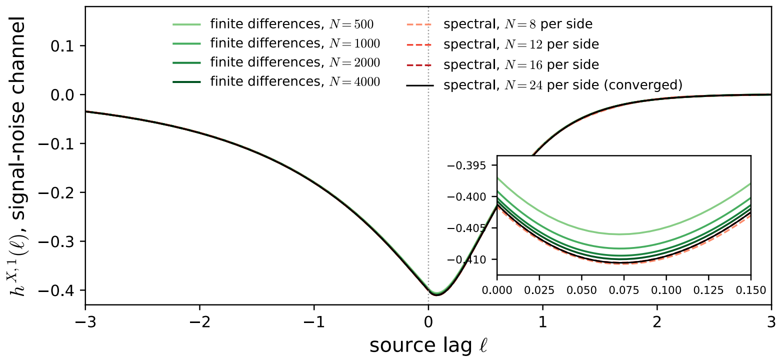

Figure 3.5: The signal-noise piece of $h^{X,1}$: grid solutions on successively finer grids (green) and the spectral solution with 24 points per side (black); spectral solutions with 8, 12, and 16 points (dashed) are indistinguishable from it at this scale. Inset: the grid solutions converging to the spectral one at the kink.

matrix is Toeplitz except in the rows for the shortest ages, and the conjugate gradient method with a circulant preconditioner solves it [30]. Only the update of the policy kernels remains iterative.

The grid scheme reads the state at the midpoint of each cell, which is an error of first order in the spacing. The solutions on grids from $N = 500$ to $N = 4000$ points still move by a few parts in a thousand between refinements. The reported figures therefore come from a second discretization. The kernels are smooth on each side of lag zero, so each one-sided kernel is represented by its values at Chebyshev points on $[0, L]$, and each two-sided kernel by two such sets on $[-L, 0]$ and $[0, L]$. The projection, the adjoint transport, and the wedge integrals become small dense matrices built from polynomial interpolation and Gauss–Legendre quadrature, and the equilibrium is found by Newton's method on the policy values. With 24 points per side the solution is converged to 10^{-13}, and a full equilibrium takes about ten milliseconds on a desktop. The grid solutions converge toward it as the spacing shrinks (Figure 3.5), and the depth of the kink agrees with an independent solution by the Wiener–Hopf method [91]. The spectral solution reproduces the closed-form filter of the uncontrolled problem to 3×10^{-8}, which is the size of the truncation at L.

Neural networks were also tried. At the fitting points, the networks reduced equation residuals below 2×10^{-3}, yet still misrepresented the future-lag side of the fundamental

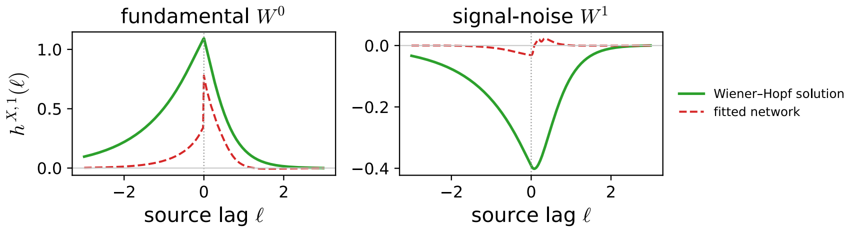

Figure 3.6: A neural-network solution compared with the Wiener–Hopf solution of the same system, with equation residuals below 2×10^{-3} at the fitting points.

piece and the entire kink at lag zero (Figure 3.6). Small residuals at the fitting points do not show that a network solution is correct, so any such solution should be checked through the finite-deviation identity of Chapter 1.

The solvers, refinement studies, and figures in this chapter are the work of Claude Fable 5, a large language model built by Anthropic [7], working under the author's direction. Interactive versions of the dissertation's computations are collected at sbabichenko.com/noisestate.

Chapter 4

The Information Wedge in a Stationary Kyle–Back Market

In a Kyle–Back market an order is at once a trade and a signal. It earns or loses money at execution, and it changes what the market maker and the other traders infer from order flow. This chapter prices that second effect for a market with finitely many strategic traders, in a stationary, ρ-discounted setting. Time runs over the whole real line, with no reset date; viewed from any anchor time, the state is each observer's stationary beliefs about past primitive shocks together with the current mispricing $\widehat{V}_t^i - P_t$. Once the initial date is sent to $-\infty$ the fundamental value and the price level need not be finite on their own, but their difference and the price moves caused by a marginal order are. A marginal order perturbs every other observer's beliefs; call that perturbation a *seed*. The chapter's two main objects are the deterministic map that propagates a seed forward into future prices and the backward adjoint system that prices it. The main result is a condition under which a stationary candidate profile is a best response for one trader. When it and the second-order condition hold for every trader, the profile is a stationary equilibrium within the noise-state-linear class, whose strategies are linear in beliefs about past shocks.

4.1 Introduction

Kyle's model has one risk-neutral informed trader, noise traders, and a competitive market maker [71]. The asset's liquidation value is $v \sim \mathcal{N}(p_0, \Sigma_0)$, and the noise traders submit orders independent of v. The informed trader observes v and submits an order x; the market maker observes only the combined flow $y = x + u$ and prices it. An equilibrium pairs a trading rule X and a pricing rule P so that the informed trader maximizes expected profit and the market maker breaks even,

$$X(v) \in \arg\max_x \mathbb{E}[(v - P(x+u))\, x \mid v]\,, \qquad P(y) = \mathbb{E}[\, v \mid x + u = y\,].$$

In a single round this has a unique linear solution $X(v) = \beta(v - p_0)$, $P(y) = p_0 + \lambda y$, with $\beta = (\sigma_u^2/\Sigma_0)^{1/2}$ and $\lambda = \frac{1}{2}(\Sigma_0/\sigma_u^2)^{1/2}$, in which half the informed trader's private information is impounded and $1/\lambda$ measures market depth.

Kyle's dynamic version runs N such auction rounds over a trading day $[0, 1]$. The equilibrium is recursive and linear, with price increments $\Delta p_n = \lambda_n(\Delta x_n + \Delta u_n)$ and a residual variance Σ_n that falls auction by auction, so the informed trader, to avoid revealing too much at once, trades gradually and information enters the price over the day. As the auctions are made frequent, the sequence converges to a continuous-time limit with constant price impact $\lambda = (\Sigma_0/\sigma_u^2)^{1/2}$ and linearly declining residual variance, $\Sigma_t = (1 - t)\Sigma_0$. Back proved this continuous-time equilibrium for general value distributions and a general strategy space [11].

The market maker remains competitive throughout this chapter. Its conditional-expectation pricing rule is a consistency condition of equilibrium, and nothing here tests the market maker's own deviations. Monitored deviations by a quoting market maker need the tools of Chapter 6 together with the lag coordinates of Chapter 3.

With several informed traders, each must forecast what the others infer from the order flow they jointly generate, the "forecasting the forecasts of others" problem of

Townsend and Sargent [104, 109], which in principle lets the state grow without bound, as traders forecast forecasts of forecasts. The canonical multi-trader extensions stay tractable through extreme symmetry in the traders' information: every trader's signal is drawn from the same distribution, with identical variance and pairwise covariance, and each trader receives exactly one signal, at time 0. Under that symmetry the average signal is a sufficient statistic for the price and the regress collapses to a finite state. Holden and Subrahmanyam [56] take the symmetry to its extreme: every trader observes the same v. Competition becomes a "rat race" in which even two traders reveal almost all of v at once, so that, unlike the single-trader case, the market becomes infinitely deep in the high-frequency limit. Foster and Viswanathan [38] keep the symmetry but let the common signal split into correlated pieces of equal quality. As traders learn the others' pieces from the order flow, the conditional correlation between their residual signals falls and can turn negative, so an early "rat race" gives way to a "waiting game" in which each trader holds back and hopes the others move the price. Competition shifts regime over the trading day. Back et al. [13] carry this imperfect competition into continuous time.

A parallel line reaches the single-insider equilibrium through enlargement of filtrations and dynamic Markov bridges [24, 29]. The construction is exact and general in the value distribution, but its scope is the same corner: one insider, one static endowment, one scalar quantity to be revealed. The obstruction is a filtering one. Estimating X from $X + \varepsilon_X$ and Y from $Y + \varepsilon_Y$ are standard problems; estimating both from $X + Y + \varepsilon_X + \varepsilon_Y$ is not the sum of the two. The order flow filtered by the market maker is exactly such a superposition, built from every trader's response to their own shock history and arriving continuously. Most models descended from Kyle remove the superposition problem through the symmetry above. The exceptions are partial: one of two traders knowing more [37], general information in a single round [72], continuous arrival to a single insider [12, 31]. With continuous signal flows of different quality no single variable pins the price, and each observer's estimate needs its own kernel and its own belief price. For the same reason, price impact is trader-specific. Each trader's kernels enter the superposition differently, so although the instantaneous impact of anonymous order flow is common, the priced dynamic response differs by trader; the symmetric models carry a single impact coefficient in both senses.

The model takes the Foster–Viswanathan branch to a continuous-time, infinite-horizon setting and relaxes its symmetry. In their model each trader i holds one private signal s_i at $t = 0$, of fixed correlation and equal quality. The market maker prices the order flow $y = \sum_i x_i + u$ at the conditional expectation while tracking $t_i = \mathbb{E}[s_i \mid \text{order flow}]$. The residual $s_i - t_i$, what trader i knows that the market maker does not, is the finite sufficient statistic that holds the regress in check.

Time runs over the whole real line rather than a fixed interval; signals arrive continuously rather than once; their quality differs across traders; the fundamental is a process V_t rather than a single terminal value v; and the fundamental and the noise-trader flow load on mutually orthogonal factor families. The public flow Z is Foster–Viswanathan's order flow y, and Σ_Z carries the noise-trader variance σ_u^2. Each Y^j is the flowing counterpart of the private signal s_j, a noisy reading of the gap $V_t - P_t$.

Once signals flow, differ in quality, and read a moving fundamental, the symmetry that reduced the state to a finite sufficient statistic is gone. The residual $s_i - t_i$ gives way to observer j's noise-state over an unbounded shock history, with filtering encoded by the deterministic kernel $F_t^j(r, u)$.

Trader i values future profit by the discounted analogue of the Foster–Viswanathan objective, with the finite horizon replaced by the discount rate ρ and, as in those models, the trader's estimate $\widehat{V}^i$ in place of v under the conditional expectation. The per-period price impact λ_n becomes a price-impact map $\Lambda(t)$, and a marginal order moves the price contemporaneously only in its range.

Two feedback loops matter. First, each trader forms a posterior from its private price signal and trades on it; the same trade earns profit from the mispricing and enters the public order flow Z. The market maker and the traders then observe public order flow, so exploiting information also changes the information environment (Figure 4.1). Second, a marginal deviation has an immediate market-maker channel and an intertemporal trader-feedback channel.

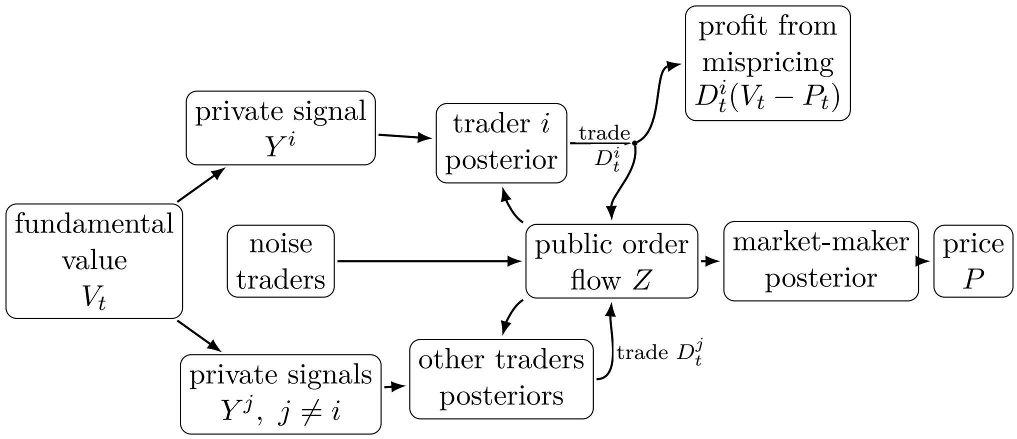

Figure 4.1: Information feedback around public order flow. The trader trades on its posterior. The same trade is a payoff action, earning $D_t^i(V_t - P_t)$, and a signal, because it enters public order flow. Other strategic trades and noise-trader demand also enter Z. The market maker updates its posterior on public order flow and sets the price from it; the traders update on the same flow.

The belief-price method of Chapter 1 prices that second, intertemporal channel: here the beliefs a deviation manipulates are the ones the market maker and the other traders form from public order flow rather than from an abstract state signal. In Foster–Viswanathan a trader who has deviated carries one extra scalar, the gap between the equilibrium-path price and the realized price, to summarize past off-equilibrium play; here the deviation is priced at the margin. Perturb trader i's order at time t by a marginal spike z. Because the order enters the public flow Z, it changes every other observer's signal history and so their noise-state estimates, in two parts: a shift in the old-history coordinates ($r < t$) and a fresh coordinate born at the current date (the birth term). A deterministic forward-response map carries the seed through the observers' filters, the other traders' policy kernels, and the pricing rule. Its output is the future price response $p_i(\tau, t; z)$. The adjoint system, the discounted dual of this map, prices the seed by integrating $(D_\tau^i)^\top p_i(\tau, t; z)$, the future execution cost the price response creates, with discount rate ρ.

4.2 The stationary Kyle–Back market

Dictionary from Chapter 1. The market notation below specializes the baseline calculus. The physical state X contains the fundamental, price, and mispricing coordinates, while D^i is trader i's order rate. The market maker is observer 0, and the belief price $H^{k,i}$ prices movements of observer k's posterior. The only persistent new market object is the price-impact map Λ, the instantaneous map from an order innovation to a price innovation.

Spike terminology. A *spike* is the marginal order deviation. Its seed is the immediate state and filter displacement created at the spike time, and its forward response propagates that seed forward. Movements of coordinates already present before the spike are old-history effects.

Fix a discount rate $\rho > 0$ and a deviating trader i. Viewed from a finite anchor time t, all past shocks $r < t$ enter the state only through the observers' filter kernels

$$F_t^j(r,u), \qquad r,u < t, \qquad j \in \{0,1,\dots,n\}.$$

Here 0 is the market maker and $1,\dots,n$ are traders. For a deviation by trader i, the shifted observer blocks are

$$\mathcal{O}_i := \{0\} \cup \{j : 1 \le j \le n,\ j \ne i\}. \tag{4.2.1}$$

4.2.1 Primitive shocks and covariance selectors

One standard d-dimensional two-sided Brownian noise W generates all randomness. Economic shocks are fixed linear combinations of the increments dW:

$$\begin{aligned} dV_t &= \Sigma_V^{1/2} E^V \, dW_t, \\ dZ_t &= \sum_{j=1}^{n} D_t^j \, dt + \Sigma_Z^{1/2} E^Z \, dW_t, \\ dY_t^j &= B_X^{Y,j}(V_t - P_t) \, dt + \Sigma_{Y,j}^{1/2} E^j \, dW_t, \qquad j = 1, \dots, n. \end{aligned} \tag{4.2.2}$$

Here $V_t \in \mathbb{R}^q$ and $Z_t \in \mathbb{R}^q$ are levels relative to an arbitrary finite origin; changing the origin adds a common formal level that cancels from M_t^i (Section 4.2.3). The market maker prices competitively and each trader values the asset from its own filtration,

$$P_t = \mathbb{E}[V_t \mid \mathcal{F}_t^0], \qquad \widehat{V}_t^i = \mathbb{E}[V_t \mid \mathcal{F}_t^i], \tag{4.2.3}$$

relative to the same origin. The first equality is imposed as a consistency rule. The matrices E^V, E^Z, E^j are selector maps into primitive-shock channels. Throughout this chapter their rows are normalized,

$$E^V (E^V)^\top = I, \qquad E^Z (E^Z)^\top = I, \qquad E^j (E^j)^\top = I, \tag{4.2.4}$$

with identities of the appropriate dimensions. The matrices $\Sigma_V^{1/2}$, $\Sigma_Z^{1/2}$, and $\Sigma_{Y,j}^{1/2}$ therefore carry the covariance scale, so

$$\mathrm{Cov}(dV_t) = \Sigma_V \, dt, \qquad \mathrm{Cov}(\Sigma_Z^{1/2} E^Z \, dW_t) = \Sigma_Z \, dt, \qquad \mathrm{Cov}(\Sigma_{Y,j}^{1/2} E^j \, dW_t) = \Sigma_{Y,j} \, dt. \tag{4.2.5}$$

Cross-covariances enter through

$$\Sigma_Y^{jk} := \Sigma_{Y,j}^{1/2} E^j (E^k)^\top \Sigma_{Y,k}^{1/2}. \tag{4.2.6}$$

The lack of correlation between fundamental factors and noise-trader factors is the orthogonality condition

$$E^V (E^Z)^\top = 0. \tag{4.2.7}$$

In the baseline version the private-signal noise blocks are also orthogonal to the fundamental and noise-trader blocks:

$$E^V (E^j)^\top = 0, \qquad E^Z (E^j)^\top = 0, \qquad j = 1, \dots, n. \tag{4.2.8}$$

A raw order-flow spike therefore has no direct private-signal noise component; any effect on Y^j comes through the drift $B_X^{Y,j}(V_t - P_t)$.

4.2.2 Standing assumptions

Assumption 4.1 (Stationary admissible environment)**.** The following conditions stand throughout the chapter.

(i) The discount rate satisfies $\rho > 0$.

(ii) The covariance Σ_Z is inverted only on the traded order-flow support. If Σ_Z is singular, $\Sigma_Z^{-1/2}$ denotes the corresponding pseudoinverse square-root map.

(iii) Policy kernels are predictable and have zero diagonal trace: $D_t^j(t) = 0$.

(iv) The notation distinguishes noise-state and primitive coordinates. The kernel $D_t^j(u)$ multiplies trader j's noise-state increment, while D_t^j and $D_{W,t}^j(u)$ denote raw demand and its primitive-shock response kernel.

(v) The stationary kernels are square-integrable enough for the displayed stochastic integrals, convolutions, and conditional expectations.

(vi) The discounted belief prices satisfy the transversality condition in Definition 4.5 (Section 4.5).

(vii) If $\Lambda(t)$ of (4.3.17) is singular, read all first-order conditions as range equations. The pseudoinverse selects a representative only on the instantaneous impacted subspace.

4.2.3 Discounted objective and finite mispricing

Write D_t^i for trader i's realized raw-demand control. Trader i evaluates future profit from time t by

$$J_t^i(D^i; D^{-i}) := \mathbb{E}\left[\int_t^\infty e^{-\rho(\tau-t)}(D_\tau^i)^\top \left(\widehat{V}_\tau^i - P_\tau\right) d\tau \;\middle|\; \mathcal{F}_t^i\right]. \tag{4.2.9}$$

Only the mispricing

$$M_t^i := \widehat{V}_t^i - P_t \tag{4.2.10}$$

will enter the first-order condition. At zero trading cost the best response is not unique. The computations of Section 4.8 therefore add a quadratic trading cost, which selects one best response and leaves the reported quantities converging as the cost vanishes.

Before expanding the filter response, write $p_i(\tau, t; z)$ for the price movement at future time τ caused by a raw spike $z \in \mathbb{R}^q$ submitted by trader i at time t. By submitting a spike of length h in direction z, trader i earns $z^\top M_t^i$ now, to first order. The spike also changes the prices at which trader i executes later, and the first variation of the continuation payoff is

$$z^\top M_t^i - \mathbb{E}\left[\int_t^\infty e^{-\rho(\tau-t)}(D_\tau^i)^\top p_i(\tau, t; z)\, d\tau \,\middle|\, \mathcal{F}_t^i\right]. \tag{4.2.11}$$

At a best response this expression is zero for every $\mathcal{F}_t^i$-measurable spike direction z, so the unexpanded spike stationarity condition is

$$z^\top M_t^i = \mathbb{E}\left[\int_t^\infty e^{-\rho(\tau-t)}(D_\tau^i)^\top p_i(\tau,t;z)\,d\tau \,\middle|\, \mathcal{F}_t^i\right], \qquad z\in\mathbb{R}^q. \tag{4.2.12}$$

Lemma 4.2 (Finite mispricing under an infinite past)**.** *Introduce a finite lower cutoff $a<t$ and set*

$$V_t^{(a)} := \int_a^t \Sigma_V^{1/2} E^V\, dW_s.$$

Let $\widehat{V}_t^{i,(a)}$ and $P_t^{(a)}$ be the corresponding trader-i estimate and market-maker price. If the difference between the trader and market-maker filter kernels is square integrable on $(-\infty,t)$, then

$$M_t^i = \lim_{a\downarrow-\infty}\left(\widehat{V}_t^{i,(a)} - P_t^{(a)}\right) \tag{4.2.13}$$

exists in L^2. The individual levels $\widehat{V}_t^{i,(a)}$ and $P_t^{(a)}$ may fail to converge as $a\downarrow-\infty$.

Proof. For each cutoff, both $\widehat{V}_t^{i,(a)}$ and $P_t^{(a)}$ are linear functionals of the primitive shock on (a,t). Their difference is the same primitive shock integrated against the difference of the two deterministic filter kernels. The stated square-integrability condition makes this family Cauchy in L^2 by Itô isometry. □

4.3 The life of a trader's order

A marginal order does three things at the instant it arrives. It re-weights everything the observers already believed, moving the old coordinates of their noise-states; it creates a new coordinate born at the current date; and it moves the price through the market maker's filter.

4.3.1 Order-flow block and filter kernels

The market maker observes order flow in normalized units, using $\Sigma_Z^{-1/2} Z$. Trader $j \geq 1$ observes this same order-flow channel plus their private signal. The selector

$$\chi^0 := I_q, \qquad \chi^j := \begin{pmatrix} I_q \\ 0_{p_j - q, q} \end{pmatrix}, \qquad j \geq 1, \tag{4.3.1}$$

embeds a normalized order-flow increment into the Z-block of trader j's stacked observation, of dimension p_j, and adds zero to the private-signal block.

For any trader j, the same realized raw-demand process has two impulse-response representations, following the convention of Chapter 1:

$$D_t^j = \bar{D}_t^j + \int_{-\infty}^t D_t^j(u)\, d\widehat{W}_t^j(u) = \bar{D}_t^j + \int_{-\infty}^t D_{W,t}^j(u)\, dW_u. \tag{4.3.2}$$

The first integral writes the policy in trader j's noise-state coordinates, $\widehat{W}_t^j(u) = \mathbb{E}[W_u \mid \mathcal{F}_t^j]$, the second the same raw demand in primitive-shock coordinates. The total primitive-coordinate order-flow drift is

$$D_{W,t}^{\text{tot}}(u) := \sum_{j=1}^n D_{W,t}^j(u), \qquad u < t. \tag{4.3.3}$$

For each trader j, the normalized drift of the private price signal is $\Sigma_{Y,j}^{-1/2} B_X^{Y,j}(V_t - P_t)$. The price is in $\mathcal{F}_t^j$, a linear functional of the order-flow history the trader sees, so it drops out of the innovation on and off the path, and only the fundamental's kernel enters; write

$$\Sigma_{Y,j}^{-1/2} B_X^{Y,j} V_t = \overline{C}_t^{Y,j} + \int_{-\infty}^t C_t^{Y,j}(u)\, dW_u, \qquad j = 1, \ldots, n. \tag{4.3.4}$$

Its unresolved part is the bottom block of (4.3.11). The observation-drift kernels are

$$C_t^0(u) = \Sigma_Z^{-1/2} D_{W,t}^{\text{tot}}(u), \tag{4.3.5}$$

for the market maker, and

$$C_t^j(u) = \begin{pmatrix} \Sigma_Z^{-1/2} D_{W,t}^{\mathrm{tot}}(u) \\ C_t^{Y,j}(u) \end{pmatrix}, \qquad j \geq 1, \tag{4.3.6}$$

for trader j.

The filter equations use the primitive objects each observer has not yet resolved. Let

$$\Pi^0 := E^{Z,\top} E^Z, \qquad \Pi^j := E^{Z,\top} E^Z + E^{j,\top} E^j, \qquad j \geq 1. \tag{4.3.7}$$

Each Π^k collects the channels observer k resolves at once: order flow for everyone, a trader's own signal. The unresolved aggregate order-flow drift for observer $k \in \{0, 1, \ldots, n\}$ is

$$\widetilde{D}_{W,t}^{\mathrm{tot},k}(u) := D_{W,t}^{\mathrm{tot}}(u)(I - \Pi^k) - \int_{-\infty}^t D_{W,t}^{\mathrm{tot}}(r) F_t^k(r,u)\, dr. \tag{4.3.8}$$

The unresolved value kernel is

$$\widetilde{V}_t^k(u) := \Sigma_V^{1/2} E^V (I - \Pi^k) - \int_{-\infty}^t \Sigma_V^{1/2} E^V F_t^k(r,u)\, dr. \tag{4.3.9}$$

Under the maintained orthogonality between fundamental shocks and contemporaneous observation-noise blocks, the first term in (4.3.9) is $\Sigma_V^{1/2} E^V$.

The unresolved observation-drift kernels are then

$$\widetilde{C}_t^0(u) = \Sigma_Z^{-1/2} \widetilde{D}_{W,t}^{\mathrm{tot},0}(u), \tag{4.3.10}$$

and, for trader $j \geq 1$,

$$\widetilde{C}_t^j(u) = \begin{pmatrix} \Sigma_Z^{-1/2} \widetilde{D}_{W,t}^{\mathrm{tot},j}(u) \\ \Sigma_{Y,j}^{-1/2} B_X^{Y,j} \widetilde{V}_t^j(u) \end{pmatrix}. \tag{4.3.11}$$

The top block is the order-flow channel, the bottom the private price signal.

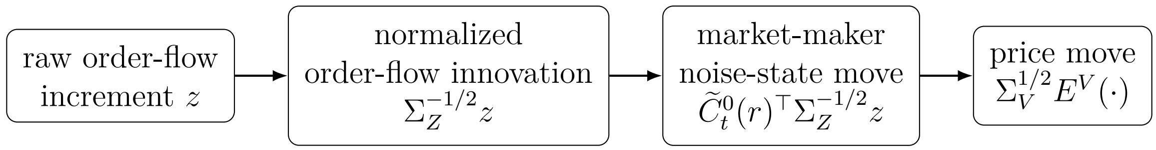

Figure 4.2: Order-flow news first becomes a normalized order-flow innovation, then a market-maker noise-state movement, and then a price movement.

The compact old-history filter equation is

$$\partial_t F_t^k(r,u) = \widetilde{C}_t^k(r)^\top \widetilde{C}_t^k(u), \qquad k = 0, 1, \ldots, n, \quad r, u < t. \tag{4.3.12}$$

The market-maker filter evolves by

$$\partial_t F_t^0(r,u) = \widetilde{D}_{W,t}^{\mathrm{tot},0}(r)^\top \Sigma_Z^{-1} \widetilde{D}_{W,t}^{\mathrm{tot},0}(u), \qquad r, u < t, \tag{4.3.13}$$

with birth traces, the kernel's boundary values at the coordinate born at t,

$$F_t^0(r,t) = \widetilde{D}_{W,t}^{\mathrm{tot},0}(r)^\top \Sigma_Z^{-1/2} E^Z, \qquad F_t^0(t,u) = E^{Z,\top} \Sigma_Z^{-1/2} \widetilde{D}_{W,t}^{\mathrm{tot},0}(u). \tag{4.3.14}$$

For trader $j \geq 1$, the old-history equation is

$$\begin{aligned} \partial_t F_t^j(r,u) = {} & \widetilde{D}_{W,t}^{\mathrm{tot},j}(r)^\top \Sigma_Z^{-1} \widetilde{D}_{W,t}^{\mathrm{tot},j}(u) \\ & + \widetilde{V}_t^j(r)^\top (B_X^{Y,j})^\top \Sigma_{Y,j}^{-1} B_X^{Y,j} \widetilde{V}_t^j(u), \qquad r, u < t. \end{aligned} \tag{4.3.15}$$

Its birth traces are

$$\begin{aligned} F_t^j(r,t) &= \widetilde{D}_{W,t}^{\mathrm{tot},j}(r)^\top \Sigma_Z^{-1/2} E^Z + \widetilde{V}_t^j(r)^\top (B_X^{Y,j})^\top \Sigma_{Y,j}^{-1/2} E^j, \\ F_t^j(t,u) &= E^{Z,\top} \Sigma_Z^{-1/2} \widetilde{D}_{W,t}^{\mathrm{tot},j}(u) + E^{j,\top} \Sigma_{Y,j}^{-1/2} B_X^{Y,j} \widetilde{V}_t^j(u). \end{aligned} \tag{4.3.16}$$

4.3.2 From an order to a price move

The price-impact map is

$$\Lambda(t) := \left(\int_{-\infty}^{t} \Sigma_V^{1/2} E^V \widetilde{C}_t^0(r)^\top \, dr \right) \Sigma_Z^{-1/2}. \tag{4.3.17}$$

An order moves the price only by moving the market maker's innovations (Figure 4.2). By the factor orthogonality $E^V E^{Z,\top} = 0$ the birth term of the transpose contraction $\mathfrak{B}^{0\top}$ of Chapter 1, (1.4.9), vanishes, and $\Lambda(t) = \Sigma_V^{1/2} E^V \, \mathfrak{B}^{0\top} \, \Sigma_Z^{-1/2}$. A raw order-flow spike in a direction $z \in \operatorname{Ker} \Lambda(t)$, equivalently with $\int_{-\infty}^{t} \widetilde{C}_t^0(r)^\top \Sigma_Z^{-1/2} z \, dr \in \operatorname{Ker}(\Sigma_V^{1/2} E^V)$, produces no instantaneous market-maker price response; its price effects are future ones, carried by the forward response and belief-price terms below.

A strategic unit spike by trader i places the old-history seeds

$$\varphi_i^0(t,r) := \widetilde{C}_t^0(r)^\top \Sigma_Z^{-1/2}, \tag{4.3.18}$$

and, for trader observers $j \neq i$,

$$\varphi_i^j(t,r) := \widetilde{C}_t^j(r)^\top \chi^j \Sigma_Z^{-1/2}, \qquad j = 1, \ldots, n, \quad j \neq i. \tag{4.3.19}$$

The same spike creates the birth seeds

$$\psi_i^0(t) := E^{Z,\top} \Sigma_Z^{-1/2}, \tag{4.3.20}$$

and, for trader observers $j \neq i$,

$$\psi_i^j(t) := E^{Z,\top} \Sigma_Z^{-1/2}, \qquad j = 1, \ldots, n, \quad j \neq i. \tag{4.3.21}$$

The equality follows because the spike enters trader j's stacked observation only through the common order-flow block; its private-signal noise component is zero.

Remark 4.3 (Order-flow martingales and primitive drift). Back's continuous-time equilibrium makes total order flow a martingale in the market-maker filtration. In the notation here this martingale property is imposed after conditioning on the market maker's information. The kernel C^0 is the primitive-shock impulse response of the normalized order-flow drift, while $\widetilde{C}^0$ is the part of that drift not yet resolved by the market-maker filter. The market-maker innovation can then be a martingale even though C^0 is nonzero. In the pure risk-neutral Kyle–Back benchmark, with no inventory or risk-premium motive for the market maker, the predictable component of public order flow is exactly the part subtracted by the filter. Section 4.8.1 works out the special case $\widetilde{C}^0 = C^0$, in which no part of the primitive order-flow drift is already resolved by the market maker.

4.3.3 Spike seed measure and immediate response

For a raw spike direction $z \in \mathbb{R}^q$ submitted by trader i, combine the old-history and birth parts into the seed measure on $(-\infty, t]$

$$d\eta_i^0(t, u; z) = \varphi_i^0(t, u) z \, du + \psi_i^0(t) z \, \delta_t(du), \tag{4.3.22}$$

$$d\eta_i^j(t, u; z) = \varphi_i^j(t, u) z \, du + \psi_i^j(t) z \, \delta_t(du), \qquad j = 1, \ldots, n, \; j \neq i. \tag{4.3.23}$$

The other traders' immediate change in demand. The old-history part of a spike generates the immediate opponent demand response

$$\delta D_{\Sigma,t}^{(-i)}(z) = R^i(t) z, \qquad R^i(t) = \sum_{\substack{j=1 \\ j \neq i}}^{n} \int_{-\infty}^{t} D_t^j(u) \widetilde{C}_t^j(u)^\top \chi^j \Sigma_Z^{-1/2} \, du. \tag{4.3.24}$$

$R^i(t)$ leaves out the birth atom, because the policy kernel is predictable and has zero trace at $u = t$, so no instantaneous term $D_t^j(t) \psi_i^j z$ appears.

4.4 Forward response

Write $\delta w^j_\tau(r)$ for the perturbation of observer j's noise-state coordinate r at future time τ, in density form: $d_r\,\delta\widehat{W}^j_\tau(r) = \delta w^j_\tau(r)\,dr$. The change in the other traders' total demand at time τ is

$$\delta D^{(-i)}_{\Sigma,\tau} = \sum_{\substack{j=1\\ j\neq i}}^{n} \int_{-\infty}^{\tau} D^j_\tau(u)\,\delta w^j_\tau(u)\,du. \tag{4.4.1}$$

The old-history perturbations satisfy

$$\partial_\tau \delta w^0_\tau(r) = \widetilde{C}^0_\tau(r)^\top \left[\Sigma_Z^{-1/2}\delta D^{(-i)}_{\Sigma,\tau} - \int_{-\infty}^{\tau} C^0_\tau(u)\delta w^0_\tau(u)\,du\right], \tag{4.4.2}$$

$$\partial_\tau \delta w^j_\tau(r) = \widetilde{C}^j_\tau(r)^\top \left[\chi^j\Sigma_Z^{-1/2}\delta D^{(-i)}_{\Sigma,\tau} - \int_{-\infty}^{\tau} C^j_\tau(u)\delta w^j_\tau(u)\,du\right], \tag{4.4.3}$$

for $r < \tau$ and $j = 1, \ldots, n$, $j \neq i$. The coordinate born at the current date is

$$\delta w^0_\tau(\tau) = E^{Z,\top}\left[\Sigma_Z^{-1/2}\delta D^{(-i)}_{\Sigma,\tau} - \int_{-\infty}^{\tau} C^0_\tau(u)\delta w^0_\tau(u)\,du\right], \tag{4.4.4}$$

$$\delta w^j_\tau(\tau) = \begin{pmatrix} E^{Z,\top} & E^{j,\top}\end{pmatrix}\left[\chi^j\Sigma_Z^{-1/2}\delta D^{(-i)}_{\Sigma,\tau} - \int_{-\infty}^{\tau} C^j_\tau(u)\delta w^j_\tau(u)\,du\right], \tag{4.4.5}$$

for $j = 1, \ldots, n$, $j \neq i$. Only the market maker's noise-state is read directly into a price movement:

$$\delta P_\tau = \int_{-\infty}^{\tau} \Sigma_V^{1/2} E^V\,\delta w^0_\tau(s)\,ds. \tag{4.4.6}$$

This perturbation is finite by the argument of Lemma 4.2.

Linearity lets one propagate a generic initial perturbation of an observer's noise-state and contract that response against the seed generated by a raw order-flow spike.

For $m, k \in \mathcal{O}_i$, define $G^{m\leftarrow k}_{\tau\leftarrow t}(s \leftarrow u)$ as the distributional kernel that maps an initial perturbation in observer k's noise-state coordinate $u \leq t$ at time t into the perturbation of observer m's coordinate $s \leq \tau$ at time τ:

$$(G^{m\leftarrow k}_{\tau\leftarrow t}\eta)(s) := \int_{(-\infty,t]} G^{m\leftarrow k}_{\tau\leftarrow t}(s\leftarrow u)\,d\eta(u), \qquad \tau \geq t. \tag{4.4.7}$$

The initial condition is $G^{m\leftarrow k}_{t\leftarrow t} = \mathbf{1}_{\{m=k\}}\mathrm{Id}$. Equivalently, $G^{m\leftarrow m}_{t\leftarrow t}(s \leftarrow u) = \delta(s-u)$ in the distributional sense. Read all identities involving the Dirac initial kernel $\delta(s-u)$ weakly, after pairing with the finite seed measure above.

For $s < \tau$, the forward-response kernels satisfy the same old-history equations as the perturbations. For the market maker's noise-state,

$$\partial_\tau G^{0\leftarrow k}_{\tau\leftarrow t}(s \leftarrow u) = \widetilde{C}^0_\tau(s)^\top \Bigg[\Sigma_Z^{-1/2} \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{\tau} D^j_\tau(r) G^{j\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr - \int_{-\infty}^{\tau} C^0_\tau(r) G^{0\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr\Bigg]. \tag{4.4.8}$$

For the noise-state of a trader $m \neq i$,

$$\partial_\tau G^{m\leftarrow k}_{\tau\leftarrow t}(s \leftarrow u) = \widetilde{C}^m_\tau(s)^\top \Bigg[\chi^m \Sigma_Z^{-1/2} \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{\tau} D^j_\tau(r) G^{j\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr - \int_{-\infty}^{\tau} C^m_\tau(r) G^{m\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr\Bigg]. \tag{4.4.9}$$

The coordinate born at $s = \tau$ is

$$G^{0\leftarrow k}_{\tau\leftarrow t}(\tau \leftarrow u) = E^{Z,\top} \Bigg[\Sigma_Z^{-1/2} \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{\tau} D^j_\tau(r) G^{j\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr - \int_{-\infty}^{\tau} C^0_\tau(r) G^{0\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr\Bigg], \tag{4.4.10}$$

for the market maker. For the noise-state of a trader $m \neq i$, split the stacked residual into its common order-flow and private-signal components:

$$\mathcal{E}^{m\leftarrow k}_{Z,\tau\leftarrow t}(u) := \Sigma_Z^{-1/2} \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{\tau} D^j_\tau(r) G^{j\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr - \int_{-\infty}^{\tau} C^0_\tau(r) G^{m\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr, \tag{4.4.11}$$

$$\mathcal{E}^{m\leftarrow k}_{Y,\tau\leftarrow t}(u) := -\int_{-\infty}^{\tau} C^{Y,m}_\tau(r) G^{m\leftarrow k}_{\tau\leftarrow t}(r \leftarrow u)\, dr. \tag{4.4.12}$$

Then the born coordinate is

$$G^{m\leftarrow k}_{\tau\leftarrow t}(\tau \leftarrow u) = E^{Z,\top}\mathcal{E}^{m\leftarrow k}_{Z,\tau\leftarrow t}(u) + E^{m,\top}\mathcal{E}^{m\leftarrow k}_{Y,\tau\leftarrow t}(u), \qquad m \neq i. \tag{4.4.13}$$

The future price movement at τ caused by the raw spike is the sum of the price responses generated by each observer whose beliefs the spike moves:

$$p_i(\tau, t; z) := \sum_{k\in\mathcal{O}_i} \int_{(-\infty,t]} p^k(\tau, t, u)\, d\eta^k_i(t, u; z). \tag{4.4.14}$$

The channel price-response density is

$$p^k(\tau, t, u) := \int_{-\infty}^{\tau} \Sigma_V^{1/2} E^V G^{0\leftarrow k}_{\tau\leftarrow t}(s \leftarrow u)\, ds, \qquad k \in \mathcal{O}_i. \tag{4.4.15}$$

Proposition 4.4 (Instantaneous price impact)**.** *For every raw spike direction z,*

$$p_i(t, t; z) = \Lambda(t)z. \tag{4.4.16}$$

Proof. At $\tau = t$ the forward response is the identity on the seed's own block and zero elsewhere, so only the market maker's seed reaches the price: $G^{0\leftarrow 0}_{t\leftarrow t}(s \leftarrow u) = \delta(s - u)$. Contracting this with the old-history market-maker seed (4.3.18) gives $\Lambda(t)z$. The birth atom in the market-maker block would contribute $\Sigma_V^{1/2} E^V E^{Z,\top} \Sigma_Z^{-1/2} z$, zero by the fundamental/order-flow orthogonality condition. □

4.5 Discounted belief prices and mispricing drift

Definition 4.5 (Discounted belief prices)**.** For the market maker and for each trader observer $j \neq i$, define $H^{j,i}$ by

$$\int_t^\infty e^{-\rho(\tau-t)}(D^i_\tau)^\top \delta P_\tau \, d\tau = \int_{-\infty}^t H^{0,i}_t(r)\eta^0(r)\, dr + \sum_{\substack{j=1 \\ j\neq i}}^n \int_{-\infty}^t H^{j,i}_t(r)\eta^j(r)\, dr, \tag{4.5.1}$$

where δw^η solves the forward perturbation equations with initial perturbation η at time t. The belief price $H^{j,i}_t(r)$ is the marginal discounted future execution cost to trader i of shifting observer j's noise-state at r. The infinite-horizon replacement for the finite terminal condition is the transversality condition

$$\lim_{\tau\to\infty} e^{-\rho(\tau-t)} \left(\int_{-\infty}^\tau H^{0,i}_\tau(r)\delta w^0_\tau(r)\, dr + \sum_{\substack{j=1 \\ j\neq i}}^n \int_{-\infty}^\tau H^{j,i}_\tau(r)\delta w^j_\tau(r)\, dr \right) = 0 \tag{4.5.2}$$

for admissible perturbations.

Substituting the total price response, channel by channel, into the discounted future execution cost isolates the recurring bracket

$$H^{k,i}_t(u) = \int_t^\infty e^{-\rho(\tau-t)}(D^i_\tau)^\top p^k(\tau,t,u)\, d\tau, \qquad k \in \mathcal{O}_i. \tag{4.5.3}$$

Equation (4.5.3) has the same row/column convention as Definition 4.5; the endpoint value $H^{k,i}_t(t)$ is the discounted value of the born coordinate.

Contracting $H^{k,i}_t$ against the unresolved drift kernel $\widetilde{C}^k_t$ on old history and against the birth selector at $u = t$ gives the stationary form of the transpose contraction (1.4.9) of Chapter 1. For the market maker,

$$H^{0,i}_t \mathfrak{B}^{0\top} := \int_{-\infty}^t H^{0,i}_t(s)\widetilde{C}^0_t(s)^\top\, ds + H^{0,i}_t(t)E^{Z,\top}. \tag{4.5.4}$$

For a trader observer $j \neq i$,

$$H_t^{j,i}\mathfrak{B}^{j\top} := \int_{-\infty}^{t} H_t^{j,i}(s)\widetilde{C}_t^{j}(s)^\top \, ds + H_t^{j,i}(t)\begin{pmatrix} E^{Z,\top} & E^{j,\top} \end{pmatrix}. \tag{4.5.5}$$

The market maker observes only order flow, so its row is a pure Z block. For trader observers, split $H_t^{j,i}\mathfrak{B}^{j\top} = \left((H_t^{j,i}\mathfrak{B}^{j\top})_Z,\, (H_t^{j,i}\mathfrak{B}^{j\top})_Y\right)$ into order-flow and private-signal blocks. The value of one normalized aggregate order-flow innovation is

$$\beta_t^i := H_t^{0,i}\mathfrak{B}^{0\top} + \sum_{\substack{j=1 \\ j\neq i}}^{n} \left(H_t^{j,i}\mathfrak{B}^{j\top}\right)_Z. \tag{4.5.6}$$

Theorem 4.6 (Discounted backward system for the belief prices)**.** *Under the transversality condition in Definition 4.5, the market maker's belief price solves*

$$(\rho - \partial_t)H_t^{0,i}(r) = \underbrace{(D_t^i)^\top \Sigma_V^{1/2} E^V}_{\text{price move}} - \underbrace{H_t^{0,i}\mathfrak{B}^{0\top}\Sigma_Z^{-1/2} D_{W,t}^{\text{tot}}(r)}_{Z\text{ self-correction}}. \tag{4.5.7}$$

For every trader observer $j = 1, \ldots, n$, $j \neq i$,

$$(\rho - \partial_t)H_t^{j,i}(r) = \underbrace{\beta_t^i \Sigma_Z^{-1/2} D_t^j(r)}_{\text{active demand}} - \underbrace{\left(H_t^{j,i}\mathfrak{B}^{j\top}\right)_Z \Sigma_Z^{-1/2} D_{W,t}^{\text{tot}}(r)}_{Z\text{ self-correction}} - \underbrace{\left(H_t^{j,i}\mathfrak{B}^{j\top}\right)_Y C^{Y,j}(t,r)}_{Y\text{ self-correction}}. \tag{4.5.8}$$

There is no equation for $H^{i,i}$, because trader i knows its own order, so the order does not move trader i's beliefs.

No $\widetilde{C}^0$ appears in the market-maker source, because the seed (4.3.18) already carries it. An opponent's order does not move the price directly; it enters the common Z observation, the market maker and the other traders update on it, and the market maker's update moves the price.

Remark 4.7 (The filter subtracts C, not another $\widetilde{C}$). The forward filter term has the form

$$\widetilde{C}^{j\top}\left[\cdots-\int C^j\delta w^j\right].$$

$H^{j,i}\mathfrak{B}^{j\top}$ already absorbs the factor $\widetilde{C}^j$. The remaining kernel is the observation drift C^j, which is why the self-correction terms in Theorem 4.6 contain $\Sigma_Z^{-1/2}D_W^{\mathrm{tot}}$ and $C^{Y,j}$.

A marginal spike inserts the old-history seeds φ_i^j and the birth seeds ψ_i^j into the observer blocks, so the first-order condition is

$$\begin{aligned}(M_t^i)^\top = \mathbb{E}\Bigg[&\int_{-\infty}^t H_t^{0,i}(r)\varphi_i^0(t,r)\,dr + H_t^{0,i}(t)\psi_i^0(t)\\ &+\sum_{\substack{j=1\\ j\neq i}}^n\left(\int_{-\infty}^t H_t^{j,i}(r)\varphi_i^j(t,r)\,dr + H_t^{j,i}(t)\psi_i^j(t)\right)\Bigg|\,\mathcal{F}_t^i\Bigg].\end{aligned}\tag{4.5.9}$$

Differentiating the stationary kernel representation and using (4.3.12),

$$dM_t^i = \left\{\int_{-\infty}^t \Sigma_V^{1/2}E^V\int_{-\infty}^t\left(\widetilde{C}_t^i(u)^\top\widetilde{C}_t^i(s)-\widetilde{C}_t^0(u)^\top\widetilde{C}_t^0(s)\right)du\,dW_s\right\}dt + dM_t. \tag{4.5.10}$$

The martingale M collects the increment at the boundary $s=t$. The other boundary term of the inner integral, at $u=t$, would pair the contemporaneous observation noise with $\Sigma_V^{1/2}E^V$; it vanishes because the selectors are orthogonal, $E^VE^{Z\top}=E^VE^{i\top}=0$. The term in which $\widetilde{C}^i$ multiplies trader i's own innovation vanishes after conditioning on $\mathcal{F}_t^i$, because $\widetilde{C}^i$ is the innovation gain of trader i's own filter, so its old-history component is orthogonal to trader i's information at time t. Hence

$$\mathrm{drift}_{\mathcal{F}_t^i}(M^i) = -\int_{-\infty}^t K_t^0(s)\,d\widehat{W}_t^i(s),\qquad K_t^0(s):=\Sigma_V^{1/2}E^V\int_{-\infty}^t\widetilde{C}_t^0(u)^\top\widetilde{C}_t^0(s)\,du. \tag{4.5.11}$$

Lemma 4.8 (The predictable drift lies in the price-impact range)**.** *For every old-history* *coordinate* $s < t$*,*

$$K_t^0(s) = \Lambda(t)\Sigma_Z^{1/2}\widetilde{C}_t^0(s). \tag{4.5.12}$$

Consequently the predictable drift in (4.5.11) *lies in* Range$(\Lambda(t))$*, with raw-demand preimage* $\Sigma_Z^{1/2}\widetilde{C}_t^0(s)$*.*

Proof. Substitute the definition of $\Lambda(t)$ in (4.3.17) and use $\Sigma_Z^{-1/2}\Sigma_Z^{1/2}$ as the identity on the traded order-flow support. □

Proof of Theorem 4.6. Let

$$\delta D_t := \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{t} D_t^j(u)\delta w_t^j(u)\,du,$$

and

$$\mathcal{P}_t(H,\delta w) := \int_{-\infty}^{t} H_t^{0,i}(r)\delta w_t^0(r)\,dr + \sum_{\substack{j=1 \\ j\neq i}}^{n} \int_{-\infty}^{t} H_t^{j,i}(r)\delta w_t^j(r)\,dr.$$

On old history coordinates $r < t$, differentiating $\mathcal{P}_t$ gives the interior terms and the contribution of the coordinates born at the current date

$$H_t^{0,i}(t)\delta w_t^0(t) + \sum_{\substack{j=1 \\ j\neq i}}^{n} H_t^{j,i}(t)\delta w_t^j(t).$$

After substituting the old-coordinate equations (4.4.2)–(4.4.3) and the birth conditions (4.4.4)–(4.4.5), the value assigned to an observation residual is the full contraction (4.5.4)–(4.5.5). The active-demand part contributes

$$H_t^{0,i}\mathfrak{B}^{0\top}\Sigma_Z^{-1/2}\delta D_t + \sum_{\substack{j=1 \\ j\neq i}}^{n} \left(H_t^{j,i}\mathfrak{B}^{j\top}\right)_Z \Sigma_Z^{-1/2}\delta D_t = \beta_t^i\Sigma_Z^{-1/2}\delta D_t, \tag{4.5.13}$$

and the self-correction part contributes

$$- H_t^{0,i}\mathfrak{B}^{0\top}\int_{-\infty}^{t} \Sigma_Z^{-1/2} D_{W,t}^{\mathrm{tot}}(u)\delta w_t^0(u)\,du$$

$$-\sum_{\substack{j=1\\j\neq i}}^{n}\left[(H_t^{j,i}\mathfrak{B}^{j\top})_Z\int_{-\infty}^{t}\Sigma_Z^{-1/2}D_{W,t}^{\text{tot}}(u)\delta w_t^j(u)\,du+(H_t^{j,i}\mathfrak{B}^{j\top})_Y\int_{-\infty}^{t}C_t^{Y,j}(u)\delta w_t^j(u)\,du\right].$$

(4.5.14)

Equations (4.5.7)–(4.5.8) cancel these terms. The remaining terms are

$$\frac{d}{dt}\mathcal{P}_t(H,\delta w)=\rho\mathcal{P}_t(H,\delta w)-(D_t^i)^\top\int_{-\infty}^{t}\Sigma_V^{1/2}E^V\delta w_t^0(r)\,dr=\rho\mathcal{P}_t(H,\delta w)-(D_t^i)^\top\delta P_t. \tag{4.5.15}$$

Equivalently,

$$\frac{d}{dt}\left(e^{-\rho(t-s)}\mathcal{P}_t(H,\delta w)\right)=-e^{-\rho(t-s)}(D_t^i)^\top\delta P_t$$

for any fixed anchor s. Integrating from s to ∞ and using transversality gives (4.5.1). Since (4.5.1) determines H uniquely as a functional on seeds, this solution is the belief price. □

4.6 Moving the spike date

Dictionary from Chapter 3. The analysis runs in that chapter's stationary lag coordinates, and one convention differs. Chapter 3 splits each stationary object into a deterministic mean and a deterministic kernel against past shocks, as in (3.2.4); here the random object is kept with its calendar subscript. So $h_t^{k,i}(b)$ of (4.6.3) is that chapter's random belief price $H_t^{k,i}(b)$ in age coordinates, not its deterministic kernel $h^{k,i}(b,\ell)$. The unresolved innovation kernels $\tilde{c}^0,\tilde{c}^j$ are its $\tilde{c}^{k,Z},\tilde{c}^{k,Y}$, and the total-policy kernel $D_{W,t}^{\text{tot}}(u)$ is the sum over traders of its control kernels d_W^m at age $b=t-u$.

The raw-spike condition (4.2.12) holds at every spike date, and moving the spike date moves the seed it creates. Proposition 4.9 records the result; paired with the mispricing drift (4.5.11), it gives the first-order condition of Section 4.7.

Write the raw-spike FOC (4.2.12) as a pairing of belief prices with the seed:

$$z^\top M_t^i = \mathbb{E}\left[S_t^i(z)\,\middle|\,\mathcal{F}_t^i\right], \tag{4.6.1}$$

where $h_t^{k,i}(b)$ is the belief price of observer k at shock age b, and $\phi_i^k(b)$ and ψ_i^k are the old-history and birth seeds trader i's spike leaves in observer k's filter, all three in the stationary lag coordinates set out in (4.6.3)–(4.6.4) below. In those coordinates,

$$S_t^i(z) := \sum_{k\in\mathcal{O}_i}\left(\int_0^\infty h_t^{k,i}(b)\phi_i^k(b)z\,db + h_t^{k,i}(0)\psi_i^k z\right). \tag{4.6.2}$$

To differentiate this identity in the spike date, use stationary lag coordinates. In a stationary solution, the old-history derivative of the seed is a derivative in the age of the shock. Set $b = t-u$, the coordinate age of the lag convention (3.1.5) of Chapter 3, and write

$$H_t^{k,i}(t-b) = h_t^{k,i}(b), \qquad \varphi_i^k(t,t-b) = \phi_i^k(b), \qquad k\in\mathcal{O}_i. \tag{4.6.3}$$

The old-history seed densities and birth matrices are

$$\phi_i^0(b) = \tilde{c}^0(b)^\top\Sigma_Z^{-1/2}, \qquad \phi_i^j(b) = \tilde{c}^j(b)^\top\chi^j\Sigma_Z^{-1/2}, \quad \psi_i^0 = \psi_i^j = E^{Z,\top}\Sigma_Z^{-1/2}. \tag{4.6.4}$$

For old-history coefficients,

$$C_t^0(t-b) = c^0(b), \qquad C_t^j(t-b) = c^j(b), \qquad \widetilde{C}_t^k(t-b) = \tilde{c}^k(b), \tag{4.6.5}$$

where

$$c^0(b) = \Sigma_Z^{-1/2}d_W^{\mathrm{tot}}(b), \qquad c^j(b) = \begin{pmatrix}\Sigma_Z^{-1/2}d_W^{\mathrm{tot}}(b)\\ c^{Y,j}(b)\end{pmatrix}, \quad j\neq i. \tag{4.6.6}$$

Here $d_W^{\mathrm{tot}}(b) = D_{W,t}^{\mathrm{tot}}(t-b)$, $d^j(b) = D_t^j(t-b)$, and $c^{Y,j}(b) = C_t^{Y,j}(t-b)$. Each identity below holds as an $\mathcal{F}_t^i$-conditional expectation; read it as holding for the mean and for every coefficient of $d\widehat{W}_t^i(s)$ in the representation (4.7.3). In lag coordinates the adjoint equations of Theorem 4.6 read $(\rho - \partial_t)H_t^{k,i}(t-b) = s_t^{k,i}(b)$, with source terms

$$s_t^{0,i}(b) := (D_t^i)^\top \Sigma_V^{1/2} E^V - H_t^{0,i} \mathfrak{B}^{0\top} c^0(b), \tag{4.6.7}$$

for the market-maker block, and

$$s_t^{j,i}(b) := \beta_t^i \Sigma_Z^{-1/2} d^j(b) - H_t^{j,i} \mathfrak{B}^{j\top} c^j(b), \qquad j \neq i, \tag{4.6.8}$$

the second term of each being the drift the observer's filter already predicted, in the stacked form (4.6.6).

Moving the spike date also moves the seed. Its density transports in age, its support gains the point $b = 0$, and the birth atom ages. At a fixed belief price the three contribute

$$\sum_{k \in \mathcal{O}_i} \Big(\int_0^\infty h_t^{k,i}(b)\, \partial_b \phi_i^k(b)\, z\, db + h_t^{k,i}(0) \phi_i^k(0) z - \partial_b h_t^{k,i}(0)\, \psi_i^k z \Big). \tag{4.6.9}$$

Proposition 4.9 (Drift of the spike condition as the spike date moves)**.** *For every raw direction z,*

$$\text{drift}\ \mathbb{E}\left[S_t^i(z) \,\middle|\, \mathcal{F}_t^i \right] = \rho\, z^\top M_t^i - z^\top \Lambda(t)^\top D_t^i + z^\top \mathcal{V}_t^i, \tag{4.6.10}$$

where the wedge, the market form of the information wedge v^i of (3.2.6), is

$$\begin{aligned} z^\top \mathcal{V}_t^i = \mathbb{E}\Bigg[&\sum_{k \in \mathcal{O}_i} H_t^{k,i} \mathfrak{B}^{k\top} \int_0^\infty c^k(b) \phi_i^k(b) z\, db - \beta_t^i \Sigma_Z^{-1/2} R^i(t) z \\ &+ \sum_{k \in \mathcal{O}_i} \Big(\int_0^\infty h_t^{k,i}(b)\, \partial_b \phi_i^k(b)\, z\, db + h_t^{k,i}(0) \phi_i^k(0) z - \partial_b h_t^{k,i}(0)\, \psi_i^k z \Big) \Bigg| \mathcal{F}_t^i \Bigg]. \end{aligned} \tag{4.6.11}$$

The first line is the value of the drift each observer's filter already predicted and subtracts from the seed less the value of the immediate opponent response; the second is the seed's own motion (4.6.9).

Proof. $S_t^i(z) = \langle H_t, \eta_t \rangle$ pairs the belief prices with the seed measure (4.3.22)–(4.3.23), and both factors move with t. Pathwise, Theorem 4.6 gives $\partial_t H_t^{k,i}(u) = \rho H_t^{k,i}(u) - s_t^{k,i}(u)$ at every fixed coordinate $u \le t$; paired with the seed this is $\rho S_t^i(z) - \langle s, \eta_t \rangle$. By (4.6.7)–(4.6.8), (4.3.17) and (4.3.24),

$$\langle s, \eta_t \rangle = z^\top \Lambda(t)^\top D_t^i + \beta_t^i \Sigma_Z^{-1/2} R^i(t) z - \sum_{k \in \mathcal{O}_i} H_t^{k,i} \mathfrak{B}^{k\top} \int_0^\infty c^k(b) \phi_i^k(b) z \, db,$$

the atom contributing nothing, since $E^V E^{Z,\top} = 0$ and the policy kernels have zero trace at age 0. The seed's motion at a fixed belief price is (4.6.9), with $\partial_u H_t^{k,i}(u)|_{u=t} = -\partial_b h_t^{k,i}(0)$. Summing, $dS_t^i(z)/dt = \rho S_t^i(z) - \langle s, \eta_t \rangle +$ (4.6.9) pathwise. Conditioning, $\mathbb{E}[S_{t+dt} \mid \mathcal{F}_{t+dt}^i] - \mathbb{E}[S_t \mid \mathcal{F}_t^i] = \mathbb{E}[S_{t+dt} - S_t \mid \mathcal{F}_{t+dt}^i] + (\mathbb{E}[S_t \mid \mathcal{F}_{t+dt}^i] - \mathbb{E}[S_t \mid \mathcal{F}_t^i])$, and the second bracket is an $\mathcal{F}^i$-martingale increment, so the $\mathcal{F}_t^i$-drift of $\mathbb{E}[S_t^i(z) \mid \mathcal{F}_t^i]$ is the conditional expectation of the pathwise derivative. Finally $\mathbb{E}[S_t^i(z) \mid \mathcal{F}_t^i] = z^\top M_t^i$ by (4.6.1), which turns ρS into $\rho z^\top M_t^i$. □

4.7 Stationary best response and equilibrium

A trader's order, priced at impact, splits into a trade on the information gap, an impatience term, and the wedge. The information gap is the trader's estimate of the unresolved total order flow.

Theorem 4.10 (Stationary best response to a stationary profile). *Fix a stationary candidate profile of the other traders' primitive-coordinate policy kernels and associated filter kernels. Under Assumption 4.1, trader i's best response satisfies the level condition (4.5.9) and its $\mathcal{F}_t^i$-drift form*

$$\Lambda(t)^\top D_t^i - \Lambda(t) \int_{-\infty}^t \Sigma_Z^{1/2} \widetilde{C}_t^0(s) \, d\widehat{W}_t^i(s) - \rho M_t^i = \mathcal{V}_t^i, \tag{4.7.1}$$

with $\mathcal{V}_t^i$ as in (4.6.11). The minimum-norm primitive-coordinate representative is

$$D_t^i = (\Lambda(t)^\top)^\dagger \Big[\Lambda(t) \int_{-\infty}^{t} \Sigma_Z^{1/2} \widetilde{C}_t^0(s)\, d\widehat{W}_t^i(s) + \rho M_t^i + \mathcal{V}_t^i\Big]. \tag{4.7.2}$$

For a single asset, or whenever $\Lambda(t)$ is symmetric, this is $D_t^i = \int_{-\infty}^{t} \Sigma_Z^{1/2} \widetilde{C}_t^0(s)\, d\widehat{W}_t^i(s) + \Lambda(t)^\dagger \big(\rho M_t^i + \mathcal{V}_t^i\big)$. The three terms are the information-gap trade, impatience, and the wedge.

Policy kernels. Write $\mathcal{V}_t^i$ and the mispricing in trader i's noise-state coordinates,

$$\mathcal{V}_t^i = \overline{\mathcal{V}}_t^i + \int_{-\infty}^{t} \mathcal{V}_t^i(s)\, d\widehat{W}_t^i(s), \qquad M_t^i = \bar{M}_t^i + \int_{-\infty}^{t} M_t^i(s)\, d\widehat{W}_t^i(s). \tag{4.7.3}$$

If

$$D_t^i = \bar{D}_t^i + \int_{-\infty}^{t} D_t^i(s)\, d\widehat{W}_t^i(s), \tag{4.7.4}$$

then the primitive mean and noise-state policy kernel are

$$\bar{D}_t^i = (\Lambda(t)^\top)^\dagger \Big[\rho \bar{M}_t^i + \overline{\mathcal{V}}_t^i\Big], \qquad D_t^i(s) = (\Lambda(t)^\top)^\dagger \Big[\Lambda(t) \Sigma_Z^{1/2} \widetilde{C}_t^0(s) + \rho M_t^i(s) + \mathcal{V}_t^i(s)\Big]. \tag{4.7.5}$$

A stationary profile gives D_W^{tot}, and with it the observation kernels C^0, C^j and unresolved kernels $\widetilde{C}^0, \widetilde{C}^j$. The adjoint equations determine $H^{0,i}$ and $H^{j,i}$, and these give $\mathcal{V}_t^i$. $\widetilde{C}$, H, and $\mathcal{V}^i$ all depend on the stationary order-flow law generated by the profile, so (4.7.2) is a fixed-point equation for the profile.

Corollary 4.11 (Stationary equilibrium condition)**.** *If a stationary admissible profile satisfies Theorem 4.10 for every trader i, with the first-order conditions read as range equations as in Assumption 4.1, then the profile is stationary for every trader's deviation problem; if in addition each trader's deviation quadratic form is positive semidefinite on the admissible class, the profile is a stationary equilibrium within the noise-state-linear class. Conversely, any stationary noise-state-linear equilibrium satisfying the standing* *assumptions obeys the system in Theorem 4.10 for each trader.*

The corollary is the market analogue of Proposition 3.1 of Chapter 3. For $\rho > 0$ the drift form (4.7.1) implies the level condition, since their difference solves $dN = \rho N\,dt + dm$ with m a martingale and stays bounded only if it vanishes; at $\rho = 0$ the level condition is imposed on the means. Section 4.8 solves the level condition directly and checks the positivity condition on the discrete class.

Proof of Theorem 4.10. The level condition is (4.2.12). For its drift form, take the $\mathcal{F}_t^i$-predictable drift of (4.6.1). The drift of the left side is

$$z^\top \operatorname{drift}_{\mathcal{F}_t^i}(M_t^i) = -z^\top \Lambda(t) \int_{-\infty}^{t} \Sigma_Z^{1/2} \widetilde{C}_t^0(s)\, d\widehat{W}_t^i(s), \tag{4.7.6}$$

by (4.5.11) and (4.5.12). The drift of the right side is (4.6.10). Therefore, for every raw direction z,

$$-z^\top \Lambda(t) \int_{-\infty}^{t} \Sigma_Z^{1/2} \widetilde{C}_t^0(s)\, d\widehat{W}_t^i(s) = \rho\, z^\top M_t^i - z^\top \Lambda(t)^\top D_t^i + z^\top \mathcal{V}_t^i.$$

Rearranging and using that the identity holds for every z gives (4.7.1). $\square$

4.8 Numerical solutions

4.8.1 Unbiasedness and the regularized objective

In the risk-neutral Kyle–Back benchmark, the local first-order condition (4.7.1) can partly collapse into a condition on the aggregate order-flow law. To see this, suppose the normalized order-flow drift is not already resolved by the market maker, so

$$\widetilde{C}_t^0(s) = C_t^0(s) = \Sigma_Z^{-1/2} D_{W,t}^{\mathrm{tot}}(s). \tag{4.8.1}$$

Then the raw preimage of the local information-gap drift is

$$\int_{-\infty}^{t} \Sigma_Z^{1/2} \widetilde{C}_t^0(s)\, d\widehat{W}_t^i(s) = \int_{-\infty}^{t} D_{W,t}^{\mathrm{tot}}(s)\, d\widehat{W}_t^i(s) = \mathbb{E}\Big[D_t^{\mathrm{tot}} - \bar{D}_t^{\mathrm{tot}} \mid \mathcal{F}_t^i\Big]. \tag{4.8.2}$$

Since $D_t^{\mathrm{tot}} = D_t^i + \sum_{j\neq i} D_t^j$, the condition can look like a restriction on the opponents' trades, or on total public order flow, rather than a formula that isolates D_t^i pointwise. This gives the continuous-time Kyle–Back unbiasedness phenomenon in noise-state coordinates. The market maker sees public order flow as an innovation martingale, while a better-informed trader can assign a nonzero drift to the same public process. Back's continuous-time equilibrium pins down the pricing rule and the law of total order flow in this way, but the risk-neutral insider has many optimal trading paths [11].

The coupled system is a stationary fixed point in kernels, and the rest of this section solves it numerically. The computations replace the objective (4.2.9) by its regularization

$$J_t^{i,\epsilon}(D^i; D^{-i}) := \mathbb{E}\left[\int_t^{\infty} e^{-\rho(\tau-t)}\Big((D_\tau^i)^\top\Big(\widehat{V}_\tau^i - P_\tau\Big) - \epsilon\,|D_\tau^i|^2\Big)\, d\tau \;\Big|\; \mathcal{F}_t^i\right], \qquad \epsilon > 0, \tag{4.8.3}$$

whose first-order condition is that of Theorem 4.10 with M_t^i replaced by $M_t^i - 2\epsilon D_t^i$. With this cost the system has locally unique strict equilibria for one, two, and three traders. At zero cost the second-order form is only semidefinite. Under permanent linear impact a trader who buys and later sells the same quantity pays nothing [59], so such round trips span a null space of the form and the best response is not unique.

Remark 4.12 (Singular instantaneous impact directions). Zero immediate price impact is not zero impact. An order the market maker does not price today can still teach the other traders and come back through their future trades, so whether it pays depends on the component of $\rho M_t^i + \mathcal{V}_t^i$ along z, the right side of (4.7.1). If a direction has neither instantaneous nor discounted future price response while a trader perceives mispricing in that direction, a finite best response fails. If the future feedback response offsets the current perceived mispricing, a finite stationary solution can exist even when the immediate impact is zero.

4.8.2 Primitives

The stationary system has been solved numerically for one, two, and three traders. The fundamental is a random walk. Each trader observes a signal $dY^i = B_X^{Y,i}(V - P)\,dt + dW^i$ with scalar gain $B_X^{Y,i}$. The trading cost ϵ of the regularized objective (4.8.3) is 10^{-3}, standing in for zero, with $\Sigma_Z = 1$ and unit signal gains; the objective is average cost, $\rho = 0$. The undiscounted monopolist has no stationary solution, so the one-trader kernels are computed at discount rate $\rho = 0.5$. The scalars computed on the lag window of Appendix 4.A (λ, total price profits, and the split between channels) are the same at $\rho = 0$ and at every $\rho > 0$ tried, so the reported values are at $\rho = 0$ throughout (Appendix 4.A). The level first-order condition (4.5.9) is evaluated with the market maker's pricing rule and each opponent's strategy adjusting as in Section 4.4. A second-order condition has also been checked at every reported equilibrium. The exact quadratic form of Chapter 1's finite-deviation expansion is positive definite on the strategy space, with smallest eigenvalue 2ϵ, so no deviation improves on the reported policy.

4.8.3 Profit accounts

Nothing is liquidated in a stationary market, so profit needs a convention. A trade pushes the quote as it fills and pays the average price, half its own impact: orders fill by *walking the book*. Write $Q_t^i = \int^t D_s^i\,ds$ for trader i's position. *Execution profits* are made at the fill, the gap between the new quote and the price paid. *Position profits* are made afterward, from quote changes on shares already held:

$$\text{execution}^i = \tfrac{1}{2}\,\mathbb{E}[\,dQ_t^i\,dP_t]/dt, \qquad \text{position}^i = \mathbb{E}[Q_{t-}^i\,dP_t]/dt.$$

Their sum, the price-profit rate, is zero-sum across the market. A finite trading rate pays no execution: it trades $D\,dt$ against its own impact $\lambda D\,dt$, second order. Only the diffusive noise-trader flow carries a first-order term, a gain, since a fill at mid sits below the post-trade mark. On a grid of cell h the charge is $\frac{1}{2}\lambda h\,\mathbb{E}[D^2]$, so the zero is a continuum statement. It is tied to the trading cost: as $\epsilon \to 0$ the age-zero loading grows like $\epsilon^{-1/2}$ toward an impulse, which would pay.

The *book gap* $\mathbb{E}[Q_t^i(V_t - P_t)]$ is the average undervaluation of the shares held; a book can be infinite with a finite gap: the mispricing has short memory, so the covariance truncates itself. Writing $G_t = Q_t^i(V_t - P_t)$,

$$\Delta G_t = \underbrace{\Delta Q_t^i\,(V_t - P_t)}_{\text{gap purchases}} - \underbrace{Q_{t-}^i\,\Delta P_t}_{\text{conversion into position profits}} + \underbrace{Q_{t-}^i\,\Delta V_t}_{\text{news, mean zero}} :$$

buying below fundamental value adds to the gap, a quote move toward the fundamental converts it into position profits, and news adds nothing on average. In the stationary state the gap is constant, so gap purchases equal position profits. Churning earns nothing: the execution gain on a buy is lost in position when the sell pushes the quote back. The market maker makes no position profits, since its inventory cannot forecast its own innovations, and carries no gap, since its pricing error is orthogonal to what it knows. The objective (4.8.3) is the price-profit rate net of the trading cost.

4.8.4 Symmetric markets

Competition speeds up trading and revelation [56] (Figure 4.3). A single trader at discount rate 0.5 trades its current signal at age zero and its filtered signal afterwards. Its loading on the fundamental peaks at 0.83 near age 0.2 and is below 0.02 by age four, and half of its profit accrues by age 0.42, its half-profit age. With three traders every kernel is negligible beyond age two, each trader's largest loading is on its own raw signal noise, concentrated near age zero, and half the profit accrues by age 0.28. Price impact is $\lambda = 1.000$ with one,

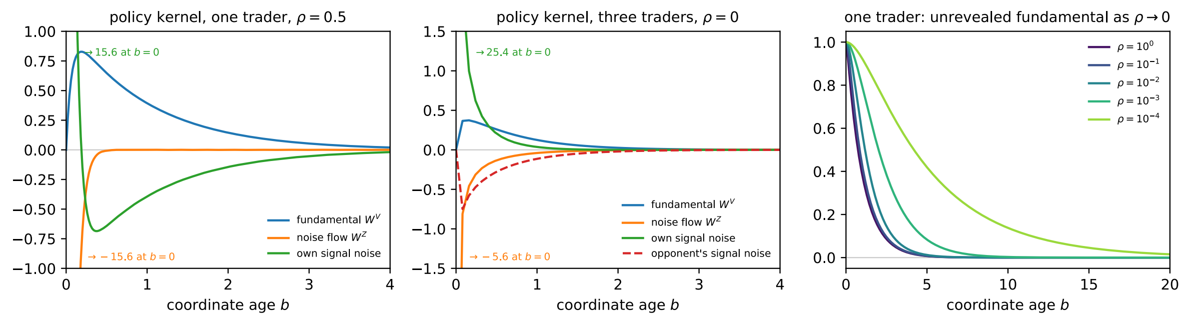

Figure 4.3: Equilibria of the stationary system at trading cost 10^{-3}. Left: one trader's policy kernel at discount rate $\rho = 0.5$; center: the policy kernel of one of three identical traders at $\rho = 0$; right: the single trader's unrevealed fundamental, the share of a fundamental shock not yet in the price by shock age, at discount rates from 1 to 10^{-4}. As $\rho \to 0$ the trader trades later and later and the curves do not converge: the undiscounted monopolist has no stationary solution, while the competitive cases exist at $\rho = 0$. The loadings on own signal noise and on noise-trader flow leave the frame near $b = 0$; their values there are marked at the exit points. This is the impulse of the profit-accounts subsection, whose loading grows like $\epsilon^{-1/2}$ as the cost vanishes.

	price profits	execution	position	book gap, $N = 1, 2, 3$
traders, together	$+\lambda\Sigma_Z$	0	$+\lambda\Sigma_Z$	0.06, 0.36, 0.29
noise traders	$-\lambda\Sigma_Z/2$	$+\lambda\Sigma_Z/2$	$-\lambda\Sigma_Z$	-0.06, -0.36, -0.29
market maker	$-\lambda\Sigma_Z/2$	$-\lambda\Sigma_Z/2$	0	0

Table 4.1: Profit accounts of the symmetric markets ($\rho = 0$, $\epsilon = 10^{-3}$). Rates are flows per unit time; the book gap is a level. Each of the N traders takes $\lambda\Sigma_Z/N$ of the traders' row.

two, or three traders. Per-trader price profits, gross of the trading cost, fall from 1.000 to 0.333, and their total is $\lambda\Sigma_Z$ in every case. Competition divides the noise traders' losses $\lambda\Sigma_Z$ without changing them.

The unrevealed fundamental, the share of a fundamental shock not yet in the price by the shock's age, is 0.39 at age one with one trader and 0.25 with three. A monopolist with a random-walk fundamental has no reason to trade its information by any particular date; its horizon comes from discounting, or, in Kyle's model, from the terminal date. With an opponent, information not traded now is traded by the opponent, so the two- and three-trader solutions exist at $\rho = 0$ and reveal the shock within two ages.

Symmetry forces the accounts (Table 4.1): each trader takes $\lambda\Sigma_Z/N$, all of it position, and the noise traders take half of what they pay back as execution. The monopolist's

	$\gamma = 0.25$	0.5	1	2	4
half-profit age	1.36	0.68	0.34	0.18	0.10
fundamental book	+0.81	+0.40	+0.20	+0.10	+0.05
noise-trader book	+0.46	+0.23	+0.12	+0.06	+0.03
own-noise book	−1.37	−0.68	−0.34	−0.17	−0.08
opponent-noise book	+0.82	+0.41	+0.20	+0.10	+0.05
net book gap	0.72	0.36	0.18	0.09	0.05

Table 4.2: One of two symmetric traders' book gap by channel as the common signal gain varies ($\rho = 0$, $\epsilon = 10^{-3}$, window scaled with the trading speed). Every channel and the net scale together with the half-profit age, as $1/\gamma$. The competitors' gap is the edge withheld from opponents. Per-trader price profits are 0.50 at every gain, and the gap, unlike the monopolist's $\sqrt{2\epsilon/\rho}$, does not shrink with the trading cost.

gap is the trading-cost residual, $\sqrt{2\epsilon/\rho}$ to leading order: its undervalued fundamental positions, worth +0.50 at every cost, net exactly at zero cost against the overvalued positions it briefly puts on by trading its own signal noise. The competitors' gaps do not shrink with the cost. The monopolist buys no gap on its own signal noise. A competitor trades the current signal immediately, buys 0.34 of overvaluation on its own noise, and recovers 0.20 on the others'. Table 4.2 varies the common signal gain of two traders.

Competition also changes whose flow pays the trader: the monopolist's position profits come entirely from its own flow moving the price, while with opponents each trader earns 0.72 from its own flow and loses 0.22 to its opponents' (spectral values at $\epsilon = 10^{-3}$). On the shared channels, the fundamental and the noise-trader flow, an opponent's flow pays a book exactly as the trader's own does (Table 4.3). The tax falls on the private channels: a trader marks up its own noise-driven positions with its later orders, and its opponent, which does not share the noise, trades against exactly those price moves. Opponents charge a trader for the errors its noise puts into the price; the monopolist escapes because its own-noise purchases are zero.

4.8.5 Two assets

The two-asset equilibria, each with two traders at trading cost 2×10^{-3} and discount rate $\rho = 0.1$, are computed in three configurations: signal gains $(1, 1)$ for both traders

	own flow	opponents' flow
fundamental	+0.125	+0.125
noise-trader flow	+0.125	+0.125
own signal noise	+0.235	−0.441
opponent's signal noise	+0.235	−0.029
total	+0.720	−0.220

Table 4.3: Position profits of one of two symmetric traders, by shock channel and by whose flow moves the price (spectral solution, $\epsilon = 10^{-3}$). The shared channels pay from either flow equally; the opponents' tax is confined to the private noise channels.

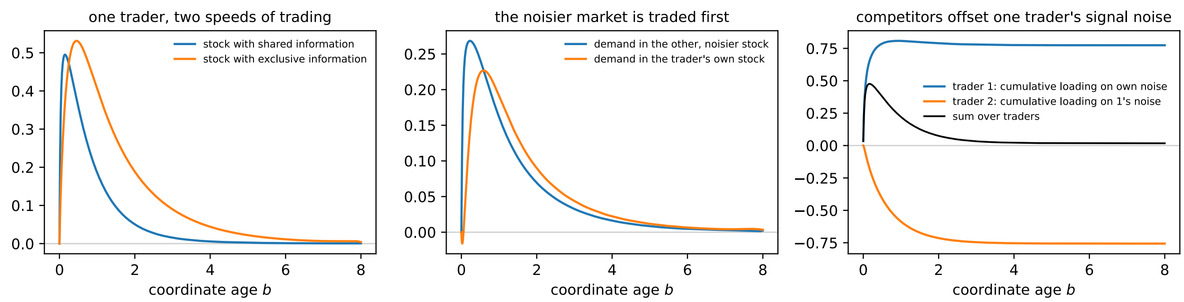

Figure 4.4: Two-asset equilibria at trading cost 2×10^{-3}. Left: one trader, two market structures, two speeds of trading. Center: the trade in the market with more noise-trader flow comes first. Right: competitors' trades offset one trader's signal noise, so the sum over traders returns to zero.

in the symmetric case with $\Sigma_V = I$; gains $(1, 1)$ and $(1, 0)$ in the generalist–specialist market; gains $(1, 0)$ and $(0, 1)$ with noise-trader variance $\mathrm{diag}(1.8, 0.2)$ in the specialist case; fundamental correlation 0.6 in the last two. A trader holding exclusive information in one market and shared information in the other trades them at different speeds, and a specialist whose own market has little noise-trader flow trades its information first in the other market where that flow is larger (Figure 4.4). The instantaneous impact Λ is a matrix across assets: in the symmetric case it is 1.001 on the diagonal with no cross-impact, while in the specialist case the off-diagonal entries reach 0.42, so an order in one asset moves the other's price.

4.8.6 Deviation impact

The impact of a deviation is not the pricing kernel, the market maker's permanent price response to a unit of unexpected order flow. Figure 4.5 splits the price response to a unit

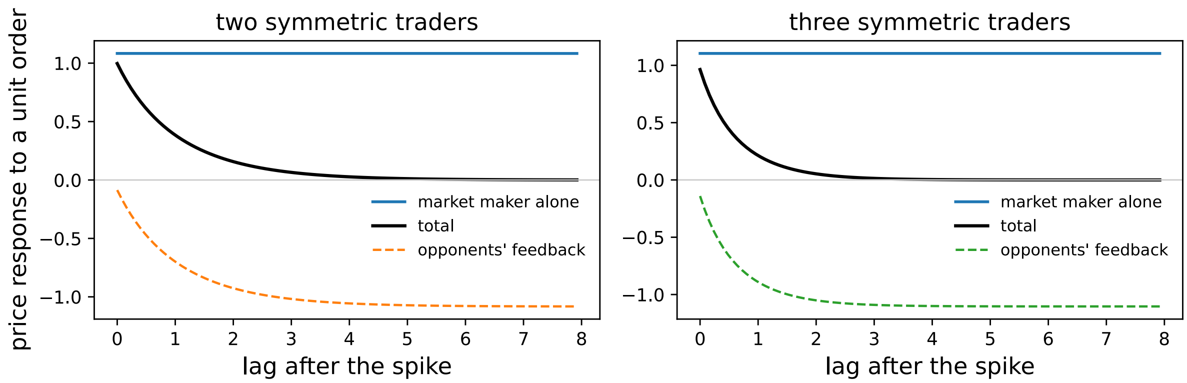

Figure 4.5: Price response to a unit order by one trader at age zero, decomposed by re-running the forward response with the opponents' policies switched off ($\epsilon = 0.1$, $\rho = 0.1$; the large cost smooths the kernels, and the decomposition is about shape). The market maker alone moves the price permanently; the opponents trade against the order and unwind the move, so the total response is transitory even though the pricing kernel is not.

order at age zero, the kernel $p_i(\tau, t; z)$ of (4.4.14), into the market maker's own update and the opponents' feedback, by re-running the forward response of Section 4.4 with the opponents' policies switched off (two and three symmetric traders, $\epsilon = 0.1$, $\rho = 0.1$). The market maker alone moves the price permanently, by 1.08 per unit. The opponents see order flow that is not theirs and no news in their own signals, so they sell against it. On impact they absorb 8 percent of the move with one opponent and 13 percent with two. Afterwards they unwind all of it, half by age 0.72 and 0.48 and everything by age eight. The instantaneous impact is λ, but the discounted execution cost in the first-order condition, $\int_t^\infty e^{-\rho(\tau-t)} D_\tau\, p_i(\tau, t; z)\, d\tau$, is a fraction of the permanent one, and the fraction falls with competition. In the generalist–specialist market a unit order in the shared stock is absorbed 14 percent on impact and unwound on the same clock.

4.8.7 Asymmetric noise

Three traders with correlated fundamentals and asymmetric noise-trader flow separate two roles that the symmetric cases merge. Table 4.4 lists the parameters; Tables 4.5 and 4.6 report price profits and impact. Price discovery does not depend on ζ: the unrevealed fundamental at ages one and two is 0.20 and 0.04, at every ζ. Price profits do. As stock

fundamentals	two, correlation 0.6
signal gains (stock one, stock two)	trader one $(1.5, 0.5)$, trader two $(0.5, 1.5)$, generalist $(1, 1)$
noise-trader variances	$\mathrm{diag}(1+\zeta, 1-\zeta)$, $\zeta = 0, 0.5, 0.8$
discount rate, trading cost	$\rho = 0$, $\epsilon = 2 \times 10^{-3}$ (halving it moves the numbers by about two percent)
window, collocation	$L = 16$, 64 Chebyshev points

Table 4.4: Parameters of the three-trader asymmetric-noise market.

	$\zeta = 0$	$\zeta = 0.5$	$\zeta = 0.8$
noisy-stock leader	0.66 (0.57)	0.76 (0.70)	0.81 (0.77)
quiet-stock leader	0.66 (0.09)	0.52 (0.12)	0.39 (0.14)
generalist	0.53 (0.26)	0.51 (0.33)	0.48 (0.36)

Table 4.5: Price profits net of the trading cost in the three-trader market with noise-trader variances $\mathrm{diag}(1+\zeta, 1-\zeta)$; in parentheses, the part earned in the noisy stock.

two's noise-trader flow shrinks, the quiet-stock leader's price profits fall from 0.66 to 0.39 and the noisy-stock leader's rises from 0.66 to 0.81; the generalist loses a little, 0.53 to 0.48. These price profits are net of the trading cost, and Table 4.7's totals are gross, 0.02 higher for the heavier traders. The quiet stock's leader takes its information to the noisy stock: at $\zeta = 0.8$ it trades its own factor in the noisy stock with integrated loading 0.19 against 0.28 in its own, up from 0.07 against 0.64 at $\zeta = 0$. Impact follows the noise: at $\zeta = 0.8$ the own impact is 0.73 in the noisy stock and 2.01 in the quiet one. The cross-impact is asymmetric because the quiet stock's flow is the more informative: an order in the quiet stock moves the noisy stock's price by 0.45, while an order in the noisy stock moves the quiet stock's by 0.31. Every trader runs two clocks: its loading on its own factor peaks by age 0.24, and its hedge against the correlated factor, traded in the opposite direction, peaks between ages 0.6 and 1.

Symmetry no longer forces the accounts; Table 4.7 gives them at $\zeta = 0.8$. On the home books the clock is common: the ratio of book gap to position profits is 0.28–0.32 for each trader's home stock, the generalist, and the noise traders, at every ζ, so the noise asymmetry redistributes who buys undervaluation, not how long the price takes to pay it out. The specialists' small cross-stock books turn over faster, 0.19–0.26. Opponents tax the generalist hardest: the specialists lose a quarter to three tenths of their own-flow position profits

	$\zeta = 0$	$\zeta = 0.5$	$\zeta = 0.8$
own impact, noisy stock	0.94	0.79	0.73
own impact, quiet stock	0.94	1.31	2.01
quiet-stock order on the noisy price	0.31	0.36	0.45
noisy-stock order on the quiet price	0.31	0.30	0.31

Table 4.6: Entries of the matrix Λ in the same market.

	noisy stock execution	noisy stock position (own/opponents')	noisy stock gap	quiet stock execution	quiet stock position (own/opponents')	quiet stock gap	total
noisy-stock leader	0.00	+0.79 (+1.08/ − 0.28)	+0.23	0.00	+0.04 (+0.06/ − 0.02)	+0.01	+0.83
quiet-stock leader	0.00	+0.14 (+0.24/ − 0.10)	+0.03	0.00	+0.25 (+0.34/ − 0.08)	+0.08	+0.40
generalist	0.00	+0.38 (+0.60/ − 0.22)	+0.11	0.00	+0.12 (+0.18/ − 0.07)	+0.04	+0.49
noise traders	+0.66	−1.31	−0.37	+0.20	−0.40	−0.12	−0.86
market maker	−0.66	−0.01	0.00	−0.20	0.00	0.00	−0.86

Table 4.7: Profit accounts in the three-trader market at $\zeta = 0.8$, by stock (spectral solution, $\epsilon = 2 \times 10^{-3}$). Position profits split into the part earned by the trader's own flow and the part its opponents' flow takes back; the total is price profits over both stocks. The gap is a level, not a flow, so it does not enter the total. Rates are flows per unit time. Each column is zero-sum, and the market maker's position profits and book gap are zero stock by stock, up to the trading cost.

to opponents' trades, the generalist more than a third, its errors being visible to both specialists. The quiet stock's leader earns seven tenths as much from its own flow in the noisy stock as at home, up from a fifth at $\zeta = 0$. The noise traders in the quiet stock lose 2.01 in position per unit of their variance against 0.73 in the noisy one, the net impacts Λ_{22} and Λ_{11}.

4.8.8 The three terms of the order

Theorem 4.10 splits each trader's order into the information-gap trade, impatience, and the wedge (Figure 4.6, computed on the lag grid). The information-gap trade carries 92 percent of the order, in L^2 norm over the policy kernels, in every case computed. Under the unbiasedness that the solutions satisfy, it is $\mathbb{E}[D_t^{\text{tot}} \mid \mathcal{F}_t^i]$, so the wedge is the price-impact value of the opponents' expected current orders with the sign reversed, $-\Lambda\,\mathbb{E}[D_t^{-i} \mid \mathcal{F}_t^i]$ (Λ is symmetric in these runs), less impatience. A trader shades its order by what it expects the others to send. The wedge is 13 percent of the order and lives at ages under two; impatience is 4 percent in the $\rho = 0.1$ case and absent at $\rho = 0$.

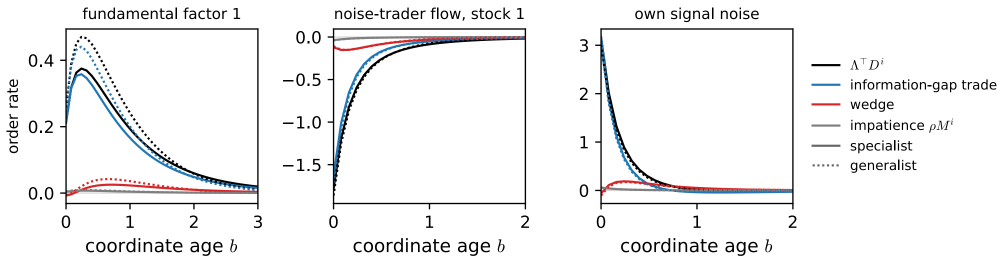

Figure 4.6: The three terms of Theorem 4.10 for each trader's order in stock one, by coordinate age and shock channel, in the generalist–specialist market at $\rho = 0.1$: specialist solid, generalist dotted. The right panel shows each trader's own signal-noise channel.

4.8.9 Numerical identities

The solutions satisfy, up to the size of the trading cost, several identities that the equations do not impose directly. Each trader's policy kernel has zero projection onto the market maker's information (relative norm below 5×10^{-3} at $\epsilon = 10^{-3}$). Summing over traders, total order flow is a Brownian motion in the market maker's own filtration, which extends the unbiasedness that Back [11] proved for a single trader to every case computed here. The pricing kernel is constant across ages to within 0.2 percent, so each unexpected unit of order flow moves the price permanently and by the same amount at every age. The sum over ages of total trader demand's loading on each signal noise is zero. Competitors fully offset one another's noise-driven positions, while a monopolist has no one to offset it and is left with a net position of 10^{-3}, the size of the cost. Total price profits gross of the trading cost equal the noise traders' losses $\lambda\Sigma_Z$ to 2×10^{-4} for one, two, and three traders, and do not depend on the lag truncation L of Appendix 4.A. The net price profits fall short by exactly the cost, 0.006 per trader with two or three traders. As the cost vanishes, each competitor's loading on its own signal noise at age zero grows like $\epsilon^{-1/2}$ while its integral stays fixed, so the loading is concentrated on ages of order $\epsilon^{1/2}$ and the net shortfall falls like $\epsilon^{1/2}$. In the limit a competitor trades its current signal noise as an impulse, and total price profits equal $\lambda\Sigma_Z$ exactly. Computed down to $\epsilon = 5 \times 10^{-5}$, total price profits are 0.998. For one trader λ and the price profits are 1.000, the values of Back's continuous-time equilibrium. Proofs of these identities are open.

4.A Computing the Kyle–Back equilibria

The equilibria of Section 4.8 are fixed points of the joint best-response map over all traders' kernels, computed by Newton's method on the difference between a profile and its best response, with the cost ϵ continued downward from a large value at which the map is well conditioned. Damped best-response iteration fails here. The equilibrium is unstable under it, with the dominant eigenvalue of the best-response map above one. More trading raises price impact, higher impact reduces trading, and less trading lowers impact again, so the iteration oscillates instead of converging. Newton iteration does not depend on this stability and converges to residuals near 10^{-10} in a few steps.

The computations use the average-cost objective ($\rho = 0$) with lags truncated at $L = 8$ as in Chapter 3; the window replaces the transversality of Assumption 4.1, and its effect is measured below. The kernels are represented by their values at 64 Chebyshev points on $[0, L]$ by the spectral collocation of Appendix 3.A. The three-trader scalars are unchanged to four digits at 96 and 128 points. The two-asset equilibria use the same solver in its two-asset form; the three-trader asymmetric-noise market runs with $L = 16$ and $\epsilon = 2 \times 10^{-3}$, and its profit accounts are computed from the node values by Chebyshev quadrature, with positions as spectral antiderivatives. The window must scale with the half-profit age, so the gain ladder of Table 4.2 uses windows up to $L = 32$. The loading on own signal noise carries an impulse of width $\epsilon^{1/2}$ at age zero, and at 64 points the collocation solution rings behind it, a node-to-node alternation of 3×10^{-3} in the tail. At 128 points the alternation is 2×10^{-4}, and the three-trader kernels of Figure 4.3 are drawn from that solution. With two or more traders the fundamental is essentially fully revealed well before age L and the reported quantities do not depend on it. The two- and three-trader kernels agree at $L = 8$ and $L = 16$ to four digits beyond age one. A single undiscounted trader has no such solution. Its kernels scale with L, half of its profit accrues after age $0.19L$, and the unrevealed fundamental falls to zero exactly at L. The truncation plays the part of Kyle's terminal date: the undiscounted computation converges on a

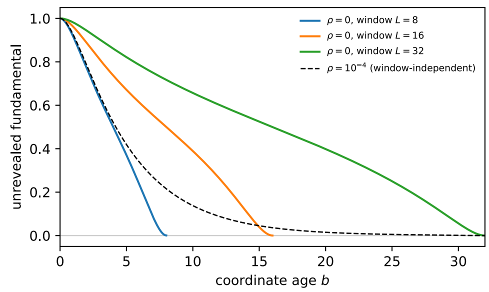

Figure 4.7: The window artifact for one undiscounted trader: the unrevealed fundamental by shock age computed at $\rho = 0$ on windows $L = 8$, 16, and 32 falls linearly to zero exactly at L, so each window gives a different solution. The discounted solution at $\rho = 10^{-4}$ (dashed) is the same on every window.

window only because the window supplies it. Figure 4.7 shows the artifact. With a discount rate $\rho > 0$ the one-trader solution is the same at $L = 8$, 16, and 32. As ρ falls (Figure 4.3, right) the horizon, measured as one over the share revealed by age one, grows like $\rho^{-1/2}$ (1.63, 1.75, 2.37, 4.88, 13.1 at $\rho = 1, 10^{-1}, \ldots, 10^{-4}$). The scalars are unaffected: λ, total price profits, and the channel split are the same at every ρ. The second-order checks are made on the average-cost system and do not rely on the $\rho > 0$ hypothesis of Assumption 4.1. The market maker's belief update is computed as in Appendix 3.A.

Every reported solution has passed the second-order check of Section 4.8 and a comparison between solvers. The spectral solver reproduces the grid solver's extrapolations to zero spacing.

On a lag grid, how price impact is timed within an interval matters more than how fine the spacing is. If the scheme realizes the price impact at the midpoint of the interval, a trader can buy and later sell the same quantity at a profit that comes entirely from the half-interval delay between an order and its own impact. The second-order condition fails below a cost threshold proportional to the spacing. Realizing the impact within the same interval as the order removes that profit. The spectral solver therefore reports the small-cost equilibria.

The solvers and figures of Section 4.8 are the work of Claude Fable 5 [7], working under the author's direction. Interactive versions are at sbabichenko.com/noisestate.

Chapter 5

Local Strategic Information in Symmetric Network Economies

5.1 Local strategic information in a large economy

A common large-economy model places many individually negligible agents around a small set of aggregate prices or population statistics. Mean-field methods make this structure tractable by replacing the rest of the population with the distribution of their states [26, 58, 73]. This is natural when interaction is anonymous and the feedback runs through an average state. It is less natural when economic activity is organized into many separate local environments, each containing only a few strategically important players.

A regional credit market may contain only a few major lenders. A local labor market may be dominated by a few employers. A product niche, hospital system, supply-chain cluster, or dealer network may likewise be small as a strategic environment even though each participant is negligible relative to the national economy. Such players are atomistic globally but oligopolistic locally. Their actions can reveal private information to a small set of strategically consequential observers, so each player may alter its action because of what those observers will infer.

The corresponding large-economy limit is large by *replication*, not by local anonymity. One first solves the finite-player information game inside each local environment and then aggregates its equilibrium outcome across many environments. Aggregating the primitives first can eliminate the local action-to-belief channel that determines behavior inside each market.

This distinction also changes the usual interpretation of local interaction. Centered local shocks may wash out across many markets, but the incentives governing how agents respond to those shocks are repeated. If structurally similar players all attenuate or otherwise distort investment, lending, hiring, or trading because their actions reveal information, the common object is an equilibrium coefficient, not a common private signal. The resulting local distortion can change the conditional cross-sectional mean and the aggregate response to fundamentals as well as the spatial variance.

Figure 1.3 in Chapter 1 shows the equilibrium mean control shifting with signal precision, so information and incentives do not separate. If informational asymmetry is correlated with incentives (sellers of a good being likelier to be experts on it), the correlation shows up in aggregate means.

The symmetry results of Appendix 5.A support this local-to-aggregate architecture. The purchase-order market of Section 5.3 only needs to be solved from the perspective of a representative firm. Chapter 1 removes the hierarchy of forecasts by representing beliefs in noise-state coordinates, the estimated primitive shocks. Graph symmetry reduces the number of unknown kernels: $d_t(w, u)$ takes one value per equivalent position of w (Proposition 5.7), while each player still observes only its own channels and moves its neighbors' posteriors.

The finite-player noise-state equations have several player labels. A policy kernel maps a shock attached to one player into another player's action response; a filter kernel maps that shock into another player's estimate; and a belief price prices the effect of changing another player's posterior. When players occupy the vertices of a homogeneous graph and the model is invariant under graph automorphisms, the relabelings of vertices that preserve the edges, many of these labels describe copies of the same object.

The results of Appendix 5.A concern a *finite* vertex-transitive graph, one whose symmetries can carry any vertex to any other. They show that the best-response map commutes with graph automorphisms, that a unique equilibrium inherits the graph symmetry, and that policy, filtering, and adjoint kernels take equal values at positions that look alike from the reference player's seat (Appendix 5.A). Cycles and finite Cayley graphs, graphs whose vertices form a group, admit displacement coordinates. Distance-transitive graphs, whose symmetries can carry any pair of vertices to any equally distant pair, admit distance-shell coordinates.

Dynamic models on financial networks motivate the construction. A close precedent is Feng et al. [36], who solve linear-quadratic stochastic differential games on a directed chain, building on the directed-chain diffusions of Detering et al. [35]. Related network dynamics include interacting diffusions for inter-bank lending [39], systemic-risk mean-field games [27, 28], and contagion with defaults [61]. Here the noise-state, filtering, and information-wedge objects of Chapter 1 inherit the graph's symmetries. Section 5.3 builds a purchase-order market on the cycle. The strategic effect is signal jamming [40], a shift in mean orders several times the filtering gap.

5.2 From local equilibrium to aggregate behavior

5.2.1 The spatial mean

On a vertex-transitive graph, a local kernel gives every vertex the same row up to a rearrangement of its entries, so constants map to constants. The spatial mean of $\mathcal{K}x$ is therefore the spatial mean of x multiplied by the sum of a single row, $\sum_r k(r)$.

Proposition 5.1 (The spatial mean of a local kernel)**.** *Let* Γ *be a finite group, let*

$x : \Gamma \to \mathbb{R}^p$, and let

$$(\mathcal{K}x)(v) = \sum_{r \in \Gamma} k(r) x(vr),$$

where $k(r) \in \mathbb{R}^{q \times p}$. Define

$$\bar{x} = \frac{1}{|\Gamma|} \sum_{v \in \Gamma} x(v), \qquad \overline{\mathcal{K}x} = \frac{1}{|\Gamma|} \sum_{v \in \Gamma} (\mathcal{K}x)(v).$$

Then

$$\overline{\mathcal{K}x} = \left(\sum_{r \in \Gamma} k(r) \right) \bar{x}. \tag{5.2.1}$$

When Γ is abelian, so that the order of two displacements does not matter, $\sum_r k(r)$ is the spatial Fourier transform of the kernel k evaluated at frequency zero.

Proof. For each fixed r, the map $v \mapsto vr$ permutes Γ, so

$$\frac{1}{|\Gamma|} \sum_{v \in \Gamma} (\mathcal{K}x)(v) = \sum_{r \in \Gamma} k(r) \left(\frac{1}{|\Gamma|} \sum_{v \in \Gamma} x(vr) \right) = \left(\sum_{r \in \Gamma} k(r) \right) \bar{x}.$$

□

On a group, after collecting the nonspatial state, source-time, and observer indices, each linear part of the state, filter, policy, and adjoint equations that treats every vertex alike has this form. Proposition 5.1 separates two roles of the equilibrium kernels. Their deviations from the spatial mean govern dispersion, clustering, and propagation across distance. Their spatial mean governs the response of the cross-sectional mean. The first-order condition of Chapter 1 reads schematically as a direct control response plus an information-wedge correction. On a symmetric graph, both pieces generate local kernels in primitive coordinates, the D_W of Chapter 1, the form Proposition 5.1 needs. Write k^{exo} for the policy kernel of the exogenous-signal equilibrium (Section 5.4.2); the aggregate loading is

$$\sum_r k^{\text{exo}}(r) + \sum_r \Big(k(r) - k^{\text{exo}}(r) \Big).$$

The kernel shift affects only nonaggregate variation in the special case

$$\sum_r \Big(k(r) - k^{\text{exo}}(r)\Big) = 0.$$

There is no reason for this cancellation to hold. Players in the same network position share the same incentive to reveal, conceal, exaggerate, or delay, so their distortions point the same way and add up in the mean rather than cancelling. For example, a common productivity component may generate less average investment because each local firm knows that investment reveals information to a few important competitors.

5.2.2 Replication across local economies

Consider M local economies indexed by m, each containing the same finite graph G but possibly different local shocks. Let

$$\bar{D}_t^{m,*} = \frac{1}{|V|} \sum_{v \in V} D_t^{m,v,*}$$

be the equilibrium average action in local economy m. The macro average is

$$\bar{D}_t^{(M)} = \frac{1}{M} \sum_{m=1}^{M} \bar{D}_t^{m,*}.$$

If, conditional on a common macro history $\mathcal{C}_t$, the local economies are identically distributed and conditionally independent with finite second moments, then the conditional law of large numbers gives

$$\bar{D}_t^{(M)} \xrightarrow[M \to \infty]{\mathbb{P}} \mathbb{E}[\bar{D}_t^{1,*} \mid \mathcal{C}_t]. \tag{5.2.2}$$

The local noises $W^{0,v}$ average out across local economies, leaving the row sum $\sum_r d_{W,t}(r,u)$ of the equilibrium kernels applied to the common coordinates.

The order of operations matters. Schematically,

$$\text{aggregate}\Big(\text{local equilibrium}\Big) \neq \text{equilibrium}\Big(\text{aggregated local primitives}\Big)$$

in general.

5.3 A purchase-order market on the cycle

The model of this section is heavily inspired by Woodford [118]. There, firms set prices on the basis of noisy private signals about nominal spending. Because each firm must also forecast what the others will forecast, prices adjust slowly to a monetary shock and the shock has persistent real effects. This section keeps that structure and changes one thing. The signal about nominal spending is no longer an abstract private observation but the market itself, what a firm can see of its own supplier's quote and its own customer's order. Those signals are other firms' actions, the observed actions of Chapter 1, and they travel on the cycle of Appendix 5.B. Firms $v \in \mathbb{Z}/N\mathbb{Z}$ each buy an input from firm $v-1$, sell to firm $v+1$ and to households, and set prices in advance (Figure 5.2). Every variable is a log deviation from a steady state. If x_t is a price or quantity and $\bar{x}$ its value on the constant path the economy would follow without shocks, the variable in the model is $\log(x_t/\bar{x})$. For small deviations this is the percentage deviation, so a value of 0.01 means one percent above the steady state. Prices and quantities that multiply in levels add in logs, which makes the relations below linear. Figure 5.1 shows one price in both units. The primitives are quadratic, which keeps the game in the class of Chapter 1.

Primitives

Nominal spending q_t, the total money households spend per unit time, is a random walk. Firm v has a productivity a_t^v, the output it gets from a unit of input, and faces a household taste shock η_t^v, a shift in how much households want its product at a given price. Both

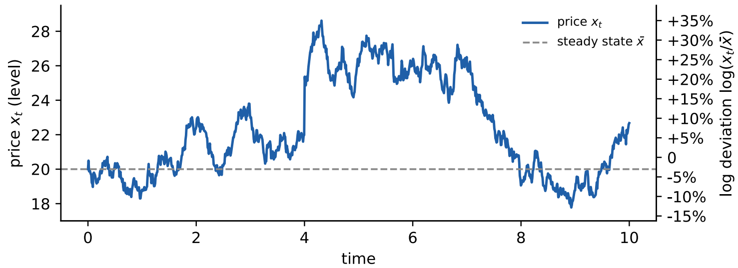

Figure 5.1: One price in levels (left axis) and as a log deviation from its steady state $\bar{x}$ (right axis), the unit used for every variable in this section.

are Ornstein–Uhlenbeck, mean-reverting rather than permanent:

$$dq_t = \sigma_u \, dW_t^q, \qquad da_t^v = -\theta_a a_t^v \, dt + \sigma_a \, dW_t^{a,v}, \qquad d\eta_t^v = -\theta_\eta \eta_t^v \, dt + \sigma_\eta \, dW_t^{\eta,v}, \quad (5.3.1)$$

with W^q, $W^{a,v}$, $W^{\eta,v}$ independent Brownian motions, so the shocks are independent across firms and of the aggregate. These $1 + 2N$ Brownian motions and the observation noises of (5.3.6) below are the only sources of randomness. At each date firm v chooses two controls: a quote P_t^v, the log price at which it will deliver at $t + \tau$, and an order o_t^v, the log quantity of input it asks firm $v - 1$ to deliver at $t + \tau$. The delivery lag $\tau > 0$ turns each control into a delayed row of the state in the sense of (2.1.5):

$$\pi_t^v = P_{t-\tau}^v, \qquad i_t^v = o_{t-\tau}^v, \qquad d_t^v = o_{t-\tau}^{v+1}, \quad (5.3.2)$$

the price in force, the input arriving, and the deliveries owed to the customer. The price index, the average of the prices in force, and household demand, the quantity households buy from firm v, are algebraic in these,

$$P_t = \frac{1}{N} \sum_v \pi_t^v, \qquad y_t^{H,v} = q_t - P_t - \theta \, (\pi_t^v - P_t) + \eta_t^v, \quad (5.3.3)$$

and firm v's output at t is $a_{t-\tau}^v + i_t^v$, since productivity is known when the order is placed. Household demand falls by θ times the firm's price relative to the others', so θ measures how readily households switch between firms.

Objective

Firm v sets its price in advance and would like it, once in force, to match a target that moves with the price index, with real activity, and with its own input cost. It sells at a markup, so every unit sold to the customer or to households adds to profit; this is the only reason a seller wants its customer to order more. Output at t was fixed by the order placed at $t-\tau$, but deliveries at t are whatever the customer and the households ask for then, so the two can differ. A shortfall has to be bought at short notice while a surplus is dumped. Orders take effort to place. The input is paid at delivery at the price quoted now, which makes an order dearer when the supplier's quote is high relative to the prices that will prevail at delivery. Firm v minimizes the expected discounted integral of the flow loss,

$$
\begin{aligned}
L_t^v = {} & \underbrace{\left(\pi_t^v - \pi_t^{v*}\right)^2}_{\text{missed price target}} \underbrace{-\,2\kappa\left(d_t^v + y_t^{H,v}\right)}_{\text{markup earned on sales}} + \underbrace{m\left(a_{t-\tau}^v + i_t^v - d_t^v - y_t^{H,v}\right)^2}_{\text{output minus deliveries}} \\
& + \underbrace{r\,(o_t^v)^2}_{\text{cost of ordering}} + \underbrace{2c\,o_t^v\left(P_t^{v-1} - P_{t+\tau}\right)}_{\text{order times real input price}},
\end{aligned} \tag{5.3.4}
$$

with target price

$$
\pi_t^{v*} = \underbrace{P_t}_{\text{price index}} + \underbrace{\xi\,(q_t - P_t)}_{\text{demand pressure}} + \underbrace{\zeta\left(\pi_t^{v-1} - P_t - a_{t-\tau}^v\right)}_{\text{real cost of the input}}. \tag{5.3.5}
$$

The target is the price the firm would choose if it knew everything. A firm that sells a differentiated product prices at its cost plus a markup, and in log deviations from the steady state the markup is a constant that drops out, so the target moves with cost. Cost here is the input: the supplier's price relative to the index, $\pi_t^{v-1} - P_t$, less the firm's own productivity $a_{t-\tau}^v$, since a more productive firm needs less input per unit of output. The index enters one for one because a firm whose competitors have all raised their prices can raise its own without losing customers. The output gap $q_t - P_t$, spending in excess of prices, is demand pressure. When households are buying more than the economy produces at current prices, each firm can charge more, and ξ is how much. This is the pricing rule of Woodford [118], with $\xi < 1$ so that firms move their prices less than one for one with demand, plus the input-cost term. The second term is the first-order part of revenue at markup κ, and the third gives the order its curvature. The last is the second-order part of the input bill. The product of order and real cost is bilinear in the firm's order and its neighbors' current quotes, which enter the order's first-order condition as an exogenous linear process. The price in force $\pi_t^v = P_{t-\tau}^v$ is a lagged own control and is carried as a state coordinate. The order's first-order condition contributes $c\,\mathbb{E}[P_t^{v-1} - P_{t+\tau} \mid \mathcal{F}_t^v]$. The buyer orders less when it perceives a high real input price and more when it expects the price index to rise over the lag, so its inflation motive is a term of its own. The buyer must decompose the supplier's quote into inflation and productivity before acting on it. The loss is convex in the quote through the two squares at $t + \tau$ and in the order through r and the mismatch term.

Information

Firm v observes exactly its own productivity, its own actions, and its own prices in force. It observes four noisy channels,

$$dY_t^{v,k} = S_t^{v,k}\,dt + s_k\,dW_t^{v,k}, \qquad k = 1, \dots, 4, \tag{5.3.6}$$

with

$$S_t^{v,1} = y_t^{H,v}, \qquad S_t^{v,2} = P_t^{v-1}, \qquad S_t^{v,3} = o_t^{v+1}, \qquad S_t^{v,4} = o_t^{v-1} :$$

its own household sales; the transaction price on its input, whose noise stands for idiosyncratic freight and discount terms; its order book, whose noise is the orders of small buyers; and its supplier's upstream order, seen through industry data. The $W^{v,k}$ are Brownian motions independent of everything else, and the noise levels s_k are fixed.

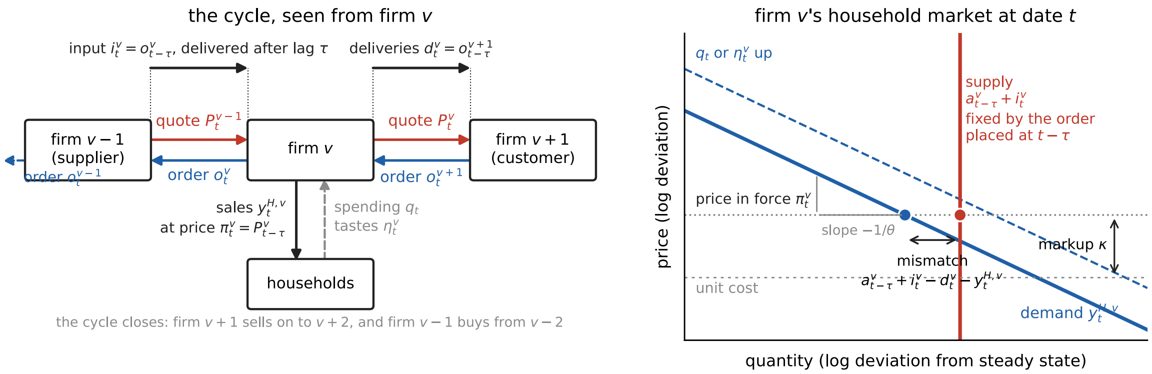

Figure 5.2: The purchase-order market. Left: the cycle seen from firm v. Goods move to the right; quotes (red) travel with the goods and orders (blue) against them; households buy from each firm. Firm v sees its own productivity and actions exactly and the four channels of (5.3.6) through noise. Right: firm v's household market at one date. Demand is the line (5.3.3); supply is fixed by the order placed a lag earlier. At the price in force the gap between the two is the mismatch term of the loss (5.3.4).

The last three channels are observations of controls, the rows $B_m^Y D_{t-}^m$ of (3.1.2). No firm observes another firm's action exactly, and no channel's noise is another firm's strategy. Latent to firm v are q, P, η^v, the neighbors' productivities, and the neighbors' quotes and orders. Firm v's information $\mathcal{F}_t^v$ is generated by the four channels, a^v, and its own past actions, and a stationary strategy is a linear rule on it,

$$u_t^v = \bar{u} + \sum_{k=1}^{4} \int_0^L g_k(l)\, dY_{t-l}^{v,k} + \int_0^L \hat{g}(l)\, dW_{t-l}^{a,v}, \tag{5.3.7}$$

for each of the two controls, with lags cut off at L.

At t the $t-\tau$ contracts are delivered; nothing decided at t touches them. The firm reads its channels and its productivity, infers spending and the average of the quotes being made now, and chooses P_t^v and o_t^v. Each action carries an information wedge at t and a state price at $t+\tau$, the two terms of the delayed policy line of Chapter 3.

Goods move one way, so the graph is the directed cycle, whose automorphism group is $\mathbb{Z}/N\mathbb{Z}$, the rotations, and the displacement reduction of Appendix 5.A applies with signed displacements, with two controls per firm, and with the supplier and the customer in distinct position classes because the cycle is directed. The price index P is the spatial mean of the price-in-force kernels, and $q - P$ is the aggregate object of Proposition 5.1.

5.4 Numerical solutions of the purchase-order market

The market of Section 5.3 is solved on a cycle of three firms, inside the rotation-symmetric class of Remark 5.6. The example does for the cycle what the two-player game of Chapter 1 did for the bilateral case.

5.4.1 Primitives

Unless otherwise stated there are $N = 3$ firms, the delivery lag is $\tau = 0.5$, and the objective is the long-run average cost, as in the numerical benchmark of Chapter 3. The shocks (5.3.1) have $\sigma_u = 1$, $\theta_a = \theta_\eta = 0.5$ and $\sigma_a = \sigma_\eta = 1$, so productivity and tastes have stationary standard deviation 1 and half-life 1.4, about three delivery lags. In the loss (5.3.4) and target (5.3.5): $\theta = 4$ is the elasticity of household demand to the relative price, $\xi = 0.15$ the weight of the output gap and $\zeta = 0.5$ the weight of the real input cost in the target price, $\kappa = 0.3$ the markup on sales, $m = 1$ the penalty on the gap between output and deliveries, $r = 0.2$ the effort cost of orders, and $c = 0.2$ the coefficient of the input bill. The channel noises in (5.3.6) are $s_1 = 2.5$ on own sales, $s_2 = 0.3$ on the transaction price, $s_3 = 0.3$ on the order book, and $s_4 = 2$ on the supplier's upstream order, so a seller sees its customer's order precisely and a buyer sees its supplier's order poorly. Strategies take the form (5.3.7) with $L = 24$. These values are not calibrated. They are chosen so that every channel is active and the equilibrium can be inspected, so the signs and orderings below should be trusted and the magnitudes should not.

5.4.2 Three information structures

The same shocks and losses are solved under three information structures. Under perfect information every firm observes every primitive shock; there is nothing to filter and

	mean quote $\bar{P}$	mean order $\bar{o}$
perfect information	-0.260	1.298
market information, exogenous signal	-0.252	1.261
market information, strategic	-0.274	1.036
strategic shift (strategic $-$ exogenous)	-0.022	-0.225

Table 5.1: Mean quote and mean order at $N = 3$ under the three information structures.

no action carries information. Under market information a firm observes only its four channels and its own productivity. The market-information equilibrium is computed twice. In the exogenous-signal version the information wedge is set to zero in every firm's first-order condition. In this counterfactual, each firm optimizes as though its actions did not change its neighbors' beliefs. Corollary 1.9 reaches that conclusion when the neighbors' observations are genuinely unaffected by the firm's control; here they are not, so the wedge is switched off as a counterfactual rather than vanishing on its own. Firms still filter the same channels; they just cannot teach. In the strategic version the wedge is kept and the neighbors learn from the actions. The strategic shift is the difference between the strategic and the exogenous-signal solutions. The difference between perfect information and the exogenous-signal solution is the filtering gap, the cost of dispersed information without any strategic response to it. Any perfect-information model can be run through the same three steps.

The cycle length is the one parameter with a direct empirical counterpart. On the cycle each firm is one stage of a production chain that closes on itself, so N is the number of stages. Across 426 US industries the mean upstreamness, an industry's average number of stages from final use, 1 meaning direct sale, is 2.09 with a standard deviation of 0.85 and a maximum of 4.65 [8], so chains of two to five stages cover the bulk of production. Table 5.2 reports the equilibrium over that range and beyond. The order shift is largest at five or six stages, and the quote shift stays between -0.012 and -0.022. The filtering gap grows slowly with N, from 3 to 6 percent, as spending has more firms to be inferred from. The example uses $N = 3$, the shortest cycle with a distinct supplier and customer. Nothing below changes qualitatively at other lengths, and the shift at $N = 3$ is the

N	perfect information	exogenous signal	strategic	gap (exo. − perfect)	shift (strat. − exo.)
3	1.298	1.261	1.036	−0.037	−0.225
4	1.340	1.306	1.018	−0.034	−0.289
5	1.364	1.323	1.020	−0.041	−0.303
6	1.379	1.330	1.027	−0.049	−0.303
8	1.397	1.336	1.043	−0.060	−0.293
12	1.414	1.340	1.062	−0.074	−0.278
24	1.430	1.343	1.079	−0.087	−0.264
48	1.438	1.346	1.090	−0.092	−0.256

Table 5.2: Mean order $\bar{o}$ on cycles of N firms under the three information structures of Section 5.4.2.

smallest in the range.

5.4.3 Mean actions and the strategic shift

Table 5.1 reports the stationary means. Removing perfect information but keeping the exogenous signal lowers the mean order by 3 percent. Letting the neighbors observe the actions lowers it by a further 18 percent. As in Chapter 1, movement of the mean action with the information structure is the direct sign that separation fails. Under a separation principle the three rows would coincide. Here the strategic shift is six times the filtering gap.

The order shift comes from the buyer's side of each link. A firm's order enters its supplier's order book, the supplier quotes higher the more demand it infers, and the firm pays more for its input. Each firm therefore orders less than it would if its order were not observed. The mechanism is signal jamming [40, 100]. It is the level distortion of limit pricing [84], in its signaling form, and of career concerns [57], where the distortion fools no one and persists. The noise-state calculus of Chapter 1 adds the magnitude and its comparative statics, which need the filtering dynamics. The quote shift is a tenth the size. The seller's incentive from Chapter 1, to quote high in the hope that the buyer attributes the quote to inflation rather than to low productivity, is present, but at these parameters the transaction price is a precise signal and a firm already learns most of what it needs to

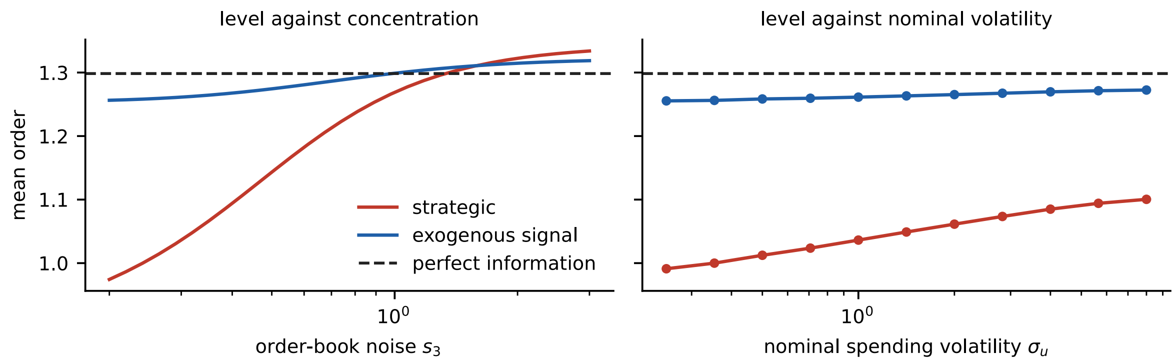

Figure 5.3: Mean order under the three information structures, against the order-book noise s_3 (left, log scale) and against the volatility of nominal spending σ_u (right, log scale).

know about the aggregate from its own sales, so the buyer hardly revises its belief when the seller shades its quote.

Figure 5.3 plots the mean order against the order-book noise and against the volatility of nominal spending, with the perfect-information level as a reference. With $s_3 = 0.2$ the order shift is -0.28; it crosses zero near $s_3 = 1.6$ and stays within 0.015 of zero up to $s_3 = 3$. A large s_3 means the identifiable customer's order is a small part of the supplier's order book, the situation of a supplier with many small buyers. In that limit the strategic and exogenous-signal solutions coincide. The model therefore contains the continuum case as a limit. The exogenous-signal level barely moves across either panel, so firms that only filter do not shift the mean; firms that also manage what their neighbors see do, by an amount that falls with nominal volatility, from -0.264 at $\sigma_u = 0.25$ to -0.172 at $\sigma_u = 8$. A supplier reads its order book as a mixture of its customer's demand and spending, weighting each by its variance. When nominal shocks dominate, a large order is read as inflation and the supplier raises its quote along with the index, which leaves the buyer's real input price alone, so the buyer has less reason to hold back. The strategic shift depends on the nominal regime in the sense of Lucas [77, 78]. The delivery lag matters much less. Between $\tau = 0.1$ and $\tau = 1$ the order shift moves from -0.20 to -0.24 and the quote shift stays at -0.022. The shift comes from observed actions, not from prices being set in advance.

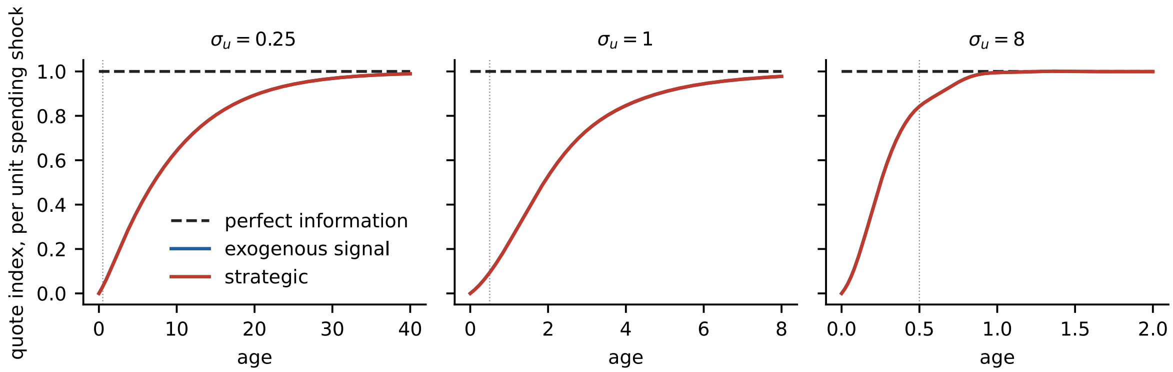

Figure 5.4: Response of the quote index to a unit shock in nominal spending at three volatilities of spending. Dashed: perfect information. Blue: market information with the exogenous signal. Red: strategic. The dotted vertical line marks the delivery lag.

5.4.4 Impulse responses

The equilibrium kernels are impulse responses to primitive shocks. Figure 5.4 plots the response of the quote index $P_{t+\tau}$ to a shock to nominal spending, per unit of the shock, at three levels of nominal volatility and under the three information structures.

Under perfect information the spending shock raises the quote index by one at once and the price in force by one at lag τ. Household demand is up by one until the prices in force catch up and zero afterwards, so money is neutral apart from the delivery lag. Under market information the index rises slowly, to 0.09 at τ, 0.53 at 4τ and 0.98 at 16τ, and demand decays over the same horizon. The innovation reaches each firm through channels that also carry its idiosyncratic shocks, and firms raise their prices only as fast as they can separate the two. This is the inertia of Woodford [118], produced here by the same higher-order expectations, except that the signals are the neighbors' quotes and orders. The speed of adjustment is set by the variance of nominal shocks relative to idiosyncratic ones, as in Lucas [77]. At $\sigma_u = 0.25$ the index is at 0.55 after 16τ and 0.93 after 46τ; at $\sigma_u = 8$ it is fully adjusted within two lags. In every panel the strategic kernels lie close to the exogenous-signal kernels. The largest difference is 1.3×10^{-3} on a unit shock, and it is not a numerical residual. It is unchanged under a tighter fixed-point tolerance and a finer grid, grows to 2.1×10^{-3} at $s_3 = 0.2$, and changes sign at $N = 6$.

The spatial mean of the kernel shift, the object of Proposition 5.1, is present in this economy but small. It is smaller than the response it perturbs by a factor of about 800, and smaller than the strategic shift in the mean action, -0.225, by a factor of about 170. The proposition places no bound on its size, and nothing here shows it is small in other economies of the class.

On the spending shock household demand responds by one and output by only 0.045. Output is $a^v_{t-\tau} + i^v_t$, fixed by orders placed a lag earlier, so a nominal shock cannot raise it on impact. It raises demand against a fixed supply, and the gap enters the mismatch term of the loss. The real effect of money in this market is a mismatch between deliveries and demand, and it lasts as long as the price adjustment.

On the productivity shock the two market-information solutions do differ (Figure 5.5). Firm 0's quote falls by 0.38 on impact, against 0.52 under perfect information, and the index by a third of that; demand and output rise by about 1.1 once the lower price is in force. The order rises by 0.18 on impact and 0.14 at τ under perfect information, by 0.11 and 0.06 under the exogenous signal, and by 0.14 and 0.08 in the strategic solution, which also has a second hump at 3τ when the customer's order returns through the order book. The strategic effect is concentrated in the responses to idiosyncratic innovations, the ones a neighbor cannot distinguish from aggregate conditions without the signal, and there the firm responds more, not less. The gap between the strategic and exogenous-signal responses is present at every level of nominal volatility and largest at intermediate levels, where the supplier finds an order hardest to read and the buyer finds it most worth shaping.

Table 5.3 collects the size of the kernel shift, the largest difference between the strategic and exogenous-signal kernels, on the two shocks across ten parameter variants. It is largest when the input bill's cross term c is large or the demand elasticity θ is small, the two places where a nominal level enters the loss other than through a relative price. On the productivity shock it rises with the cost pass-through ζ and with c and falls when own sales are precise. The sales incentive κ leaves every kernel unchanged and moves only the means.

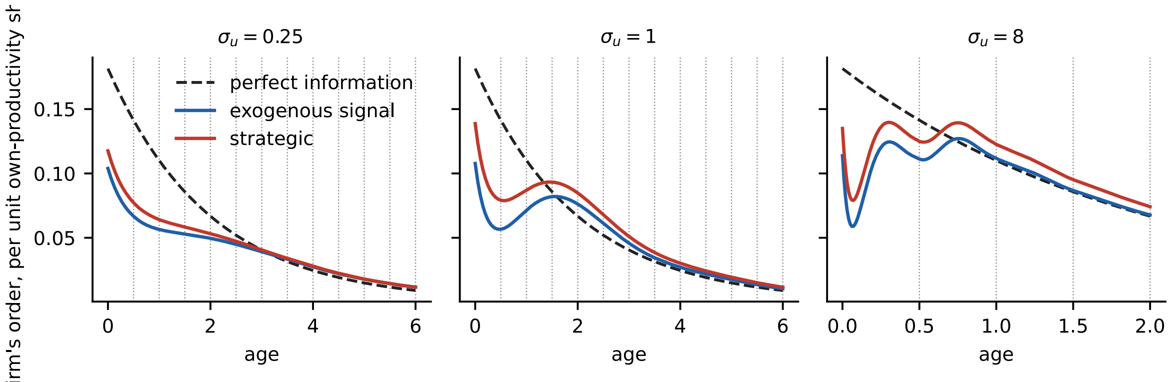

Figure 5.5: Response of firm 0's order to a unit shock in its own productivity at three volatilities of nominal spending. Dashed: perfect information. Blue: market information with the exogenous signal. Red: strategic. Dotted verticals at multiples of τ.

	spending	own productivity
base	1.3	31
delivery lag $\times 2$ ($\tau = 1$)	1.5	36
quote on q ($\xi = 0.5$)	1.9	16
pass-through $\times 2$ ($\zeta = 1$)	3.6	55
cross term $\times 3$ ($c = 0.6$)	9.7	50
sales incentive $\times 3$ ($\kappa = 1$)	1.3	31
precise sales ($s_1 = 1$)	1.3	3.3
inelastic demand ($\theta = 1$)	6.7	21
calm spending ($\sigma_u = 0.25$)	0.98	14
volatile spending ($\sigma_u = 8$)	0.80	21

Table 5.3: Largest difference between the strategic and exogenous-signal kernels on the spending shock and on firm 0's productivity shock, by parameter variant, in units of 10^{-3} per unit shock. Each variant changes one parameter from the base case $N = 3$, $\tau = 0.5$, $\theta = 4$, $\xi = 0.15$, $\zeta = 0.5$, $\kappa = 0.3$, $m = 1$, $r = 0.2$, $c = 0.2$, channel noise $(s_1, s_2, s_3, s_4) = (2.5, 0.3, 0.3, 2)$, $\sigma_u = \sigma_a = \sigma_\eta = 1$, $\theta_a = \theta_\eta = 0.5$.

In Figure 5.6 the supplier reads its customer's order for many lags, which is the channel the order shift runs through, while the buyer reads the transaction price at once and forgets it within one lag. The kinks at multiples of τ in Figure 5.5 and in the policy kernels of Figure 5.6 are produced by the delivery lag. The firm may already have acted on an observation older than τ, in a price now in force, and the optimal weight on it changes slope at that age. Each further multiple of τ is one more round of the same effect through the neighbors. The two resolutions in Figure 5.6 are indistinguishable, so the kinks are properties of the equilibrium and not of the discretization.

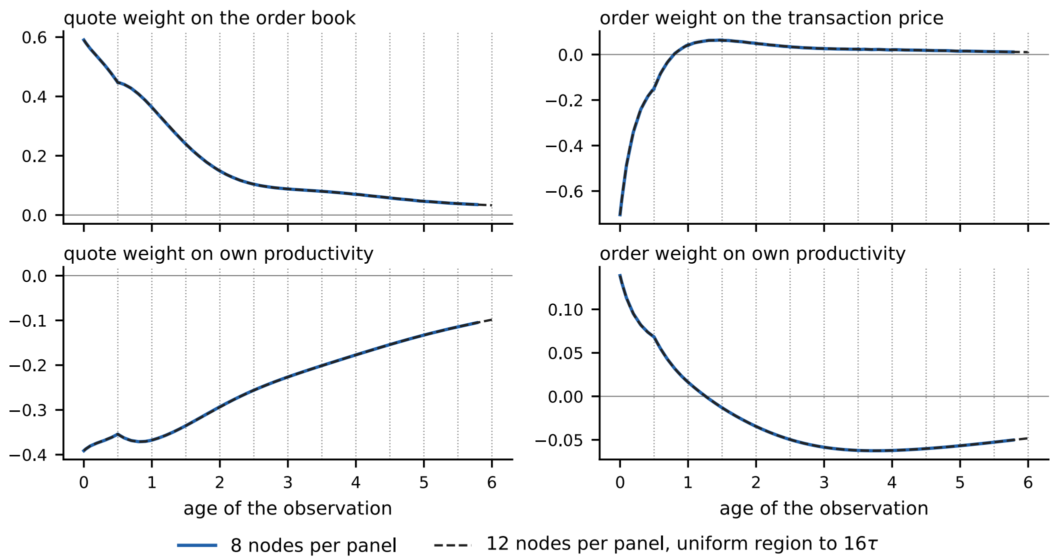

Figure 5.6: Four policy kernels of the strategic solution, by the age of the observation, at two resolutions of the age grid: the supplier's quote weight on its order book, the buyer's order weight on the transaction price, and the quote and order weights on the firm's own productivity. Dotted verticals at multiples of τ.

5.4.5 Computation

The solver is the stationary solver of Chapter 3 extended with the delayed-action and observed-action terms of Section 5.3. Kernels are represented by their values at Chebyshev points on panels in age, with a panel boundary at every multiple of τ up to 8τ and panels of doubling width beyond, out to $L = 24$. A delay is then an exact shift between panels, and the cost integrals and convolutions are Gauss quadratures on the pieces between breakpoints. The kernels of delayed processes are never resampled onto the grid, which keeps the kinks at multiples of τ exact. The equilibrium is computed as a fixed point of best responses in the control-free parametrization of Chapter 1, in which each best response is a single linear solve, and the fixed point is found by Anderson acceleration. One equilibrium takes about 30 best responses and two seconds. The cyclic symmetry is used as in Appendix 5.A. The system with every firm on the common strategy is block circulant (each block row is the previous one shifted by one firm), so it is inverted as N independent smaller systems, and the firm whose best response is being computed enters as a low-rank correction.

Recomputing every equilibrium with 12 nodes per panel and the uniform region extended to 16τ changes the means in the sixth digit and does not move the plotted kernels. Doubling the window to $L = 48$ changes the baseline means in the sixth digit as well. The one exception is the low-volatility case $\sigma_u = 0.25$, whose price adjustment is slow enough that $L = 24$ truncates it. Its results are computed at $L = 48$, where they agree with $L = 96$ to three digits. Random starting points converge to the same equilibrium. The leading eigenvalue of the best-response map at the equilibrium has modulus about 0.5.

The solvers and figures of this section are the work of Claude Fable 5, a large language model built by Anthropic [7], working under the author's direction. Interactive versions of the dissertation's computations are collected at sbabichenko.com/noisestate.

5.5 Aggregate interaction and the infinite-network limit

Reducing by graph symmetry keeps who observes whom and how an action changes a particular neighbor's posterior. Incomplete-information macroeconomics needs this detail, since aggregate dynamics can depend on how dispersed information is converted into actions [5].

Common aggregate effects can enter the vertex equations without changing the spatial indexing. A common factor C_t may enter every vertex equation,

$$\mathrm{d}X_t^v = (\cdots + A_c C_t)\,\mathrm{d}t + \sigma_c\,\mathrm{d}C_t + \sigma\,\mathrm{d}W_t^{0,v}, \tag{5.5.1}$$

or every observation equation. It acts as common noise for the local graph problem. The relative-position kernels then become conditional on the common history.

Mean-field effects can also sit *across* local economies. For example, a graph-symmetric state equation may include

$$A_m m_t, \qquad m_t = \mathrm{Law}\Big(X_t^o, D_t^o, \widehat{W}_t^o(\cdot,\cdot) \mid C_{[0,t]}\Big), \tag{5.5.2}$$

or a finite-dimensional statistic of that law. The consistency condition comes after the local vertex-transitive reduction: solve the finite strategic-information problem given (C, m), then require the law of a single local economy to reproduce m.

The finite-cycle formulas suggest an extension from C_N to $\mathbb{Z}$, but the finite proof does not itself establish an infinite-player game. Formally one would write

$$dX_t^v = \Big(aX_t^v + \beta X_t^{v-1} + \beta X_t^{v+1} + B^X D_t^v\Big)\,dt + \sigma\,dW_t^{0,v}, \tag{5.5.3}$$

$$D_t^v = \bar{D}_t + \sum_{r\in\mathbb{Z}} \int_0^t d_t(r,u)\,d\widehat{W}_t^v(r,u). \tag{5.5.4}$$

Turning these displays into a theorem first requires a state space for the spatial process. Spatially homogeneous independent noise is generally not $\ell^2(\mathbb{Z})$-valued, so the construction needs a weighted space or stationary-random-field formulation even before the strategic fixed point. One must also impose enough decay for the policy and wedge sums to converge and formulate admissible controls and unilateral deviations for infinitely many players. The remaining step is to prove equilibrium existence or convergence of the finite-cycle equilibria. The joint age–displacement representation above suggests one route: obtain estimates uniform in the cycle size for the spatial Fourier transform of the kernels and then invert it.

Conjecture 5.2 (Finite-cycle approximation). *Suppose the best-response map is a contraction over short horizons, uniformly in N, that interactions have finite range, and that the filtering, policy, and adjoint kernels decay exponentially in the distance between the two vertices. Then the symmetric equilibria on C_N converge locally as $N \to \infty$ to a translation-equivariant equilibrium random field on $\mathbb{Z}$.*

5.6 Economic implications and conclusion

The local-to-aggregate architecture changes which local quantities count as macroeconomic parameters. Local concentration, observation precision, disclosure delays, and the identity of observable actions can all change the kernel shift. When the same local structure repeats across many markets, these institutional details change aggregate elasticities and means even if technology and preferences are unchanged. In the purchase-order market, the strategic reduction in mean orders is several times the filtering gap, while responses to aggregate shocks barely move.

The adjustment-cost or sluggish-response coefficient that a representative-agent model would fit to the aggregate impulse response is not structural here. It changes when a policy alters local observability or market concentration. This is an information-network version of the usual warning [79] against treating reduced-form aggregate coefficients as policy invariant.

The row sum $\sum_r k(r)$ sets the aggregate impulse response and the individual entries $k(r)$ set the covariance between neighbors, so the two together over-identify k. The kernel shift, $k - k^{\text{exo}}$, moves with the information parameters (Table 5.3), which no technological adjustment cost contains.

An intervention need not target a nationally large player to have a macroeconomic effect. A disclosure, privacy, antitrust, or reporting rule can change behavior modestly in every local economy.

The noise-state closes the local hierarchy of beliefs, graph symmetry identifies equivalent positions, and aggregation comes after solving each finite local equilibrium. On a finite vertex-transitive graph, best responses commute with relabeling (Proposition 5.4), a unique equilibrium is invariant under it (Corollary 5.5), and the policy, filtering and adjoint kernels depend on position only relative to the reference player (Proposition 5.7). Local shocks may diversify, but a systematically repeated information incentive can remain in the aggregate response through the kernel's spatial mean (Proposition 5.1).

5.A Symmetry reduction on a vertex-transitive graph

5.A.1 Automorphism-invariant games

Let $G = (V, E)$ be a finite vertex-transitive graph, let $\mathcal{I} = V$ be the player set, and fix a reference vertex $o \in V$. Write

$$K = \mathrm{Stab}(o) = \{g \in \mathrm{Aut}(G) : go = o\}.$$

The stabilizer K collects the relabelings that leave the reference vertex where it is. For $g \in \mathrm{Aut}(G)$, let U_g denote the relabeling action on vertex-indexed objects,

$$(U_g x)(v) = x(g^{-1}v).$$

Relabeling moves each vertex's object along with the vertex. The same notation applies componentwise to controls, observations, filtrations, strategy profiles, and deterministic kernels.

Assumption 5.3 (Automorphism-invariant game)**.** For every $g \in \mathrm{Aut}(G)$, relabeling all vertex-indexed primitive shocks, initial conditions, state coefficients, observation channels, costs, delays, and admissible strategy sets by U_g leaves the game unchanged. If D is admissible, then $U_g D$ is admissible, and player v's objective under D equals player gv's objective under $U_g D$ after the same relabeling of primitive randomness.

Figure 5.7 illustrates the position classes on two graphs. On a cycle, displacement from the reference vertex indexes the vertices. On a distance-transitive graph such as the Petersen graph, the equivalent positions are the distance shells.

cycle C_{10}

relative positions

Petersen graph

distance shells

Figure 5.7: Position classes on two graphs. Vertices with the same shade occupy equivalent positions relative to the filled reference player.

5.A.2 Relabeling commutes with best response

Equivariant means relabeling before the map matches relabeling after. For a strategy profile D, write $BR(D)$ for the profile whose vth component is player v's unique best response to D^{-v}. When individual best responses are not unique, read BR as a correspondence and the equality below as equality of sets.

Proposition 5.4 (Equivariance of the best-response map)**.** *Under Assumption 5.3,*

$$BR(U_g D) = U_g BR(D), \qquad g \in \mathrm{Aut}(G). \tag{5.A.1}$$

Proof. Fix $g \in \mathrm{Aut}(G)$ and a profile D. Let $\widetilde{D}^v$ be any admissible deviation for player v. Assumption 5.3 gives a bijection between deviations of player v against D^{-v} and deviations of player gv against $(U_g D)^{-gv}$, namely $\widetilde{D}^v \mapsto U_g \widetilde{D}^v$. The bijection preserves the objective values. The deviation $\widetilde{D}^v$ then minimizes player v's objective against D^{-v} if and only if $U_g \widetilde{D}^v$ minimizes player gv's objective against $(U_g D)^{-gv}$. Applying this statement at every vertex gives (5.A.1). □

Corollary 5.5 (Symmetry of a unique equilibrium)**.** *Under Assumption 5.3, if the Nash* *equilibrium D^* is unique, then*

$$U_g D^* = D^*, \qquad g \in \mathrm{Aut}(G).$$

Proof. Because $D^* = BR(D^*)$, Proposition 5.4 implies

$$U_g D^* \in U_g BR(D^*) = BR(U_g D^*).$$

The profile $U_g D^*$ is then also a Nash equilibrium. Uniqueness gives $U_g D^* = D^*$. $\square$

Remark 5.6 (Solving among symmetric profiles)*.* Even without global uniqueness, Proposition 5.4 implies that the subspace of automorphism-invariant profiles is mapped to itself by the best-response map. Any fixed point obtained there is a Nash equilibrium of the original finite game whenever the best-response closure of Chapter 1 applies against arbitrary admissible deviations. Symmetry reduces the deterministic fixed-point problem, but it does not restrict an individual deviating player to symmetric deviations.

Proposition 5.4 says that relabeling the local economy and then solving it gives the same equilibrium as solving first and relabeling afterward. This holds for the endogenous filtering and belief maps as well as for the state dynamics.

5.A.3 Kernels indexed by relative position

Fix a graph-symmetric equilibrium. At the reference vertex, write a policy in noise-state coordinates as

$$D_t^o = \bar{D}_t + \sum_{w \in V} \int_0^t d_t(w,u)\, d\widehat{W}_t^o(w,u), \qquad \widehat{W}_t^o(w,u) = \mathbb{E}[W_u^w \mid \mathcal{F}_t^o]. \tag{5.A.2}$$

The kernel $d_t(w,u)$ weights the reference player's estimate of the shock that hit vertex w at time u.

Proposition 5.7 (Stabilizer-orbit reduction)**.** *Under Assumption 5.3, let D^* be an automorphism-invariant equilibrium, $U_g D^* = D^*$. Then every deterministic one-point kernel attached to the reference player is constant on K-orbits. Explicitly,*

$$d_t(gw, u) = d_t(w, u), \qquad g \in K. \tag{5.A.3}$$

Here $F_t(w, w'; u, s)$ is the reference player's filter kernel from the shock at vertex w' and time s to its estimate of the shock at vertex w and time u, and $H_t(w, w'; u)$ is its belief price on player w's estimate of the shock at vertex w' and time u. A two-point filter kernel and a two-point belief prices are invariant when both vertices are moved by the same automorphism,

$$F_t(gw, gw'; u, s) = F_t(w, w'; u, s), \tag{5.A.4}$$

$$H_t(gw, gw'; u) = H_t(w, w'; u), \tag{5.A.5}$$

for all $g \in K$.

Proof. Fix $g \in K$. Since $go = o$, the relabeling U_g leaves the reference player fixed while permuting every source vertex. The equilibrium is symmetric, $U_g D^* = D^*$. Apply this identity to the noise-state representation (5.A.2). The blueprint of Chapter 1 is the linear map with deterministic kernel from the primitive shocks dW to the reference player's estimated shocks $d\widehat{W}^o$; expanding each estimate through it writes D^o on the primitive coordinates, where the deterministic L^2 representation is unique because the primitive coordinates are orthogonal. The relabeling permutes those coordinates, so the kernel in primitive-shock coordinates at w and gw agree. The blueprint commutes with U_g for $g \in K$ by Assumption 5.3, so composing the primitive-coordinate equality with it gives the same equality for the noise-state kernels on its range, which is (5.A.3). The linear map from source coordinates to the reference player's conditional estimates defines the filter kernel, so applying the same relabeling simultaneously to its two spatial indices gives (5.A.4). The adjoint equations are the transpose linearization of this equivariant forward system with an invariant quadratic cost source; uniqueness of their solution gives (5.A.5). □

Corollary 5.8 (Displacement and distance-shell coordinates)**.** *On a finite Cayley graph, group displacement from the reference vertex may index the kernels. On the cycle C_N, this gives*

$$D_t^v = \bar{D}_t + \sum_{r \in \mathbb{Z}/N\mathbb{Z}} \int_0^t d_t(r,u)\, d\widehat{W}_t^v(r,u), \qquad \widehat{W}_t^v(r,u) = \mathbb{E}[W_u^{v+r} \mid \mathcal{F}_t^v].$$

If the model is also reflection invariant, then $d_t(r,u) = d_t(-r,u)$. On a distance-transitive graph, every one-point kernel attached to the reference vertex depends only on graph distance $\operatorname{dist}(o,w)$*.*

Proof. For a Cayley graph, left translation, multiplying every vertex by a fixed group element, carries any vertex to the reference vertex and converts an ordered pair of vertices into their group displacement. This uses only invariance under the translations; when the model is also reflection invariant, Proposition 5.7 identifies r with $-r$. In a distance-transitive graph, the symmetries fixing o carry any vertex of a distance shell to any other, so a kernel constant on position classes depends only on $\operatorname{dist}(o,w)$. □

The information wedge inherits the same reduction. If $\mathcal{V}_t(w)$ denotes the contribution of opponent w to the reference player's wedge, then

$$\mathcal{V}_t^o = \sum_{w \neq o} \mathcal{V}_t(w), \qquad \mathcal{V}_t(gw) = \mathcal{V}_t(w), \quad g \in K.$$

The sum can then be grouped by position class, with each class's contribution multiplied by the number of positions in it. The number of distinct relative positions controls the number of equilibrium objects, not the number of ordered player pairs.

5.B A finite-cycle market with delayed observation

Take $V = \mathbb{Z}/N\mathbb{Z}$. A nearest-neighbor LQG state is

$$dX_t^v = \left(aX_t^v + \beta X_t^{v-1} + \beta X_t^{v+1} + B^X D_t^v\right) dt + \sigma\, dW_t^{0,v}, \tag{5.B.1}$$

with local observation

$$dY_t^v = B_X^Y X_t^v\, dt + dW_t^v. \tag{5.B.2}$$

All indices are modulo N. Any bounded automorphism-equivariant matrix can replace the nearest-neighbor term: the adjacency matrix, the graph Laplacian, a finite-range convolution. Corollary 5.8 reduces all deterministic one-point kernels to displacement coordinates. Two-point filters and adjoints retain two relative indices because they compare two source locations.

Let each vertex w emit a local signal

$$\mathrm{d}Z_t^w = CX_t^w \mathrm{d}t + \mathrm{d}W_t^{Z,w}. \tag{5.B.3}$$

Vertex v observes Z^w with delay $\Delta(v, w)$. The delay rule is graph symmetric when

$$\Delta(gv, gw) = \Delta(v, w), \qquad g \in \mathrm{Aut}(G). \tag{5.B.4}$$

On the cycle this becomes

$$\Delta(v, w) = \Delta_{w-v}.$$

At vertex 0, the filtration contains

$$Z_s^r \quad \text{for} \quad s \le t - \Delta_r. \tag{5.B.5}$$

The filtering equation has birth surfaces

$$u = t - \Delta_r, \tag{5.B.6}$$

indexed by relative position. Each surface marks the source time at which a delayed signal first enters the filtration. In the finite-player notation these are delayed observation rows. An observation delay small in each local economy can change the equilibrium coefficient at a given relative position and so alter the aggregate response in (5.2.2).

The delayed rows are the local counterpart of Chapter 2, where an observation delay changes the birth and predictable-correction terms of the filter inside each local economy. On a stationary finite cycle, shock age and displacement can jointly index policy, filtering, and adjoint kernels. Spatial Fourier transforms and the temporal transforms of Chapter 3 can then split the linear filtering and adjoint steps into N independent smaller problems. The Kyle–Back market of Chapter 4 is a setting of this kind: a locally informative action also moves an aggregate price.

Chapter 6

Monitored Deviations: Naive Distortions and Privy Responses

6.1 Introduction

Chapter 4 retains the classical Kyle–Back division of roles, in which informed traders choose strategic controls while the competitive market maker's pricing rule is pinned down by conditional-expectation consistency. The price is then not a control the market maker can deviate from unilaterally. A strategic market maker requires a different equilibrium test. Its quote must be an admissible control, and changing the quote unilaterally must not improve its payoff.

Such a deviation is not hidden inside aggregate order flow. In the transparent market, traders observe the information needed to recognize a departure from the market maker's quote rule. The deviation propagates through the other players' strategies before returning to the market maker's objective. Pricing this feedback requires a coordinate for each player who sees the deviation and responds to it, and an adjoint valuing the resulting chain of responses. The finite-horizon response calculus below defines these objects. Applying it to a strategic stationary market maker additionally requires specifying the market maker's objective, solving the market equilibrium, and translating the gains, one for each pair of a deviating player and a player who responds to it, into the lag coordinates of Chapter 3.

The preceding chapters vary what players know about nature, the projection $\widehat{W}^i$, while every deviation stays hidden inside others' filtering. This chapter varies what players know about the origin of an off-equilibrium deviation.

In Chapter 1 the deviation's origin is known exclusively to the deviating player. The deviating player knows the drift it inserted, so its own innovation is invariant, while every other player attributes the resulting movement to primitive shocks and filters it accordingly. The information wedge $\mathcal{V}^i$ prices this unrecognized-deviation channel. It does not record how another player changes its strategy after recognizing who deviated.

Here a deviation may instead be *monitored*. A player who knows its origin, a privy player, treats it as a labeled, perfectly observed, zero-variance shock and responds through a separate strategy coordinate. A player who does not know its origin, a naive player, continues to filter its effect as ordinary primitive shock. The same seed can generate privy responses for some players and naive filtering for others. The convention of Chapter 1 is the corner in which no opponent monitors the deviation, while a publicly attributable market-maker quote lies in the opposite corner in which every relevant opponent is privy. When every opponent is privy, a deviation moves no opponent's beliefs, the information wedge vanishes, and equilibrium play is the separation principle applied to the closed-loop strategies of the perfect-information game [16, 25]. In the classical perfect-information game every player sees the shocks and is privy. The monitoring relation (Definition 6.1) moves the game between these corners; Proposition 6.13 recovers the two endpoints exactly. The naive players here differ from the deceived agents of the adversarial-deception literature [67, 119]. Their strategies remain equilibrium best responses, and their inference is Bayes-correct given their information (Remark 6.3); they lack only the origin label. Table 6.1 locates the four regimes. This chapter mixes the two columns: the same deviation can be observed by some players and unobserved by others.

	deviations	
shocks	unobserved	observed
observed	open-loop Nash	closed-loop Nash
filtered	Chapter 1	all-privy corner

Table 6.1: Four information regimes. The monitoring relation moves the game along the bottom row, and Proposition 6.13 recovers its two corners exactly; the top row is the classical pair.

Mathematically, monitoring enlarges the forward system. A labeled deviation gives every player who sees it a response coordinate, in addition to any shifts in ordinary noise-state estimates, so merely removing privy players from the wedge sum would omit their strategic responses. The response is the linear perturbation implied by the equilibrium policy kernels, not a re-solution of the game after the deviation. For each responder–origin pair one backward equation, (6.5.16), gives a gain $H^{j\leftarrow i}$, and the response to a seed at any time t is read off by (6.5.23), so no forward–backward system is solved per seed time.

Index roles are fixed throughout this chapter. The pair $j \leftarrow i$ always reads "responder j, origin i." Symbols carrying both k or k' and $j \leftarrow i$ identify the ordinary noise-state row or column inside that responder–origin map.

6.2 Baseline noise-state setting

The objects of Chapter 1 carry over. Let $\mathcal{I} = \{1, \dots, n\}$ be the player set and let the primitive uncertainty be a Brownian shock path W. The state and player i's observation evolve as

$$dX_t = \left(B_X^X(t)X_t + \sum_{m=1}^{n} B_m^X(t)\, D_t^m\right)dt + E^0\, dW_t, \qquad dY_t^i = B_X^{Y,i}(t)\, X_t\, dt + E^i\, dW_t, \tag{6.2.1}$$

so that player m's control enters the state through the control-to-state map B_m^X, and $B_X^{Y,i}$ is the state-to-observation map; E^0 and E^i load the common shocks W into the state and into player i's observation. Player j observes a filtration $\mathcal{F}_t^j$ and carries the noise-state

$$\widehat{W}_t^j(u) := \mathbb{E}[W_u \mid \mathcal{F}_t^j], \qquad 0 \le u \le t, \tag{6.2.2}$$

and its equilibrium action is noise-state linear,

$$D_t^j = \bar{D}_t^j + \int_0^t D_t^j(u)\, d\widehat{W}_t^j(u). \tag{6.2.3}$$

Here "linear" refers to linearity in the represented shock coordinates, not to linearity of the kernels, themselves determined by nonlinear fixed-point equations.

Each player's running cost is quadratic in the joint physical-state and control variable, as in (1.3.2), with Hessian

$$G^j(s) = \begin{pmatrix} G^{XX,j}(s) & G^{XD,j}(s) \\ G^{DX,j}(s) & G^{DD,j}(s) \end{pmatrix}, \qquad G^{DX,j}(s) = (G^{XD,j}(s))^\top. \tag{6.2.4}$$

Throughout, $G^j(s)$ satisfies the Schur condition (1.3.3) of Chapter 1.

6.3 Deviation origins and monitoring

6.3.1 The monitoring relation

The monitoring relation is directed. Every player knows the origin of its own deviation, so the relation is reflexive and the set of *privy players* includes the deviating player itself. The blip-of-irrationality convention of Section 6.4.2, a single forced departure after which the deviating player resumes equilibrium play, fixes what this means for the deviating player's own continuation.

Definition 6.1 (Monitoring relation)**.** Write

$$j \trianglerighteq i \tag{6.3.1}$$

if player j can resolve deviations whose origin is player i. Equivalently, when player i creates a deviation seed, the impulse that starts the deviation, player j observes its origin perfectly and treats it as a separate deviation shock rather than filtering it as part of the primitive uncertainty. The relation is reflexive, so $i \trianglerighteq i$ for every i. Throughout, a deviation is a control spike z of player i, and $v_i = B_i^X(t)z$ is the state impulse it produces; the seed carried by the forward system is v_i, not z. It is the state impulse alone; Chapter 4's seed was the filter displacement a spike leaves.

For each origin i, define

$$\mathcal{P}_i := \mathrm{Privy}(i) := \{j \in \mathcal{I} : j \trianglerighteq i\}, \qquad \mathcal{N}_i := \mathrm{Naive}(i) := \mathcal{I} \setminus \mathcal{P}_i. \tag{6.3.2}$$

The set $\mathcal{P}_i$ contains every player who observes i's deviation seed, the deviating player among them, while $\mathcal{N}_i$ contains the players who do not.

A common source of the relation is information inclusion. If every information source available to player i is also available to player j (for example $\mathcal{F}_t^i \subseteq \mathcal{F}_t^j$, with observation of i's policy rule and action channel), it is natural to set $j \trianglerighteq i$. Inclusion is sufficient: j can compute i's equilibrium action, and any gap between it and the realized action is i's deviation. It is not necessary: j may observe i's action directly, a posted quote for instance, without observing i's signals, as in Section 6.8. The calculus below requires only the relation $\trianglerighteq$; the modeler specifies the institutional or informational mechanism that generates it.

6.3.2 Transitivity

Transitivity is assumed:

$$k \trianglerighteq j \trianglerighteq i \quad \Longrightarrow \quad k \trianglerighteq i. \tag{6.3.3}$$

The condition is substantive. Suppose j is privy to i and so responds to i's seed. If k is privy to j but not to i, then k sees j's response while lacking the coordinate needed to represent the cause of that response. The response chain is then ill-typed: k can identify the proximate actor but not the root seed, a murder mystery.

Contrapositively, if k is naive to i and j is privy to i, then k must also be naive to j's response to i:

$$k \ntrianglerighteq i,\ j \trianglerighteq i \quad \Longrightarrow \quad k \ntrianglerighteq j \text{ for every } j \text{ responding to the } i\text{-seed.} \tag{6.3.4}$$

Transitivity is the monitoring analogue of partial nestedness in team theory [55]. There, an agent affected by another's action observes that other's information; here, a player who sees a response also sees its origin.

Lemma 6.2 (Naive sets)**.** *Under transitivity* (6.3.3)*, $\mathcal{N}_i \subseteq \mathcal{N}_j$ for every privy player $j \in \mathcal{P}_i$. A player naive to the deviating player is naive to every response by a privy player $j \in \mathcal{P}_i$. Section 6.5 reads this as a containment of index sets: the rows of every gain, indexed by $\mathcal{N}_j$, contain its columns, indexed by $\mathcal{N}_i$.*

Proof. If $k \in \mathcal{N}_i$ and $j \in \mathcal{P}_i$, then (6.3.4) gives $k \ntrianglerighteq j$, so $k \in \mathcal{N}_j$. □

Remark 6.3 (No smoke without fire)*.* Consider the ill-typed case above, in which k observes player j's response to an i-origin seed but not the origin i. Seeing only the response move the state, k could not treat it as a labeled coordinate. It would have to attribute the movement by Bayes' rule, forming a posterior over which origin produced the response, supported on every deviation k cannot rule out. That posterior is not a free belief but an equilibrium object. Transitivity is exactly the condition that forbids this case. On the privy side the attribution never arises: a privy player's response kernel $D^{j \leftarrow i}$ is indexed by the originating seed i, not by an inferred mixture over possible origins. To rule out this attribution problem, no player may see smoke without seeing the fire that caused it.

With reflexivity and transitivity, $\trianglerighteq$ is a preorder. If $i \trianglerighteq j$ and $j \trianglerighteq i$, then i and j are mutually informed about each other's deviation origins.

Assumption 6.4 (Perfect transitive monitoring)**.** The relation $\trianglerighteq$ is reflexive and transitive, and monitoring is perfect: a player either observes a deviation origin exactly or not at all, with no intermediate noisy observation. Formally, $j \trianglerighteq i$ means that player j's strategies may depend on the coordinate $z^i \in \mathcal{Z}_i$ of Remark 6.5.

This assumption is maintained for the rest of the chapter; Remark 6.12 returns to the role of perfect observation and to what fails when it is dropped.

6.4 Naive versus privy coordinates

Monitoring breaks a convention every other chapter relies on. The control is noise-state linear, (6.2.3), and everything a player could react to reaches it through $\widehat{W}^j$, an opponent's deviation included, since a deviation enters the observation stream and is filtered like any other movement of the state. A monitored seed is not like that. It moves the physical state, and it moves the beliefs of the players who do not recognize it, but for a player who does recognize it there is nothing to filter. The seed has zero variance under the equilibrium law, so it is not a coordinate of any noise-state, and a strategy written as a function of $\widehat{W}^j$ alone cannot respond to it. The response has to live somewhere else. A strategy must therefore split into two parts. The on-path part is the policy kernel of the other chapters. The response part is a kernel on the labeled seed coordinate, zero on the equilibrium path and pinned only by what the player would do off it. Section 6.4.3 makes the decomposed object the strategy of the monitored game.

6.4.1 Naive and privy players

A naive player $k \in \mathcal{N}_i$ does not recognize the deviation seed $v_i(s)$ as a deviation. Its physical effect enters k's observation stream, and k filters that effect as if it were evidence about the primitive shock path. The ordinary noise-state changes:

$$\delta w_t^{k \leftarrow i}(u; s) \neq 0, \qquad k \in \mathcal{N}_i. \tag{6.4.1}$$

Throughout, δw is the noise-state perturbation in density form, $d_u\, \delta \widehat{W}_t^k(u) = \delta w_t^k(u)\, du$; the singular part $\Pi^k\, dW_u$ does not vary. For a privy player $j \in \mathcal{P}_i$, the ordinary shock estimate does not change when i deviates:

$$\delta w_t^j(u) = 0 \quad \text{for the ordinary fundamental-shock coordinate.} \tag{6.4.2}$$

Instead, j receives a separate input coordinate, the deviation seed $v_i(s)$ (Remark 6.5). The seed enlarges player j's sufficient statistic from $\widehat{W}^j$ to

$$\Big(\widehat{W}^j,\ \{v_i(s) : j \trianglerighteq i\}\Big). \tag{6.4.3}$$

This labeled coordinate needs a response kernel of its own.

Remark 6.5 (Deviations on their own factor). Deviations live on their own factor. Let $\mathcal{Z}_i$ be player i's admissible deviations, the L^2 control paths of Chapter 1, and $\mathcal{Z} = \prod_i \mathcal{Z}_i$. Player j's strategy is measurable in $\mathcal{F}_s^j$ and in the factors $\mathcal{Z}_i$ with $j \trianglerighteq i$, its own among them, and constant in the others. Under the equilibrium law the seed has zero variance, so a naive player's filter carries no coordinate for it and never anticipates it; a privy player reads it as a known input. At fixed equilibrium gains the displaced state, filters, and responses are affine in $z \in \mathcal{Z}$, so the coefficients below are exact for finite deviations, and two deviations, from one origin or several, displace by the sum. The seed affects conditional means but not conditional covariances or the ordinary filter for the primitive shocks. A monitored deviation can move a privy player's action but cannot manipulate that player's ordinary noise-state. That perfect observation keeps the privy response affine in the seed.

As in the Chapter 1 spike variation, v denotes the *induced state impulse*, the state movement a control spike produces. The forward variational system starts from

$$\delta_i X_t = v. \tag{6.4.4}$$

If the underlying control-space spike is z, then $v = B_i^X(t)z$, and the deviating player's first-order condition pulls the state price back to control space through $(B_i^X(t))^\top H_t^{X,i}$. The kernel $D_t^{j\leftarrow i}(s)$ represents a privy player's response to an observed state impulse, normalized per unit state impulse; equivalently a control-space spike first maps into the state impulse and then propagates.

6.4.2 Blip of irrationality

The convention that fixes how privy players continue off equilibrium is a *blip of irrationality*, in which a deviation is a single forced departure from i's best response, after which i continues through its own response kernel $D^{i\leftarrow i}$. The other players do not read the blip as evidence about i's type or as a loss of reputation. Continuation play treats the deviation as bygone, as in Markov perfect equilibrium [83]. The responders price the mechanical effect of the seed and carry on as though i is again rational.

The seed spawns no independent root seeds. It induces responses that change the state and feed naive filters and other controls, but each response, the deviating player's own included, is a deterministic linear function of $v_i(s)$ within the per-origin system.

Chapter 1 needed no such convention. No opponent observes a deviation's origin, so how the deviating player's own future departures are grouped into seeds is invisible to every other player, and only the total control path matters. Once some players are privy, the grouping matters. A responder reacts to each labeled seed it observes, so what counts as one seed with its continuation, rather than several separate seeds, determines the responses. The blip convention pins this down. Given it, the deviating player's own deviation calculus does not depend on how a departure is decomposed.

Lemma 6.6 (Blip and frozen continuations)**.** *Fix an origin i. For a seed time t in the horizon $[0, T]$ of Chapter 1, let e_t denote the control spike at t followed by the frozen continuation, as in Corollary 1.11. Let b_t denote the same spike followed by the blip continuation, in which i resumes its equilibrium response. Then*

$$b_t = e_t + \int_t^T K(s,t)\, e_s\, ds, \tag{6.4.5}$$

where K is the bounded strictly causal kernel generated by the deviating player's response to its own seed, and $\mathrm{Id} + K$ is invertible on the space of admissible L^2 deviations, the control class of Corollary 1.11. Consequently the families $\{e_t\}$ and $\{b_t\}$ span the same space of admissible deviations, and the deviating player's first-order condition (6.5.36) is the same under either convention. The deviation quadratic forms of Corollary 1.11, written in the two bases as Q_e and Q_b, are related by the invertible change of variable $(\mathrm{Id} + K)$, with $Q_b = (\mathrm{Id} + K)^* Q_e (\mathrm{Id} + K)$, so one is positive definite exactly when the other is.

Proof. The two continuations agree on $[0, t]$ and differ on $(t, T]$ by the response path that the blip convention appends, the linear image of the seed under the deviating player's own response kernels propagated through (6.5.8). This difference is an admissible control path supported on $(t, T]$, and every such path is a superposition of frozen spikes at times $s > t$; reading off the coefficient of e_s defines $K(s, t)$. The kernels are bounded on the compact horizon (admissibility), so K is a bounded Volterra operator and $\mathrm{Id} + K$ is invertible. The inverse is given by the series $\mathrm{Id} - K + K^2 - \cdots$, because K is strictly causal, feeding a spike at t only into times after t. Both families parameterize the same space. A linear functional vanishes on a spanning family exactly when it vanishes on the space, which gives the first-order claim. Substituting $b = (\mathrm{Id} + K)e$ into the deviation quadratic form gives $Q_b =$

$(\mathrm{Id}+K)^*Q_e(\mathrm{Id}+K)$, and an invertible change of variable preserves positive definiteness. $\square$

The exact finite-deviation identity of Corollary 1.11 transfers through $(\mathrm{Id} + K)$, so verifying a blip-convention profile against all admissible deviations reduces to the Chapter 1 identity.

Monitoring creates two belief freedoms that the all-naive setting of Chapter 1 lacks. A player who sees a movement may need to know who caused it, and a player who sees a seed may need to know what the deviating player will do next. Transitivity closes the first, so no attribution belief is ever needed (Remark 6.3), and the blip convention closes the second by fixing the continuation.

6.4.3 The monitored game and its equilibrium

The primitives of the monitored game are the baseline game of Section 6.2, the monitoring preorder of Definition 6.1 under Assumption 6.4, and the blip convention. A strategy for player j has two parts. On the equilibrium path it reduces to the noise-state linear map (6.2.3). Off the path, player j's strategy may depend on the coordinates z^i with $j \trianglerighteq i$ (Remark 6.5), and it specifies a response kernel pair $(D^{j\leftarrow i,X}, D^{j\leftarrow i,W^k})$ for each origin i with $j \trianglerighteq i$. A player's *full feedback map* is the pair of on-path kernels and response kernels.

Definition 6.7 (Monitored-deviation equilibrium)**.** A profile of full feedback maps is a *monitored-deviation equilibrium* if

(i) for every player i, no admissible control deviation improves i's payoff against the other players' full feedback maps. The deviation's consequences are generated by the monitoring structure: each privy player responds through its response kernels and each naive player filters the deviation's effects with its equilibrium filter;

(ii) for every origin i, seed, and privy player $j \in \mathcal{P}_i$, player j's response is optimal for j conditional on the observed seed, against the other players' full feedback maps.

Condition (ii) is sequential rationality, optimality conditional on events that equilibrium play never produces. A seed occurs with probability zero, so (ii) changes no player's ex-ante payoff. It pins the response kernels. For $j = i$ it pins the deviating player's own post-blip continuation. At first order, (ii) is the privy player's first-order condition (6.5.9), with the response kernels (6.5.23), and (i) is the deviating player's condition (6.5.36); both come out of the rectangular backward system of Proposition 6.9. Lemma 6.6 lets one verify condition (i) beyond first order. The finite-deviation identity of Corollary 1.11, transported to the blip basis, expresses the payoff change of an arbitrary admissible deviation as its first-order term plus a quadratic form in the deviation. Verification requires this quadratic form to be positive semidefinite.

Existence is a separate question. At the corners, Proposition 6.13 reduces the concept to systems with known existence theory, the coupled Riccati equations of the closed-loop game [16, 26] and the finite-horizon results of Chapter 1. Away from the corners the profile and the response kernels form one fixed-point problem, the one the market of Section 6.8 poses; a general existence theorem is open.

6.5 The seed-dependent system

Fix a deviating player i and a seed time t, with privy set $\mathcal{P}_i$ and naive set $\mathcal{N}_i$ as in (6.3.2). The state dynamics are (6.2.1), with player m's control entering through B^X_m. We first derive the forward equations and the privy player's response to the seed. We then write the backward system in operator form and in coefficient blocks, before treating the deviating player's response to its own seed.

6.5.1 Forward deviation equations

The initial condition for a state impulse by player i at time t is

$$\delta_i X_t = v, \qquad \delta_i w_t^k(u) = 0 \quad (k \in \mathcal{N}_i,\ 0 \le u \le t). \tag{6.5.1}$$

For $s \ge t$, the displaced state is

$$\Big(\delta_i X_s,\ \{\delta_i w_s^k(u) : k \in \mathcal{N}_i,\ 0 \le u \le s\}\Big). \tag{6.5.2}$$

The state equation is

$$\boxed{\frac{d}{ds}\delta_i X_s = B_X^X(s)\delta_i X_s + \sum_{k\in\mathcal{N}_i}\int_0^s B_k^X(s)\, D_s^k(u)\, \delta_i w_s^k(u)\, du + \sum_{j\in\mathcal{P}_i} B_j^X(s)\, \delta_i D_s^j.} \tag{6.5.3}$$

For each naive player $k \in \mathcal{N}_i$, define the naive error variation

$$\delta_i \widetilde{X}_s^k := \delta_i X_s - \int_0^s X(s,r)\, \delta_i w_s^k(r)\, dr. \tag{6.5.4}$$

The kernel $X(s,r)$ is the impulse response of the state at s to a primitive shock at r. It carries no player index, since the dynamics are common; integrating it against the truth gives the state, and against a player's estimates gives that player's estimate of the state. Player k's unresolved-state kernel $\widetilde{X}_s^k(u)$ of Chapter 1 is the part of that response player k has not yet resolved, and its innovation kernel is

$$\widetilde{C}_s^{k,Y}(u) := B_X^{Y,k}(s)\widetilde{X}_s^k(u). \tag{6.5.5}$$

The naive filter variation has an interior revision equation

$$\boxed{\partial_s \delta_i w_s^k(u) = \Big(\widetilde{C}_s^{k,Y}(u)\Big)^{\top} B_X^{Y,k}(s)\, \delta_i \widetilde{X}_s^k, \qquad 0 \le u < s,} \tag{6.5.6}$$

and a birth term

$$\boxed{\delta_i w^k_s(s) = E^{k\top} B^{Y,k}_X(s)\, \delta_i \widetilde{X}^k_s.} \tag{6.5.7}$$

Substituting (6.5.4) into (6.5.6), the coefficient of $\delta_i w^k_s(r)$ in $\partial_s \delta_i w^k_s(u)$ is

$$-\widetilde{C}^{k,Y}_s(u)^\top B^{Y,k}_X(s) X(s,r).$$

This is the belief–belief block of the generator, and it is not symmetric in (u,r): the lag u being revised enters through the error kernel, the lag r of the shock being estimated through the state kernel. That asymmetry makes the generator non-self-adjoint.

6.5.2 A privy player's response

Fix a privy player $j \in \mathcal{P}_i$. To compute player j's response to the origin-i seed, one tests whether j would profitably deviate from its proposed response. This introduces a second, hypothetical, j-origin spike. The i-seed and this test j-seed are distinct labels. In the gain defined below, the rows are indexed by $\mathcal{N}_j$, the players naive to j, because they come from player j's own deviation problem; the columns by $\mathcal{N}_i$, the players naive to i, because they come from the already-realized i-origin displacement.

Let the closed-loop coefficients for an a-origin deviation be

$$\begin{aligned}
A^{XX}_a(s) &:= B^X_X(s) + \sum_{q\in\mathcal{P}_a} B^X_q(s)\, D^{q\leftarrow a,X}_s, \\
A^{XW^p}_a(s,u) &:= B^X_p(s)\, D^p_s(u) + \sum_{q\in\mathcal{P}_a} B^X_q(s)\, D^{q\leftarrow a,W^p}_s(u), \qquad p \in \mathcal{N}_a.
\end{aligned} \tag{6.5.8}$$

These coefficients use the response kernels appropriate to the origin label a. When $a = i$ they propagate the observed i-seed; when $a = j$ they are the row-side coefficients that price a hypothetical deviation by j.

The i-variation of player j's closed-loop adjoint has rows $\delta_i H^{X,j}_s$ and $\delta_i H^{W^{k'},j}_s(u)$ for $k' \in \mathcal{N}_j$, dual to the row state S^j. The next two subsections write it as one rectangular gain, first in operator form and then in coefficient blocks, with source $G_{\text{eff}}^{j\leftarrow i}$.

Player j's first-order condition, conditioned on its own information $\mathcal{F}_s^j$, gives its response. Player j's strategies are measurable in the factor $\mathcal{Z}_i$ (Remark 6.5), so the origin-i seed is a known input to j. Subtracting the condition player j satisfies on the equilibrium path from the one it satisfies after the i-seed gives

$$0 = G^{DD,j}(s)\,\delta_i D_s^j + G^{DX,j}(s)\,\delta_i X_s + \mathbb{E}\Big[(B_j^X(s))^\top \delta_i H_s^{X,j} \mid \mathcal{F}_s^j\Big]. \tag{6.5.9}$$

Player j observes the i-seed, and the forward and adjoint systems are linear with deterministic coefficients, so the conditional expectation collapses:

$$\boxed{\delta_i D_s^j = -\Big(G^{DD,j}(s)\Big)^{-1}\Big[G^{DX,j}(s)\,\delta_i X_s + (B_j^X(s))^\top \delta_i H_s^{X,j}\Big].} \tag{6.5.10}$$

Remark 6.8 (Two labels, no cross term). The forward system is affine in the seeds, so the i- and j-labeled forward effects add, and the test uses the same gains as the ordinary j-origin deviation problem. The i-seed shifts the value of player j's adjoint while leaving that gain system intact.

6.5.3 Operator form and source

The operator notation compresses the seed-dependent physical and noise-state equations above. Stacked symbols collect Chapter 1 coordinates, and the rectangular gain $H^{j\leftarrow i}$ collects their adjoint responses into one backward solve.

An i-origin seed moves the noise-states of $\mathcal{N}_i$, while a hypothetical deviation by j moves those of $\mathcal{N}_j$. There is then no single displaced state, so the origin label indexes it. For any origin label a, define the state attached to that origin by

$$S_s^a := \Big(X_s,\ \{\widehat{W}_s^p(u) : p \in \mathcal{N}_a,\ 0 \le u < s\}\Big). \tag{6.5.11}$$

The superscript a is the origin label, and the noise-state fields are those of the players naive to origin a.

For an a-origin deviation, write

$$\delta_a S_s^a := \left(\delta_a X_s,\ \{\delta_a w_s^p(u) : p \in \mathcal{N}_a\}\right). \tag{6.5.12}$$

The deviation dynamics are

$$\frac{d}{ds}\delta_a S_s^a = A_a(s)\delta_a S_s^a, \qquad A_a(s) = A_a^{\mathrm{tr}}(s) + A_a^{\mathrm{fil}}(s), \tag{6.5.13}$$

where A_a^{tr} is the controlled-state transport part and A_a^{fil} is the filter birth-and-revision part. Throughout, A_a is the full generator and A_a^{XX} its physical-state block.

In coordinates, the X row is (6.5.3) with the coefficients (6.5.8) at label a, and the filter rows are (6.5.6)–(6.5.7) for $p \in \mathcal{N}_a$. The same non-self-adjoint generator template is used for every origin label. On the row side of a Sylvester equation it is A_a^*; on the column side, A_a.

Fix an observed seed of origin i and a privy player $j \in \mathcal{P}_i$. The i-origin state S^i supplies the columns and S^j the rows, because player j's adjoint lies in the dual of S^j. The rectangular gain is then one operator

$$H_s^{j\leftarrow i} : S_s^i \longrightarrow (S_s^j)^*, \qquad \delta_i H_s^{S^j} = H_s^{j\leftarrow i}\delta_i S_s^i. \tag{6.5.14}$$

In block form,

$$H_s^{j\leftarrow i} = \begin{pmatrix} H_s^{XX,j\leftarrow i} & \{H_s^{XW^k,j\leftarrow i}(r)\}_{k\in\mathcal{N}_i} \\ \{H_s^{W^{k'}X,j\leftarrow i}(u)\}_{k'\in\mathcal{N}_j} & \{H_s^{W^{k'}W^k,j\leftarrow i}(u,r)\}_{k'\in\mathcal{N}_j,\,k\in\mathcal{N}_i} \end{pmatrix}. \tag{6.5.15}$$

At the operator level, the backward equation is the rectangular Sylvester equation

$$-\frac{d}{ds}H_s^{j\leftarrow i} = G_{\mathrm{eff},s}^{j\leftarrow i} + A_j(s)^* H_s^{j\leftarrow i} + H_s^{j\leftarrow i} A_i(s), \qquad H_T^{j\leftarrow i} = \Pi_X^{j,\top} G^{XX,j}(T)\Pi_X^i. \quad (6.5.16)$$

The source is player j's ordinary running cost, differentiated once along j's own deviation direction and once along the i-seed direction, the Hessian G^j evaluated on the two directions. Let $\Pi_X^a : S^a \to X$ be the physical-state projection. Let $D^{j\leftarrow j} : S^j \to D^j$ be player j's response to its own seed, and let $D^{j\leftarrow i} : S^i \to D^j$ be player j's response to an i-origin seed. Then

$$\boxed{G_{\mathrm{eff},s}^{j\leftarrow i} = \begin{pmatrix} \Pi_X^j \\ D^{j\leftarrow j} \end{pmatrix}^\top G^j(s) \begin{pmatrix} \Pi_X^i \\ D^{j\leftarrow i} \end{pmatrix}.} \quad (6.5.17)$$

For $j = i$, the row and column maps coincide and $G_{\mathrm{eff}}^{i\leftarrow i}$ is symmetric. For $j \neq i$, $G_{\mathrm{eff}}^{j\leftarrow i}$ is generally rectangular and nonsymmetric.

Write the response kernels as

$$\begin{aligned} \delta_i D_s^j &= D_s^{j\leftarrow i,X}\delta_i X_s + \sum_{k\in\mathcal{N}_i}\int_0^s D_s^{j\leftarrow i,W^k}(r)\delta_i w_s^k(r)\,dr, \\ \delta_j D_s^j &= D_s^{j\leftarrow j,X}\delta_j X_s + \sum_{k'\in\mathcal{N}_j}\int_0^s D_s^{j\leftarrow j,W^{k'}}(u)\delta_j w_s^{k'}(u)\,du. \end{aligned} \quad (6.5.18)$$

Then the source blocks of (6.5.17) are

$$\begin{aligned} G_{\mathrm{eff},s}^{XX,j\leftarrow i} &= G^{XX,j}(s) + G^{XD,j}(s)D_s^{j\leftarrow i,X} + (D_s^{j\leftarrow j,X})^\top G^{DX,j}(s) \\ &\quad + (D_s^{j\leftarrow j,X})^\top G^{DD,j}(s)D_s^{j\leftarrow i,X}, && (6.5.19) \\ G_{\mathrm{eff},s}^{XW^k,j\leftarrow i}(r) &= G^{XD,j}(s)D_s^{j\leftarrow i,W^k}(r) + (D_s^{j\leftarrow j,X})^\top G^{DD,j}(s)D_s^{j\leftarrow i,W^k}(r), && (6.5.20) \\ G_{\mathrm{eff},s}^{W^{k'}X,j\leftarrow i}(u) &= (D_s^{j\leftarrow j,W^{k'}}(u))^\top G^{DX,j}(s) + (D_s^{j\leftarrow j,W^{k'}}(u))^\top G^{DD,j}(s)D_s^{j\leftarrow i,X}, && (6.5.21) \\ G_{\mathrm{eff},s}^{W^{k'}W^k,j\leftarrow i}(u,r) &= (D_s^{j\leftarrow j,W^{k'}}(u))^\top G^{DD,j}(s)D_s^{j\leftarrow i,W^k}(r). && (6.5.22) \end{aligned}$$

The response kernels themselves are

$$\boxed{\begin{aligned} D_s^{j\leftarrow i,X} &:= -\big(G^{DD,j}(s)\big)^{-1}\big[G^{DX,j}(s) + (B_j^X(s))^\top H_s^{XX,j\leftarrow i}\big], \\ D_s^{j\leftarrow i,W^k}(r) &:= -\big(G^{DD,j}(s)\big)^{-1}(B_j^X(s))^\top H_s^{XW^k,j\leftarrow i}(r), \qquad k\in\mathcal{N}_i. \end{aligned}} \tag{6.5.23}$$

Together, (6.5.16), (6.5.17), and (6.5.23) form a rectangular Riccati–Sylvester fixed point.

6.5.4 Coefficient blocks

The gain equations are the coefficient form of the seed-dependent system. The four block equations below share one backward template, specialized to a row coefficient and a column coefficient.

Proposition 6.9 (Rectangular gain reduction)**.** *Fix an origin i and a privy player $j\in\mathcal{P}_i$. Under Assumption 6.4, suppose the variations $\delta_i H^{X,j}$ and $\delta_i H^{W^{k'},j}$, $k'\in\mathcal{N}_j$, are linear in $(\delta_i X, \delta_i w^k)$ with bounded deterministic coefficients. Then they take the rectangular form*

$$\delta_i H_s^{X,j} = H_s^{XX,j\leftarrow i}\delta_i X_s + \sum_{k\in\mathcal{N}_i}\int_0^s H_s^{XW^k,j\leftarrow i}(r)\,\delta_i w_s^k(r)\,dr, \tag{6.5.24}$$

$$\delta_i H_s^{W^{k'},j}(u) = H_s^{W^{k'}X,j\leftarrow i}(u)\delta_i X_s + \sum_{k\in\mathcal{N}_i}\int_0^s H_s^{W^{k'}W^k,j\leftarrow i}(u,r)\,\delta_i w_s^k(r)\,dr. \tag{6.5.25}$$

The coefficient functions solve (6.5.28)–(6.5.32), with source blocks (6.5.19)–(6.5.22). Equation (6.5.23) then recovers player j's response to any realized origin-i seed without solving a new forward-backward system for each seed time.

The filter contraction. The filter generator enters the adjoint through one operator and its transpose. For a naive player k', the filter contraction of Chapter 1, equation (1.4.9), acts here on a noise-state lag field M by

$$\boxed{\mathfrak{B}^{k'}M := E^{k'}\,M(s) + \int_0^s \widetilde{C}_s^{k',Y}(r)\,M(r)\,dr,} \tag{6.5.26}$$

$$
\begin{array}{r|c|c|}
 & \text{column term } HA_i \text{ (origin } i\text{)} & \\
 & \begin{gathered} + H^{\cdot X} A_i^{XX} \\ + \textstyle\sum_k H^{\cdot W^k} \mathfrak{B}^{k\top} B_k \\ X \end{gathered} & \begin{gathered} + H^{\cdot X} A_i^{XW^k} \\ - H^{\cdot W^k} \mathfrak{B}^{k\top} B_k \, X(s,r) \\ W^k \ (k \in \mathcal{N}_i) \end{gathered} \\
\hline
\text{row term } A_j^* H \text{ (responder } j\text{)} \quad \begin{gathered} + A_j^{XX\top} H^{X\cdot} \\ + \textstyle\sum_{k'} B_{k'}^\top \mathfrak{B}^{k'} H^{W^{k'}\cdot} \end{gathered} \quad X & -\dot{H}^{XX} = G_{\text{eff}}^{XX} + \cdots & -\dot{H}^{XW^k}(r) = G_{\text{eff}}^{XW^k}(r) + \cdots \\
\hline
\begin{gathered} + A_j^{XW^{k'}\top} H^{X\cdot} \\ - X^\top B_X^{Y,k'\top} \mathfrak{B}^{k'} H^{W^{k'}\cdot} \end{gathered} \quad \begin{gathered} W^{k'} \\ k' \in \mathcal{N}_j \end{gathered} & -\dot{H}^{W^{k'}X}(u) = G_{\text{eff}}^{W^{k'}X}(u) + \cdots & -\dot{H}^{W^{k'}W^k}(u,r) = G_{\text{eff}}^{W^{k'}W^k}(u,r) + \cdots \\
\hline
\end{array}
$$

Figure 6.1: Reconstructing the backward system $-\dot{H}^{\bullet\bullet} = G_{\text{eff}}^{\bullet\bullet} + A_j^* H + HA_i$. Row terms come from $A_j^* H$ (left), column terms from HA_i (top); the $W^{k'}$-row revision and the W^k-column residual term carry the opposite sign. Each cell is completed by adding the two terms on its row margin and the two on its column margin; the dot is the free index, filled by the cell's column on the row margins and by its row on the column margins.

the filter birth $E^{k'}$ at lag s together with the revision $\widetilde{C}^{k',Y}$ at interior lags. Its transpose acts on a field N from the right,

$$
N \, \mathfrak{B}^{k'\top} := N(s) \, E^{k'\top} + \int_0^s N(r) \left(\widetilde{C}_s^{k',Y}(r) \right)^\top dr. \tag{6.5.27}
$$

Contracting the responder's $W^{k'}$-row on the left, $\mathfrak{B}^{k'} H_s^{W^{k'}\,\cdot}$ for $k' \in \mathcal{N}_j$, gives the cost to player j of the players naive to it treating its own hypothetical deviation as noise, the sum over $\mathcal{N}_j$ in (6.5.28). Contracting the deviating player's W^k-column on the right, $H_s^{\cdot\, W^k} \mathfrak{B}^{k\top}$ for $k \in \mathcal{N}_i$, gives the term through which the deviating player's seed enters the naive filters, the sum over $\mathcal{N}_i$ there. Each contraction leaves the other lag as the displayed argument. The observation map $B_X^{Y,k}(s)$ multiplies the column contraction on the right in the gain equations below.

Figure 6.1 assembles the four equations as a block. The state X and the naive coordinates $W^{k'}$ ($k' \in \mathcal{N}_j$) and W^k ($k \in \mathcal{N}_i$) index the rows and columns, and $B_a := B_X^{Y,a}(s)$.

Rows. Matching the $\delta_i X_s$ coefficient gives

$$\boxed{\begin{aligned}-\frac{d}{ds}H_s^{XX,j\leftarrow i} &= G_{\mathrm{eff},s}^{XX,j\leftarrow i} + A_j^{XX}(s)^\top H_s^{XX,j\leftarrow i} + H_s^{XX,j\leftarrow i}A_i^{XX}(s)\\ &\quad + \sum_{k'\in\mathcal{N}_j} B_X^{Y,k'}(s)^\top \mathfrak{B}^{k'} H_s^{W^{k'}X,j\leftarrow i} + \sum_{k\in\mathcal{N}_i} H_s^{XW^k,j\leftarrow i}\mathfrak{B}^{k\top}B_X^{Y,k}(s), \qquad H_T^{XX,j\leftarrow i} = G^{XX,j}(T).\end{aligned}} \tag{6.5.28}$$

Here and below the terminal condition appears blockwise; all non-XX terminal blocks are zero, matching the operator terminal $\Pi_X^{j,\top}G^{XX,j}(T)\Pi_X^i$ of (6.5.16).

For each $k\in\mathcal{N}_i$, matching the coefficient of $\delta_i w_s^k(r)$ gives

$$\boxed{\begin{aligned}-\frac{d}{ds}H_s^{XW^k,j\leftarrow i}(r) &= G_{\mathrm{eff},s}^{XW^k,j\leftarrow i}(r) + A_j^{XX}(s)^\top H_s^{XW^k,j\leftarrow i}(r) + H_s^{XX,j\leftarrow i}A_i^{XW^k}(s,r)\\ &\quad - H_s^{XW^k,j\leftarrow i}\mathfrak{B}^{k\top}B_X^{Y,k}(s)X(s,r)\\ &\quad + \sum_{k'\in\mathcal{N}_j} B_X^{Y,k'}(s)^\top\mathfrak{B}^{k'}H_s^{W^{k'}W^k,j\leftarrow i}(r), \qquad H_T^{XW^k,j\leftarrow i}(r) = 0.\end{aligned}} \tag{6.5.29}$$

The $X(s,r)$ term inherits its sign from the residual convention

$$\delta_i\widetilde{X}_s^k = \delta_i X_s - \int_0^s X(s,r)\delta_i w_s^k(r)\,dr. \tag{6.5.30}$$

For each row $k'\in\mathcal{N}_j$, matching the $\delta_i X_s$ coefficient gives

$$\boxed{\begin{aligned}-\frac{d}{ds}H_s^{W^{k'}X,j\leftarrow i}(u) &= G_{\mathrm{eff},s}^{W^{k'}X,j\leftarrow i}(u) + A_j^{XW^{k'}}(s,u)^\top H_s^{XX,j\leftarrow i} - X(s,u)^\top B_X^{Y,k'}(s)^\top\mathfrak{B}^{k'}H_s^{W^{k'}X,j\leftarrow i}\\ &\quad + H_s^{W^{k'}X,j\leftarrow i}(u)A_i^{XX}(s) + \sum_{k\in\mathcal{N}_i} H_s^{W^{k'}W^k,j\leftarrow i}(u)\mathfrak{B}^{k\top}B_X^{Y,k}(s), \qquad H_T^{W^{k'}X,j\leftarrow i}(u) = 0.\end{aligned}} \tag{6.5.31}$$

For each row $k'\in\mathcal{N}_j$ and column $k\in\mathcal{N}_i$, matching the coefficient of $\delta_i w_s^k(r)$ gives

$$\boxed{\begin{aligned}-\frac{d}{ds}H_s^{W^{k'}W^k,j\leftarrow i}(u,r) &= G_{\mathrm{eff},s}^{W^{k'}W^k,j\leftarrow i}(u,r) + A_j^{XW^{k'}}(s,u)^\top H_s^{XW^k,j\leftarrow i}(r)\\ &\quad - X(s,u)^\top B_X^{Y,k'}(s)^\top\mathfrak{B}^{k'}H_s^{W^{k'}W^k,j\leftarrow i}(r) + H_s^{W^{k'}X,j\leftarrow i}(u)A_i^{XW^k}(s,r)\\ &\quad - H_s^{W^{k'}W^k,j\leftarrow i}(u)\mathfrak{B}^{k\top}B_X^{Y,k}(s)X(s,r), \qquad H_T^{W^{k'}W^k,j\leftarrow i}(u,r) = 0.\end{aligned}} \tag{6.5.32}$$

Together these are the rectangular backward system. It is nonlinear because the response kernels in (6.5.23) and the source blocks (6.5.19)–(6.5.22) depend algebraically on the same coefficients.

6.5.5 The deviating player's own response and its first-order condition

By reflexivity, the deviating player is one of the privy players. Under the blip-of-irrationality convention (Section 6.4.2), the deviating player's post-blip continuation is its equilibrium continuation law, so its response to its own seed is the diagonal response $D^{i\leftarrow i}$. The diagonal $j = i$ solves a backward equation of the form $\dot{H}^{i\leftarrow i} = -G^{i\leftarrow i}_{\text{eff}} - A_i^* H^{i\leftarrow i} - H^{i\leftarrow i} A_i$, a source plus the generator acting on each side. It is the rectangular Sylvester equation (6.5.16) at $j = i$, expanded into the blocks (6.5.28)–(6.5.32). It is special in two ways. First, the deviating player's continuation control moves with the seed, so its own response is charged in the second variation rather than dropped. Second, the row and column maps coincide. The source is i's running-cost Hessian evaluated on its own response kernel,

$$G^{i\leftarrow i}_{\text{eff},s} = \begin{pmatrix} \Pi^i_X \\ D^{i\leftarrow i} \end{pmatrix}^{\top} G^i(s) \begin{pmatrix} \Pi^i_X \\ D^{i\leftarrow i} \end{pmatrix}. \tag{6.5.33}$$

Here Π^i_X is the physical-state projection of (6.5.17). Contracting the $j = i$ case of (6.5.19)–(6.5.22) against the seed gives, on the X row,

$$\left[G^{i\leftarrow i}_{\text{eff},s} \delta_i S^i_s\right]^X = G^{XX,i\leftarrow i}_{\text{eff},s} \delta_i X_s + \sum_{k\in\mathcal{N}_i} \int_0^s G^{XW^k,i\leftarrow i}_{\text{eff},s}(r) \delta_i w^k_s(r)\, dr, \tag{6.5.34}$$

and, on each belief row,

$$\left[G^{i\leftarrow i}_{\text{eff},s} \delta_i S^i_s\right]^{W^k}(u) = G^{W^k X,i\leftarrow i}_{\text{eff},s}(u) \delta_i X_s + \sum_{k'\in\mathcal{N}_i} \int_0^s G^{W^k W^{k'},i\leftarrow i}_{\text{eff},s}(u,r) \delta_i w^{k'}_s(r)\, dr. \tag{6.5.35}$$

The deviating player's first-order condition is the coefficient of an arbitrary $\mathcal{F}^i_t$-

measurable control spike. The condition is

$$\boxed{0 = \mathbb{E}\Big[G_t^{DD,i} D_t^i + G_t^{DX,i} X_t + (B_i^X(t))^\top H_t^{X,i} \mid \mathcal{F}_t^i\Big]}\,, \tag{6.5.36}$$

or equivalently

$$\boxed{D_t^i = -\Big(G_t^{DD,i}\Big)^{-1} \mathbb{E}\Big[G_t^{DX,i} X_t + (B_i^X(t))^\top H_t^{X,i} \mid \mathcal{F}_t^i\Big]\,.} \tag{6.5.37}$$

Read as a response kernel rather than a level, this condition is the $j = i$ case of (6.5.10).

Remark 6.10 (Symmetry of the own-response equation without a self-adjoint generator). The diagonal equation has the form

$$\dot{H}^{i\leftarrow i} = -G_{\text{eff}}^{i\leftarrow i} - A_i^* H^{i\leftarrow i} - H^{i\leftarrow i} A_i. \tag{6.5.38}$$

The full generator A_i is not self-adjoint, by the belief–belief asymmetry noted after (6.5.6), so the row and column equations need not look term-by-term symmetric. What is symmetric is the diagonal source and terminal condition, so the solution preserves the adjoint symmetry

$$H^{XW^k,i\leftarrow i}(r) = \Big(H^{W^k X,i\leftarrow i}(r)\Big)^\top, \qquad H^{W^{k'} W^k,i\leftarrow i}(u,r) = \Big(H^{W^k W^{k'},i\leftarrow i}(r,u)\Big)^\top, \tag{6.5.39}$$

provided the usual uniqueness conditions hold.

6.6 Superposition under perfect monitoring

For a fixed origin i, the privy players in $\mathcal{P}_i$ all see the same seed and solve one self-consistency problem, coupled only through the shared forward state. The monitoring relation orders who sees whose seeds. Each origin produces its own gain system of the form (6.5.28)–(6.5.32); the systems take different values because the naive/privy split depends on the origin.

Write $R_t^{\leftarrow i}(s)$ for the per-unit privy response to an origin-i seed at time s, and R_t for the response to a configuration of seeds.

Proposition 6.11 (Superposition of monitored seeds)**.** *Under Assumption 6.4, the forward system is affine in the seeds, so the response to several monitored seeds is the sum of the single-seed responses,*

$$R_t\Big(\sum_m v_{i_m}(s_m)\Big) = \sum_m R_t^{\leftarrow i_m}(s_m)\, v_{i_m}(s_m). \tag{6.6.1}$$

If player i blips, then the response is the same regardless of whether some other player j blipped sooner or not. Without superposition, verifying would mean thinking about an infinite tree of deviations rather than treating them one at a time. The gains themselves are solved in order rather than independently, since the source (6.5.19) for $j \neq i$ contains $D^{j\leftarrow j}$, player j's response to its own seed, so each $j \leftarrow i$ solve uses the $j \leftarrow j$ solve first. The pairs $j \leftarrow i$ with $j \neq i$ for one origin are coupled through A_i^{XX}, which contains every $D^{q\leftarrow i,X}$, $q \in \mathcal{P}_i$.

Remark 6.12 (The role of perfect observation). The linear theory relies on binary monitoring. A player either observes a deviation seed perfectly, as a known zero-variance input not jointly estimated with the fundamentals, or not at all, in which case it is folded into the ordinary filter with the equilibrium filter fixed. Keep two relaxations apart. In the first, a player observes the seed through an additive noisy channel. A noisy reading of a point mass at zero carries no information, so this channel is informative only if the seed carries a nondegenerate prior. Deviations then occur with positive probability on the equilibrium path, and the object being solved is no longer a deviation calculus. In the second, transitivity fails and a player sees a response whose origin it cannot resolve. The response discrepancy is identically zero on the equilibrium path, so the observation is a null event, and attribution conditional on it is a mixture over the origins the player cannot rule out. The mixture weights play the role that off-path beliefs play in sequential equilibrium [68]. Given the weights the theory remains linear, but superposition fails, since attribution couples the per-origin blocks into one system.

6.7 The two corners

For a fixed origin i, solve (6.5.28)–(6.5.32) for each privy player $j \in \mathcal{P}_i$. By Lemma 6.2 the columns $\mathcal{N}_i$ are contained in the rows $\mathcal{N}_j$.

The variations $\delta_i H^{X,j}$ and $\delta_i H^{W^{k'},j}$ in (6.5.24)–(6.5.25) are the marginal values of the seed to player j. The responder's own first-order condition pins each response kernel rather than assuming it. This is the consistency required of conjectured reactions in oligopoly theory [22], and the selection that closed-loop differential games otherwise leave open [16].

The two corners of the monitoring relation recover the two classical regimes.

Proposition 6.13 (Corner reductions)**.** *Maintain Assumption 6.4.* (i) (All privy.) *Fix an origin i with $\mathcal{N}_i = \varnothing$. Then the diagonal $j = i$ of* (6.5.28)–(6.5.32) *retains only its XX block, a Riccati equation in $H^{XX,i\leftarrow i}$ with row and column coefficient A_i^{XX}. If $\mathcal{P}_a = \mathcal{I}$ for every origin a, so every naive set is empty, then every solve retains only its XX block and its solution does not depend on the origin, $H_s^{XX,j\leftarrow i} = H_s^{XX,j}$ for all i. Writing $D_s^{j,X}$ for the common response kernel $D_s^{j\leftarrow i,X}$, the pairs $(H^{XX,j}, D^{j,X})$ solve the coupled Riccati system of the closed-loop Nash equilibrium of the perfect-information game [16, 26],*

$$
\begin{aligned}
-\dot{H}_s^{XX,j} &= G^{XX,j}(s) + G^{XD,j}(s)D_s^{j,X} + (D_s^{j,X})^\top G^{DX,j}(s) \\
&\quad + (D_s^{j,X})^\top G^{DD,j}(s)D_s^{j,X} + A^{XX}(s)^\top H_s^{XX,j} + H_s^{XX,j}A^{XX}(s), \\
H_T^{XX,j} &= G^{XX,j}(T), \qquad\qquad (6.7.1) \\
D_s^{j,X} &= -\Big(G^{DD,j}(s)\Big)^{-1}\Big[G^{DX,j}(s) + (B_j^X(s))^\top H_s^{XX,j}\Big], \\
A^{XX}(s) &= B_X^X(s) + \sum_m B_m^X(s)D_s^{m,X}.
\end{aligned}
$$

No seed moves any player's ordinary filter. (ii) (All naive.) *Fix an origin i with $\mathcal{P}_i = \{i\}$. Then only the diagonal $j = i$ remains, and its rows and columns are both indexed by $\mathcal{N}_i = \mathcal{I} \setminus \{i\}$. The system (6.5.28)–(6.5.32) is the coefficient form of the i-variation of the adjoint system of Theorem 1.8 of Chapter 1, written in the blip basis of Lemma 6.6, and the exact verification identity of Corollary 1.11 transfers through the basis map.*

Proof. (i) With $\mathcal{N}_i = \varnothing$ the diagonal has no rows or columns beyond X, so only (6.5.28) remains, its source the $j = i$ case of (6.5.19) and its row and column coefficient A_i^{XX}. With every naive set empty the W rows and columns are absent for every pair. The candidate $H^{XX,j\leftarrow i} = H^{XX,j}$ matches the terminal condition for every pair. Substituting it into (6.5.23) gives $D^{j\leftarrow i,X} = D^{j,X}$ independent of the origin, so $A_a^{XX} = A^{XX}$ for every origin label a in (6.5.8), and the source (6.5.19) becomes the first four terms of (6.7.1). The right sides of (6.5.28) and (6.7.1) then coincide, so the candidate solves the corner system, and it is the solution whenever the closed-loop Riccati system (6.7.1) has a unique solution. That no seed moves any ordinary filter is (6.4.2) applied to every player.

(ii) With $\mathcal{P}_i = \{i\}$ only the diagonal pair remains, with rows and columns both indexed by the deviating player's naive set. Expanding the i-variation of (1.4.7)–(1.4.8) of Chapter 1 in the coefficients of $(\delta_i X, \delta_i w^k)$ and comparing with (6.5.28)–(6.5.32), the terms of (6.5.19) in $D^{i\leftarrow i}$ and the closed-loop part of the row and column coefficient A_i^{XX} differ from the frozen-continuation expansion by

$$
\begin{aligned}
&G^{XD,i} D_s^{i\leftarrow i,X} + (D_s^{i\leftarrow i,X})^\top G^{DX,i} + (D_s^{i\leftarrow i,X})^\top G^{DD,i} D_s^{i\leftarrow i,X} \\
&\qquad + H_s^{XX,i\leftarrow i} B_i^X D_s^{i\leftarrow i,X} + (D_s^{i\leftarrow i,X})^\top (B_i^X)^\top H_s^{XX,i\leftarrow i} = -(D_s^{i\leftarrow i,X})^\top G^{DD,i} D_s^{i\leftarrow i,X},
\end{aligned}
$$

by the $j = i$ case of (6.5.23) and its transpose, and blockwise the same for the W rows and columns, so the two systems differ by the operator $-(D^{i\leftarrow i})^* G^{DD,i} D^{i\leftarrow i}$. The two conventions are the two continuations of Lemma 6.6, whose change of variable $Q_b = (\mathrm{Id} + K)^* Q_e (\mathrm{Id} + K)$ carries this term, and the lemma transfers the verification identity. ☐

Remark 6.14 (Lyapunov and Riccati forms). The frozen continuation holds the deviating player's control fixed after the spike, so its backward equation is linear in H, a Lyapunov equation. The blip continuation lets the deviating player re-optimize, and substituting the gain (6.5.23) makes the equation quadratic in H, a Riccati equation. The term $-(D^{i\leftarrow i,X})^\top G^{DD,i} D^{i\leftarrow i,X}$ is the value of that re-optimization.

6.8 A strategic market maker: transparent and opaque markets

In the Kyle–Back markets of Chapter 4 the market maker is not a player. It posts the conditional expectation of the fundamental given order flow, and the only strategic question is how the informed traders trade against that rule. Here the market maker chooses its quote and cares about its inventory, and whether the trader can see its deviations depends on whether order flow is published. Publishing the flow turns the same market, with the same information on the equilibrium path, from a naive corner into a privy one. The market maker observes the order flow, so its observation loads on the trader's control rather than on the state, which puts the market outside the baseline (6.2.1) and inside Chapter 4. The origin sets and the blip convention are those of this chapter; the trader's response kernels come from Chapter 4's adjoint system.

6.8.1 The market

One asset, one informed trader (player 1), one market maker (player 0), a finite horizon T. The fundamental V_t is a Brownian motion with variance rate Σ_V, the market maker posts a quote P_t, and both players' payoffs depend on the two only through the mispricing

$$M_t := V_t - P_t, \tag{6.8.1}$$

the gap between what the asset is worth and what it is quoted at. The trader observes the signal flow

$$dY_t = M_t\,dt + dW_t^Y \tag{6.8.2}$$

of Chapter 4 and submits an order rate D_t; noise traders submit $\Sigma_Z^{1/2} dW_t^Z$; the market maker sees the total order flow

$$dZ_t = D_t\,dt + \Sigma_Z^{1/2} dW_t^Z \tag{6.8.3}$$

and absorbs it into its inventory,

$$dQ_t = -dZ_t. \tag{6.8.4}$$

The trader is risk neutral with a quadratic trading cost and earns the mispricing on its orders,

$$\mathbb{E}\int_0^T \Big[M_t D_t - \epsilon D_t^2\Big]\,dt, \tag{6.8.5}$$

and the market maker pays it on the flow it absorbs and pays for inventory,

$$\mathbb{E}\int_0^T \Big[-M_t D_t - \gamma\, Q_t^2\Big]\,dt, \qquad \gamma > 0. \tag{6.8.6}$$

The mispricing is both the market maker's control and something it does not know. It sets M_t through its quote, but only up to its own estimation error: with $\widehat{V}_t^0 = \mathbb{E}[V_t \mid \mathcal{F}_t^0]$,

$$M_t = M_t^0 + \Big(V_t - \widehat{V}_t^0\Big), \qquad M_t^0 := \mathbb{E}[M_t \mid \mathcal{F}_t^0] = \widehat{V}_t^0 - P_t, \tag{6.8.7}$$

and M_t^0 is its control; $M_t^0 > 0$ is a quote below the belief. The trader likewise acts on its own estimate $M_t^1 := \mathbb{E}[M_t \mid \mathcal{F}_t^1] = \widehat{V}_t^1 - P_t$, which for the trader is a state of its filter. (M_t, Q_t) is the state, of which only Q_t has dynamics of its own, (W^V, W^Z, W^Y) are the shocks, and V and P appear only through M_t. At $\gamma = 0$ the first-order condition returns $M_t^0 = 0$, the competitive rule of Chapter 4; with $\gamma > 0$ the market maker quotes off it, and γ stays fixed across the two cases below.

Two markets. The market maker's information is the order flow in both cases, $\mathcal{F}_t^0 = \sigma(Z_s, s \le t)$, and it cannot see the trader's order apart from the noise. Only the trader's information differs.

1. *Transparent market.* The trader sees the quote and the order flow, $\mathcal{F}_t^1 = \sigma(Y_s, P_s, Z_s, s \le t)$.

2. *Opaque market.* The trader sees the quote only, $\mathcal{F}_t^1 = \sigma(Y_s, P_s, s \le t)$.

On the equilibrium path the two coincide. The quote is a function of the flow history, and when that function is invertible the trader recovers the flow from the price. Off the path they differ. In the transparent market the trader can compare the posted quote with the equilibrium rule applied to the flow it sees. Any gap is the market maker's, so the trader is privy to its deviations, $1 \trianglerighteq 0$. In the opaque market a quote off the rule looks like a quote on the rule after an unusual run of noise trades. The trader filters the gap as noise, $1 \ntrianglerighteq 0$. In both markets the trader's order is hidden in the flow, $0 \ntrianglerighteq 1$. For a trader origin, $\mathcal{N}_1 = \{0\}$ in both markets: the trader keeps the information wedge of Chapter 4 against the market maker's filter. For a market-maker origin, $\mathcal{P}_0 = \{0, 1\}$ and $\mathcal{N}_0 = \varnothing$ in the transparent market, while $\mathcal{N}_0 = \{1\}$ in the opaque one.

What a quote deviation does. A quote below the rule draws extra buying now, since the trader sees a larger mispricing. The extra buying lowers the market maker's inventory, and the inventory outlives the quote: once the market maker is back on its rule, the quote keeps moving with the inventory and the trader keeps trading against it. The market maker therefore weighs the margin it gives up on the flow it already expects against the inventory the extra flow removes. The trader weighs the mispricing it earns against the inventory its order leaves behind and, where the market maker cannot see the order, against the beliefs it moves.

The seed. A control spike z of width h is a quote z below the rule. It raises the mispricing M by z during the blip, and the trader's estimate M^1 with it in both markets. All information about V reaches the market through the trader, so the quote carries none and a quote seed moves no belief about the fundamental. With $D_t^{1\leftarrow 0,M}$ the trader's contemporaneous loading on the mispricing, the trader raises its order by $D_t^{1\leftarrow 0,M} z$, which lowers the inventory drift by as much and over the blip displaces the inventory by

$$\delta_0 Q = -D_t^{1\leftarrow 0,M}\, z\, h, \tag{6.8.8}$$

of order h and permanent. Every origin-0 system below has this displacement as its state column.

The market maker's cost. The market maker maximizes (6.8.6), so its running cost is the inventory penalty plus the mispricing paid on the flow, $\gamma Q_t^2 + M_t D_t$. The market maker sees neither factor of the product, so it minimizes the expectation given $\mathcal{F}_t^0$. That is the product of the expectations plus a covariance, and the expectation of the mispricing is its control,

$$\mathbb{E}\Big[\gamma Q_t^2 + M_t D_t \,\Big|\, \mathcal{F}_t^0\Big] = \gamma Q_t^2 + M_t^0\, \mathbb{E}[D_t \mid \mathcal{F}_t^0] + \mathrm{Cov}\Big(M_t, D_t \,\Big|\, \mathcal{F}_t^0\Big). \tag{6.8.9}$$

The covariance is the adverse-selection cost, what the trader earns from knowing more. Since the quote is known to the market maker it equals $\mathrm{Cov}(V_t, D_t \mid \mathcal{F}_t^0)$, an equilibrium object that the quote does not move. With the trader's loadings substituted, the market maker's estimate of the order is $\mathbb{E}[D_t \mid \mathcal{F}_t^0] = D_t^{1\leftarrow 0,M} M_t^0 + D_t^{1\leftarrow 0,Q} Q_t$. The running cost is then a quadratic form in the market maker's own mispricing and its inventory, the block form of (6.2.4),

$$\ell_t^0 = D_t^{1\leftarrow 0,M}\Big(M_t^0\Big)^2 + D_t^{1\leftarrow 0,Q} M_t^0 Q_t + \gamma Q_t^2 + \mathrm{Cov}\Big(M_t, D_t \,\Big|\, \mathcal{F}_t^0\Big), \tag{6.8.10}$$

so $G^{MM,0} = 2D_t^{1\leftarrow 0,M}$, $G^{MQ,0} = D_t^{1\leftarrow 0,Q}$, $G^{QQ,0} = 2\gamma$. The trader's reaction supplies the curvature of the market maker's cost in its own control. The trader's running cost is $-M_t D_t + \epsilon D_t^2$, so $G^{DD,1} = 2\epsilon$. Writing the mispricing through the quote rule, $M_t = (V_t - \widehat{V}_t^0) - P_t^Q Q_t$, separates the part the trader knows better from the part the inventory has already moved. It gives the two cross blocks $G^{VD,1} = -1$ and $G^{QD,1} = P_t^Q$, with $P_t^Q := \partial P_t / \partial Q_t$ the quote's inventory loading found below.

6.8.2 The transparent market

A quote deviation: the trader's response. The trader is privy, so its estimates do not move and it responds through the two coordinates the seed moves,

$$\delta_0 D_s = D_s^{1\leftarrow 0,M}\, \delta_0 M_s + D_s^{1\leftarrow 0,Q}\, \delta_0 Q_s. \tag{6.8.11}$$

The first loading is the trader's on-path loading on its own mispricing; a quote shift is, for the trader, a movement of M like any other. The second is its loading on the market maker's inventory with the quote held fixed,

$$D_s^{1\leftarrow 0,Q} = \frac{1}{2\epsilon}\Big[H_s^{QQ,1\leftarrow 0} - \Big(H_s^{W^0Q,1\leftarrow 0}\mathfrak{B}^{0\top}\Big)\Sigma_Z^{-1/2}\Big], \tag{6.8.12}$$

where $H^{QQ,1\leftarrow 0}$ is the value to the trader of a unit of market-maker inventory and $H^{W^0Q,1\leftarrow 0}(u)$ is that unit's effect on the price the trader puts on the market maker's beliefs. The contraction against the market maker's filter is $N\mathfrak{B}^{0\top} := N(s)E^{Z\top} + \int_0^s N(r)\, \widetilde{C}_s^0(r)^\top dr$, as in Chapter 4. A departure from this loading is ordinary flow to the market maker. Adding the trader's reaction to the quote's own response to inventory gives the closed-loop drift of an inventory displacement,

$$a_Q(s) = -D_s^{1\leftarrow 0,Q} + D_s^{1\leftarrow 0,M} P_s^Q, \tag{6.8.13}$$

negative when the displacement unwinds; $-a_Q$ is the order's total loading on inventory. When the market maker's inventory aversion outweighs the trader's loading, $P^Q_s < 0$, and a long market maker quotes below its belief. The two gains solve

$$-\frac{d}{ds}H^{QQ,1\leftarrow 0}_s = a_Q(s)\Big(H^{QQ,1\leftarrow 0}_s - P^Q_s\Big), \tag{6.8.14}$$

$$-\partial_s H^{W^0Q,1\leftarrow 0}_s(u) = -a_Q(s)\Sigma_V^{1/2}E^V - \Big(H^{W^0Q,1\leftarrow 0}_s \mathfrak{B}^{0\top}\Big)\Sigma_Z^{-1/2}D_{W,s}(u) + a_Q(s)H^{W^0Q,1\leftarrow 0}_s(u), \tag{6.8.15}$$

with zero terminal conditions. A unit of inventory is worth, to the trader, the orders it will draw times the quote's inventory loading net of the inventory gain, accumulated along the unwinding. It moves the price the trader puts on the market maker's beliefs by the same orders times the quote's lift per unit of belief, $\Sigma_V^{1/2}E^V$, which the market maker then filters away. On the path no belief moves, but the loading the market maker faces depends on how it would misfilter a departure from it.

A quote deviation: the forward system. With $\mathcal{N}_0 = \varnothing$ there are no naive noise-states, so the filter equations (6.5.6)–(6.5.7) are absent and the displaced state of Section 6.5 is the inventory alone. The seed is the state impulse (6.5.1) through the market maker's control-to-state map,

$$\delta_0 M_t = z, \qquad \delta_0 Q_t = 0, \tag{6.8.16}$$

a quote z below the rule and an inventory not yet moved. After the blip the market maker's control returns to its rule, so the only mispricing left is the quote's response to the displaced inventory, and the trader answers through both kernels,

$$\begin{aligned} \delta_0 M_s &= -P^Q_s\,\delta_0 Q_s, \\ \delta_0 D_s &= D^{1\leftarrow 0,M}_s\,\delta_0 M_s + D^{1\leftarrow 0,Q}_s\,\delta_0 Q_s = \Big(D^{1\leftarrow 0,Q}_s - D^{1\leftarrow 0,M}_s P^Q_s\Big)\,\delta_0 Q_s. \end{aligned} \tag{6.8.17}$$

Substituting into (6.5.3) closes the system on the inventory,

$$\begin{aligned} \frac{d}{ds}\delta_0 Q_s &= \Big(-D_s^{1\leftarrow 0,Q} + D_s^{1\leftarrow 0,M} P_s^Q\Big)\,\delta_0 Q_s, \\ \delta_0 Q_s &= -D_t^{1\leftarrow 0,M}\, z\, h\, \exp\!\bigg(\int_t^s \Big(-D_r^{1\leftarrow 0,Q} + D_r^{1\leftarrow 0,M} P_r^Q\Big)\, dr\bigg), \end{aligned} \tag{6.8.18}$$

so the seed's whole forward trace is one exponential at the closed-loop drift, and it decays exactly when that drift is negative.

A quote deviation: what the market maker does after its blip. The market maker's own continuation is the diagonal $j = i = 0$ on the inventory coordinate, with the blocks of (6.8.10) and its own response kernel $D_s^{0\leftarrow 0,Q} = -P_s^Q$ substituted,

$$G_{\text{eff},s}^{QQ,0\leftarrow 0} = 2\gamma - 2D_s^{1\leftarrow 0,Q} P_s^Q + 2D_s^{1\leftarrow 0,M}\Big(P_s^Q\Big)^2, \tag{6.8.19}$$

and the row and column maps coincide on the diagonal (Section 6.5.5), both being the closed-loop drift (6.8.13). The covariance in (6.8.10) is dropped, since the quote does not move it. Here $H_s^{QQ,0}$ is the marginal value of inventory per unit of inventory. The quote's inventory loading is the trader's direct loading on inventory less the inventory price, divided by the curvature of profit in the quote,

$$M_t^0 = -P_t^Q Q_t, \qquad P_t^Q = \frac{D_t^{1\leftarrow 0,Q} - D_t^{1\leftarrow 0,M} H_t^{QQ,0}}{2D_t^{1\leftarrow 0,M}}, \qquad \text{so} \qquad P_t = \widehat{V}_t^0 + P_t^Q Q_t, \tag{6.8.20}$$

Read backward from T, where inventory has no value, $H^{QQ,0}$ accumulates three things: the cost of holding inventory net of what the trader's loading on it earns, the drift of inventory at that loading, and the saving from quoting against it,

$$-\frac{d}{ds}H_s^{QQ,0} = 2\gamma - \frac{\Big(D_s^{1\leftarrow 0,Q}\Big)^2}{2D_s^{1\leftarrow 0,M}} - D_s^{1\leftarrow 0,Q} H_s^{QQ,0} - \frac{D_s^{1\leftarrow 0,M}}{2}\Big(H_s^{QQ,0}\Big)^2, \qquad H_T^{QQ,0} = 0, \tag{6.8.21}$$

the closed-loop Riccati equation of corner (i) on the inventory coordinate. The inventory price in the first-order condition is $H_t^{Q,0} = H_t^{QQ,0} Q_t$, a public quantity. There is no belief row and no belief column, since no seed of the market maker's moves any filter and its own filter is public. At $\gamma = 0$ the pair $D^{1\leftarrow 0,Q} \equiv 0$, $H^{QQ,0} \equiv 0$ solves (6.8.20)–(6.8.21) and gives $P = \widehat{V}^0$. The trader loads on inventory only because the market maker quotes against it, and the market maker quotes against inventory only because the trader loads on it, so the competitive rule is a fixed point of that loop. Whether the loop has a second, self-fulfilling fixed point at $\gamma = 0$ is not settled here.

An order deviation: the trader's own condition. Only the market maker is naive to a trader seed, $\mathcal{N}_1 = \{0\}$, so it indexes both the rows and the columns. The diagonal $j = i = 1$ is the coefficient form of Chapter 4's adjoint system, corner (ii), with the inventory added. Two entries are new. Because the market maker quotes against its inventory, the trader's cost acquires the cross term $D_s P_s^Q Q_s$, the derivative of $-M_s D_s$ in Q_s through the quote. Q_t is $\mathcal{F}_t^0$-measurable, so the market maker's filter has no Q coordinate and its observation loads on the trader's order rather than on Q. The Q row of the trader's state price therefore has no filter coupling and involves only $H^{Q,1}$ and the on-path order D_s. With a single trader there is no other trader whose reaction to the inventory the deviating player faces, so the row has no drift term. The market maker's quoting against inventory sits in the trader's cost, through M, not in the drift of Q. The trader's inventory price is its own future orders, each weighted by how much the quote it will face moves per unit of inventory,

$$-\frac{d}{ds} H_s^{Q,1} = P_s^Q D_s, \qquad H_T^{Q,1} = 0, \qquad \text{so} \qquad H_t^{Q,1} = \int_t^T P_s^Q D_s \, ds. \tag{6.8.22}$$

Against the competitive market maker $P^Q \equiv 0$ and the row vanishes.

First-order conditions. The market maker's condition is corner (i), with B_0^Q replaced by the closed-loop map $-D^{1\leftarrow 0,M}$, since the quote moves the inventory only through the trader's order. Quoting a unit further below the belief does three things. It gives up a unit of margin on the flow already expected, $\mathbb{E}[D_t \mid \mathcal{F}_t^0]$. It attracts new flow $D_t^{1\leftarrow 0,M}$, on which the current margin M_t^0 is lost. And it lowers the inventory by that new flow, which is worth the inventory price. At the optimum the three cancel,

$$0 = \mathbb{E}[D_t \mid \mathcal{F}_t^0] + D_t^{1\leftarrow 0,M} M_t^0 - D_t^{1\leftarrow 0,M}\, H_t^{QQ,0} Q_t. \tag{6.8.23}$$

There is no wedge. With $\mathbb{E}[D_t \mid \mathcal{F}_t^0]$ written out this is (6.8.20): a long market maker quotes below its belief to attract the buying that unwinds the position. The trader's is the deviating player's condition with the market maker as its naive set, corner (ii), Theorem 4.10 of Chapter 4 on the enlarged state. Besides earning the mispricing, an order moves the market maker's inventory, priced by the inventory row, and its beliefs, priced by the belief prices as in Chapter 4. The trader sets its estimated mispricing equal to the expected cost of the two,

$$M_t^1 = 2\epsilon D_t^1 + \mathbb{E}\Big[- H_t^{Q,1} + \int_{-\infty}^{t} H_t^{W^0,1}(r)\, \varphi_1^0(t,r)\, dr + H_t^{W^0,1}(t)\, \psi_1^0(t) \,\Big|\, \mathcal{F}_t^1 \Big]. \tag{6.8.24}$$

The belief terms are those of (4.5.9) with the adjoints now propagating through the quote rule (6.8.20). Differentiating (6.8.24) in M_t^1 with the inventory and belief terms held fixed, since a privy trader's filter does not move and the inventory has not yet, gives the contemporaneous loading $D_t^{1\leftarrow 0,M} = 1/(2\epsilon)$. In the transparent market only the trader filters.

6.8.3 The opaque market

What changes. The trader now sees the quote but not the flow, so it cannot separate a quote off the rule from a quote on the rule after unusual noise. That holds for a quote shaded smoothly, which is absolutely continuous against the diffusing rule quote. A blip is a jump and would be seen; it enters only as the basis from which the response to a smooth shading is assembled by linearity. For an origin-0 seed the naive set is $\mathcal{N}_0 = \{1\}$. The displaced state acquires a second column beside the inventory displacement, the displacement of the trader's estimates of every past shock,

$$\Big(\delta_0 Q_s,\ \{\delta_0 w_s^1(u) : 0 \le u \le s\}\Big), \tag{6.8.25}$$

which evolves by (6.5.3)–(6.5.7) with one difference in where the seed enters. The quote is an observation of the trader's. Knowing its own order, the trader reads the noise flow exactly from the quote's martingale part $(\Lambda_t - P_t^Q)\Sigma_Z^{1/2} dW_t^Z$, with Λ the market maker's loading on order flow, whenever $\Lambda_t \neq P_t^Q$, so a quote seed enters the trader's filter through that channel rather than through a state displacement. Because P is observed directly, its movement in the signal drift $M_t = V_t - P_t$ cancels from the trader's innovation, as in Chapter 4. The trader reads a shading δP as the noise flow that would have produced it, $\delta_0 w_s^Z(u) = \delta P_u / \Big((\Lambda_u - P_u^Q)\Sigma_Z^{1/2}\Big)$, a density born on the window, which the birth term (6.5.7) carries. Its perceived inventory and its perceived market-maker belief then move by the on-path maps, $-\Sigma_Z^{1/2}\int \delta_0 w^Z$ and $\Lambda\Sigma_Z^{1/2}\int \delta_0 w^Z$, while the actual inventory has not moved: the trader expects an unwinding that is not coming. The trader no longer responds through the response kernel; it applies its policy kernel to the displaced estimates, its estimate of the inventory among them,

$$\delta_0 D_s = \int_0^s D_s^1(u)\, \delta_0 w_s^1(u)\, du, \tag{6.8.26}$$

and the market maker's own gain system acquires the rows dual to the new column, $H_s^{W^1 Q,0}(u)$ and $H_s^{W^1 W^1,0}(u,r)$, computed below on the lag window. Paired with the seed's belief displacement they give the market maker a wedge of the kind the trader has against it, the value of the belief displacement a shading leaves in the trader's filter,

$$\mathcal{V}_t^0 = \int_{-\infty}^t H_t^{W^1,0}(u)\, \varphi_0^1(t,u)\, du + H_t^{W^1,0}(t)\, \psi_0^1(t), \tag{6.8.27}$$

quote's inventory loading	P^Q	-0.33
trader's quote-fixed inventory loading	$D^{1\leftarrow 0,Q}$	-0.70
inventory's closed-loop drift	a_Q	-0.12
belief correction to the inventory loading	κ	$+0.21$
loading on order flow	Λ	1.057
net price impact	$\Lambda - P^Q$	1.38

Table 6.2: The stationary transparent market at $\gamma = 0.1$ ($\rho = 0.5$, $\epsilon = 0.2$, unit signal gain and noise, $L = 8$, 60 cells). At $\gamma = 0$ the loadings on inventory vanish and $\Lambda = 1.053$.

and the market maker's first-order condition is the transparent one (6.8.23) with the inventory price projected onto the market maker's information and the wedge subtracted,

$$0 = \mathbb{E}[D_t \mid \mathcal{F}_t^0] + D_t^{1\leftarrow 0,M} M_t^0 - D_t^{1\leftarrow 0,M}\, \mathbb{E}\Big[H_t^{Q,0} \,\Big|\, \mathcal{F}_t^0\Big] - \mathbb{E}\Big[\mathcal{V}_t^0 \,\Big|\, \mathcal{F}_t^0\Big]. \tag{6.8.28}$$

The inventory price is projected now, since the market maker's belief is no longer public. In the last term the market maker shades its quote to make the trader believe noise traders have been active, and is paid for it through the trader's later orders. The trader's condition (6.8.24) keeps its form, its origin sets being unchanged, but every kernel in it is different, because the market maker it faces now sets its quote with the wedge in its first-order condition.

6.8.4 Computation and comparison

The transparent market, computed. The stationary form of the transparent market is solved on the lag window of Chapter 4, with the trader's problem exactly as there and the market maker's quote loading P^Q fixed by the origin-0 system above in its stationary form. At discount rate ρ the inventory gain is $H^{QQ,1\leftarrow 0} = a_Q P^Q/(a_Q - \rho)$, the market maker's Riccati equation is algebraic, and the belief row is a linear equation in age with one nonlocal scalar $\kappa = (H^{W^0 Q,1\leftarrow 0}\mathfrak{B}^{0\top})\Sigma_Z^{-1/2}$, the correction to the trader's inventory loading from how the market maker would filter a departure from it; P^Q is then a root of a scalar equation, found by tracking the root in γ upward from the competitive rule. The average-cost limit is degenerate: at $\rho = 0$ the trader's inventory value equals P^Q, its inventory loading responds to P^Q with the same slope as the market maker's rule, and the rule cannot pin P^Q. At $\rho = 0.5$, $\epsilon = 0.2$, unit signal gain and unit noise, a quarter of the impact a trader faces at $\gamma = 0.1$ is inventory, not information: the market maker's loading on order flow barely moves, from 1.053 to 1.057, while the net price impact of an order, $\Lambda - P^Q$, rises from 1.05 to 1.38. Table 6.2 collects the equilibrium at $\gamma = 0.1$; over the range computed the quote loading is linear in the inventory weight, $P^Q \approx -3.3\gamma$ for $\gamma \leq 0.15$ (Figure 6.2). At $\gamma = 0$ the computation returns the competitive kernels of Chapter 4 to 3×10^{-11}, and refining the grid from 60 to 120 cells moves P^Q by 0.2 percent. The gain equations admit a second self-consistent gain at $P^Q = 0$, with $a_Q \approx -2.3$ and the trader unwinding inventory it has no reason to touch; the exact best response rules it out, so the competitive rule is the only fixed point of the loop on this branch. Above $\gamma \approx 0.2$ the trader's second-order form goes indefinite on this grid at $\epsilon = 0.2$, the same boundary as in Chapter 4, and a larger trading cost is needed to certify larger inventory weights.

The opaque market, computed. On the lag window the market maker's problem after its blip, with the trader reading the shading as noise flow, is an ordinary linear-quadratic problem. Its state is the inventory and its own shading history, since the trader's displaced estimates are a function of the shadings; its control is today's shading; and the trader's response to a shading path is its on-path kernel applied to the noise flow it infers by inverting the quote rule, with the signal drift shifted to match. A stationary Riccati equation in $n+1$ unknowns gives the quote's inventory loading, and $H^{W^1Q,0}$ and $H^{W^1W^1,0}$ are its inventory–history and history–history blocks. With the trader's privy response $-1/(2\epsilon)$ in place of the inference, the same equation returns the transparent loading to 0.4 percent at 120 cells. At $\gamma = 0.1$ the opaque loading is $P^Q = -0.19$ against -0.33: the market maker loads its quote on inventory forty percent less when the trader cannot see the flow, and the ratio is 0.6 over the whole range, $P^Q \approx -1.9\gamma$ (Figure 6.2). The reason is the trader's response to a shading. Read as noise flow, a quote one unit off the rule draws 0.58 of the order it draws when seen for what it is, and afterwards

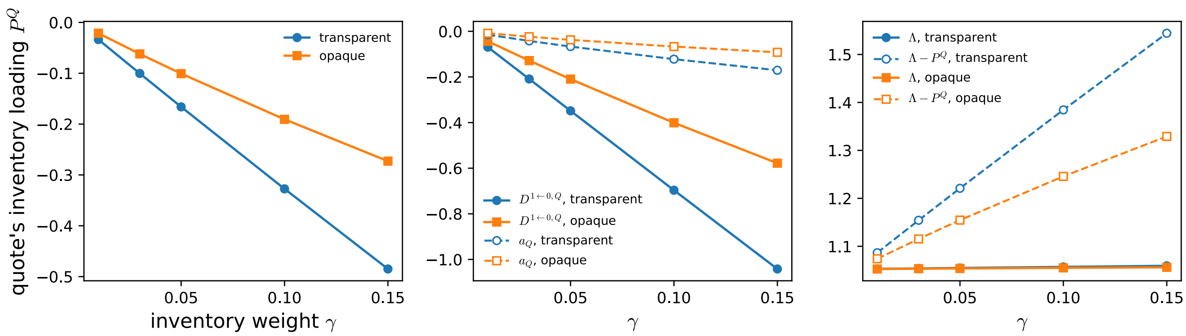

Figure 6.2: The stationary transparent and opaque markets against the inventory weight γ ($\rho = 0.5$, $\epsilon = 0.2$, unit signal gain and noise, $L = 8$, 60 cells). Left: the quote's inventory loading P^Q. Center: the trader's quote-fixed inventory loading and the inventory's closed-loop drift. Right: the market maker's loading on order flow and the net price impact $\Lambda - P^Q$.

the trader trades the other way, $+0.22$ in total over the window, toward the unwinding it expects and that is not coming. Shading is a weaker tool for managing inventory in the opaque market, so the market maker uses less of it: the trader's inventory loading is -0.40 against -0.70, the inventory unwinds at -0.076 against -0.12, and the net price impact is 1.25 against 1.38. The market maker also loads its optimal quote on its own past shadings, -0.04 on the most recent and decaying, a loading that never fires on the equilibrium path.

The solvers and figures of this section are the work of Claude Fable 5, a large language model built by Anthropic [7], working under the author's direction. Interactive versions of the dissertation's computations are collected at sbabichenko.com/noisestate.

What the pair identifies. Any difference between the two equilibria is the price of monitoring. Publishing order flow switches the market maker's wedge off, but it does not discipline the market maker: it loads its quote more on inventory in the transparent market, because a trader who can see the flow can be steered. Most of the gain from being monitored is real inventory efficiency, and the trader's orders carry the same amount of the fundamental into the price in both markets.

Table 6.3 gives the accounts at $\gamma = 0.1$. The noise traders buy at quotes their own flow has pushed up and sell at quotes it has pushed down. The information concession is identical in the two markets, and the inventory concession carries the whole difference.

	transparent	opaque
trader's price profits	1.07	1.01
noise traders' loss	1.25	1.13
information concession	0.96	0.95
inventory concession	0.30	0.17
market maker's total loss	0.28	0.63
$\mathbb{E}[Q^2] = \Sigma_Z/(2\lvert\bar{a}_Q\rvert)$	4.7	7.5
expected unwinding rate $\bar{a}_Q$	-0.106	-0.067

Table 6.3: The two markets at $\gamma = 0.1$, at the parameters of Table 6.2. Price profits are gross of a trading cost of 0.13 in both markets. The information concession is paid through the loading on order flow and the inventory concession through the quote's inventory loading. The market maker's loss is flow plus inventory penalty.

Transparency transfers 0.13 from the noise traders, split evenly between the two strategic players, and separately saves the market maker 0.28 of real inventory cost, since a trader who can see the flow unwinds the inventory for it.

Which relation holds is an institutional assumption; it does not follow from quotes being actions. A specialist who posts a schedule is monitored; a market maker who prices after seeing the flow is not, and the opaque case is the natural one for the batch auction of Kyle [71]. Existence is known only at the corners; both equilibria above pass the trader's second-order check of Chapter 4 at each computed γ, the opaque one is the fixed point of the market maker's Riccati equation, and both rents vanish continuously as $\gamma \to 0$. The nearest benchmarks are the monopolist specialist of Glosten [43] for the transparent market, the transparency comparisons of Madhavan [80] and Pagano and Röell [93], and the manipulation of Allen and Gale [3], with the roles reversed, for the opaque one. The comparison is an old policy fight: the London Stock Exchange moved block-trade publication from immediate to next-day and then to a ninety-minute delay between 1987 and 1992, dealers arguing that delayed publication protected the unwinding of their inventory [42].

6.9 Conclusion

Allowing some players to monitor a deviation's origin extends the Chapter 1 deviation calculus to an augmented forward system with two channels. In naive filtering, a naive player misattributes the deviation to primitive shocks and contributes an information wedge. In privy response, a privy player observes the deviation as a labeled zero-variance coordinate and responds through a new kernel. Under perfect transitive monitoring (Assumption 6.4) the seed responses are linear and superposable (Proposition 6.11), and the seed-dependent forward-backward system reduces to a single backward gain sweep that serves all seed times. Section 6.8 works out what monitoring means in a Kyle–Back market.

One direct extension combines this chapter with the delayed public signals of Chapter 2. Suppose an i-origin deviation at time s is initially hidden in an aggregate signal but is identified by a reporting rule at time $s+\Delta$. Before the report, player j is naive to the deviating player, so the seed moves j's ordinary noise-state estimate and is priced through the information wedge. At the report, the new information does more than add another measurement row. It attributes an already propagated seed to player i. In the resulting "naive until disclosure, privy afterward" system, origin attribution makes the same delayed transition from hidden influence to disclosure that Chapter 2 derives for signal manipulation.

Two other combinations concern the scale of the calculation. First, translating the rectangular gain into the stationary lag coordinates of Chapter 3 is the bridge to the strategic-market-maker extension; Section 6.8 does the scalar case. Second, if the dynamics, costs, and transitive monitoring preorder on a finite network are invariant under the graph automorphisms of Chapter 5, the responder–origin gains should inherit an orbit reduction. The stationary rectangular system and the equivariant monitored-response reduction remain open.

The blip convention also makes the chapter a step toward modeling irrationality. A blip is an accidental deviation of variance zero. It is priced, but it occurs with probability zero, so equilibrium behavior never anticipates it. Treating irrationality formally would mean giving blips a variance, which raises the questions of what law the blips follow, whether the naive and privy channels survive when attribution is probabilistic, and how the equilibrium deforms as the blip variance grows from zero.

Chapter 7

Conclusion: The Information Wedge Across Environments

Every equilibrium in this dissertation is a set of kernels against the same primitive shocks. Under perfect information the kernels are the ordinary impulse responses. The chapters changed what players could see, and the coordinates never moved.

The games have two feedback channels: the physical one, and the one that runs through what other players learn.

The noise-state representation keeps the second channel by writing private information as conditional estimates of the primitive shocks while retaining endogenous signals, private histories, and individual impact. In the LQG case every object of the calculus is a deterministic kernel, and the information wedge, a term in the state-price equation, prices the effect of changing another player's posterior.

The chapters vary what players see of the shocks. Monitoring varies what they know about a deviation's origin. A movement a player observes may be filtered as noise or recognized as another player's deviation. Chapter 6 prices both ways of responding, and the classical closed-loop game returns at the corner where every player sees the shocks and every player is privy.

The wedge and the accompanying equilibrium results come from the finite-horizon benchmark: Theorem 1.8 derives the belief-price representation, Corollary 1.11 gives closure over the full admissible class, and Proposition 1.12 solves unilateral best responses globally. Short-horizon equilibrium follows in Theorem A.7. The two-player benchmark separates statistical benefits from strategic costs and makes allocating signal precision a policy instrument.

The same object appears outside the baseline game. In identical-interest games the wedge prices the effect of shifting teammates' posteriors and so their future controls (Section 1.4.1). Delayed public signals add to it a birth term and a jump in the belief price at the date the news arrives, and the stationary equations carry the same terms in lag coordinates. In Kyle–Back trading the wedge becomes a belief price on the market maker's and the other traders' filters, since order flow changes current prices and the future filtering environment (Theorems 4.6 and 4.10). On finite symmetric local markets, the graph symmetry reduces policy and filter kernels to equivalent-position coordinates (Proposition 5.7). Asymmetric monitoring splits the feedback into a naive channel, in which the naive player filters the deviation as noise (this is the wedge), and a privy channel, in which the privy player responds to the deviation directly, knowing who made it (Definition 6.7, Proposition 6.9).

7.1 Directions beyond the dissertation

7.1.1 Closure

A noise-state linear fixed point is already a Nash equilibrium against arbitrary admissible deviations, but whether every equilibrium takes this form is open. Proving it means allowing strategies with random martingale-representation coefficients, writing the induced nonlinear filters through the Kushner-Stratonovich equations, and showing that equilibrium forces those coefficients to collapse to deterministic kernels. The continuous-time martingale representation theorem gives a handle that is absent in the discrete setting where Witsenhausen's counterexample [116] applies; a Lyapunov argument in the variance of the representation coefficients is the natural route.

7.1.2 Computation and estimation

One route to computing the kernels adapts high-dimensional PDE methods, in the spirit of Han's work [50], to the kernel HJB or the forward-backward kernel system. A different route approximates the noise-state itself. The filtering equations reweight old sources through kernels and project them into a current decision, the operation a linear attention layer performs. Linear transformers are a natural approximating class, with the noise-state as the target filter the trained network should recover. On the stationary system of Chapter 3, however, networks trained on the continuum residuals converged to spurious kernels while reporting small residuals (Figure 3.6).

On the dynamic-programming side, lifting the lag dynamics of Section A.2.3 to a fixed Hilbert space and taking a quadratic ansatz should recover the belief prices as off-diagonal Riccati blocks, and would likely give the CARA extension of the appendices a cleaner positivity condition than its operator inversions.

For empirical work the problem becomes estimation. Once the kernels are known the observed signal law has a Gaussian likelihood. But the kernels are endogenous equilibrium objects, so the estimator must either fit parameters while solving the kernel fixed point or learn a fast surrogate for that fixed point.

There is an easier road than fitting kernels. Across the chapters the strategic effects live in the means: the pooling gains of Chapter 1 are mostly mean-path, in Chapter 5 firms hold back their orders by several times the filtering gap while responses to aggregate shocks barely move, and the quote's inventory loading of Chapter 6 is a level. The near-invariance of the responses to aggregate shocks is camouflage: an econometrician who fit those responses would match them and miss a level distortion, one that moves whenever policy changes who sees whose actions. The natural design is levels against observability changes, post-trade transparency [20, 45], platform visibility, disclosure mandates, rather than kernel comparisons.

7.1.3 Sources

The most immediate extension is endogenous information acquisition. Signal technologies are treated as fixed throughout, and Appendix A.1 takes a first step by isolating the obstruction. One clean version uses a Brownian sheet of independent measurement channels, with each player selecting a portion of the sheet. Acquiring a source then changes the equilibrium policy kernels, and so changes what other players expect the holder to do with what it learns.

Compound Poisson signal sources are a larger step than a change of noise distribution. Throughout the dissertation each player knows the precision of their channels, so a source's informativeness is fixed and the randomness lies in its realizations. A Poisson source with an uncertain arrival rate breaks this. Individual jumps may carry messages of known quality, but if the rate is unknown the value of the source as a whole is unknown. Signal quality itself becomes an inferred state, which is closer to how news, expert reports, and disclosures actually work. The filter is then generally infinite-dimensional, but that is not obviously the binding constraint; the question is whether the filter over primitive sources stays stable and remains easier to approximate than the hierarchy it replaces.

7.1.4 Networks

The most immediate extension of the graph chapter is the large-cycle limit of Conjecture 5.2. Replicated markets can also be linked endogenously, with a mean-field consistency condition on the distribution of market-level outcomes and possibly a few major markets of finite aggregate impact. Heterogeneous local environments would be aggregated over market types. In every case the aggregated object is a solved finite strategic environment, which distinguishes the construction from a graphon limit where local informational impact has already vanished.

7.1.5 Markets

Chapter 6 makes the market maker strategic and computes the stationary market both ways, transparent and opaque. Existence remains open there, away from the corner cases and away from the $\rho \to 0$ degeneracy in which the market maker's first-order condition cannot pin down how much the quote loads on inventory.

The delayed-public and monitored-deviation chapters suggest a separate extension in which attribution arrives after the action, with a trader's order initially hidden in aggregate flow and a reporting rule naming it later. How players should respond on the day the order is named remains open.

In Rebonato's parable [99], two identical volatility traders facing the same mispriced quote behave differently depending on whether the risk function marks their book to their own model or to the market consensus. In the Kyle–Back setting the consensus is the market maker's price. A quote recognized as the market maker's action is one the trader responds to directly, the way a privy player responds to an observed deviation in Chapter 6, while the trader's private valuation remains a projection of primitive shocks through its noise-state. The gap between the trader's valuation and the consensus mark is the performance measure in a principal–agent contract between the trader as agent and risk management as principal, with Sannikov-type continuation values encoding it [103]. The belief prices of Chapter 4 separately price order flow that moves the consensus. The contract, risk penalty, and strategic channel then jointly blunt how hard the trader trades on its private information.

7.1.6 Noise traders with momentum

The strategic market maker above has a balance sheet it cares about, and Chapter 6 lets its quote be tested against deviations because every trader knows the quote is the market maker's action. On its own this adds an inventory charge to the price. With Brownian noise that charge depends only on the inventory already held, since the noise adds variance but does not move the market maker's forecast of future flow.

Noise traders with momentum do not need a strategic market maker. In Kyle–Back the informed traders already trade with long memory, since staying under the radar means trading slowly, and it is the noise flow that is Brownian. Retail herding has momentum. Fractional Brownian motion with $H > 1/2$ is the natural way to write this [81].

With both, the net order flow is a state the market maker manages, so a run of trend-chasing buying is a forecastable inventory problem rather than noise to be absorbed. The future of the noise now matters on its own, and everyone holds beliefs about it where Brownian noise gave them nothing to forecast. I expect the market maker's conditional expectation of value and the price it posts to come apart, with the price reflecting inventory management against the flow the market maker expects.

A trader who can forecast the coming run, and the charge the market maker will add against it, buys ahead and sells into that charge. They buy to sell later. De Long et al. [34] have rational speculators front-running trend-chasing traders; here the trend-chasers' persistence is a primitive. An order now moves the market maker's forecast of flow as well as its estimate of value, so the trader manipulates through a second channel, and the wedge prices both. The round trip pays only if the market maker's charge overshoots. The market maker forecasts the selling to come and prices it in, the traders buying ahead forecast the market maker, and how much of the rise survives would be a fixed point in the same kernels. The theorem to aim for is the market maker's inventory cost against competition among the traders buying ahead.

As markets shift toward passive and flow-driven trading, the part of order flow that is not about fundamentals grows [41] and the market maker's inventory problem becomes a larger part of the price. The model at the end has a strategic market maker who manages inventory, noise traders with momentum, and informed traders whose signals keep arriving.

Afterword

My background was originally in cognitive neuroscience. It started from looking at optical illusions. Weird patterns shown to our eyes, our brain playing tricks on us. As the theory goes, optical illusions serve a purpose. Our brain is not powerful enough to process the world in its entirety, and so resorts to certain practicalities. It decides to tell us just enough to act on, guided and stabilized by billions of years of natural selection towards choosing the proper realities. Natural selection is a statistical tool present wherever there is continuous heredity. Our laws [54], economies [107], languages [32], traditions [85], and even our sense of fairness [17] are guided and stabilized by the principle of natural selection. After many years of guidance, these aspects of our world have stabilized enough to be predictable through the lens of optimality. But our world is far from stationary. Whether we look backwards ten, a hundred, or a thousand years, much of our past looks foreign. We still have not reached the end, if there will ever be one. To what extent are our optical illusions optimal responses, and to what extent are they relics of a less than optimal past? And how much of our sense of fairness?

To understand our irrationality, we must first understand optimality. Ariely[1], Loewenstein, and Prelec [9], following Kahneman and Tversky [110], ran an experiment that stuck with me. Test subjects bid on an object after seeing the last two digits of their social security number. Those with a higher number bid higher. In a world where our mind must process many things at once, it cannot afford the luxury of properly classifying

[1] Better known for other work [74].

every source of information. In a decentralized world where our sense of value is not fully grounded, we look to what others do for guidance. What is optimal? What is a relic?

In my teens, I was enraptured by Leonard Read's "I, Pencil" and Hayek's "The Use of Knowledge in Society". Could economics, a field that studies optimality in decentralized societies, arm me with the frame to answer my two simple questions? After reading further, I kept finding papers that came close in description, but kept falling short of what I wanted. Each introduction would have beautiful prose stating, clearly, that the author shares much of the motivation that I do. And yet whenever I looked deeper, I kept encountering a barrier that stopped the rest of the text from matching the introductions' promises. Economists, to me, did not have what they needed to properly write down optimality. Mathematics, the tool that let Isaac Newton write down the motion of celestial bodies, my joy I had not yet found a use for, became my answer.

The unknown beyond what we can see. If we do not know it, or simply cannot understand it, can we tell it apart from randomness? The proper mathematics describes optimality in a random world, a society more complex than we can understand.

More than a decade later, this dissertation's only goal is to arm those who care with better tools to understand this world, and help me start again where I left off and answer my two simple questions.

Appendix A

Additional Finite-Horizon Extensions

A.1 Endogenous signal precision

This appendix identifies the obstruction to optimizing precision in the multi-player setting.

A.1.1 Why choosing precision is hard

In single-agent LQG, the separation principle makes precision easy to optimize: the policy gain $S(t)$ does not depend on signal precision, so the value of information can be computed independently of the control problem. With multiple agents and endogenous signals, separation fails (Remark 1.13), because the policy kernel $D_t^i(u)$ depends on precision through the information wedge, the shadow price of moving another player's noise-state.

The filtrations generated by $B_X^{Y,i} X\, dt + dW^i$ and $(B_X^{Y,i} + \delta B_X^{Y,i}) X\, dt + dW^i$ are generically incomparable: neither contains the other. This breaks the envelope argument that makes single-agent rational inattention tractable, because the player at $P^i + \delta P^i$ cannot implement the control they would have chosen at P^i.

There is also a more basic problem with calling $\delta B_X^{Y,i}$ a change in precision. The dW^i in the reduced observation equation need not be the error from one literal measurement. It can already be the noise left after several measurements have been collapsed to a sufficient statistic. For example, if

$$dY^m = b_m X\,dt + dV^m, \qquad m \in S,$$

with independent unit-variance measurement noises, then the informative part of these measurements can be written

$$d\bar{Y}^S = \sqrt{P_S}\,X\,dt + d\bar{V}^S, \qquad P_S = \sum_{m\in S} b_m^2, \qquad d\bar{V}^S = \frac{1}{\sqrt{P_S}}\sum_{m\in S} b_m\,dV^m.$$

Adding another measurement increases P_S, but it also changes $d\bar{V}^S$. At the level of the actual measurements, more precision therefore does not usually mean increasing the gain while holding the same measurement noise fixed. The latter is a perfectly good perturbation of the reduced-form filtering problem, including for the calculations below, but it needs some care before it is interpreted as acquiring more information.

Correlation makes this more important rather than less. If the measurement errors have covariance R, the corresponding precision is $b^\top R^{-1} b$ and the sufficient statistic weights the measurements by $R^{-1}b$. Adding a source can then change the weights on measurements the player already had, and correlations with other information the player observes have to be carried along as well. None of this makes the filtering mathematics difficult; it means that an economically meaningful deviation in precision should come from a specified change in the underlying measurements, with the new reduced-form gain and noise derived from that change.

The incomparability disappears when an independent source adds precision (the Brownian-sheet formulation sketched in Section 7.1.3). Monitoring an additional channel strictly enlarges the filtration, and the marginal value of the source is the value of optimally exploiting its independent increment. The full rational-inattention game in which players choose source portfolios requires this formulation and is left to later work. Whether opponents are privy or naive to the precision change, in the sense of Chapter 6, depends on whether they see it. If they do not see it, they filter its effects as primitive shock. This is coherent: the change reaches them through the deviating player's control, so it moves the drift of their observations and not the volatility, and treating a deviation as primitive shock breaks down only when the deviation changes the volatility. The channels described in Section A.1.2 are what a naive opponent filters as noise. If they see it, through a data purchase or disclosed research, they are privy and respond through the gains of Chapter 6. In the naive case the best response to deterministic precisions is deterministic, because fixed gains make the conditional variances deterministic, so the precision game can be solved over deterministic gain paths. Unfortunately my intuition does not stretch far enough for players with both privy and naive opponents, although when all opponents are privy it should still apply.

A.1.2 Two channels from a precision perturbation

Despite the incomparability, one can compute the first-order effect of a gain perturbation on equilibrium objects by linearizing the forward–backward system.

Fix an equilibrium profile and spike-perturb player i's gain at time t_0. By the envelope theorem, δD^i has no first-order cost effect, because the FOC is satisfied. But the primitive-shock control $D^i_{W,\tau}(r) = D^i_\tau(r)\Pi^i + \int_0^\tau D^i_\tau(u)\, F^i_\tau(u,r)\, du$ responds at first order through δF^i, changing which part of the fixed policy kernel the player's own information resolves.

The variation of the unresolved adjoint decomposes as

$$\delta\widetilde{H}^{X,i}_t(s) = \underbrace{-\int_s^t H^{X,i}_t(u)\,\delta F^i_t(u,s)\,du}_{\text{(a) filtering channel}} + \underbrace{\int_{t_0}^T \Phi^{HD,i}(t,\tau)\,\widetilde{\delta D}^i_{W,t}(\tau,s)\,d\tau}_{\text{(b) equilibrium channel}}, \tag{A.1.1}$$

where $\Phi^{HD,i}$ is the deterministic kernel through which a control variation at time τ enters the time-t adjoint, and $\widetilde{\delta D}^i_W$ is the unresolved component of the control variation induced by δF^i.

Through channel (a) the player resolves more of the fixed adjoint from the perturbed signal, the standard value-of-information effect present in single-agent problems. Channel (b) arises because the same δF^i changes the projection, leaving part of $D^i_{W,\tau}$ unresolved by the player's current information. The equilibrium channel operates through the information wedge and vanishes when $\mathcal{V}^i_t = 0$ (exogenous signals, Corollary 1.9). Multi-player rational inattention couples the filtering and control problems that separate in single-agent settings: the marginal value of precision depends on the equilibrium through the wedge. Optimizing attention requires solving for the equilibrium first.

Local profile-space notation. In the remainder of this appendix, bold symbols denote stacked policy profiles, and calligraphic and sans-serif symbols denote maps on the profile space. These are local functional-analytic abbreviations and do not carry into the substantive chapters.

A.2 Well-posedness of the deterministic impulse-response fixed point

This appendix records the a priori bounds that follow from the projection structure of the filtering map and separates three well-posedness questions. A unilateral best response against fixed linear opponent maps is globally well posed for every finite horizon, and the joint policy update is a contraction on short horizons. For arbitrary horizons, existence and uniqueness follow under a weighted monotonicity condition that one can also check numerically near a computed equilibrium.

Notation. C denotes a generic constant depending on (n, d, L, λ); $C(M)$ when it also depends on a policy-norm bound M.

A.2.1 Standing assumptions

Assumption A.1. *There exist $L > 0$ and $\lambda > 0$ such that, for all $t \in [0,T]$ and all i, $\|B^X_X(t)\| + \|B^X_i(t)\| + \|\sigma(t)\| + \|P^i(t)\| + \|G^{XX,i}(t)\| + \|\bar{G}^{X,i}_t\| + \|G^{X,i}(t,\cdot)\|_{L^\infty} + \|G^{DX,i}(t)\| + \|G^{XX,i}(T)\| + \|\bar{G}^{X,i}_T\| + \|G^{X,i}(T,\cdot)\|_{L^\infty} \le L$, the joint running Hessian is positive semidefinite, and the Schur condition (1.3.3) holds with constant λ (which implies $G^{DD,i}(t) \succeq \lambda I_d$, since $S^i \preceq G^{DD,i}$).*

A.2.2 The kernel state space and its dynamics

At time t the kernel state is

$$\mathcal{Z}_t := \Big(\bar{X}_t,\ X_t(\cdot),\ \{F^j_t\}_{j=1}^n\Big) \ \in\ \mathbb{R}^d \times L^2([0,t];\mathbb{R}^{d\times(n+1)d}) \times \prod_{j=1}^n L^2_{\mathrm{sym}}([0,t]^2;\mathbb{R}^{(n+1)d\times(n+1)d}). \tag{A.2.1}$$

The unresolved response $\widetilde{X}^j$ is not a state variable, but follows algebraically from $(\mathcal{Z}_t)$ by

$$\widetilde{X}^j_t(u) = X_t(u)(I - \Pi^j) - \int_0^t X_t(v)\, F^j_t(v,u)\, dv. \tag{A.2.2}$$

It is the part of X_t that player j has not resolved from its own observations.

Under a noise-state linear profile $\mathbf{D} = \{(\bar{D}^j, D^j)\}_{j=1}^n$, the kernel state evolves by

$$\dot{\bar{X}}(t) = B^X_X(t)\bar{X}_t + \textstyle\sum_k B^X_k(t)\bar{D}^k_t, \tag{A.2.3}$$

$$\dot{X}_t(s) = B^X_X(t)X_t(s) + \textstyle\sum_k B^X_k(t)D^k_{W,t}(s), \qquad X_s(s) = \sigma(s)E^0, \tag{A.2.4}$$

$$\partial_t F^j_t(u,s) = \widetilde{X}^j_t(u)^\top P^j(t)\, \widetilde{X}^j_t(s), \qquad F^j_0 = 0, \tag{A.2.5}$$

with E^0 the block selector for the state noise W^0, boundary $F^j_t(u,t) = \widetilde{X}^j_t(u)^\top B^{Y,j}_X(t)^\top E^j$, and primitive-shock control $D^k_{W,t}(s) := D^k_t(s)\Pi^k + \int_0^t D^k_t(u)\, F^k_t(u,s)\, du$.

The policy $\mathbf{D}$ enters (A.2.3)–(A.2.4) but does *not* enter (A.2.5) directly, since F^j depends on $\mathbf{D}$ only through X by (A.2.2).

A.2.3 Dynamic-programming interpretation of a fixed-profile best response

The deterministic kernels in (A.2.1) are useful for computing equilibrium, but they are not themselves a state observed by any player. Their domains also grow with time and create the birth terms at the current date derived in Appendix 1.B. For these reasons, the well-posedness results below do not rely on a Hamilton–Jacobi–Bellman equation on $(\bar{X}, X, F)$.

Dynamic programming instead becomes available after fixing the opponents' complete linear strategy maps. For player i, let S_t^i denote the augmented hidden state containing the physical state and the opponents' noise-state coordinates that enter their future controls, lifted to a fixed Hilbert space so that its generator includes transport and coordinate births. The player-specific separated state is

$$m_t^i := \mathbb{E}[S_t^i \mid \mathcal{F}_t^i].$$

Under the frozen linear profile the problem is linear Gaussian, the conditional covariance is deterministic, and the separated state has the formal innovation representation

$$dm_t^i = \Big(\mathsf{A}_t^i m_t^i + \mathsf{B}_t^i u_t^i + b_t^i\Big)dt + \mathsf{K}_t^i \, dI_t^i.$$

The control u_t^i moves the drift, and the innovation dI_t^i carries what the player learns. For the value function normalized as one half of the continuation cost, the corresponding formal HJB equation is

$$\begin{aligned} 0 = \partial_t V^i(t,m) + \inf_u \Bigg\{ &\frac{1}{2}\bar{\ell}^i(t,m,u) + \Big\langle DV^i(t,m), \mathsf{A}_t^i m + \mathsf{B}_t^i u + b_t^i \Big\rangle \\ &+ \frac{1}{2}\operatorname{Tr}\Big[\mathsf{K}_t^i (\mathsf{K}_t^i)^* D^2 V^i(t,m)\Big] \Bigg\}. \end{aligned} \tag{A.2.6}$$

The first-order condition is

$$u_t^{i,*} = -(G^{DD,i}(t))^{-1}\Big[G^{DX,i}(t)\,\Pi_X m_t^i + (\mathsf{B}_t^i)^* DV^i(t, m_t^i)\Big], \tag{A.2.7}$$

where $\Pi_X m_t^i = \mathbb{E}[X_t \mid \mathcal{F}_t^i]$ is the physical component of the separated state. A quadratic operator ansatz for V^i gives an operator Riccati system. Its off-diagonal blocks, coupling the physical component of m^i to the opponents' noise-state coordinates, are the dynamic-programming counterparts of the belief prices and the information wedge. A rigorous HJB treatment would require specifying the fixed-space lift and the unbounded transport and birth operators.

A.2.4 A priori bounds from projection structure

The filter kernel F_t^j parametrizes a portion of the conditional projection $(\mathcal{P}_t^j f)(u) := \Pi^j f(u) + \int_0^t F_t^j(u,s)\, f(s)\, ds$, player j's blueprint, an orthogonal projection on $L^2([0,t];\mathbb{R}^{(n+1)d})$ (Theorem 1.5 and the projection interpretation following it).

Lemma A.2 (Projection bounds)**.** *For each j and $t \geq 0$: (i) $\|\mathcal{P}_t^j\|_{\mathrm{op}} = 1$; (ii) both $\mathcal{P}_t^j$ and $I - \mathcal{P}_t^j$ are L^2-contractions; (iii) $\|\mathcal{P}_t^j - \Pi^j\|_{\mathrm{op}} \leq 2$.*

Proof. Standard properties of orthogonal projections; (iii) is the triangle inequality. □

Lemma A.3 (Dissipation)**.** *$\{\mathcal{P}_t^j\}_{t\geq 0}$ is non-decreasing in the positive semidefinite order, and $\|\widetilde{X}^j(t,\cdot)\|_{L^2([0,t])} \leq \|X(t,\cdot)\|_{L^2([0,t])}$ for all t.*

Proof. $\mathcal{P}_t^j$ is the orthogonal projection onto the $\mathcal{F}_t^j$-measurable linear functionals of W (Lemma 1.7); since $\mathcal{F}_s^j \subseteq \mathcal{F}_t^j$ for $s \leq t$, these projections are nested and $\mathcal{P}_t^j$ is non-decreasing. Since $\widetilde{X}^j(t,\cdot) = (I - \mathcal{P}_t^j)X(t,\cdot)$ by (A.2.2), Lemma A.2(ii) gives the bound. □

These bounds hold uniformly in the precision paths, the policy profile, and the horizon T. As $\mathcal{P}_t^j$ grows toward I, the complementary projection contracts, driving

$\widetilde{X}^j \to 0$ and decelerating F^j, so the quadratic right-hand side of (A.2.5) is self-limiting. This prevents finite-time blowup despite the quadratic nonlinearity, just as the Kalman filter covariance remains bounded.

A.2.5 Well-posedness

Definition A.4. Define $\mathbb{B}_T^M := \Big\{\mathbf{D} = \{(\bar{D}^i, D^i)\}_{i=1}^n : \bar{D}^i \in C([0,T];\mathbb{R}^d),\ D^i \in C(\Delta_T),\ \|\mathbf{D}\| \le M\Big\}$, where $\|\mathbf{D}\| := \max_i(\|\bar{D}^i\|_C + \|D^i\|_{L^\infty})$. With continuous data the Volterra structure of (1.A.5)–(1.A.6) and of the adjoint equations propagates continuity to F^i and to the belief prices, so the birth traces $F_t^i(u,t)$, $F_t^i(t,u)$ and $H_t^{k,i}(t)$ are defined pointwise.

Definition A.5 (Forward–backward policy update)**.** Define

$$\mathcal{T}(\mathbf{D}) := \Big\{(-(G^{DD,i})^{-1}[G^{DX,i}\bar{X}^{\mathbf{D}} + \bar{H}^{X,i,\mathbf{D}}],\ -(G^{DD,i})^{-1}[G^{DX,i}X^{\mathbf{D}} + H^{X,i,\mathbf{D}}])\Big\}_{i=1}^n,$$

where the adjoints solve the backward system of Theorem 1.8 at the forward environment induced by $\mathbf{D}$. Fixed points are Nash equilibria (Corollary 1.11).

Proof of Proposition 1.12. Use the control-free reference filtration of Proposition 1.6 and let

$$\mathcal{U}_i := L^2_{\text{prog}}(\Omega \times [0,T], \mathbb{F}^i; \mathbb{R}^d).$$

The exact system for the difference between two controls makes the physical state and all induced opponent controls affine and continuous functions of the realized input $u_i \in \mathcal{U}_i$. The cost then has the form

$$J^i(u_i; D^{-i}) = \frac{1}{2}\langle u_i, A_i u_i\rangle_{\mathcal{U}_i} + \langle b_i, u_i\rangle_{\mathcal{U}_i} + c_i,$$

where A_i is bounded (Proposition 1.6), deterministic, self-adjoint, and $\langle u_i, A_i u_i\rangle \ge \lambda\|u_i\|^2$ by Lemma 1.10. There is then a unique minimizer.

It remains to show that the minimizer is noise-state linear, without assuming in advance that the optimum lies in that class. Split every square-integrable functional of the reference noise into its affine part, a constant plus a linear functional of the noise, and a remainder orthogonal to all affine functionals. The frozen linear causal response maps send affine functionals to affine functionals and remainders to remainders (they are built from deterministic kernels and from conditional expectations given filtrations generated by linear functionals of the noise), so A_i does the same, b_i is affine, and the cost has no cross term between the two parts. Any remainder in a control contributes only its nonnegative quadratic cost and cannot occur in the unique minimizer. Since $\mathcal{F}_t^i$ is generated by linear functionals of the reference noise, conditioning on it also maps affine functionals to affine functionals and remainders to remainders, so the affine part of an adapted control is adapted and the minimization over $\mathcal{U}_i$ may be carried out part by part. The minimizer is consequently affine Gaussian and adapted. By the Gaussian-span statement in Lemma 1.7, it has a deterministic noise-state representation. The corresponding deterministic kernels satisfy the stationarity system. $\square$

Fixed coordinates for arbitrary-horizon equilibrium. The monotonicity result below is stated on a fixed Hilbert space of complete linear causal strategy maps. Let

$$\mathcal{S}_T := \prod_{i=1}^{n} \mathcal{S}_T^i$$

be such a realization, using a profile-independent observation or reference- innovation coordinate system. If profile-dependent noise-state coordinates are used computationally, the coordinate-change terms are included when the best-response derivative is formed. Proposition 1.12 defines the exact best-response map

$$\mathcal{B}(\mathbf{s}) := \Big(\mathcal{B}_i(\mathbf{s}^{-i})\Big)_{i=1}^{n}, \qquad \mathbf{s} \in \mathcal{S}_T,$$

and the equilibrium residual

$$\mathcal{R}(\mathbf{s}) := \mathbf{s} - \mathcal{B}(\mathbf{s}). \tag{A.2.8}$$

A zero of $\mathcal{R}$ is a fixed point of the exact linear-strategy best-response map, and so a Nash equilibrium (Corollary 1.11).

The forward estimates below follow by Grönwall arguments on the forward system (A.2.3)–(A.2.5), using the a priori bounds of Lemmas A.2–A.3 to control the quadratic nonlinearity in (A.2.5). The adjoint estimates follow by backward Grönwall on the linear adjoint system (1.B.19)–(1.B.20), using Cauchy–Schwarz on the belief-price coupling terms, including the birth term $B_X^{Y,k}(t)^\top E^k \bar{H}_t^{k,i}(t)$ and the drift-inference integral $\int_0^t P^k(t)\widetilde{X}_t^k(r)\,\bar{H}_t^{k,i}(r)\,dr$.

Proposition A.6 (Bounds and Lipschitz continuity)**.** *Under Assumption A.1, for* $\mathbf{D} \in \mathbb{B}_T^M$*:*

(i) The forward system has a unique global solution with $\|\bar{X}\|_C + \|X\|_{L^\infty} \le K_{\text{fwd}}$*,* $\|\widetilde{X}^j(t,\cdot)\|_{L^2} \le K_{\text{fwd}}\sqrt{T}$*,* $\|\mathcal{P}_t^j\|_{\text{op}} = 1$*.*

(ii) The backward system has a unique solution with $\|\bar{H}^{X,i}\|_C + \|H^{X,i}\|_{L^\infty} \le K_{\text{bwd}}$*, and time derivatives bounded by* $C(M)K_{\text{bwd}}$*.*

(iii) Both maps $\mathbf{D} \mapsto (\bar{X}, X, \widetilde{X}^j)$ *and* $\mathbf{D} \mapsto (\bar{H}^{X,i}, H^{X,i})$ *are Lipschitz with constants* $C_{\text{fwd}}(T,M)$ *and* $C_{\text{bwd}}(T,M)$ *respectively, both bounded by* $C(M)\,T\,e^{C(M)T}$*.*

Here $K_{\text{fwd}}, K_{\text{bwd}} \le C(M)e^{C(M)T}$*.*

Proof sketch. For the forward system, $(\bar{X}, X)$ solve linear equations with bounded coefficients once D_W^k is controlled. The operator bound $\|\mathcal{P}_t^k - \Pi^k\|_{\text{op}} \le 2$ (Lemma A.2(iii)) on the filter kernel closes the system in L^2 by Grönwall, ruling out finite-time blowup. The L^∞ closure follows by bootstrapping. With $\phi(t) := \|X\|_{L^\infty(\Delta_t)} + \max_k \|\widetilde{X}^k\|_{L^\infty(\Delta_t)}$, the explicit formula (1.A.10) gives $\|F_t^k(\cdot,s)\|_{L^2([0,t])} \le C(M)(1+\phi(t))$, so Cauchy–Schwarz on $D_{W,t}^k(s) = D_t^k(s)\Pi^k + \int_0^t D_t^k(u)F_t^k(u,s)\,du$ gives $|D_{W,t}^k(s)| \le C(M)(1+\phi(t))$, and ϕ satisfies $\phi(t) \le C(M) + C(M)\int_0^t \phi(r)\,dr$.

For the Lipschitz bound, the variation δD_W^k decomposes into a direct term bounded by $C(M)\|\delta \mathbf{D}\|$ and an indirect term through δF^k, which couples to $\delta\widetilde{X}^k$ and back to δX by (A.2.2). A Grönwall argument on the aggregate $\Psi(t) := \|\delta X\|_{L^\infty(\Delta_t)} + \max_j \sup_{r\le t}\|\delta\widetilde{X}^j(r,\cdot)\|_{L^2}$ gives $\Psi(T) \le C(M)Te^{C(M)T}\|\delta\mathbf{D}\|$; the factor of T arises because the policy enters through integrals of length $\le T$.

The backward estimates follow from linearity of the adjoint system, Cauchy–Schwarz on coupling terms, and backward Grönwall. $\square$

Theorem A.7 (Short-horizon equilibrium)**.** *Under Assumption A.1, there exists $T^* > 0$ depending only on (n, d, L, λ) such that for $T \in (0, T^*]$, $\mathcal{T}$ is a contraction on $\mathbb{B}_T^{M_0}$ for a suitable M_0. There is then a unique noise-state linear equilibrium in $\mathbb{B}_T^{M_0}$, Picard iterates converge geometrically, and it is Nash over the full admissible L^2 class.*

Proof. Write $\mathcal{T}_T$ for the update on horizon T. By Proposition A.6(ii), there are $T_0 > 0$ and $C_0 < \infty$ such that

$$\sup_{0<T\le T_0}\|\mathcal{T}_T(0)\| \le C_0.$$

Set $M_0 := 2C_0$. Proposition A.6(iii) gives

$$\mathrm{Lip}\Big(\mathcal{T}_T;\mathbb{B}_T^{M_0}\Big) \le \lambda^{-1}C(M_0)Te^{C(M_0)T}.$$

Choose $T^* \le T_0$ so that the right-hand side is at most $1/2$ for $T \le T^*$. Then, for $\mathbf{D} \in \mathbb{B}_T^{M_0}$,

$$\|\mathcal{T}_T(\mathbf{D})\| \le \|\mathcal{T}_T(0)\| + \frac{1}{2}\|\mathbf{D}\| \le C_0 + \frac{1}{2}M_0 = M_0.$$

The map $\mathcal{T}_T$ is then a self-map and a contraction on $\mathbb{B}_T^{M_0}$. Banach's theorem gives existence, uniqueness, and geometric convergence in $\mathbb{B}_T^{M_0}$. $\square$

Proposition A.8 (Weighted Jacobian criterion)**.** *Let $\mathsf{W} : \mathcal{S}_T \to \mathcal{S}_T$ be bounded, self-*

adjoint, and uniformly positive, and write

$$\langle h, k\rangle_{\mathsf{W}} := \langle \mathsf{W}h, k\rangle, \qquad \|h\|_{\mathsf{W}}^2 := \langle h, h\rangle_{\mathsf{W}}.$$

Let $K \subset \mathcal{S}_T$ be convex, and suppose that the exact residual $\mathcal{R}$ is continuously Fréchet differentiable on K. If there is $\mu > 0$ such that

$$\frac{1}{2}\left(\mathsf{W}D\mathcal{R}(\mathbf{s}) + D\mathcal{R}(\mathbf{s})^*\mathsf{W}\right) \succeq \mu\mathsf{W} \qquad \textit{for every } \mathbf{s} \in K, \tag{A.2.9}$$

then $\mathcal{R}$ is μ-strongly monotone on K:

$$\langle \mathcal{R}(\mathbf{s}) - \mathcal{R}(\widetilde{\mathbf{s}}), \mathbf{s} - \widetilde{\mathbf{s}}\rangle_{\mathsf{W}} \geq \mu\|\mathbf{s} - \widetilde{\mathbf{s}}\|_{\mathsf{W}}^2.$$

Proof. Set $\Delta = \mathbf{s} - \widetilde{\mathbf{s}}$. Convexity of K and the fundamental theorem of calculus give

$$\mathcal{R}(\mathbf{s}) - \mathcal{R}(\widetilde{\mathbf{s}}) = \int_0^1 D\mathcal{R}(\widetilde{\mathbf{s}} + \theta\Delta)\Delta\, d\theta.$$

Taking the W-inner product with Δ and using (A.2.9) proves the claim. □

Theorem A.9 (Certified equilibrium on a strategy ball)**.** *Fix a reference profile $\mathbf{s}_0 \in \mathcal{S}_T$ and $R > 0$, and set*

$$K_R(\mathbf{s}_0) := \{\mathbf{s} \in \mathcal{S}_T : \|\mathbf{s} - \mathbf{s}_0\|_{\mathsf{W}} \leq R\}.$$

Suppose that the exact best-response map is well defined on $K_R(\mathbf{s}_0)$ and that its residual $\mathcal{R} = I - \mathcal{B}$ is μ-strongly monotone and L-Lipschitz there. If

$$\|\mathcal{R}(\mathbf{s}_0)\|_{\mathsf{W}} < \mu R, \tag{A.2.10}$$

then $K_R(\mathbf{s}_0)$ contains a unique profile $\mathbf{s}^*$ satisfying

$$\mathcal{R}(\mathbf{s}^*) = 0.$$

The profile $\mathbf{s}^*$ is a Nash equilibrium over the full admissible L^2 control class. For every

$$0 < \eta < \frac{2\mu}{L^2},$$

the projected residual iteration

$$\mathbf{s}^{m+1} = P^{\mathsf{W}}_{K_R(\mathbf{s}_0)} \left[\mathbf{s}^m - \eta \mathcal{R}(\mathbf{s}^m)\right] \tag{A.2.11}$$

converges geometrically to $\mathbf{s}^*$.

Proof. For $\mathbf{s}, \widetilde{\mathbf{s}} \in K_R(\mathbf{s}_0)$, nonexpansiveness of metric projection gives

$$\begin{aligned}
&\left\| P^{\mathsf{W}}_{K_R(\mathbf{s}_0)}(\mathbf{s} - \eta\mathcal{R}(\mathbf{s})) - P^{\mathsf{W}}_{K_R(\mathbf{s}_0)}(\widetilde{\mathbf{s}} - \eta\mathcal{R}(\widetilde{\mathbf{s}})) \right\|_{\mathsf{W}}^2 \\
&\quad \leq \left(1 - 2\eta\mu + \eta^2 L^2\right) \|\mathbf{s} - \widetilde{\mathbf{s}}\|_{\mathsf{W}}^2.
\end{aligned}$$

The stated range of η makes the projected map a contraction. Banach's fixed-point theorem supplies a unique fixed point $\mathbf{s}^*$ of that map. The projection characterization gives

$$\langle \mathcal{R}(\mathbf{s}^*), \mathbf{y} - \mathbf{s}^* \rangle_{\mathsf{W}} \geq 0 \qquad \text{for every } \mathbf{y} \in K_R(\mathbf{s}_0).$$

If $\|\mathbf{s}^* - \mathbf{s}_0\|_{\mathsf{W}} = R$, choosing $\mathbf{y} = \mathbf{s}_0$ gives

$$\langle \mathcal{R}(\mathbf{s}^*), \mathbf{s}^* - \mathbf{s}_0 \rangle_{\mathsf{W}} \leq 0.$$

Strong monotonicity relative to $\mathbf{s}_0$ instead gives

$$\begin{aligned}\langle \mathcal{R}(\mathbf{s}^*), \mathbf{s}^* - \mathbf{s}_0 \rangle_{\mathsf{W}} &\geq \mu R^2 + \langle \mathcal{R}(\mathbf{s}_0), \mathbf{s}^* - \mathbf{s}_0 \rangle_{\mathsf{W}} \\ &\geq \mu R^2 - \|\mathcal{R}(\mathbf{s}_0)\|_{\mathsf{W}} R > 0,\end{aligned}$$

a contradiction. The profile $\mathbf{s}^*$ then lies in the interior of the ball, and testing the variational inequality in both directions gives $\mathcal{R}(\mathbf{s}^*) = 0$. Strong monotonicity shows that two zeros in the ball must coincide. □

Remark A.10 (A two-player check). For two players and scalar weights $r_1, r_2 > 0$, a sufficient form of (A.2.9) on a convex set K is

$$\sup_{\mathbf{s} \in K} \left\| \frac{1}{2} \left[\sqrt{\frac{r_1}{r_2}}\, D_2 \mathcal{B}_1(\mathbf{s}) + \sqrt{\frac{r_2}{r_1}}\, D_1 \mathcal{B}_2(\mathbf{s})^* \right] \right\|_{\mathrm{op}} < 1.$$

Only the symmetrically reinforcing part of reciprocal best responses enters. Monotonicity can then hold even when ordinary Picard iteration is not a contraction because the dominant feedback is antisymmetric or rotational.

Corollary A.11 (Numerical certification near a computed profile)**.** *Let* $\widehat{\mathbf{s}}$ *be a computed strategy profile for a finite- dimensional discretization, let*

$$J_{\mathcal{R}} \approx D\mathcal{R}(\widehat{\mathbf{s}})$$

be the computed Jacobian of the best-response residual at $\widehat{\mathbf{s}}$*, with* $\|J_{\mathcal{R}} - D\mathcal{R}(\widehat{\mathbf{s}})\|_{\mathsf{W}} \leq \varepsilon_J$ *in the weighted operator norm* $\|A\|_{\mathsf{W}} := \|\mathsf{W}^{1/2} A \mathsf{W}^{-1/2}\|_{\mathrm{op}}$*, and define the local weighted margin*

$$\widehat{\mu} := \lambda_{\min} \left[\mathsf{W}^{-1/2} \frac{\mathsf{W} J_{\mathcal{R}} + J_{\mathcal{R}}^{\mathsf{T}} \mathsf{W}}{2} \mathsf{W}^{-1/2} \right]. \tag{A.2.12}$$

Suppose that, throughout $K_R(\widehat{\mathbf{s}})$*,*

$$\|D\mathcal{R}(\mathbf{s}) - D\mathcal{R}(\widehat{\mathbf{s}})\|_{\mathsf{W}} \le \Lambda R.$$

Then $\mathcal{R}$ *is strongly monotone on that ball with constant*

$$\mu_R := \widehat{\mu} - \varepsilon_J - \Lambda R.$$

If

$$\mu_R > 0, \qquad \|\mathcal{R}(\widehat{\mathbf{s}})\|_{\mathsf{W}} < \mu_R R, \tag{A.2.13}$$

then the ball contains a unique equilibrium.

Proof. The Jacobian-variation bound and the variational characterization of the smallest eigenvalue imply

$$\frac{1}{2}\left(\mathsf{W}D\mathcal{R}(\mathbf{s}) + D\mathcal{R}(\mathbf{s})^{*}\mathsf{W}\right) \succeq \mu_R \mathsf{W}$$

throughout the ball. Proposition A.8 and Theorem A.9 then apply with center $\widehat{\mathbf{s}}$. □

Numerical implementation. After discretizing the deterministic policy kernels, the margin (A.2.12) is the smallest generalized eigenvalue of

$$\frac{1}{2}\left(\mathsf{W}J_{\mathcal{R}} + J_{\mathcal{R}}^{\mathsf{T}}\mathsf{W}\right)\psi = \mu \mathsf{W}\psi.$$

Take the Jacobian for the exact best-response residual, computing it from tangent equations, automatic differentiation, complex-step differentiation, or centered finite differences. A positive $\widehat{\mu}$ alone certifies local monotonicity and local uniqueness, but the variation and residual bounds in (A.2.13) turn it into a rigorous certificate on a whole ball. Report separately

$$\widehat{\mu}, \qquad \sigma_{\min}(J_{\mathcal{R}}), \qquad \rho(D\mathcal{B}(\widehat{\mathbf{s}})).$$

The first measures the monotonicity margin, the second local invertibility, and the third local convergence of undamped Picard iteration. Then $\widehat{\mu} > 0$ together with $\rho(D\mathcal{B}) > 1$ describes a locally unique equilibrium for which ordinary Picard iteration is unstable but sufficiently damped residual descent remains convergent.

Remark A.12 (What remains model-specific)*.* Theorem A.9 is an arbitrary-horizon result, but one must verify its hypotheses on the chosen strategy ball. In the noise-state system this requires differentiating the deterministic forward filter and backward belief-price equations with respect to the fixed linear strategy maps. The resulting tangent equations are linear Volterra–backward systems. Proposition A.6 supplies the zeroth-order bounds; a corresponding first-order estimate gives $D\mathcal{B}$, while a second- order estimate or validated numerical bound gives Λ in Corollary A.11.

Remark A.13 (Stochastic coefficients)*.* If primitives depend on a public factor process ξ_t, the kernel state augments to $(\mathcal{Z}_t, \xi_t)$. The projection bounds of Lemmas A.2–A.3 are algebraic and hold pathwise, so forward well-posedness extends between jumps of ξ. At each jump the filtering gains, strategies, and the information wedge shift discontinuously, while $\mathcal{Z}$ is continuous across jumps.

Bibliography

[1] Vadim M. Adamyan, Damir Z. Arov, and Mark G. Krein. Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur–Takagi problem. *Mathematics of the USSR-Sbornik*, 15(1):31–73, 1971.

[2] George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

[3] Franklin Allen and Douglas Gale. Stock-price manipulation. *Review of Financial Studies*, 5(3):503–529, 1992.

[4] Manuel Amador and Pierre-Olivier Weill. Learning from prices: Public communication and welfare. *Journal of Political Economy*, 118(5):866–907, 2010.

[5] George-Marios Angeletos and Chen Lian. Incomplete information in macroeconomics: Accommodating frictions in coordination. In John B. Taylor and Harald Uhlig, editors, *Handbook of Macroeconomics*, volume 2, pages 1065–1240. Elsevier, 2016.

[6] George-Marios Angeletos and Alessandro Pavan. Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142, 2007.

[7] Anthropic. Claude Fable 5. Large language model, 2026. https://www.anthropic.com.

[8] Pol Antràs, Davin Chor, Thibault Fally, and Russell Hillberry. Measuring the upstreamness of production and trade flows. *American Economic Review*, 102(3):412–416, 2012.

[9] Dan Ariely, George Loewenstein, and Drazen Prelec. "coherent arbitrariness": Stable demand curves without stable preferences. *Quarterly Journal of Economics*, 118(1):73–105, 2003.

[10] Sam Babichenko. Forecasting and Manipulating the Forecasts of Others, 2026. URL https://arxiv.org/abs/2603.12140.

[11] Kerry Back. Insider trading in continuous time. *The Review of Financial Studies*, 5(3): 387–409, 1992.

[12] Kerry Back and Hal Pedersen. Long-lived information and intraday patterns. *Journal of Financial Markets*, 1(3–4):385–402, 1998.

[13] Kerry Back, C. Henry Cao, and Gregory A. Willard. Imperfect competition among informed traders. *The Journal of Finance*, 55(5):2117–2155, 2000.

[14] Alan Bain and Dan Crisan. *Fundamentals of Stochastic Filtering*, volume 60 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2009.

[15] Abhijit V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992.

[16] Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. SIAM, Philadelphia, 2nd edition, 1999.

[17] Nicolas Baumard. *The Origins of Fairness: How Evolution Explains Our Moral Nature*. Oxford University Press, Oxford, 2016. Translated by Paul Reeve.

[18] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

[19] Dirk Bergemann, Tibor Heumann, and Stephen Morris. Information and volatility. *Journal of Economic Theory*, 158:427–465, 2015.

[20] Hendrik Bessembinder, William Maxwell, and Kumar Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288, 2006.

[21] Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5): 992–1026, 1992.

[22] Timothy F. Bresnahan. Duopoly models with consistent conjectures. *American Economic Review*, 71(5):934–945, 1981.

[23] Markus K. Brunnermeier. *Investor Behaviour, Financial Markets and the International Economy*. PhD thesis, London School of Economics and Political Science, 1999.

[24] Luciano Campi, Umut Çetin, and Albina Danilova. Dynamic Markov bridges motivated by models of insider trading. *Stochastic Processes and their Applications*, 121(3):534–567, 2011.

[25] René Carmona. *Lectures on BSDEs, Stochastic Control, and Stochastic Differential Games with Financial Applications*, volume 1 of *Financial Mathematics*. SIAM, Philadelphia, 2016.

[26] René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018.

[27] René Carmona, Jean-Pierre Fouque, and Li-Hsien Sun. Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4):911–933, 2015.

[28] René Carmona, Jean-Pierre Fouque, Seyyed Mostafa Mousavi, and Li-Hsien Sun. Systemic risk and stochastic games with delay. *Journal of Optimization Theory and Applications*, 179(2):366–399, 2018.

[29] Umut Çetin and Albina Danilova. *Dynamic Markov Bridges and Market Microstructure*, volume 90 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018.

[30] Raymond H. Chan and Michael K. Ng. Conjugate gradient methods for Toeplitz systems. *SIAM Review*, 38(3):427–482, 1996.

[31] Minh Chau and Dimitri Vayanos. Strong-form efficiency with monopolistic insiders. *Review of Financial Studies*, 21(5):2275–2306, 2008.

[32] Morten H. Christiansen and Nick Chater. Language as shaped by the brain. *Behavioral and Brain Sciences*, 31(5):489–509, 2008.

[33] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

[34] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance*, 45(2):379–395, 1990.

[35] Nils Detering, Jean-Pierre Fouque, and Tomoyuki Ichiba. Directed chain stochastic differential equations. *Stochastic Processes and their Applications*, 130(4):2519–2551, 2020.

[36] Yichen Feng, Jean-Pierre Fouque, and Tomoyuki Ichiba. Linear–quadratic stochastic differential games on directed chain networks. *Journal of Mathematics and Statistical Science*, 7(2):25–67, 2021.

[37] F. Douglas Foster and S. Viswanathan. Strategic trading with asymmetrically informed traders and long-lived information. *Journal of Financial and Quantitative Analysis*, 29(4): 499–518, 1994.

[38] F. Douglas Foster and S. Viswanathan. Strategic trading when agents forecast the forecasts of others. *The Journal of Finance*, 51(4):1437–1478, 1996.

[39] Jean-Pierre Fouque and Tomoyuki Ichiba. Stability in a model of inter-bank lending. *SIAM Journal on Financial Mathematics*, 4:784–803, 2013.

[40] Drew Fudenberg and Jean Tirole. A “signal-jamming” theory of predation. *The RAND Journal of Economics*, 17(3):366–376, 1986.

[41] Xavier Gabaix and Ralph S. J. Koijen. In search of the origins of financial fluctuations: The inelastic markets hypothesis. Working Paper 28967, National Bureau of Economic Research, 2021.

[42] Gordon Gemmill. Transparency and liquidity: A study of block trades on the London Stock Exchange under different publication rules. *Journal of Finance*, 51(5):1765–1790, 1996.

[43] Lawrence R. Glosten. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business*, 62(2):211–235, 1989.

[44] Keith Glover. All optimal Hankel-norm approximations of linear multivariable systems and their L^∞-error bounds. *International Journal of Control*, 39(6):1115–1193, 1984.

[45] Michael A. Goldstein, Edith S. Hotchkiss, and Erik R. Sirri. Transparency and liquidity: A controlled experiment on corporate bonds. *Review of Financial Studies*, 20(2):235–273, 2007.

[46] Benjamin Golub and Matthew O. Jackson. Naïve learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.

[47] Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

[48] Yuxiang Guan, Iman Shames, and Tyler Summers. Forward-backward dynamic programming for LQG dynamic games with partial and asymmetric information, 2026.

[49] Ben M. Hambly, Renyuan Xu, and Huining Yang. Linear–quadratic Gaussian games with asymmetric information: Belief corrections using the opponents actions, 2023. SSRN working paper.

[50] Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34): 8505–8510, 2018.

[51] Lars Peter Hansen and Thomas J. Sargent. Linear rational expectations models for dynamically interrelated variables. In Robert E. Lucas, Jr. and Thomas J. Sargent, editors, *Rational Expectations and Econometric Practice*, volume 1, pages 127–156. University of Minnesota Press, Minneapolis, 1981.

[52] John C. Harsanyi. Games with incomplete information played by “Bayesian” players, I–III. Part I. the basic model. *Management Science*, 14(3):159–182, 1967.

[53] Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4): 519–530, 1945.

[54] Friedrich A. Hayek. *Law, Legislation and Liberty, Volume 1: Rules and Order*. University of Chicago Press, Chicago, 1973.

[55] Yu-Chi Ho and K’ai-Ching Chu. Team decision theory and information structures in optimal control problems—Part I. *IEEE Transactions on Automatic Control*, 17(1):15–22, 1972.

[56] Craig W. Holden and Avanidhar Subrahmanyam. Long-lived private information and imperfect competition. *The Journal of Finance*, 47(1):247–270, 1992.

[57] Bengt Holmström. Managerial incentive problems: A dynamic perspective. *The Review of*

Economic Studies, 66(1):169–182, 1999.
[58] Minyi Huang, Roland P. Malhamé, and Peter E. Caines. Large population stochastic dynamic games: Closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Communications in Information and Systems*, 6(3):221–252, 2006. doi: 10.4310/CIS.2006.v6.n3.a5.
[59] Gur Huberman and Werner Stanzl. Price manipulation and quasi-arbitrage. *Econometrica*, 72(4):1247–1275, 2004.
[60] Zhen Huo and Naoki Takayama. Rational expectations models with higher-order beliefs. *The Review of Economic Studies*, 92(5):3138–3173, 2025. doi: 10.1093/restud/rdae096.
[61] Tomoyuki Ichiba, Michael Ludkovski, and Andrey Sarantsev. Dynamic contagion in a banking system with births and defaults. *Annals of Finance*, 15(4):489–538, 2019.
[62] David H. Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control*, 18(2):124–131, 1973.
[63] Emir Kamenica and Matthew Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
[64] Kenneth Kasa. Forecasting the forecasts of others in the frequency domain. *Review of Economic Dynamics*, 3(4):726–756, 2000.
[65] Younes Kchia and Philip Protter. Progressive filtration expansions via a process, with applications to insider trading. *International Journal of Theoretical and Applied Finance*, 18(4):1550027, 2015. doi: 10.1142/S0219024915500272.
[66] John Maynard Keynes. *The General Theory of Employment, Interest and Money.* Macmillan, London, 1936.
[67] Yerin Kim, Haosheng Zhou, Alexander Benvenuti, Ruimeng Hu, and Matthew Hale. Deception in linear–quadratic control, 2026. Preprint, arXiv:2604.00227.
[68] David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, 1982.
[69] David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253–279, 1982.
[70] Finn E. Kydland and Edward C. Prescott. Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370, 1982.
[71] Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
[72] Nicolas S. Lambert, Michael Ostrovsky, and Mikhail Panov. Strategic trading in informationally complex environments. *Econometrica*, 86(4):1119–1157, 2018.
[73] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
[74] Gideon Lewis-Kraus. They studied dishonesty. Was their work a lie? *The New Yorker*, September 30 2023.
[75] Robert S. Liptser and Albert N. Shiryaev. *Statistics of Random Processes I: General Theory*, volume 5 of *Stochastic Modelling and Applied Probability.* Springer, Berlin, Heidelberg, 2nd edition, 2001.
[76] Guido Lorenzoni. A theory of demand shocks. *American Economic Review*, 99(5):2050–2084, 2009.
[77] Robert E. Lucas, Jr. Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124, 1972.
[78] Robert E. Lucas, Jr. Some international evidence on output-inflation tradeoffs. *American*

Economic Review, 63(3):326–334, 1973.

[79] Robert E. Lucas, Jr. Econometric policy evaluation: A critique. In *The Phillips Curve and Labor Markets*, volume 1 of *Carnegie-Rochester Conference Series on Public Policy*, pages 19–46. North-Holland, 1976.

[80] Ananth Madhavan. Consolidation, fragmentation, and the disclosure of trading information. *Review of Financial Studies*, 8(3):579–603, 1995.

[81] Benoit B. Mandelbrot and John W. Van Ness. Fractional Brownian motions, fractional noises and applications. *SIAM Review*, 10(4):422–437, 1968.

[82] Jacob Marschak and Roy Radner. *Economic Theory of Teams*. Number 22 in Cowles Foundation Monograph. Yale University Press, New Haven, 1972.

[83] Eric Maskin and Jean Tirole. Markov perfect equilibrium I: Observable actions. *Journal of Economic Theory*, 100(2):191–219, 2001.

[84] Paul Milgrom and John Roberts. Limit pricing and entry under incomplete information: An equilibrium analysis. *Econometrica*, 50(2):443–459, 1982.

[85] Olivier Morin. *How Traditions Live and Die*. Oxford University Press, Oxford, 2016.

[86] Stephen Morris and Hyun Song Shin. Social value of public information. *American Economic Review*, 92(5):1521–1534, 2002.

[87] John F. Muth. Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335, 1961.

[88] Ashutosh Nayyar, Aditya Mahajan, and Demosthenis Teneketzis. Optimal control strategies in delayed sharing information structures. *IEEE Transactions on Automatic Control*, 56 (7):1606–1620, 2011.

[89] Ashutosh Nayyar, Aditya Mahajan, and Demosthenis Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Transactions on Automatic Control*, 58(7):1644–1658, 2013. doi: 10.1109/TAC.2013.2239000.

[90] Kristoffer Nimark. Dynamic higher order expectations. Meeting Paper 1132, Society for Economic Dynamics, 2017.

[91] Benjamin Noble. *Methods Based on the Wiener–Hopf Technique for the Solution of Partial Differential Equations*. Pergamon Press, London, 1958.

[92] Yi Ouyang, Hamidreza Tavafoghi, and Demosthenis Teneketzis. Dynamic games with asymmetric information: Common information based perfect Bayesian equilibria and sequential decomposition. *IEEE Transactions on Automatic Control*, 62(1):222–237, 2017.

[93] Marco Pagano and Ailsa Röell. Transparency and liquidity: A comparison of auction and dealer markets with informed trading. *Journal of Finance*, 51(2):579–611, 1996.

[94] Joseph G. Pearlman and Thomas J. Sargent. Knowing the forecasts of others. *Review of Economic Dynamics*, 8(2):480–497, 2005.

[95] Edmund S. Phelps, editor. *Microeconomic Foundations of Employment and Inflation Theory*. W. W. Norton, New York, 1970.

[96] A. W. Phillips. The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957. *Economica*, 25(100):283–299, 1958.

[97] R. Radner. Team decision problems. *The Annals of Mathematical Statistics*, 33(3):857–881, 1962.

[98] Leonard E. Read. I, Pencil: My family tree as told to Leonard E. Read. *The Freeman*, 8 (12):32–37, 1958.

[99] Riccardo Rebonato. *Volatility and Correlation: The Perfect Hedger and the Fox*. Wiley, 2nd edition, 2004.

[100] Michael H. Riordan. Imperfect information and dynamic conjectural variations. *The*

RAND Journal of Economics, 16(1):41–50, 1985.

[101] Giacomo Rondina and Todd B. Walker. Confounding dynamics. *Journal of Economic Theory*, 196:105251, 2021.

[102] Paul A. Samuelson and Robert M. Solow. Analytical aspects of anti-inflation policy. *American Economic Review*, 50(2):177–194, 1960.

[103] Yuliy Sannikov. A continuous-time version of the principal–agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.

[104] Thomas J. Sargent. Equilibrium with signal extraction from endogenous variables. *Journal of Economic Dynamics and Control*, 15(2):245–273, 1991.

[105] Thomas C. Schelling. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA, 1960.

[106] Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.

[107] Adam Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*. W. Strahan and T. Cadell, London, 1776.

[108] Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973.

[109] Robert M. Townsend. Forecasting the forecasts of others. *Journal of Political Economy*, 91(4):546–588, 1983.

[110] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

[111] Pravin Varaiya and Jean Walrand. On delayed sharing patterns. *IEEE Transactions on Automatic Control*, 23(3):443–445, 1978.

[112] Deepanshu Vasal and Achilleas Anastasopoulos. Signaling equilibria for dynamic LQG games with asymmetric information. *IEEE Transactions on Control of Network Systems*, 8(3):1177–1188, 2021. doi: 10.1109/TCNS.2021.3059835.

[113] Mark Veraar. The stochastic Fubini theorem revisited. *Stochastics: An International Journal of Probability and Stochastic Processes*, 84(4):543–551, 2012.

[114] Xavier Vives. Endogenous public information and welfare in market games. *The Review of Economic Studies*, 84(2):935–963, 2017.

[115] Peter Whittle. Risk-sensitive linear/quadratic/Gaussian control. *Advances in Applied Probability*, 13(4):764–777, 1981.

[116] Hans S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):131–147, 1968.

[117] Hans S. Witsenhausen. Separation of estimation and control for discrete time systems. *Proceedings of the IEEE*, 59(11):1557–1566, 1971.

[118] Michael Woodford. Imperfect common knowledge and the effects of monetary policy. Working Paper 8673, National Bureau of Economic Research, 2001.

[119] Haosheng Zhou, Daniel Ralston, Xu Yang, and Ruimeng Hu. Adversarial decision-making in partially observable multi-agent systems: A sequential hypothesis testing approach. *IEEE Transactions on Control of Network Systems*, 13(2):1218–1229, 2026.