Producing to Signal: Private Information in Dynamic Quantity Competition

Andrew Sweeting
University of Maryland, CEPR and NBER*

Xuezhen Tao
Shanghai University of Finance and Economics

Qian Wang
Sciences Po

October 2026

Abstract

We model a finite-horizon game of repeated oligopoly competition in which sellers choose quantities of a homogeneous product with linear demand and have private information about their own marginal costs, which are serially correlated over time. Focusing on fully separating equilibria, we provide conditions for existence and uniqueness. Relative to static complete- and incomplete-information benchmarks, sellers' incentives to signal low costs can substantially affect output, productive efficiency, and welfare, even when there are ten or more sellers. We apply the model to the global oil market, and compare predictions to a model where sellers can make quantity commitments.

JEL CODES: C73, D82, L13, L71.

Keywords: oligopoly, asymmetric information, signaling, separating equilibria, dynamics, strategic incentives, oil markets.

*Sweeting (corresponding author): 3114 Tydings Hall, University of Maryland, College Park MD 20742, USA. Email: atsweet@umd.edu. This research was originally supported by US National Science Foundation Grant SES-1658670 (Sweeting) and National Natural Science Foundation of China Grant 72003120 (Tao). Xinlu Yao collaborated on this project when she was a graduate student at the University of Maryland, and we are very grateful for Xinlu's excellent work and insights, and Mario Leccese was a critical collaborator on complementary work in Leccese, Sweeting, and Tao (2025). We thank Dan Vincent for useful comments, and Gaurab Aryal for helpful discussions about Aryal and Zincenko (2024). AI (Claude) was used in the summer of 2026 to improve the performance of previously written MATLAB code. All errors are our own.

1 Introduction

Industrial Organization (I.O.) models commonly assume that oligopolists choose prices or quantities with complete information about demand and all sellers' marginal costs.[1] This assumption is tractable, but it raises two questions. Can private information alter predicted outcomes, and can the resulting differences be large enough to matter in applications?

While, as we will discuss below, the theoretical literature provides some answers to the question of whether the assumed information structure affects expected outcomes, little attention has been paid to when differences may be substantial. In this paper, we examine a model where, following Mester (1992), private information combines with repeated play to give rise to substantial effects. Mester characterizes a unique fully separating equilibrium in a three-period Cournot duopoly game where each seller's marginal cost is private information and serially correlated across periods. She shows that expected seller output in the first two periods will be higher than it would be under complete information.

The economic logic is as follows. In a separating equilibrium, sellers will infer their rivals' current costs from their current output choices, and, based on how costs evolve and their understanding of next period strategies, they will use these beliefs to calculate how much each rival is likely to produce in the next period. Given that a seller will want to produce less when it expects its rivals to produce more, and sellers will tend to produce more when they have lower costs, each seller will have an incentive to increase its current output in order to signal to its rivals that its current costs are low and it will likely produce more in the next period. When all sellers act in this way, Mester shows that, relative to complete information, signaling raises expected output in both nonterminal (i.e., the first and second) periods of the game, reflecting "perpetual signaling". However, Mester

[1] The auction literature typically assumes bidders' signals are private information. However, even when there are repeated auctions with the same bidders, the auction literature has not considered the types of signaling equilibria that we consider here.

does not compare output levels across the first two periods, discuss whether output and welfare differences will be large, or consider how the size of differences would vary with additional periods or additional sellers. We focus on these issues in our paper.

We consider models with a known linear demand curve for a homogeneous product and sellers who have quadratic cost functions where the intercepts may be private information. In one-shot games, expected equilibrium output with complete or incomplete information is the same (Vives (2002)). We are interested in what happens given fully separating Markov Perfect Bayesian Nash equilibrium strategies when play is repeated and costs are persistent, but not perfectly so, across periods, which is consistent with results in the productivity literature (Syverson (2004), Lee and Mukoyama (2015)).

We begin with a finite-horizon duopoly example, illustrating sellers' strategies – which are characterized by differential equations and boundary conditions in every period except the final one – and showing how signaling can substantially increase output, with consequences for welfare and productive efficiency relative to one-shot play. The expected quantity increases become larger as one moves to earlier periods of the game, with strategies converging, in a relatively small number of periods, to the stationary strategies that we find in an infinite-horizon version of the game with the same parameters. As one would expect, we also see that signaling effects become larger when we increase the persistence of the marginal cost intercepts, so that one period's cost signal becomes more informative about a seller's next period output.

We then illustrate that substantial quantity and welfare effects persist even when the number of sellers is fairly large: for example, many of the differences between signaling and static outcomes are almost as large when there are thirteen sellers as when there are two sellers. This contrasts both with what we have found in price-setting games (Sweeting, Tao, and Yao (2024) (STY)) and with what we believe is a common intuition that strategic investment effects tend to be large only

in games with a very small number of investing sellers. We provide an explanation for this effect by noting that in a quantity-setting linear demand-quadratic cost game, rivals' equilibrium outputs can often become more responsive to the expected output of a single firm when rivals become more numerous. The fact that effects can be large in environments with both concentrated and unconcentrated market structures suggests that researchers may want to consider the possibility of signaling in any application where sellers set quantities repeatedly and it is plausible that their costs are serially correlated and imperfectly observed by their rivals.[2]

We apply our model to the global oil market. We build on Aryal and Zincenko (2024) (AZ), who use a panel of monthly country-level production to estimate the linear demand and quadratic cost parameters of a static game where countries use Bayesian Nash equilibrium strategies and their marginal cost intercepts are assumed to be private information. We change AZ's specifications in several ways. We assume that five of the largest producers have serially correlated marginal cost intercepts and may use signaling strategies. Other producers do not signal but respond optimally to the output they expect the signaling countries to produce. Serial correlation of the intercepts is consistent with the persistence of country-level outputs, and the pattern of posterior cost intercepts that we obtain when we apply AZ's estimated model period by period. Our use of a model where some producers signal and others do not signal, but respond optimally to expected signaling outputs, also allows us to speak to ideas of leadership by Saudi Arabia or OPEC that are common in the oil industry literature (see references below).

Calibrating the marginal costs of the signaling sellers assuming static behavior, we show that signaling can significantly change output, productive efficiency and welfare relative to either a static incomplete information or a static complete information equilibrium. We also show that calibrating the marginal costs of the signaling sellers assuming signaling behavior implies quite

[2]Our examples will assume the product being sold is homogeneous, a setting where economists usually assume quantity competition. However, quantity-setting can also be an appropriate assumption when products are differentiated (Duarte, Magnolfi, Quint, Sølvsten, and Sullivan (2026)).

different levels of marginal costs, and that static and signaling models imply that seller-specific outputs and welfare may change differently in response to two stylized counterfactuals (a change in the responsiveness of non-signaling supply and the exclusion of Russia and Iran from the world oil market). We also compare the predictions of a signaling model to the predictions of a model where the same subset of producers can act as quantity leaders by making quantity commitments. For given cost parameters, both strategic models predict higher output (relative to static simultaneous Nash) by sellers that signal or lead. However, we show that they can generate different predictions for the output changes of individual sellers and different welfare effects in our counterfactuals.

Our primary contributions are twofold. First, we show that a dynamic quantity-setting game in which multiple sellers signal simultaneously can be tractable. Second, and more substantively, we show that the equilibrium effects of signaling can be substantial, both in absolute terms and relative to those identified by STY in price-setting models, even when the number of sellers is quite large. We show that a signaling model's output, welfare and comparative static predictions can differ substantially from those of static models, with or without incomplete information, and from models where strategic quantity-setting behavior comes from other assumptions, such as the ability of firms to make quantity commitments by acting as leaders.

After a review of the related literature, the paper proceeds by describing the model and equilibrium (Section 2), analyzing examples with symmetric duopoly and up to thirteen symmetric sellers (Section 3), and providing our analysis of the oil market (Section 4). Section 5 concludes.

Related Literature. An analytical literature, often motivated by the possible non-collusive effects of information exchange, compares complete- and incomplete-information equilibria in one-shot oligopoly games with price or quantity competition. Different assumptions can be made about which variables are privately known: for example, demand intercepts (Vives (1984), Gal-Or (1985))

or marginal cost intercepts (Fried (1984), Gal-Or (1986), Shapiro (1986), Vives (2002)), as well as whether firms choose prices or quantities. Vives (2002), who studies homogeneous-good quantity competition with commonly known linear demand and quadratic seller costs, provides the closest static benchmark for our paper. When every seller produces a strictly positive quantity for every cost realization, each firm's expected output is the same in the complete- and incomplete-information equilibria, while the increased variance of output raises expected consumer surplus and lowers expected total surplus. These benchmark results apply to our specifications under static behavior, and our contribution is to show that repeated signaling, which arises when incomplete information and repeated play combine, can produce large differences in both output and welfare.[3]

As noted, Mailath (1988), Mailath (1989) and Mester (1992) also consider signaling. We will make use of Mailath's characterization of separating strategies, although his application was to two-period price-setting games where cost types are fixed across periods. Mester applies similar logic to a three-period Cournot duopoly game, and shows that when marginal costs are serially correlated, rather than fixed across periods, it is possible to generate what she calls "perpetual signaling" where, in each nonterminal period, sellers expand output to persuade rivals that their current costs are low. We retain the serial correlation assumption and use homogeneous-good demand and quadratic costs, allowing us to link our analysis directly to Vives (2002), and, most importantly, we examine the magnitude of the effects and what happens when there may be substantially more sellers.[4]

STY extend the Mailath–Mester approach to multi-period price-setting games and show that signaling can produce prices well above complete-information or static Bayesian Nash levels in

[3]Vives (2002) also shows that as the market size and the number of sellers increase the deadweight loss from incomplete information dominates the loss from market power.

[4]A separate literature studies strategies and outcomes in longer games where players' payoff-relevant characteristics are fixed. For example, Milgrom and Roberts (1990), Fudenberg and Levine (1993), Kalai and Lehrer (1993) and Doraszelski, Lewis, and Pakes (2018) analyze convergence towards Nash, Toxvaerd (2017) considers signaling to deter entry, while Bonatti, Cisternas, and Toikka (2017) analyze a game where a noisy market price creates "signal-jamming" incentives to distort rivals' learning.

concentrated markets. They also show how a signaling model can explain observed changes in both price levels and price dynamics after the 2008 MillerCoors joint venture. Our results in this paper differ in two important respects. First, because in our quantity-setting examples sellers' incentives are to signal low costs, we find that signaling raises output and consumer surplus rather than increasing prices. Second, the effects in our model remain substantial with a fairly large number of sellers, whereas the effects identified in STY are very small in magnitude once there are more than three or four firms. This pattern is noteworthy because signaling can be viewed as a strategic investment, and that literature has largely focused on monopoly and duopoly settings (Tirole (1988), chapters 8 and 9). We show that, as the number of sellers increases, a seller's own output falls but combined rival output can become more sensitive to beliefs about that seller's cost, and that this explains why signaling effects can remain substantial.

Like Mester and STY, we assume that state variables (current costs, and beliefs about past costs) and actions are continuous, and that costs evolve in a first-order Markovian fashion. This allows us to focus on fully separating Markovian strategies, which are tractable because beliefs about rivals' types have a simple form (i.e., the value of sellers' costs in the last period). Other approaches are possible, such as the Experience Based Equilibrium approach of Asker, Fershtman, Jeon, and Pakes (2020), where types and actions are discrete.

Empirical I.O. typically assumes that firms choose continuous prices or quantities under complete information, either in a static Nash equilibrium or under an alternative conduct assumption. This permits researchers to evaluate first-order conditions at rivals' realized choices rather than integrate over potentially high-dimensional beliefs. Empirical models with incomplete information are rare even in static settings, although Armantier and Richard (2003), Gardete (2016) and AZ are exceptions.

Our application to the oil market builds on AZ's stylized model, using the same data and a num-

ber of their estimated parameters. Both papers abstract from resource exhaustibility, differences between short-run and long-run demand elasticities and possible adjustment costs. AZ focus on comparisons between static incomplete information and complete information equilibria, whereas we focus on the effects of dynamic signaling that arise quite naturally in an incomplete information model when the persistence of marginal costs is accounted for. Using detailed cost data, Asker, Collard-Wexler, and De Loecker (2019) empirically document significant inefficiencies in global production. Our analysis shows that signaling can also reallocate production: when relatively low-cost sellers signal, their output tends to rise and productive efficiency may improve. As discussed in Section 4.6, our comparison between a model with signaling and a model where leaders move first relates to the extensive literature on possible leadership by one or more countries in the global oil market. Much of the oil literature has been concerned with cartel behavior in the face of alternative sources of supply (e.g., Griffin (1985), Smith (2005), Almoguera, Douglas, and Herrera (2011)). For example, Hyndman (2008) develops a model where two cartel members attempt to bargain over production quotas under asymmetric information. Our application can be viewed as focusing on the exact logic of how, in an asymmetric information environment where formal commitments to output levels are not credible, large producers may act to influence, using output signals, each other's future output decisions as well as those of a set of smaller producers. We compare this model to alternatives where signaling does not happen or commitments are possible.

2 Model and Equilibrium

In this section we present our model and characterize equilibrium strategies. We make the assumptions that we use in Section 3, leaving discussion of assumptions used to study the oil market to Section 4.

2.1 Model

There are N risk-neutral sellers of a homogeneous product. The sellers play a repeated game with $t = 1, \ldots, T$ discrete time periods, where T may be finite or infinite. The discount factor is $0 < \delta < 1$, and our computations use $\delta = 0.99$ unless otherwise stated.

Demand is commonly known, linear, static and fixed across periods. Specifically, inverse demand is $P_t = u - \beta \sum_i q_{i,t}$ where $q_{i,t}$ is the output of seller i in period t and $\beta > 0$. We will assume that the market price can go negative, or that, if there is a non-negativity constraint, it never binds for any realization of sellers' costs. Seller i has costs $C(q_{i,t}) = v_{i,t}q_{i,t} + \frac{\lambda}{2}q_{i,t}^2$, so i's marginal cost is $v_{i,t} + \lambda q_{i,t}$, with $\lambda \geq 0$ (all of our calculations will assume $\lambda > 0$). The $v_{i,t}$s are continuous random variables privately observed by seller i. They have time-invariant supports $[\underline{v_i}, \overline{v_i}]$, and they evolve from period-to-period according to independent, stochastic, exogenous and time-invariant first-order Markov processes, $\psi_i : v_{i,t-1} \to v_{i,t}$.[5] We assume that the conditional PDFs $g_i(v_{i,t}|v_{i,t-1})$ have the following properties.

Assumption 1 (Type Transitions for the Continuous Type Model) *The conditional PDF* $g_i\left(v_{i,t}|v_{i,t-1}\right)$

1. *is continuous and differentiable, with appropriate one-sided derivatives at the boundaries;*
2. *is strictly positive for all* $v_{i,t}$*s and* $v_{i,t-1}$*s on the supports* $[\underline{v_i}, \overline{v_i}]$*;*
3. *for any* $v_{i,t-1}$ *there is some* v' *such that* $\frac{\partial g_i(v_{i,t}|v_{i,t-1})}{\partial v_{i,t-1}}|_{v_{i,t}=v'} = 0$ *and* $\frac{\partial g_i(v_{i,t}|v_{i,t-1})}{\partial v_{i,t-1}} < 0$ *for all* $v_{i,t} < v'$ *and* $\frac{\partial g_i(v_{i,t}|v_{i,t-1})}{\partial v_{i,t-1}} > 0$ *for all* $v_{i,t} > v'$.

The third part of the assumption implies that higher values of $v_{i,t}$ are more likely if $v_{i,t-1}$ increases, so that there is positive serial correlation in a seller's costs. Below, we will make the specific

[5]The assumption of independent exogenously-evolving serially correlated seller productivity is consistent with Olley and Pakes (1996) and most of the structural production function literature.

assumption that the vs evolve according to an AR(1) process with truncated normal innovations. The assumption also implies that a higher conditioning type shifts the next period type distribution upwards in the first-order stochastic dominance sense.[6]

Within a period t, timing is as follows. Sellers enter t with their $t-1$ types, which then evolve according to the ψs. Each seller observes its own new type, but not the previous or new types of rivals (except that we assume that types are publicly observed in a fictitious period $t = 0$ before the first period). Sellers then simultaneously choose, and commit to, their period t quantities, $q_{i,t}$. There are no adjustment costs to changing output across periods. Current quantities, once chosen, and all past quantity choices, are observable.

2.2 Equilibrium Concept

We consider fully separating Markov Perfect Bayesian equilibria (MPBE) (Toxvaerd (2008), Roddie (2012)). An MPBE specifies each seller's beliefs about rival costs given observed quantities and the application of Bayes' rule, and quantity strategies for each seller in each period that maximize the seller's discounted expected value, given the seller's current type and its beliefs, and rivals' strategies. We will note below what can be shown about existence and uniqueness of a fully separating equilibrium. Mailath (1989) and Mester (1992) also restrict consideration to fully separating equilibria. Other equilibria with partial pooling may exist, but they are hard to characterize and enumerate.

In a fully separating equilibrium, quantity choices on the equilibrium path will allow a seller's current type to be inferred exactly, and the Markov restriction on the equilibrium and the first-order Markov assumption on type transitions imply that strategies will only depend on sellers' current types and beliefs about previous period cost types (not types in earlier periods). If a seller

[6] For completeness, see Appendix A.1 for a proof.

sets a quantity below (above) the lowest (highest) quantity in the range of its equilibrium quantity strategy function, we assume that rivals will infer it to have the highest (lowest) possible $v_{i,t}$ type. If sellers are symmetric, we will only consider symmetric equilibria. Even if sellers are not symmetric, we will impose that all rival sellers will interpret a seller's quantity choice in the same way, so that they will have the same beliefs about a seller's lagged type.

2.3 Strategies in a Finite-Horizon Game

Assuming that T is finite, we derive strategies recursively from the end of the game and define values at the start of each period, before the type transitions have occurred. Assuming that equilibria are fully separating, the value of seller i will depend on the following state variables: (i) i's actual $t-1$ type, $v_{i,t-1}$; (ii) the $t-1$ type that i's rivals perceive it to have had, $\widehat{v_{i,t-1}^{-i}}$; and (iii) i's perception of its rivals' previous period types, $\boldsymbol{v}_{-i,t-1}$ (where bold indicates a vector), which i will assume will correspond to their true types given equilibrium play.[7] Quantities are chosen when i knows its period t type. Conditional on $v_{i,t}$ and $\widehat{v_{i,t-1}^{-i}}$, $v_{i,t-1}$ has no effect on i's profits or future values, so when calculating optimal strategies, the state variables are only $v_{i,t}$, $\widehat{v_{i,t-1}^{-i}}$ and $\boldsymbol{v}_{-i,t-1}$.

2.3.1 Period T

Sellers will use Bayesian Nash equilibrium strategies to maximize their final period T expected profits given rivals' quantity functions $\boldsymbol{q}_{-i,T}\left(\boldsymbol{v}_{-i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right)$. i's best response quantity choice will be the solution to

$$q_{i,T}^{*}\left(v_{i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right) = \arg\max_{q_{i,T}} \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \pi\left(v_{i,T}, q_{i,T}, \boldsymbol{q}_{-i,T}\left(\boldsymbol{v}_{-i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right)\right) \times g_{-i}\left(\boldsymbol{v}_{-i,T} \mid \boldsymbol{v}_{-i,T-1}\right) d\boldsymbol{v}_{-i,T} \tag{1}$$

[7] We will derive i's strategy using the possibility that it may unilaterally deviate by signaling a different type, but i will be assuming that other sellers signal truthfully.

where $g_{-i}(\boldsymbol{v}_{-i,T}|\boldsymbol{v}_{-i,T-1})$ is the product of rivals' transition PDFs. Denoting rivals' combined output as $Q_{-i,T}$,

$$\pi\left(v_{i,T}, q_{i,T}, \boldsymbol{q}_{-i,T}\left(\boldsymbol{v}_{-i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right)\right) = \left(u - \beta(q_{i,T} + Q_{-i,T}) - v_{i,T}\right) q_{i,T} - \frac{\lambda}{2} q_{i,T}^2. \tag{2}$$

Assuming, as we will throughout this section, that sellers choose positive quantities for all type realizations, there is an analytical formula for the unique and separating equilibrium quantity choices (Vives (2002, Proposition 1) and AZ (Lemma 1)),

$$q_{i,T}^* = \frac{1}{\lambda + (N+1)\beta}\left\{u - \frac{1}{\lambda+\beta}\left[(\lambda + N\beta)\widehat{\mu_{i,T}^{v,-i}} - \beta\sum_{j\neq i}\mu_{j,T}^v\right]\right\} - \frac{v_{i,T} - \widehat{\mu_{i,T}^{v,-i}}}{\lambda + 2\beta} \tag{3}$$

where $\mu_{j,T}^v = \mathbb{E}[v_{j,T}|v_{j,T-1}]$, i.e., the expected value of $v_{j,T}$ given j's $T-1$ type, and $\widehat{\mu_{i,T}^{v,-i}} = \mathbb{E}\left[v_{i,T}|\widehat{v_{i,T-1}^{-i}}\right]$, i.e., rivals' expectation of $v_{i,T}$ given $\widehat{v_{i,T-1}^{-i}}$.

2.3.2 Earlier Periods $(1, \ldots, T-1)$

Assumption 1 implies $\widehat{\mu_{i,T}^{v,-i}} = \mathbb{E}\left[v_{i,T}|\widehat{v_{i,T-1}^{-i}}\right]$ is strictly increasing in $\widehat{v_{i,T-1}^{-i}}$, so that equation (3) implies that rivals' period T outputs will fall if they believe seller i has a lower period $T-1$ v. Therefore, seller i has an incentive to distort its $T-1$ output to "signal" a lower $T-1$ cost type.

We follow Mailath (1988) and Mailath (1989) to characterize the strategies that will satisfy the incentive compatibility conditions. To do so, we define seller i's t-period-specific "expected signaling payoff function", which returns i's expected payoff when it signals a type $\widehat{v_{i,t}^{-i}}$, and is assumed to be continuous and differentiable (with appropriate one-sided derivatives at the boundaries). Using

$V_{i,t+1}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, \boldsymbol{v_{-i,t}}\right)$ to denote i's value function at the start of the next period,

$$\Pi^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \pi\left(v_{i,t}, q_{i,t}, \boldsymbol{q}_{-i,t}\left(\boldsymbol{v}_{-i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)\right) g_{-i}(\boldsymbol{v}_{-i,t}|\boldsymbol{v}_{-i,t-1}) d\boldsymbol{v}_{-i,t} + \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \delta V_{i,t+1}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, \boldsymbol{v_{-i,t}}\right) g_{-i}(\boldsymbol{v}_{-i,t}|\boldsymbol{v}_{-i,t-1}) d\boldsymbol{v}_{-i,t}. \tag{4}$$

Strategies. Under four conditions, seller i's separating best response function in state $\left(v_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ can be uniquely defined as the solution to a differential equation

$$\frac{dq_{i,t}^{*}(v_{i,t})}{dv_{i,t}} = -\left.\frac{\Pi_2^{i,t}\left(v_{i,t},\ \widehat{v_{i,t}^{-i}},\ q_{i,t},\ \boldsymbol{v}_{-i,t-1},\ \widehat{v_{i,t-1}^{-i}}\right)}{\Pi_3^{i,t}\left(v_{i,t},\ \widehat{v_{i,t}^{-i}},\ q_{i,t},\ \boldsymbol{v}_{-i,t-1},\ \widehat{v_{i,t-1}^{-i}}\right)}\right|_{\widehat{v_{i,t}^{-i}}=v_{i,t},\ q_{i,t}=q_{i,t}^{*}(v_{i,t})} \tag{5}$$

where $\Pi_n^{i,t}$ denotes the partial derivative of $\Pi^{i,t}$ with respect to the nth argument, and a boundary condition where the highest v type will choose a static best response, i.e., $q_{i,t}^{*}\left(\overline{v}_i\right)$ solves

$$\Pi_3^{i,t}\left(\overline{v_i}, \overline{v_i}, q_{i,t}^{*}\left(\overline{v}_i\right), \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = 0. \tag{6}$$

The numerator in equation (5) is the derivative of i's continuation value with respect to an increase in rivals' beliefs about i's type. If rivals increase their $t+1$ outputs as $\widehat{v_{i,t}^{-i}}$ increases, the numerator will be negative. The denominator is the derivative of i's period-t expected profit with respect to $q_{i,t}$, and this will also be negative if i is producing more than its static best response, so the quantity function will slope downwards.[8]

Required Conditions. This characterization depends on four conditions. Conditions 1 and 2, formally stated and shown to hold in Appendix A, are shape and type monotonicity (always more

[8]Of course, exactly at the boundary condition, the denominator in (5) is zero and $\frac{dq_{i,t}^{*}(v_{i,t})}{dv_{i,t}}$ is not defined. When solving the differential equation we start at a quantity 10^{-5} above the best response quantity (Appendix B).

profitable to produce more when i has a lower cost type) conditions on static profits.[9] On the other hand, Conditions 3 and 4 depend on continuation values given equilibrium play.

Condition 3 (Belief Monotonicity) *Given* $\boldsymbol{v}_{-i,t-1}$ *and* $\widehat{v_{i,t-1}^{-i}}$, $\Pi_2^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ *is either greater than* 0 *for all* $\left(v_{i,t}, \widehat{v_{i,t}^{-i}}\right)$ *or less than* 0 *for all* $\left(v_{i,t}, \widehat{v_{i,t}^{-i}}\right)$.[10]

Appendix A shows that, if rivals also produce less in period $t+1$ when $\widehat{v_{i,t}^{-i}}$ is lower, then this condition will hold, and equation (3) implies this is true in period $T-1$. We verify belief monotonicity numerically in each period $t < T-1$ when we compute strategies (Appendix B).

Condition 4 (Single-Crossing) *Given* $\boldsymbol{v}_{-i,t-1}$ *and* $\widehat{v_{i,t-1}^{-i}}$, $\dfrac{\Pi_3^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)}{\Pi_2^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)}$ *is a monotone function of* $v_{i,t}$*, holding* $q_{i,t}$ *and* $\widehat{v_{i,t}^{-i}}$ *fixed, at every point* $(v_{i,t}, q_{i,t})$ *in the graph of the quantity function.*

We will numerically check that single-crossing holds in each period, and we have identified failures only when the desired outputs of some sellers would be zero, or very close to zero, for some cost realizations. This contrasts with price-setting games with logit or nested logit demand, where STY find that both the belief monotonicity and single-crossing conditions can fail when signaling incentives become too large.

The four conditions are sufficient for seller i to have a unique fully separating period t *best response* function, given the quantity functions of other sellers and given beliefs about previous period types. In general, these conditions are not sufficient for an equilibrium to be unique, but Appendix A shows that, with linear demand and quadratic costs, the conditions do imply that there will be a unique fully separating period t equilibrium given continuation values. A recursive argument, combined with the uniqueness of the period T equilibrium, can be used to prove uniqueness in a finite-horizon game where the conditions hold in every period.

[9] Note that a condition on the shape of the profit function requires that we allow prices to go negative (Hurkens (2014)).

[10] $\Pi_2^{i,t}$ does not depend on $q_{i,t}$.

3 Equilibrium Strategies and Welfare Outcomes: Examples

This section uses example parameters to illustrate signaling strategies, the effects of signaling on outcomes, and what happens when we change the number of sellers.

We assume that there are N ex-ante symmetric sellers of a homogeneous product, with inverse demand $P_t = u - \beta \sum_j q_{j,t}$, where $\beta = 0.05$ and $u = 200$. Each seller has a quadratic cost function, $C_{i,t} = v_{i,t} q_{i,t} + \frac{\lambda}{2} q_{i,t}^2$ where $\lambda = 0.01$, and for all sellers $v_{i,t}$ lies on an interval $[\underline{v} = 20, \overline{v} = 30]$. $v_{i,t}$ evolves according to an exogenous truncated AR(1) process

$$v_{i,t} = \rho v_{i,t-1} + (1 - \rho)\frac{\overline{v} + \underline{v}}{2} + \eta_{i,t}, \tag{7}$$

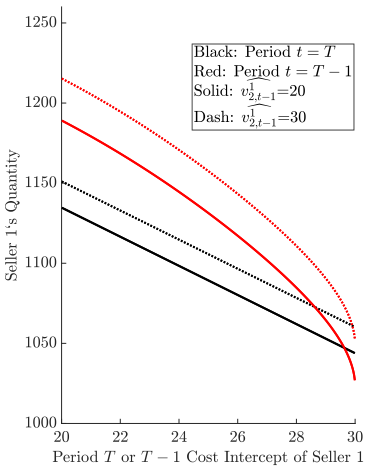
where ρ (the "AR(1) serial correlation parameter") equals 0.8, $\eta_{i,t}$ is truncated normal, with an untruncated mean and standard deviation (σ_v) of zero and 4 respectively, and lower and upper truncation points of $\underline{v} - \rho v_{i,t-1} - (1 - \rho)\frac{\overline{v}+\underline{v}}{2}$ and $\overline{v} - \rho v_{i,t-1} - (1 - \rho)\frac{\overline{v}+\underline{v}}{2}$ respectively. These assumptions imply that the probability that a seller with $v_{i,t-1}$ of 20 will have a $v_{i,t}$ of more than 25, i.e., above the unconditional median, in the next period is approximately 0.25, so we view this specification as reflecting moderately persistent cost shocks. The η innovations are independent across sellers. With these parameters, sellers' signaling strategies imply outputs, and prices, greater than zero for all possible cost realizations and for all sellers.

3.1 Duopoly

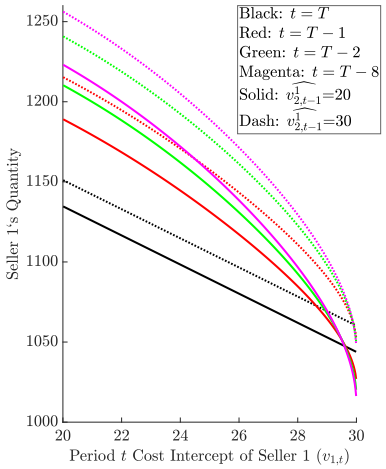
We first consider a duopoly game ($N = 2$). Appendix B details our solution method, and our verification that Conditions 3 and 4 are satisfied. We use grids with 21 equally spaced values for previous (actual and perceived) period costs and 400 values for each seller's own current cost to store strategies and values.

Figure 1: Seller 1 Strategies in a Finite Horizon Game (parameters described in text). $\widehat{v^j_{k,t-1}}$ denotes the v that rival seller j believes seller k had in the previous period. Panels (a) and (b) show strategies assuming that $\widehat{v^2_{1,t-1}} = 20$ in all cases, but varying seller 1's beliefs about seller 2. Panel (c) illustrates convergence of example strategies to those found by solving an infinite-horizon game for stationary strategies, which are represented by the corresponding dashed and dotted lines.

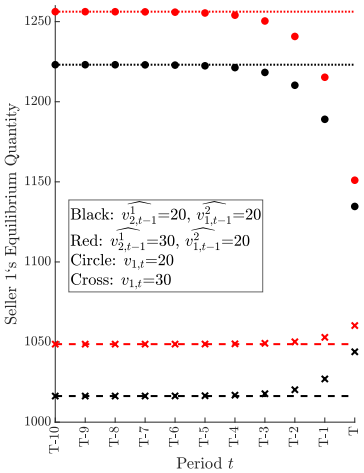
(a) T and $T-1$ Strategies for Different $\widehat{v^1_{2,t-1}}$



(b) Strategies in Multiple Periods for Alternative Beliefs



(c) Convergence of Finite Horizon Strategies



Strategies. The black lines in Figure 1(a) show seller 1's BNE strategies in the final period for two alternative values of seller 1's belief about seller 2's $T-1$ v, assuming that seller 2 believes $v_{1,T-1} = 20$. The strategies are linear in $v_{1,T}$, and, as expected, seller 1 produces more when it expects seller 2 to have a higher cost. The red lines show seller 1's strategies in $T-1$, for the same prior period beliefs. The curvature reflects the solutions to the differential equation. As seller 2 is also raising its output to signal, seller 1 produces less than in period T when $v_{1,t} = \overline{v}$, but its output increases significantly for lower $v_{1,T-1}$. The increase in seller 2's expected output, relative to its T period expected output, is larger when $\widehat{v^1_{2,T-2}}$ is small, so that the drop in seller 1's static best response quantity is larger for the solid (red) line than the dashed line. This leads to seller 1's $T-1$ quantity functions being more vertically spread out than its T functions, as well as having steeper slopes. A change in the type that a rival believes that a seller has in the previous period will therefore result in a larger change in the seller's $T-1$ expected output, compared to T, implying

that sellers will have stronger incentives to make their rival believe that their types are low in $T-2$ than they do in $T-1$.

The green lines in panel (b) show the same seller 1 strategy functions but in period $T-2$. Reflecting the stronger signaling incentives, they imply higher outputs unless $v_{1,T-2}$ is very close to $\overline{v}$, and they are also more steeply sloped and spread out than the $T-1$ functions implying stronger signaling incentives in $T-3$ than in $T-2$. The magenta lines show the strategies in $T-8$ (several periods earlier than considered by Mester (1992)), and these functions imply higher quantities, for almost all types, than in periods $T-2$, $T-1$ or T.

Convergence and Infinite-Horizon Strategies. If we go back further than $T-8$, strategies hardly change. This convergence is illustrated in more detail in panel (c), by considering each period back to $T-10$. The circles and crosses correspond to the highest and lowest points, respectively, on the quantity functions (i.e., those associated with $v_{1,t}=20$ and $v_{1,t}=30$). Black marks correspond to the strategies shown by the solid lines in panels (a) and (b), and red marks correspond to the dashed line strategies. It is notable that convergence of strategies happens in only a few periods given that the discount factor is high. This pattern is less surprising when it is recognized that a seller's period t signal only directly affects what happens in period $t+1$, as $t+1$ signals determine what happens after that. Therefore, the slope of the quantity function reflects the marginal effect of one period's inferred type on next period's expected payoff, and these marginal effects can converge in a few periods even though seller values would take hundreds of periods to converge. The dashed and dotted lines show the corresponding outputs for the stationary strategies computed when, instead of a finite-horizon game with period-specific strategies, we solve for stationary strategies in an infinite-horizon game (Appendix B). While it is plausible that multiple separating equilibria exist in the infinite-horizon game, these strategies represent the only equilibrium we have found for these

Table 1: Welfare and Output in $N = 2$ Games with Different Parameters. Inverse Demand is $P = 200 - 0.05 \sum_j q_j$ and each seller has a quadratic cost function, $C_{i,t} = v_{i,t} q_{i,t} + \frac{\lambda}{2} q_{i,t}^2$ where $\lambda = 0.01$, and $v_{i,t}$ lies on an interval $[\underline{v}, \overline{v}]$. Expected outcomes are averages over 100,000 simulations from the last 25 periods of a finite-horizon game (unless a specific t is indicated), where we draw the vs from a uniform distribution for $T - 45$ and simulate their evolution up to period T.

Parameters	(1)	(2)	(3)	(4)	(5)
$\underline{v}$	20	15	20	20	15
$\overline{v}$	30	35	30	30	35
AR(1) ρ	0.8	0.8	0.9	0.8	0.9
σ_v	4.0	4.0	4.0	2.5	2.5
Static Complete Information Nash					
Mean Total Output	2,187.5	2,187.5	2,187.5	2,187.5	2,187.5
Std. Dev. Total Output	23.6	41.9	23.8	22.0	40.0
Mean Consumer Surplus (/1000)	119.6	119.7	119.6	119.6	119.7
Mean Profits (/1000)	131.7	132.0	131.7	131.7	131.9
Mean Welfare (/1000)	251.4	251.7	251.4	251.3	251.6
Inefficiency Index	0.38%	1.23%	0.38%	0.33%	1.12%
Static Incomplete Information Nash					
Mean Total Output	2,187.5	2,187.5	2,187.5	2,187.5	2,187.5
Std. Dev. Total Output	33.4	53.5	33.4	29.2	45.8
Mean Consumer Surplus (/1000)	119.7	119.7	119.7	119.7	119.7
Mean Profits (/1000)	131.7	131.9	131.7	131.7	131.9
Mean Welfare (/1000)	251.3	251.6	251.3	251.3	251.6
Inefficiency Index	0.44%	1.36%	0.45%	0.37%	1.18%
Finite-Horizon Dynamic Incomplete Information MPBE, Period $T - 20$					
Mean Total Output	2,315.9	2,460.6	2,331.2	2,418.1	2,598.0
Std. Dev. Total Output	66.9	110.4	72.1	86.6	121.8
Mean Consumer Surplus (/1000)	134.2	151.7	136.0	146.4	169.1
Mean Profits (/1000)	123.7	113.0	122.7	116.1	100.6
Mean Welfare (/1000)	257.9	264.6	258.6	262.5	269.7
Inefficiency Index	0.33%	0.87%	0.32%	0.19%	0.43%

parameters. Differences in quantities between finite-horizon $T - 10$ strategies and infinite-horizon strategies are of the order of 10^{-6}, or around one billionth of average output.

The first column of Table 1 reports expected output and welfare measures, under three alternative equilibrium concepts, when we assume the parameters listed above:

- static complete information Nash equilibrium given realized costs in each period[11];
- static incomplete information Bayesian Nash equilibrium strategies where sellers condition

[11] These strategies would be the unique subgame perfect Nash equilibrium strategies if the game has complete information and a finite horizon.

on all sellers' previous-period types but ignore signaling incentives when choosing current output.[12] This conditioning makes the comparison informative when costs are persistent: it preserves the predictive content of lagged types while removing the incentive to manipulate future beliefs; and

- finite-horizon dynamic signaling MPBE strategies, where we report outcomes in $T - 20$.[13]

The reported outcomes include an "inefficiency index", defined as the percentage increase in total variable costs given equilibrium play relative to the minimum costs of producing the same level of total output given realized cost intercepts.[14] In these examples, the percentages are small because the sellers are ex-ante symmetric. Appendix Table B.1 also reports signaling equilibrium outcomes in $T-1$, and when we solve the infinite-horizon game (infinite-horizon outcomes are almost identical to those reported for $T - 20$).

Consistent with Vives (2002), expected outputs are the same in the static complete information and incomplete information equilibria. The static incomplete information equilibrium leads to higher consumer surplus, lower total welfare and less efficient production, but the percentage differences are small even though the standard deviation of output is about 42% larger. However, combining incomplete information and dynamics leads to larger quantitative differences. For example, expected $T - 20$ consumer surplus and total welfare in the signaling game are 12.2% and 2.6% higher than under complete information, with the total welfare increase reflecting the 6% increase in expected output. The dynamic signaling equilibrium also leads to more efficient production than the complete information equilibrium, at least for these parameters, reflecting how low-cost sellers expand their outputs to signal.

[12] These strategies would be the unique equilibrium strategies if period $t - 1$ vs were publicly revealed after $t - 1$ quantities are chosen, but before period t quantities are chosen.

[13] In specifications (1)-(4) strategies converge within ten periods; convergence takes more periods for specification (5), but strategy changes are tiny by $T - 20$.

[14] For example, if equilibrium play leads to expected costs of 206 and minimized costs would be 200, then the inefficiency index is 3%.

The remaining columns of Table 1 present the outcomes for four alternative parameterizations. Columns (2)-(4) increase $\overline{v} - \underline{v}$, increase serial correlation (ρ) and reduce the standard deviation of the innovations (σ_v) respectively. Column (5) implements all three changes together. All of these changes go in the direction of making a seller's current costs more informative about its next period costs, so that we would expect receivers to respond more to the signals they receive, and therefore the quantity increases required to achieve separation to increase. The unconditional mean of v remains the same as in column (1) in each case, and the expected surplus outcomes change little across the columns when we assume either form of static play. On the other hand, the welfare differences between the static and the dynamic signaling outcomes become significantly larger: for example, in column (5) expected $T - 20$ consumer surplus and total welfare are 41.2% and 7.2% higher in the signaling equilibrium than in the static complete information equilibrium.

3.2 Output and Welfare for Different Numbers of Sellers

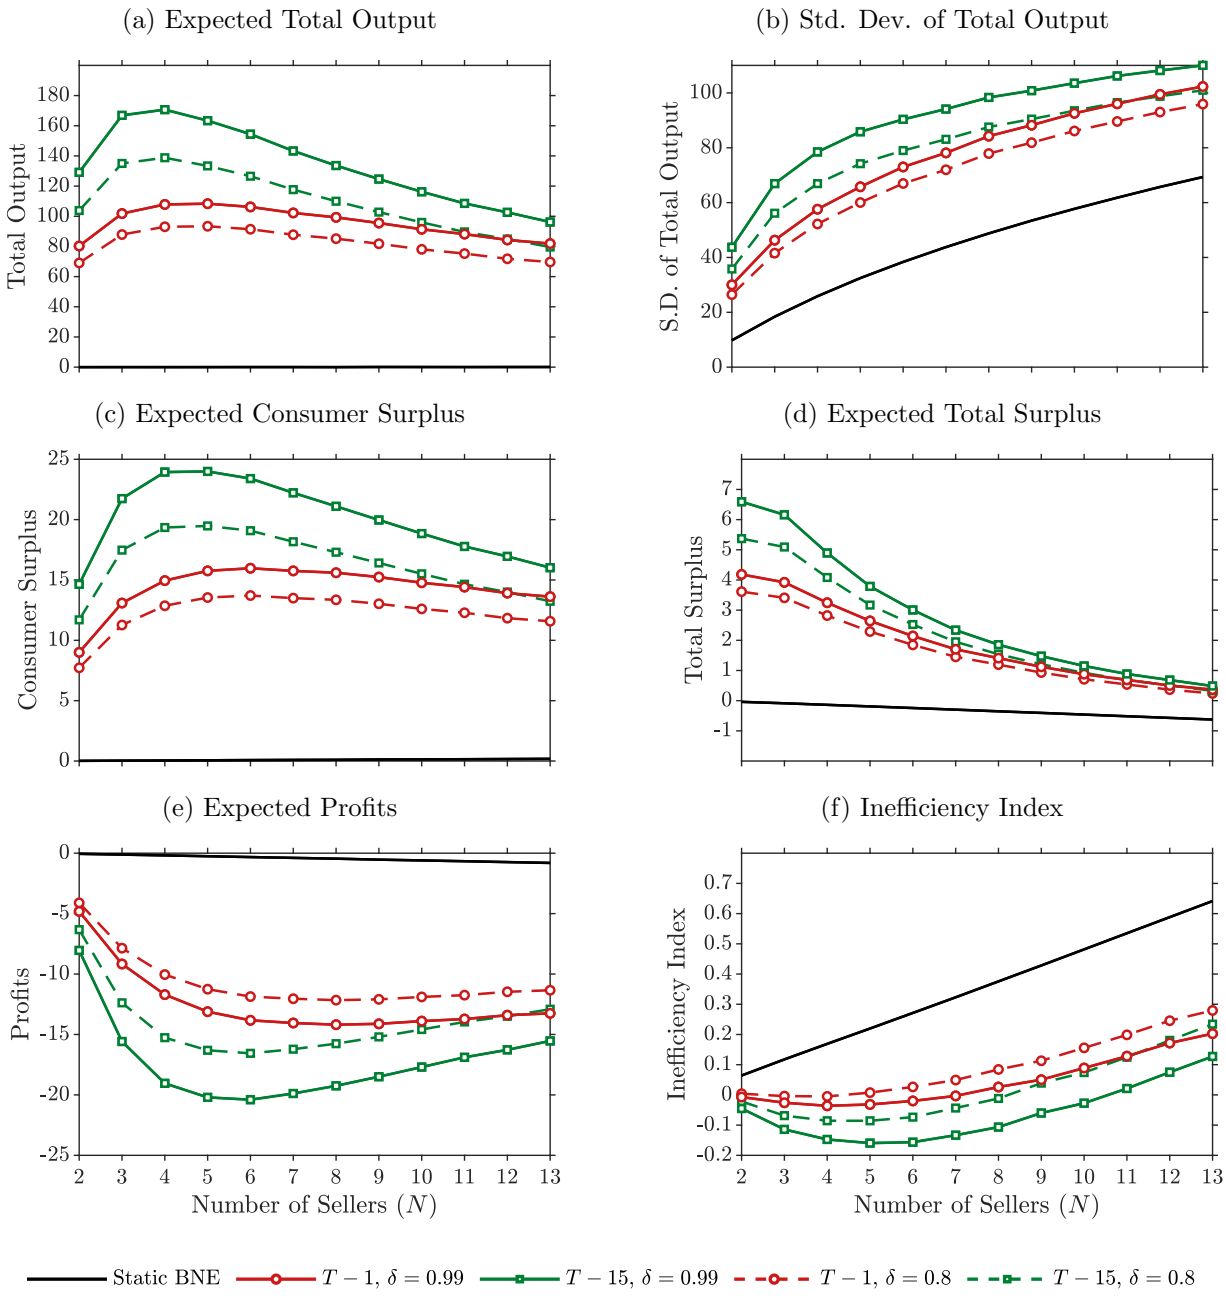
We now examine what happens when we allow more sellers. A common intuition might be that the effects of signaling should decline quite quickly as sellers are added, as STY find in the price-setting case, because sellers have fewer inframarginal units that might benefit from a reduction in rivals' output and rivals are likely to respond less to any one seller's signal.

We solve games with 2 to 13 symmetric sellers with the same demand and cost parameters as our baseline specification. The growth of the state space requires us to use only four grid points for each seller's previous period costs. Appendix Table B.2 compares computed outcomes when we use 4, 6 or 8 grid points with 2, 4 or 6 sellers, and finds differences are small.[15] For this parameterization, seller outputs are always strictly positive, and Conditions 3 and 4 are satisfied in every period that we look at.

Figure 2 presents differences in output and welfare measures from their expected values in

[15] When we compare $N = 2$ dynamic signaling outcomes computed with 4 or 21 grid points, the differences in expected outputs are around 0.03%.

Figure 2: Absolute Differences in Output and Welfare Relative to the Static Complete Information Nash Equilibrium, $N = 2$ to $N = 13$ Ex-Ante Symmetric Sellers. Linear inverse demand $P = 200 - 0.05 \sum_j q_j$, v on $[20, 30]$, $\rho = 0.8$, $\sigma_v = 4$. All panels report the absolute difference in average outcomes (based on 100,000 simulation draws) from the static complete information Nash outcome; the horizontal axis is the number of sellers in every panel, and tick labels are shown on the bottom row only. Signaling outcomes are finite-horizon MPBE strategies in periods $T-1$ and $T-15$. The units of the inefficiency index are as a percentage of minimized total costs.

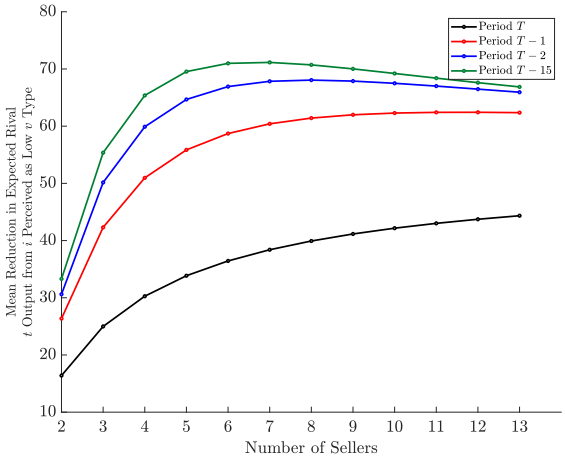


the static complete information equilibrium. The black lines indicate differences for the static incomplete information equilibrium. For dynamic signaling, we consider differences in outcomes at both $T-1$ and $T-15$ in the finite-horizon game (computational limitations prevent us from solving the infinite-horizon outcome with as many as 13 sellers). We add another comparison by solving games with discount factors of $\delta = 0.99$ and 0.8.

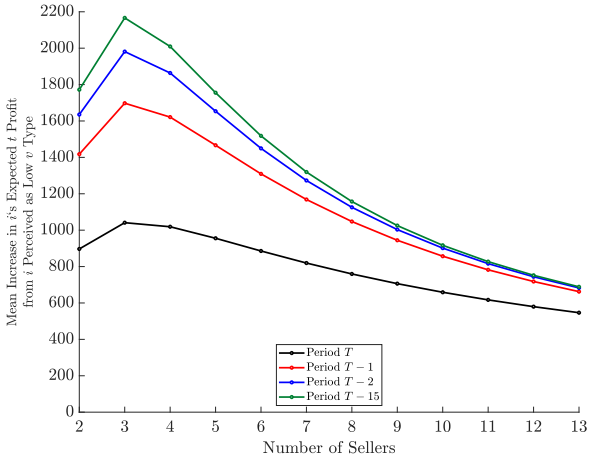
For all N, expected output (panel (a)) in the static incomplete information game is the same as in the complete information game. On the other hand, dynamic signaling raises expected outputs for all N, and the effects are large even when $N = 13$. For example, when $\delta = 0.99$, signaling raises expected $T-15$ total output, relative to complete information, by 3.0%, or 96 units, compared to 5.9%, or 129 units, when $N = 2$. Both absolute and percentage differences in output with dynamic signaling are larger when either $N = 3$ or $N = 4$ than when $N = 2$, and, when $\delta = 0.99$, the absolute difference in $T-15$ expected output is larger when $N = 8$ (134) than when $N = 2$ (129). The differences are smaller for all N and both periods if $\delta = 0.8$, but the effects of signaling remain substantial, reflecting how signaling is driven by the incentive to affect outcomes in the next period, rather than many periods in the future. The effects of signaling on consumer surplus are larger when $N = 13$ than when $N = 2$, reflecting how not only the level of output but also the standard deviation of output in the signaling equilibrium rises, relative to the complete information equilibrium. Signaling lowers expected profits and increases expected total surplus, relative to complete information, for all N, although the size of the total surplus increase becomes smaller as N rises, whereas incomplete information play without dynamics causes total surplus to fall. As noted above, signaling can increase efficiency to the extent that it corrects the tendency for low-cost sellers to produce too little in static equilibria. On the other hand, the presence of incomplete information can still lead to inefficient production, as sellers do not respond to other sellers' realized cost shocks. These effects are illustrated in panel (f) where $T-15$ signaling output

Figure 3: Measures of Seller i's Benefit from Being Perceived as a Low-Cost Type in Finite-Horizon Dynamic Signaling Games with Different N. Panel (a) shows the average (across last period cost states) reduction in expected combined rival output when $\widehat{v_{i,t-1}^{-i}} = \underline{v} = 20$ rather than $\widehat{v_{i,t-1}^{-i}} = \overline{v} = 30$. Panel (b) shows the change in i's expected period t profit. Calculations assume $\delta = 0.99$. The reported values are simple averages of the changes across the prior-period types of all sellers (including i).

(a) Reduction in Expected Combined Rival Total Output



(b) Increase in i's Expected Period t Profit



is produced more efficiently than it would be under complete information when $N \leq 10$, but it is produced less efficiently when N is larger.

Explanation for Why Dynamic Signaling Incentives Remain Substantial When N Increases. Appendix A.4, which considers whether the belief monotonicity condition holds, shows that $\frac{\partial V_{i,t+1}(v_{i,t},\widehat{v_{i,t}^{-i}},\boldsymbol{v}_{-i,t})}{\partial \widehat{v_{i,t}^{-i}}}$, i.e., the change in i's future value when rivals believe that i has a higher cost type is equal to $-\beta \times i$'s expected $t+1$ best-response output $\times$ the derivative of expected combined rival $t+1$ output with respect to $\widehat{v_{i,t}^{-i}}$.[16] The same appendix also shows that, after discounting and taking the expectation across rivals' types, this expression gives the numerator in differential equation (5), so it captures the source of a seller's incentives to signal.

[16]The derivation of the three-part expression is not immediate, and it reflects how the effect of shifting i's own quantity schedule exactly cancels the effect of rival output expansion on the units that i produces above its static best response.

The formula can be used to think through how these incentives change with N. As N increases, i's expected output will shrink, reducing signaling incentives. However, and this works against the intuition mentioned at the start of this subsection, the derivative of expected combined rival $t+1$ output with respect to $\widehat{v_{i,t}^{-i}}$ can become larger, rather than smaller, as N increases. To see this, consider the analytic derivative of expected rival output in period T:

$$\frac{\partial \sum_{j \neq i} \mathbb{E}[q_{j,T}]}{\partial \widehat{\mu_{i,T}^{v,-i}}} = \frac{\beta(N-1)}{(\lambda+\beta)\left(\lambda+(N+1)\beta\right)}. \tag{8}$$

This derivative is positive and *increasing* in N, i.e., aggregate rival output becomes *more* sensitive to the expected v of any individual seller.

The lines in Figure 3(a) show the (numerically computed) average of the expected change in rival sellers' combined period t output when rivals believe i has a $\underline{v}$ type, compared to type $\overline{v}$, in $t-1$ for periods $t = T$, $T-1$, $T-2$ and $T-15$, as a function of N. The period T black line increases monotonically with N, consistent with the formula considered above. In earlier periods, analytic formulae are not available, although the fact that the boundary condition reflects the highest cost type making a static best response means that we would still expect the effect identified for period T, where all types are best-responding, to remain relevant. We see that the responsiveness of rival output peaks at $N = 8$ for $T-2$ and $N = 7$ for $T-15$. However, the key point to note is that the change in the seller's perceived type has a larger effect on rival output when $N = 13$ than when $N = 2$ or $N = 3$. Panel (b) shows the effect of the change in i's expected period t profit so that the effects in panel (a) are combined with the decrease in i's expected output as N rises. For all periods shown, the effect on profit is maximized when $N = 3$, but, consistent with the results in Figure 2, it remains substantial when $N = 13$.

4 Application: The World Oil Market

AZ quantify the effects of incomplete information on outcomes in the world oil market. In this section we use an adapted version of their model to quantify the effects of dynamic signaling on outcomes and two stylized counterfactuals, and to compare its effects with an alternative model where there is also strategic behavior. The setting is well-suited for our context in the sense that AZ's specifications imply that rivals' outputs are sensitive to how much they expect any particular seller to produce, which, as we have shown, is a driver of signaling incentives. On the other hand, it is important to emphasize that the model abstracts from many important features of the oil market, so that the counterfactuals should be interpreted as stylized experiments with the model rather than quantitative predictions about how outcomes would change in the real world.

4.1 AZ's Static Empirical Model

AZ develop a model that they apply to detrended output data from twenty countries and deflated price data.[17] Inverse demand in period t is $P_t = u_t - \beta \sum_i q_{i,t}$ and each seller i has quadratic costs $C(q_{i,t}, v_{i,t}, w_t) = (v_{i,t} + w_t)q_{i,t} + \frac{\lambda}{2} q_{i,t}^2$. The $v_{i,t}$s are private information. AZ's model departs from the model we considered above in having common demand and cost shocks (u_t and w_t) that are assumed to be observed by all sellers in period t, and which are i.i.d. across periods. AZ assume that $v_{i,t} \sim \overline{w} \times (\text{Beta}(a_i, b_i) + 1)$ where a_i, b_i and $\overline{w} > 0$, and the a_is and b_is vary across six country groups, but the supports of the vs, $[\overline{w}, 2\overline{w}]$, are the same for all groups. The parameters are estimated using maximum likelihood assuming that sellers play a static Bayesian Nash equilibrium with linear quantity strategies. Their estimates imply that the common shocks and the i.i.d. country-level shocks have wide supports (e.g., the vs have supports from \$28.206 per barrel to \$56.412 per barrel).

[17] See AZ Assumption 4, and their subsequent discussion, for the assumptions AZ make to detrend the data.

4.2 Serial Correlation as a Feature of the Data

AZ's static specification treats all shocks as independent across months, which is natural for their static analysis. For our model, we need to specify the persistence of the private-information cost shocks. AZ's Figures 1 and 4 indicate that country-level output is highly persistent, and country-specific AR(1) regressions give an average AR(1) coefficient of 0.969, with similar averages for OPEC and non-OPEC countries. Appendix C presents an analysis of the serial correlation in the u_ts, w_ts and $v_{i,t}$s implied by the application of AZ's estimated model to the data period-by-period. This analysis supports two conclusions underpinning assumptions that we make below: the $u_t - w_t$ difference between the common shocks, which is what affects output choices, is quite stable over time, particularly during the early years of the 1992-2019 data; and the implied $v_{i,t}$ are highly persistent with changes of more than 50 cents being atypical.[18] Of course, we recognize that we are assuming that persistence reflects the persistence of a cost variable that is private information to the seller, rather than the persistence of variables that may be observable to rivals including capacity constraints.[19]

4.3 Our Specification

It is not computationally feasible to solve a signaling game with twenty asymmetric sellers. Moreover, it may not be plausible to treat all countries as signalers when some of them have production dispersed across many small and uncoordinated producers, whose individual outputs would be hard to observe. Our baseline specifications make the following assumptions. We will explain how the $[\underline{v_i}, \overline{v_i}]$ supports are chosen later.

- the monthly discount factor is $\delta = 0.99$.

[18] As a comparison, AZ's estimates and their independence assumption imply that month-to-month absolute changes in the vs would average \$4.78.

[19] On the other hand, our model would still apply if a variable that causes persistence, such as a capacity constraint, only becomes publicly observable with a lag of a couple of periods.

- inverse demand is $P_t = u - \beta \sum_i q_{i,t}$, where $\beta = 0.03329$ (AZ's estimate) and $u = 100$.

- seller i's costs are $C(q_{i,t}, v_{i,t}) = v_{i,t}q_{i,t} + \frac{\lambda}{2}q_{i,t}^2$. We use AZ's estimate of $\lambda = 0.0007$. w is therefore set equal to zero.[20]

- the supports of $v_{i,t}$ are $[\underline{v_i}, \overline{v_i}]$, with $\overline{v_i} - \underline{v_i} = 1.5$, for each seller.[21]

- Saudi Arabia, Russia, Iran, Iraq and China are "signaling sellers" with serially correlated vs (note we will identify these countries as signaling sellers even when we consider equilibria without signaling). Their vs evolve according to independent AR(1) processes, $v_{i,t} = (1 - \rho)\frac{\overline{v_i}+\underline{v_i}}{2} + \rho v_{i,t-1} + \varepsilon_{i,t}$. We assume $\rho = 0.96$ (the median of these countries' estimated ρs in Table C.1 is 0.959), and the $\varepsilon_{i,t}$ innovations are i.i.d. draws from a normal distribution, which is then truncated so the vs remain on their supports. The untruncated innovation distribution has mean zero with a standard deviation of 0.5, so that the standard deviation of the truncated draws is, on average, 0.4 (lower than the estimated standard deviation of the AR(1) residuals in Table C.1 for Saudi Arabia and Iraq, but higher than for the other signaling countries).[22]

- the remaining countries ("non-signaling" sellers) are assumed to have private information vs which are not serially correlated. They therefore have no incentives to signal their current costs, but we assume that they are sophisticated in the sense of having expectations about

[20] We have solved specifications where u_t and w_t vary i.i.d. across periods, which increases the computational burden and does not appear to affect the comparisons we are interested in, which is intuitive given that, when these shocks are i.i.d., their current values do not affect the expected future payoff from signaling private information. We choose $u_t = 100$ and $w_t = 0$, so that $u_t - w_t = 100$ for simplicity.

[21] This choice is designed to reflect how, even if some countries' vs change by several dollars over the 28 years, they change much less over shorter periods. For example, across country-years, the average within-calendar year difference between the maximum and minimum $v_{i,t}$s implied by AZ's model is \$1.06.

[22] We recognize that our calculations in Appendix C use AZ's estimated distributions as priors, and that this may tend to lower the estimated variances of the residuals in the AR(1) regressions. However, the estimated standard deviations of the innovations do not increase much (e.g., the Saudi Arabia estimate increases from 0.44 to 0.55) if we use uniform priors, and our use of a standard deviation of 0.5 is, if anything, generous since roughly 95% or more of the implied changes in the vs for most countries are less than \$0.50. If we change the standard deviation of the untruncated innovation distribution from 0.5 to 0.7, the effects of signaling are smaller but remain substantial. For example, the first four values for the combined output of the signaling sellers in Table 2 would be 1079.7, 1079.7, 1585.4 and 2186.2, so that signaling would still result in a 47% increase in signaling output.

the signaling sellers' output that are consistent with equilibrium play. The non-signaling sellers' vs are drawn, i.i.d. each period, from symmetric truncated normal distributions where the untruncated distributions have means equal to $\frac{\overline{v_i}+\underline{v_i}}{2}$, with $\overline{v_i} - \underline{v_i} = 1.5$, and standard deviations equal to 0.6 (implying the truncated vs have a standard deviation of 0.39).

We recognize that our assumption that the non-signaling sellers have uncorrelated costs is inconsistent with the large serial correlation coefficients for these countries in Appendix Table C.1. We make this assumption because we need to limit the size of the state space, which will be based on the cost types of the signaling sellers, especially once we allow for sellers to be asymmetric. Our chosen signaling countries are five large producers, with the selection also reflecting whether country-level production decisions can plausibly be interpreted as strategic. For example, we should not interpret country-level output levels in countries such as the U.S., where production decisions are taken by many independent producers, as reflecting signaling incentives.[23] One might think of the non-signaling sellers as playing the role of a "fringe" within our model, but we will assume that these sellers use static best response strategies, given the quantities they expect other sellers to choose, rather than acting competitively.[24]

4.4 Comparing Static and Dynamic Signaling Equilibrium Outcomes for Given Parameters

We begin our presentation of results by comparing outcomes when some sellers are using dynamic signaling strategies with outcomes when they are using static complete information or incomplete information Nash strategies for given parameters (the same approach as in Section 3).

In columns (1)-(4), the $\underline{v_i}$ and $\overline{v_i}$ parameters of each seller i are calibrated so that, when we assume i uses a static incomplete information best response function, and it expects other sellers to

[23] We accept that one might believe that it would be more reasonable to think of oil exporters such as U.A.E. or Kuwait as being as likely to signal as a country that consumes most of its production domestically such as China.

[24] We have experimented with specifications where these sellers do set quantities competitively, with the insights about the equilibrium effects of signaling similar to those described here.

produce their average (detrended) output (as reported in Appendix Table C.1), i's expected output equals its own observed average output.[25]

The first four columns of Table 2 report the mean $v_{i,t}$ values implied by static behavior for the signaling sellers and the U.S. Across all countries, these values range from 7.5 to 19.2, with the maximum for Ecuador, which is not shown in the table. Even though the implied average outputs are the same as in AZ's static model, the average values of the cost intercepts are different because of our different assumptions about u and w. For example, if we were to increase our assumed u to 135, the levels of our mean $v_{i,t}$s would increase by 35, and prices would increase by 35, but outputs would remain the same in every column.[26] For this reason, and also because our specification assumes that the λ parameter is the same for all sellers, whereas the cost curves in different countries would vary with reserves and production technologies, we do not believe it would be appropriate to try to "test" any of the models we consider by asking whether the estimated cost intercepts are empirically "reasonable".

The first two columns of Table 2 show simulated average outcomes in the static equilibria, with percentage changes from the complete information outcome underneath the incomplete-information entry. The solution of our static incomplete information model assumes that, for the signaling sellers, the game is played conditional on the $v_{i,t}$s in the previous period, as in Section 3. In the static specifications in columns (1) and (2), all sellers have positive outputs for all cost realizations, so the positive-output condition in Vives (2002) applies, and equal simulated output levels are therefore expected. Across the simulations, the variance of output is higher in the static BNE model

[25] i's best response function is $q_{i,t}^{BR} = \frac{u-\beta\mathbb{E}Q_{-i,t}-v_{i,t}}{\lambda+2\beta}$. Our assumption that the unconditional distribution of v_i is symmetric on a support with width 1.5 therefore leads us to calculate $\underline{v_i} = u - \beta\overline{Q_{-i}} - (\lambda + 2\beta)\overline{q_i} - 0.75$, where $\overline{q_i}$ is i's average output and $\overline{Q_{-i}}$ is the combined average output of the other sellers. Our assumption of linear quantity functions is valid because, when rivals are expected to produce their average outputs, the assumed narrow width of our supports implies that every seller's best response output is strictly positive for any cost realization.

[26] Of course, percentage changes in prices across the columns would be different. The values of the inefficiency indices (described below) would change, as they are normalized by total costs, which would increase. However, the percentage relationships between columns would be unaffected. For example, the signaling seller inefficiency index in column (4) is 59% of the value in column (3). If $u = 135$, the reported indices would fall but the second value remains 59% of the first value.

Table 2: Outcomes under Alternative Equilibrium Assumptions and Alternative Calibrations. The % differences underneath the outcomes are the percentage differences relative to the first column. Signaling outcomes are based on equilibrium strategies calculated in period $T - 20$. Outcomes are averages across 100,000 simulations. In all columns, the v_i supports of non-signaling sellers are calibrated so that their expected outputs match their average detrended outputs when they use static best responses to all other sellers producing their average outputs. "SD Δs Signaling Output" refers to the standard deviation (SD) of cross-period changes in the combined output of the signaling sellers. CI = complete information. II = incomplete information.

Calibration Assumption for Signaling Seller $[\underline{v_i}, \overline{v_i}]$s	Static Behavior				Dynamic Signaling	Leaders-Followers
Midpoint of Support of $v_{i,t}$s for Signaling Countries, plus U.S.						
Saudi Arabia	7.5	7.5	7.5	7.5	14.8	18.8
Russia	8.3	8.3	8.3	8.3	15.1	18.9
China	15.1	15.1	15.1	15.1	17.9	19.5
Iran	15.7	15.7	15.7	15.7	18.1	19.5
Iraq	15.8	15.8	15.8	15.8	18.2	19.5
U.S.	7.5	7.5	7.5	7.5	7.5	7.5
	(1)	(2)	(3)	(4)	(5)	(6)
Average Simulated Equilibrium Outcomes	Static CI Nash	Static II BNE	Dynamic Signaling	Leaders-Followers	Dynamic Signaling	Leaders-Followers
Total Output MM bbl/month	2,409.2	2,409.2 (0.0%)	2,469.4 (+2.5%)	2,500.7 (+3.8%)	2,410.5 (+0.1%)	2,409.1 (0.0%)
SD Total Output MM bbl/month	2.5	25.4 (+906%)	97.0 (+3744%)	10.9 (+332%)	59.4 (+2254%)	22.9 (+809%)
Signaling Output	1,079.7	1,079.7 (0.0%)	1,901.3 (+76.1%)	2,186.2 (+102.5%)	1,079.5 (+0.0%)	1,079.2 (-0.0%)
SD Δs Signaling Output	25.2	13.0 (-48.4%)	108.7 (+331.6%)	6.6 (-73.8%)	62.0 (+146.0%)	68.6 (+172.2%)
Saudi Arabia	362.7	362.8 (0.0%)	738.1 (+103.5%)	971.4 (+167.8%)	362.5 (-0.0%)	362.1 (-0.2%)
Russia	339.8	339.9 (0.0%)	682.8 (+100.9%)	888.2 (+161.4%)	339.9 (0.0%)	339.8 (0.0%)
China	138.3	138.3 (0.0%)	190.8 (+38.0%)	154.7 (+11.9%)	138.3 (0.0%)	137.7 (-0.4%)
Iran	121.8	121.8 (0.0%)	150.7 (+23.7%)	94.8 (-22.2%)	121.8 (0.0%)	122.5 (+0.6%)
Iraq	117.0	117.0 (0.0%)	138.9 (+18.7%)	77.1 (-34.1%)	117.1 (+0.0%)	117.1 (+0.0%)
Non-Signaling Output	1,329.5	1,329.5 (0.0%)	568.1 (-57.3%)	314.5 (-76.3%)	1,331.0 (+0.1%)	1,329.9 (+0.0%)
Price $ per barrel	19.81	19.81 (0.0%)	17.80 (-10.1%)	16.76 (-15.4%)	19.76 (-0.3%)	19.81 (0.0%)
Consumer Surplus $m. per month	96,600	96,616 (+0.0%)	101,647 (+5.2%)	104,086 (+7.7%)	96,768 (+0.2%)	96,607 (+0.0%)
Signaling Profits $m. per month	9,936	9,921 (-0.2%)	14,638 (+47.3%)	16,442 (+65.5%)	3,865 (-61.1%)	891 (-91.0%)
Non-Signaling Profits $m. per month	7,184	7,136 (-0.7%)	3,530 (-50.9%)	2,506 (-65.1%)	7,088 (-1.3%)	7,150 (-0.5%)
Total Surplus $m. per month	113,720	113,674 (-0.0%)	119,814 (+5.4%)	123,033 (+8.2%)	107,721 (-5.3%)	104,648 (-8.0%)
Inefficiency Indices, Allowing Reallocation Between						
Signaling Sellers Only	36.0%	36.1%	24.0%	14.1%	7.0%	0.6%
All Sellers	62.7%	62.9%	32.0%	17.3%	83.1%	98.7%

than the static CI model. However, a slightly different picture emerges if we look at the standard deviation of *period-to-period changes* in the combined output of the five signaling sellers. This standard deviation is 13.0 for the static incomplete information BNE model, but higher (25.2) for the static complete information Nash model. The fact that consumer surplus and profit differences are small, in percentage terms, reflects how our specifications allow limited within-country variation in the $v_{i,t}$s, although we will see that this variation is more than enough for signaling to have large effects. We also report the value of the inefficiency index. Whereas the values of these indices are small in Section 3, where sellers are symmetric, they can be much larger here where sellers' distributions of costs do not necessarily overlap. As expected, the static incomplete information equilibrium implies that production is more inefficient, but the difference is small.

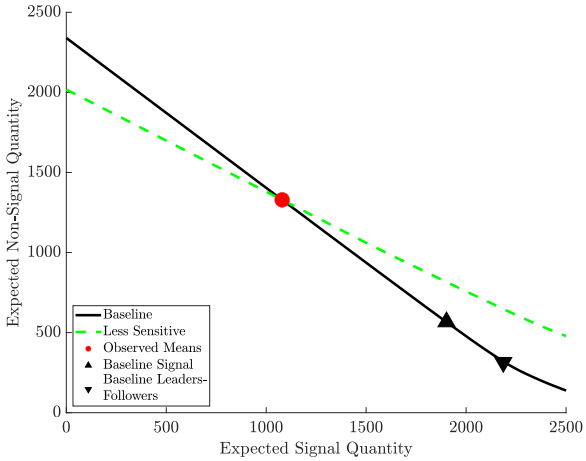
We note that AZ report quite large differences between static complete and incomplete information outcomes, including the level of output. The explanation is that AZ's estimated cost-shock distributions have very wide supports and imply that sellers will often produce zero. The strictly-positive-output assumption in Vives (2002) therefore does not hold and expected outputs can also differ across information structures.

Column (3) shows outcomes when we solve the finite-horizon signaling model (method detailed in Appendix B), and we simulate outcomes using $T - 20$ strategies (using strategies from earlier periods gives almost identical values). A feature of our method is that we pre-solve for a period-invariant function that returns the expected combined equilibrium output of the non-signaling sellers given a level of the expected output of the signaling sellers. We then jointly solve for the five signaling sellers' quantity functions and the non-signaling aggregate output consistent with that function. Individual non-signaling sellers may choose zero output for some cost realizations, although their aggregate output remains positive in the equilibria we consider.

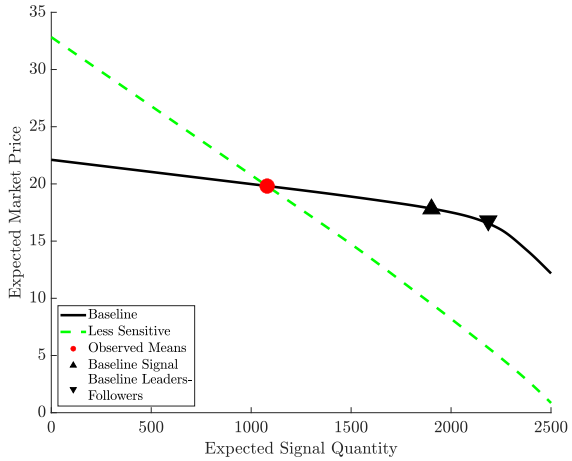
The solid black line in Figure 4(a) plots this function for our baseline specification. The red dot

Figure 4: Expected Output of the Non-Signaling Sellers and the Expected Price Under the Baseline Specification and the Specification with Less Sensitive Non-Signaling Output. The red dot corresponds to the outcome associated with the average detrended outputs, and this also matches the static BNE outcome in column (2) of Table 2. The black triangles mark average outcomes in columns (3) and (4) of the same table.

(a) Expected Non-Signaling Output



(b) Expected Price



marks the average observed output of the signaling sellers and the non-signaling sellers in the data, which is also the average outcome in the simulated static games. Unless expected signaling output is so large that many of the non-signaling sellers find production to be unprofitable for some, or all, v draws, the function implies that a one-barrel decrease in expected signaling output increases non-signaling output by approximately 0.94 barrels, so that the effect on the expected market price (panel (b)) is small. This high replacement rate also implies that, even if signaling sellers were to ignore how they might affect each other's next period outputs, they would have strong incentives to signal that their costs are low in order to reduce the next period output of the non-signaling sellers.

Comparing column (3) of Table 2 with the first two columns shows that, on average, dynamic signaling increases average total output by 2.5% and lowers prices by 10.1%. The reallocation of output across the sellers is more dramatic: Saudi Arabia and Russia double their outputs, and non-signaling sellers lower their combined output by 57%. The ▲ symbols in Figure 4 indicate

the average signaling quantity and price outcomes. A feature of our assumed specification is that, apart from the U.S., it is the lower-cost sellers that are assumed to be signaling. The reallocation of output towards these sellers means that, despite the fall in prices, total profits increase and output is produced much more efficiently. As in Section 3, signaling is associated with a very large proportional increase in the variance of total output relative to the static complete information model, with the static incomplete information model somewhere in between. The standard deviation of the changes in the output of the signaling sellers is 108.7, which is much larger than in either of the static models. The corresponding standard deviation of cross-period changes in the combined output of the five signaling sellers is 11.6 in the detrended data. In our setting, these volatility measures do not provide a clean way to discriminate between the models because we are abstracting from real-world factors, including demand shocks, adjustment costs, capacity constraints and inventory changes that undoubtedly play important roles in determining how output adjusts from month to month.[27]

We can also calibrate the signaling seller v_i supports so that their equilibrium signaling outputs match their observed outputs in the data. Column (5) reports the calibrated mean $v_{i,t}$s, and the equilibrium outcomes when we solve the calibrated model.[28] To match the observed average outputs, the calibrated vs of the signaling sellers with the largest outputs, such as Saudi Arabia, increase significantly relative to the intercepts calibrated using the static model, so that the intercepts are actually larger than those of several of the non-signaling sellers. While the exact calibrated values of the cost intercepts reflect the somewhat implausible assumption that all sellers

[27]The higher standard deviation of quantity changes across periods for a signaling seller is driven by the steepness of its quantity functions and by how much its type can change across periods. Note that if one assumes that there is a high probability that a seller's type does not change at all, then it is possible to have an equilibrium with large signaling effects on total output and small expected changes in output across periods. Therefore observed small changes in output across periods do not necessarily indicate that the equilibrium does not involve signaling.

[28]Consistent with the model, the supports of the non-signaling sellers are calibrated assuming static behavior. The calibration of signaling seller supports is done by repeatedly solving the model, and minimizing the sum of squared differences between predicted and observed outputs of the signaling sellers. As the distribution of signaling seller outputs across states changes, the average output of the non-signaling sellers, which we do not target in the calibration, changes slightly.

have quadratic cost functions with a common quadratic-cost slope parameter, the main point that we take away is that allowing for signaling can affect the level of implied costs substantially even in a market with numerous sellers. The change in intercepts naturally has the effect of decreasing implied signaling seller profits and changing how the reallocation of output between signaling sellers and between signaling sellers and non-signaling sellers affects the efficiency of production (now the planner would like the non-signaling sellers to produce more). Interestingly, the standard deviations of both the level of output and output changes implied by the signaling model fall substantially, relative to those in column (3), when we use the signaling model-calibrated parameters, although the value of the standard deviation of output changes remains well above the value in the detrended data.

4.5 Counterfactuals

We now compare how outcomes change under two stylized counterfactuals when we use either the static or the calibrated signaling models. The first counterfactual considers how equilibrium outcomes would change if the output of non-signaling sellers became less responsive to signaling output. For example, this might occur if there was a change in production technology.[29]

The results are reported in Table 3. Column (1) repeats the static incomplete information outcomes under the static calibration, and column (2) repeats the signaling outcome under the dynamic calibration. Columns (4) and (5) show the static Bayesian Nash and dynamic signaling outcomes, with percentage changes from columns (1) and (2) underneath the outcome values, when we set λ for the non-signaling sellers equal to 0.25, and we change the v_i supports of these sellers so that their observed average outputs would remain best responses. The green dashed lines in Figure 4 show how this causes the aggregate non-signaling seller output function to rotate around the pre-

[29]Sensitivity is high under the baseline specification, so we consider a reduction rather than an increase. Bornstein, Krusell, and Rebelo (2023) consider the effects on production and investment of a change in oil industry production technology in some countries, due to fracking, in a model where OPEC acts as a cartel.

Table 3: Counterfactual Effects under Alternative Specifications. Baseline columns report levels. Counterfactual columns report levels with the percentage difference from the corresponding baseline underneath. BN = static Bayesian Nash, DS = dynamic signaling, LF = leaders-followers. Midpoints of the supports of the $v_{i,t}$s for the signaling countries are as reported in Table 2, columns (2), (5) and (6). Signaling outcomes use equilibrium strategies calculated in period $T - 20$; outcomes are averages across 100,000 simulations. In the "Less Sensitive Supply" counterfactual, the λ for non-signaling sellers rises from 0.0007 to 0.2500 and their v_i supports are recalibrated so expected outputs still match average outputs in the detrended data (the implied U.S. midpoint support becomes -83.2); because that recalibration heavily affects non-signaling costs, non-signaling profits, total surplus and the all-seller inefficiency index are not reported (n.r.) for it. In the "Remove Iran and Russia" counterfactual those two countries produce zero and global demand is unaffected. Inefficiency indices are levels in every column, not changes.

	Baseline			Less Sensitive Non-Signaling Supply			Remove Iran and Russia		
	(1) Bayesian Nash	(2) Dynamic Signaling	(3) Leaders-Followers	(4) Bayesian Nash	(5) Dynamic Signaling	(6) Leaders-Followers	(7) Bayesian Nash	(8) Dynamic Signaling	(9) Leaders-Followers
Total Output MM bbl/month	2,409.2	2,410.5	2,409.1	2,409.2 (0.0%)	2,398.6 (-0.5%)	2,355.1 (-2.2%)	2,384.4 (-1.0%)	2,389.6 (-0.9%)	2,399.4 (-0.4%)
Signaling Output	1,079.7	1,079.5	1,079.2	1,079.7 (0.0%)	1,048.4 (-2.9%)	930.4 (-13.8%)	690.8 (-36.0%)	757.3 (-29.9%)	925.8 (-14.2%)
Saudi Arabia	362.8	362.5	362.1	362.8 (0.0%)	335.8 (-7.4%)	220.5 (-39.1%)	387.1 (+6.7%)	406.8 (+12.2%)	474.7 (+31.1%)
Russia	339.9	339.9	339.8	339.9 (0.0%)	316.5 (-6.9%)	215.6 (-36.6%)	0 (–)	0 (–)	0 (–)
China	138.3	138.3	137.7	138.3 (0.0%)	143.2 (+3.5%)	168.4 (+22.3%)	162.6 (+17.6%)	185.9 (+34.4%)	239.4 (+73.8%)
Iran	121.8	121.8	122.5	121.8 (0.0%)	128.2 (+5.2%)	163.7 (+33.6%)	0 (–)	0 (–)	0 (–)
Iraq	117.0	117.1	117.1	116.9 (0.0%)	124.8 (+6.6%)	162.3 (+38.6%)	141.2 (+20.6%)	164.6 (+40.6%)	211.7 (+80.8%)
Non-Signaling Output	1,329.5	1,331.0	1,329.9	1,329.5 (0.0%)	1,350.3 (+1.4%)	1,424.7 (+7.1%)	1,693.6 (+27.4%)	1,632.3 (+22.6%)	1,473.6 (+10.8%)
Price $ per barrel	19.81	19.76	19.81	19.81 (0.0%)	20.16 (+2.0%)	21.60 (+9.1%)	20.63 (+4.1%)	20.46 (+3.5%)	20.13 (+1.6%)
Consumer Surplus $m. per month	96,616	96,768	96,607	96,605 (-0.0%)	95,796 (-1.0%)	92,316 (-4.4%)	94,638 (-2.0%)	95,084 (-1.7%)	95,831 (-0.8%)
Signaling Profits $m. per month	9,921	3,865	891	9,922 (+0.0%)	4,092 (+5.9%)	2,238 (+151.2%)	6,603 (-33.4%)	3,016 (-22.0%)	909 (+2.0%)
Non-Signaling Profits $m. per month	7,136	7,088	7,150	n.r.	n.r.	n.r.	9,606 (+34.6%)	9,099 (+28.4%)	8,054 (+12.6%)
Total Surplus $m. per month	113,674	107,721	104,648	n.r.	n.r.	n.r.	110,847 (-2.5%)	107,199 (-0.5%)	104,794 (+0.1%)
Inefficiency Indices (levels), Allowing Reallocation Between									
Signaling Sellers Only	36.1%	7.0%	0.6%	36.1%	7.6%	1.9%	43.6%	8.4%	0.7%
All Sellers	62.9%	83.1%	98.7%	n.r.	n.r.	n.r.	76.6%	85.5%	98.0%

change equilibrium outcome. This exercise makes the marginal cost intercepts of the non-signaling sellers very low (for example, the mean intercept of the U.S. becomes -83.2), although their marginal costs at their observed level of outputs are similar to the baseline. We do not report non-signaling seller profits, total surplus or overall efficiency as the changes in the calibrated intercepts make these numbers uninformative.

By construction, the reported outputs predicted in column (4) are almost identical to those in column (1). In contrast, because the responsiveness of non-signaling output is a key driver of signaling incentives, outcomes in column (5) differ significantly from those in column (2): the lowest-cost signaling sellers produce less (e.g., Saudi Arabia's output falls 7.4%), and while some of the higher-cost signaling sellers raise production, which slightly lowers the efficiency of signaling output, signaling output falls overall, while non-signaling output increases and prices rise. The combined profits of the signaling sellers rise. Intuitively this reflects how signaling incentives, and the requirement to separate from higher cost types, can lead the signaling sellers to "overproduce", so that a reduction in signaling incentives can be profitable.

Our second counterfactual assumes that Russia and Iran, two countries that we have included in our set of "signaling sellers", are completely excluded from the supply side of the world oil market (for example, due to perfectly enforced sanctions). For simplicity, we assume that global oil demand is unaffected by this exclusion.[30] λ and the cost intercepts of all sellers are the same in columns (7) and (8) as in columns (1) and (2).

This counterfactual changes outcomes in both the static and the dynamic equilibria, with total output falling by around 1%, reflecting the responsiveness of non-signaling output, which rises significantly. However, there are significant differences in quantitative responses at the seller level: in the signaling equilibrium, the remaining signaling countries increase their outputs roughly twice

[30] According to Energy Institute (2024), these countries account for around 5% of global consumption, and 16% of global production, although they make up just under 20% of the average total output in AZ's twenty-country data.

as much as in the static equilibrium, so that the reduction in total surplus is smaller.

4.6 Comparison with a "Leaders-Followers" Model

One intuitive interpretation of signaling when there are two types of sellers is that the ability to signal allows the signaling sellers to gain a form of first-mover advantage over the non-signaling sellers, as they are able to give an indication of how much they will produce. To analyze whether this intuition captures what happens in the signaling equilibrium, we also calibrate and solve a model in which the five signaling sellers act as "quantity leaders". We assume that their current costs are public, and, in each period, they choose quantities simultaneously before the non-signaling sellers, who are "followers". The followers observe the leaders' commitments and then play a static Bayesian Nash equilibrium. We call this model the "leaders-followers" model (a more complete discussion of the model is included in Appendix B).

A priori, the different information and commitment assumptions could push in different directions. On the one hand, leaders can commit to whatever outputs they choose, whereas, in the signaling model, output strategies have to take a specific form to satisfy the incentive compatibility constraints. On the other hand, in our signaling model, signaling sellers can affect each other's future output choices, whereas our leaders-followers model assumption that the leaders' costs are publicly observed rules this out.

This model comparison is particularly relevant in the context of the oil industry, where one or more of the largest oil-producing countries have often been modeled as leaders (Hnyilicza and Pindyck (1976), Gilbert (1978), Crémer and Salehi-Isfahani (1991), Behar and Ritz (2017), Golombek, Irarrazabal, and Ma (2018)). In leadership models, the mechanism that allows a leader to move first, when in reality all sellers produce simultaneously, is rarely micro-founded, so it is useful to know whether signaling based on serially correlated private information generates com-

parative statics and quantitative effects similar to those of a model where leadership is assumed.[31]

Column (4) of Table 2 shows the predicted outcomes under leaders-followers play, using cost parameters calibrated assuming static simultaneous-move behavior. Relative to the static outcomes, the combined outputs of the signaling and the non-signaling sellers change in the same directions as when we assume dynamic signaling but the changes are larger. For example, signaling raises signaling seller output by 76%, whereas leadership raises the outputs of the same sellers by 102%. The effects of leadership on prices and surplus are also larger. However, signaling generates more output volatility. The standard deviation of total output is 97.0 under signaling and 10.9 under leaders–followers, while the standard deviations of cross-period changes in signaling-seller output are 108.7 and 6.6, respectively. There are also directional differences in how country outputs are predicted to change. For example, Iranian output is predicted to increase 24% with signaling but to fall 22% in the leaders-followers model, reflecting how the lowest-cost sellers, such as Saudi Arabia, increase output more dramatically when they act as leaders.

Column (6) of Table 2 shows the calibrated cost parameters, and the equilibrium outcomes, when we calibrate the signaling seller cost supports so that the leaders-followers model matches observed signaling outputs. Consistent with the lowest-cost sellers producing disproportionately more in the leaders-followers model, v-intercepts for Saudi Arabia and Russia increase to almost the same level as those of the signaling sellers that have lower output, and, consequently, signaling profits are much reduced. Interestingly, whereas using model-calibrated parameters significantly reduced the standard deviations of output and changes of signaling output in the signaling model, in the leaders-followers model these changes are in the opposite direction.

Columns (6) and (9) of Table 3 present the counterfactual results using the calibrated leaders-followers model. Under both counterfactuals, output and price changes have the same directions

[31] More broadly, one might view dynamic signaling models as potentially providing a novel micro-foundation for some of the qualitative patterns associated with dominant firm-fringe (e.g., Tasnádi (2010)) or Stackelberg models (e.g., van Damme and Hurkens (1999)).

across the three models (where the static model predicts any change), but different magnitudes. For example, in the Iran/Russia removal counterfactual, Iraqi production is predicted to increase 21% in the Bayesian Nash model, 41% in the signaling model and 81% in the leaders-followers model. In the less-sensitive-supply counterfactual, Iraqi supply is unchanged in the Bayesian Nash model, while rising 7% in the dynamic signaling model and 39% in the leaders-followers model. The changes in profits and surplus can, however, differ in direction: for example, in the Iran/Russia removal counterfactual, signaling profits rise using the calibrated leaders-followers model, but they fall using the calibrated dynamic signaling model. Therefore, while a simpler-to-compute leadership model may be able to provide some useful intuition as to how some outcomes may change in a counterfactual, such a model may predict different directional outcome changes and different magnitudes from a model where sellers have to use prior output choices to signal their future intentions.

5 Conclusion

We develop a quantitative dynamic model in which either all sellers, or a subset of them, use their output choices to signal private information about persistent marginal costs in order to affect rivals' future output choices. Both private information and the persistence of unobserved costs are plausible in many settings. We ask whether incorporating the resulting signaling incentives may yield different predictions from models that impose complete information and static behavior.

We have shown that economically meaningful differences in expected output, consumer surplus and efficiency can exist even when there are quite a few sellers and we assume model parameters for which differences between equilibrium static complete information and static incomplete information outcomes would be very small. We have also shown that dynamic signaling can matter for counterfactual predictions. We believe that the result that these differences can be substantial

in markets that are not concentrated is important, because it conflicts with what we think are common intuitions about when strategic investment incentives are likely to be important, and it suggests that the effects we identify may matter in a broad set of industries.

Once one allows for repeated play, it is also natural to ask whether sellers might be able to support some form of tacit collusion.[32] It is important to emphasize that, in a quantity-setting environment, a dynamic signaling model will tend to predict higher output, and greater consumer surplus, compared to static outcomes, whereas collusive outcomes would tend to lead to lower output and lower consumer surplus. We have also shown that signaling effects can remain substantial with many sellers and a discount factor of 0.8, an environment in which conventional repeated-game incentives to collude would be weaker. As signaling can substantially lower industry profits through costly output distortions, our model provides one explanation for why firms would want to share private cost information, if they are able to do so credibly, even if this information-sharing cannot result in a collusive outcome with close to joint-profit-maximizing outcomes (Kühn and Vives (1995) and Møllgaard and Overgaard (2008) provide policy-focused discussions of the trade-offs involved in allowing information exchanges).

Our results raise the practical question of how empirical researchers should proceed when examining repeated quantity choices. If it is plausible that marginal costs may be private information, a natural first step is to examine whether seller outputs, and the marginal costs implied by static models, are serially correlated. If persistence is limited, then signaling incentives may have limited effects unless cost shocks are very dispersed. If a high degree of serial correlation is present, a calibration of a signaling model might be used to assess its potential importance, or researchers should at least state that possibly material signaling incentives are being ignored.

Our results also highlight the importance of developing rigorous tests for whether sellers have

[32]While most of the literature on collusion assumes that sellers' costs are observable, Athey and Bagwell (2001), Athey, Bagwell, and Sanchirico (2004) and Athey and Bagwell (2008) consider the possibility that the collusive scheme will have to allow for sellers to have persistent private information cost shocks.

private information on some components of their marginal costs, whether these components are serially correlated, and, if so, whether sellers use the types of fully separating signaling strategies that are assumed here. While several researchers have developed tests for the information structure of static discrete choice games between firms (most notably Grieco (2014)), different approaches are required when choices are continuous. A test procedure would also need to allow for how real-world factors not considered here, such as adjustment costs, supply disruptions and heterogeneity in the shape of sellers' cost functions, may also contribute to either persistence or volatility in observed output choices.[33]

[33]Recent work by Hsu, Li, Liu, and Takahashi (2026) provides an approach to testing the monotonicity of strategies in a game with private information, but it does not provide a framework for testing the information structure more generally.

References

Almoguera, P. A., C. C. Douglas, and A. M. Herrera (2011): "Testing for the Cartel in OPEC: Non-Cooperative Collusion or Just Non-Cooperative?," *Oxford Review of Economic Policy*, 27(1), 144–168.

Armantier, O., and O. Richard (2003): "Exchanges of Cost Information in the Airline Industry," *RAND Journal of Economics*, 34(3), 461–477.

Aryal, G., and F. Zincenko (2024): "Empirical Framework for Cournot Oligopoly with Private Information," *RAND Journal of Economics*, 55(3), 375–402.

Asker, J., A. Collard-Wexler, and J. De Loecker (2019): "(Mis)Allocation, Market Power, and Global Oil Extraction," *American Economic Review*, 109(4), 1568–1615.

Asker, J., C. Fershtman, J. Jeon, and A. Pakes (2020): "A Computational Framework for Analyzing Dynamic Procurement Auctions: The Market Impact of Information Sharing," *RAND Journal of Economics*, 51(3), 805–839.

Athey, S., and K. Bagwell (2001): "Optimal Collusion with Private Information," *RAND Journal of Economics*, 32(3), 428–465.

——— (2008): "Collusion with Persistent Cost Shocks," *Econometrica*, 76(3), 493–540.

Athey, S., K. Bagwell, and C. Sanchirico (2004): "Collusion and Price Rigidity," *Review of Economic Studies*, 71(2), 317–349.

Behar, A., and R. A. Ritz (2017): "OPEC vs US Shale: Analyzing the Shift to a Market-Share Strategy," *Energy Economics*, 63, 185–198.

Bonatti, A., G. Cisternas, and J. Toikka (2017): "Dynamic Oligopoly with Incomplete Information," *Review of Economic Studies*, 84(2), 503–546.

Bornstein, G., P. Krusell, and S. Rebelo (2023): "A World Equilibrium Model of the Oil Market," *Review of Economic Studies*, 90(1), 132–164.

Crémer, J., and D. Salehi-Isfahani (1991): *Models of the Oil Market*. Harwood Academic Publishers, New York.

Doraszelski, U., G. Lewis, and A. Pakes (2018): "Just Starting Out: Learning and Equilibrium in a New Market," *American Economic Review*, 108(3), 565–615.

Duarte, M., L. Magnolfi, D. Quint, M. Sølvsten, and C. Sullivan (2026): "Conduct and Scale Economies: Evaluating Tariffs in the US Automobile Market," Discussion paper, University of North Carolina Chapel Hill.

Energy Institute (2024): *Statistical Review of World Energy*. Energy Institute, London, 73rd edn.

Fried, D. (1984): "Incentives for Information Production and Disclosure in a Duopolistic Environment," *Quarterly Journal of Economics*, 99(2), 367–381.

Fudenberg, D., and D. K. Levine (1993): "Steady State Learning and Nash Equilibrium," *Econometrica*, 61(3), 547–573.

Gal-Or, E. (1985): "Information Sharing in Oligopoly," *Econometrica*, 53(2), 329–343.

——— (1986): "Information Transmission—Cournot and Bertrand Equilibria," *Review of Economic Studies*, 53(1), 85–92.

Gardete, P. M. (2016): "Competing under Asymmetric Information: The Case of Dynamic Random Access Memory Manufacturing," *Management Science*, 62(11), 3291–3309.

Gilbert, R. J. (1978): "Dominant Firm Pricing Policy in a Market for an Exhaustible Resource," *Bell Journal of Economics*, 9(2), 385–395.

Golombek, R., A. A. Irarrazabal, and L. Ma (2018): "OPEC's Market Power: An Empirical Dominant Firm Model for the Oil Market," *Energy Economics*, 70, 98–115.

Grieco, P. L. E. (2014): "Discrete Games with Flexible Information Structures: An Application to Local Grocery Markets," *RAND Journal of Economics*, 45(2), 303–340.

Griffin, J. M. (1985): "OPEC Behavior: A Test of Alternative Hypotheses," *American Economic Review*, 75(5), 954–963.

Hnyilicza, E., and R. S. Pindyck (1976): "Pricing Policies for a Two-Part Exhaustible Resource Cartel: The Case of OPEC," *European Economic Review*, 8(2), 139–154.

Hsu, Y.-C., T. Li, C.-A. Liu, and H. Takahashi (2026): "Testing for Monotone Equilibrium Strategies in Games of Incomplete Information," Discussion paper, Vanderbilt University.

Hurkens, S. (2014): "Bayesian Nash Equilibrium in "Linear" Cournot Models with Private Information About Costs," *International Journal of Economic Theory*, 10(2), 203–217.

Hyndman, K. (2008): "Disagreement in Bargaining: An Empirical Analysis of OPEC," *International Journal of Industrial Organization*, 26(3), 811–828.

Judd, K. L. (1998): *Numerical Methods in Economics*. MIT Press.

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica*, 61(5), 1019–1045.

Kühn, K.-U., and X. Vives (1995): "Information Exchanges among Firms and their Impact on Competition," Discussion paper, European Commission.

Leccese, M., A. Sweeting, and X. Tao (2025): "Should We Expect Merger Synergies to Be Passed Through to Consumers?," *Journal of Industrial Economics*, 73(1), 1–30.

Lee, Y., and T. Mukoyama (2015): "Productivity and Employment Dynamics of US Manufacturing Plants," *Economics Letters*, 136, 190–193.

Mailath, G. J. (1988): "An Abstract Two-Period Game with Simultaneous Signaling: Existence of Separating Equilibria," *Journal of Economic Theory*, 46(2), 373–394.

Mailath, G. J. (1989): "Simultaneous Signaling in an Oligopoly Model," *Quarterly Journal of Economics*, 104(2), 417–427.

Mester, L. J. (1992): "Perpetual Signalling with Imperfectly Correlated Costs," *RAND Journal of Economics*, 23(4), 548–563.

Milgrom, P., and J. Roberts (1990): "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58(6), 1255–1277.

Møllgaard, H. P., and P. B. Overgaard (2008): "Information Exchange, Market Transparency, and Dynamic Oligopoly," in *Issues in Competition Law and Policy*, ed. by W. D. Collins, vol. 2, pp. 1241–1268. American Bar Association, Chicago, IL.

Olley, G. S., and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64(6), 1263–1297.

Roddie, C. (2012): "Signaling and Reputation in Repeated Games, II: Stackelberg Limit Properties," Discussion paper, University of Cambridge.

Shapiro, C. (1986): "Exchange of Cost Information in Oligopoly," *Review of Economic Studies*, 53(3), 433–446.

Smith, J. L. (2005): "Inscrutable OPEC? Behavioral Tests of the Cartel Hypothesis," *Energy Journal*, 26(1), 51–82.

Sweeting, A., X. Tao, and X. Yao (2024): "Dynamic Oligopoly Pricing with Asymmetric Information: Implications for Horizontal Mergers," *American Economic Journal: Microeconomics*, 16(3), 345–373.

Syverson, C. (2004): "Market Structure and Productivity: A Concrete Example," *Journal of Political Economy*, 112(6), 1181–1222.

Tasnádi, A. (2010): "Quantity-Setting Games with a Dominant Firm," *Journal of Economics*, 99(3), 251–266.

Tirole, J. (1988): *The Theory of Industrial Organization*. MIT Press.

Toxvaerd, F. (2008): "Strategic Merger Waves: A Theory of Musical Chairs," *Journal of Economic Theory*, 140(1), 1–26.

——— (2017): "Dynamic Limit Pricing," *RAND Journal of Economics*, 48(1), 281–306.

van Damme, E., and S. Hurkens (1999): "Endogenous Stackelberg Leadership," *Games and Economic Behavior*, 28(1), 105–129.

Vives, X. (1984): "Duopoly Information Equilibrium: Cournot and Bertrand," *Journal of Economic Theory*, 34(1), 71–94.

——— (2002): "Private Information, Strategic Behavior, and Efficiency in Cournot Markets," *RAND Journal of Economics*, 33(3), 361–376.

Online Appendices for

"Producing to Signal: Private Information in Dynamic Quantity Competition"

by Sweeting, Tao and Wang

A Proofs

A.1 Proof that Assumption 1 Implies First-Order Stochastic Dominance

Lemma A.1 (First-Order Stochastic Dominance) *Under Assumption 1, the conditional distribution of $v_{i,t}$ given $v_{i,t-1}$ is increasing in $v_{i,t-1}$ in the sense of first-order stochastic dominance. That is, if $b' > b$, then*

$$G_i(x \mid b') \leq G_i(x \mid b) \qquad \text{for all } x \in [\underline{v}_i, \overline{v}_i], \tag{9}$$

where

$$G_i(x \mid b) = \int_{\underline{v}_i}^{x} g_i(z \mid b)\, dz. \tag{10}$$

The inequality is strict for every $x \in (\underline{v}_i, \overline{v}_i)$.

Proof. Fix a conditioning value $b \in [\underline{v}_i, \overline{v}_i]$. By Assumption 1, there exists a value $v'(b)$ such that

$$\frac{\partial g_i(z \mid b)}{\partial b} < 0 \quad \text{for } z < v'(b), \qquad \frac{\partial g_i(z \mid b)}{\partial b} > 0 \quad \text{for } z > v'(b). \tag{11}$$

Consider the derivative of the conditional distribution function with respect to b (by part 1 of Assumption 1 we can differentiate under the integral sign):

$$\frac{\partial G_i(x \mid b)}{\partial b} = \int_{\underline{v}_i}^{x} \frac{\partial g_i(z \mid b)}{\partial b}\, dz. \tag{12}$$

If $x \leq v'(b)$, then

$$\frac{\partial G_i(x \mid b)}{\partial b} = \int_{\underline{v}_i}^{x} \frac{\partial g_i(z \mid b)}{\partial b}\, dz < 0 \tag{13}$$

for every $x \in (\underline{v}_i, v'(b)]$.

If $x > v'(b)$, use the fact that $g_i(\cdot \mid b)$ is a density for every b. Differentiating

$$\int_{\underline{v}_i}^{\overline{v}_i} g_i(z \mid b)\, dz = 1 \tag{14}$$

with respect to b gives

$$\int_{\underline{v}_i}^{\overline{v}_i} \frac{\partial g_i(z \mid b)}{\partial b}\, dz = 0. \tag{15}$$

It follows that

$$\frac{\partial G_i(x \mid b)}{\partial b} = -\int_{x}^{\overline{v}_i} \frac{\partial g_i(z \mid b)}{\partial b}\, dz. \tag{16}$$

Because $x > v'(b)$, the integrand on the right-hand side is positive, and hence

$$\frac{\partial G_i(x \mid b)}{\partial b} < 0 \tag{17}$$

for every $x \in (v'(b), \overline{v}_i)$.

Therefore,

$$\frac{\partial G_i(x \mid b)}{\partial b} < 0 \qquad \text{for all } x \in (\underline{v}_i, \overline{v}_i). \tag{18}$$

Consequently, for any $b' > b$,

$$G_i(x \mid b') - G_i(x \mid b) = \int_{b}^{b'} \frac{\partial G_i(x \mid s)}{\partial s}\, ds \leq 0, \tag{19}$$

with strict inequality for every interior x. Thus the conditional distribution given $b' > b$ is first-order stochastically larger than the conditional distribution given b. ■

A.2 Statement and Proof of Conditions 1 and 2

Conditions 1-4 are sufficient to characterize sellers' signaling best response functions using a differential equation and a boundary condition. As explained below, Conditions 3 and 4 have to be verified numerically when solving the game. However, with linear demand and quadratic costs, Conditions 1 and 2, which are conditions on static profit functions, can be shown to hold generally.

Condition 1 (Shape of $\Pi^{i,t}$ with Respect to $q_{i,t}$) *For any* $\left(v_{i,t}, \widehat{v_{i,t}^{-i}}\right)$, $\Pi^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ *has a unique optimum in* $q_{i,t}$*, and, for any* $q_{i,t}$ *where* $\Pi_{33}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) > 0$*, there is some* $k > 0$ *such that* $\left|\Pi_{3}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)\right| > k$.

Proof.

$$\Pi_{3}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \frac{\partial \pi\left(v_{i,t}, q_{i,t}, \boldsymbol{q}_{-i,t}(\boldsymbol{v}_{-i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}})\right)}{\partial q_{i,t}} g_{-i}\left(\boldsymbol{v}_{-i,t} \mid \boldsymbol{v}_{-i,t-1}\right) d\boldsymbol{v}_{-i,t}. \tag{20}$$

With linear demand, and assuming prices can go negative,

$$\Pi_{3}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = u - 2\beta q_{i,t} - \beta \mathbb{E}\left[Q_{-i,t} | \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right] - v_{i,t} - \lambda q_{i,t}. \tag{21}$$

Note that $\Pi_{3}^{i,t}$ only depends on the shape of the static profit function as $q_{i,t}$ does not enter the continuation value function. Assuming sellers want to produce positive quantities, i's unique optimal output is

$$q_{i,t}^{*} = \frac{u - \beta \mathbb{E} Q_{-i,t} - v_{i,t}}{2\beta + \lambda} \tag{22}$$

$\Pi_{33}^{i,t}(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}) = -(2\beta + \lambda) < 0$. Because $\Pi_{33}^{i,t} < 0$ everywhere, the second part of the condition is trivially satisfied. ■

Condition 2 (Type Monotonicity) $\Pi_{13}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) \neq 0$ *for all* $\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}\right)$ *given* $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$.

Proof. Continuing from the proof of Condition 1

$$\Pi_3^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = u - 2\beta q_{i,t} - \beta\mathbb{E}\left[Q_{-i,t}|v_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right] - v_{i,t} - \lambda q_{i,t}. \tag{23}$$

Therefore, $\Pi_{13}^{i,t}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = -1$, and the condition is satisfied. ■

A.3 Proof of the Existence of a Unique Fully Separating Equilibrium in Period $t < T$

We begin by proving a useful lemma.

Lemma A.2 (i's Separating Quantity Function in Period $t < T$) *Consider a given* $\widehat{v_{i,t-1}^{-i}}$ *and* $\boldsymbol{v}_{-i,t-1}$ *(grouping together the previous period types of other sellers). Hold fixed i's value in $t+1$ states (which will not depend on the period t strategies of any seller). Assume that i's signaling payoff function satisfies Conditions 1-4, and that, for all values of $v_{i,t}$, i wants to produce strictly positive output. Then the solution to differential equation (5) and boundary condition (6) can be written as $q_{i,t}^*(\overline{v_i}, \mathbb{E}Q_{-i,t}) + h_{i,t}(v_{i,t})$ where $q_{i,t}^*(\overline{v_i}, \mathbb{E}Q_{-i,t})$ is the static best response of i when it has the highest cost type $\overline{v_i}$ and other sellers have expected combined output $\mathbb{E}Q_{-i,t}$, and $h_{i,t}(v_{i,t})$ is a function that does not depend on the period-t strategies of other sellers (although it may depend on beliefs about other sellers' lagged types).*

Proof. i's signaling payoff function is defined in equation (4).

$$\int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \pi\Big(v_{i,t},\, q_{i,t},\, \boldsymbol{q}_{-i,t}\Big(\boldsymbol{v}_{-i,t},\, \boldsymbol{v}_{-i,t-1},\, \widehat{v}_{i,t-1}^{-i}\Big)\Big) g_{-i}(\boldsymbol{v}_{-i,t}\,|\,\boldsymbol{v}_{-i,t-1})\, d\boldsymbol{v}_{-i,t} = \Big(u - \beta q_{i,t} - \beta\,\mathbb{E}\Big[Q_{-i,t}\,\Big|\,\boldsymbol{v}_{-i,t-1}, \widehat{v}_{i,t-1}^{-i}\Big] - v_{i,t}\Big)\, q_{i,t} - \frac{\lambda}{2}\, q_{i,t}^2 \tag{24}$$

where $\mathbb{E}Q_{-i,t}$ is the expected combined output of other sellers, which, because demand is linear, is the only function of rivals' output that matters for i's expected flow profit. This function does not

depend on $\widehat{v_{i,t}^{-i}}$.

$$\int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \delta V_{i,t+1}\left(v_{i,t}, \widehat{v_{i,t}^{-i}}, \boldsymbol{v}_{-i,t}\right) g_{-i}\left(\boldsymbol{v}_{-i,t} | \boldsymbol{v}_{-i,t-1}\right) d\boldsymbol{v}_{-i,t} \tag{25}$$

does not depend on the period t quantity function of any seller, as long as those functions are separating. As Conditions 1-4 hold, i's unique fully separating best response function is determined by the differential equation,

$$\frac{dq_{i,t}^{*}(v_{i,t})}{dv_{i,t}} = -\left.\frac{\Pi_2^{i,t}\left(v_{i,t},\, \widehat{v}_{i,t}^{-i},\, q_{i,t},\, \mathbf{v}_{-i,t-1},\, \widehat{v}_{i,t-1}^{-i}\right)}{\Pi_3^{i,t}\left(v_{i,t},\, \widehat{v}_{i,t}^{-i},\, q_{i,t},\, \mathbf{v}_{-i,t-1},\, \widehat{v}_{i,t-1}^{-i}\right)}\right|_{\widehat{v}_{i,t}^{-i}=v_{i,t},\; q_{i,t}=q_{i,t}^{*}(v_{i,t})} \tag{26}$$

which determines the slope of the function, and the boundary condition,

$$\Pi_3^{i,t}\left(\overline{v_i}, \overline{v_i}, q_{i,t}^{*}\left(\overline{v}_i\right), \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right) = 0, \tag{27}$$

which determines the equilibrium quantity choice value when $v_{i,t} = \overline{v_i}$. The boundary condition implies that $q_{i,t}^{*}(\overline{v_i}, \mathbb{E}Q_{-i,t}) = \frac{u-\beta\mathbb{E}Q_{-i,t}-\overline{v_i}}{2\beta+\lambda}$.

The numerator does not depend on the period-t quantity functions. The denominator equals $u - 2\beta q_{i,t} - \beta\mathbb{E}Q_{-i,t} - v_{i,t} - \lambda q_{i,t}$, and using the boundary condition equation, this can be re-written as $(2\beta+\lambda)(q_{i,t}^{*}(\overline{v_i}, \mathbb{E}Q_{-i,t}) - q_{i,t}) + (\overline{v_i} - v_{i,t})$. Now write $q_{i,t} = q_{i,t}^{*}(\overline{v_i}, \mathbb{E}Q_{-i,t}) + h_{i,t}(v_{i,t})$. Because $q_{i,t}^{*}(\overline{v_i}, \mathbb{E}Q_{-i,t})$ does not vary with $v_{i,t}$, $\frac{dq_{i,t}^{*}(v_{i,t})}{dv_{i,t}} = \frac{dh_{i,t}(v_{i,t})}{dv_{i,t}}$, and the denominator becomes $(\overline{v_i}-v_{i,t})-(2\beta+\lambda)h_{i,t}(v_{i,t})$. Therefore, $h_{i,t}(v_{i,t})$ solves an initial value problem, with initial condition $h_{i,t}(\overline{v_i}) = 0$, in which neither the numerator nor the denominator depends on other sellers' period-t quantity functions. Therefore, i's quantity function can be written as $q_{i,t}^{*}(\overline{v_i}, \mathbb{E}Q_{-i,t}) + h_{i,t}(v_{i,t})$, where $h_{i,t}(v_{i,t})$ is a function that does not depend on the quantity functions of other sellers. ■

This result implies that a change in either the level or the slope of the period t quantity functions of other sellers will simply translate i's best response quantity function up or down as a function

of how the expected output of the other sellers has changed. Of course, the same argument applies for any other seller j. The following proposition provides our equilibrium uniqueness result.

Proposition A.1 (Unique Fully Separating Equilibrium in Period $t < T$) *Consider a given set of beliefs about all sellers' previous period cost types, given the $t+1$ value functions of each seller. Assume each seller's signaling payoff function satisfies Conditions 1-4, and that, for all type realizations, all sellers want to produce strictly positive outputs. Then there will be a unique fully separating equilibrium in period t.*

Proof. Lemma A.2 implies that the separating best response functions of each seller can be written as $q^*_{i,t}(\overline{v_i}, \mathbb{E}Q_{-i,t}) + h_{i,t}(v_{i,t})$, where the $h_{i,t}(v_{i,t})$s are functions that do not depend on the functions of the other sellers, so that interactions between sellers' strategies are captured only through the boundary intercepts $q^*_{i,t}(\overline{v_i}, \mathbb{E}Q_{-i,t})$. We show that there will be a unique vector of equilibrium intercepts, given the $h_{i,t}(v_{i,t})$s, by solving for the intercepts explicitly.

Define $\Delta_{j,t} = \mathbb{E}[h_{j,t}(v_{j,t})|v_{j,t-1}]$, i.e., the expected increment above j's best response given sellers' beliefs about j's previous cost type. The first-order condition determining $q^*_{i,t}(\overline{v_i}, \mathbb{E}Q_{-i,t})$ is

$$u - (2\beta + \lambda)q^*_{i,t} - \beta \sum_{j \neq i} (q^*_{j,t} + \Delta_{j,t}) - \overline{v_i} = 0. \tag{28}$$

Writing $y^*_{i,t} = q^*_{i,t} + \Delta_{i,t}$ for i's expected output, these N equations form the linear system $(\lambda + \beta)y^*_{i,t} + \beta \sum_{j=1}^N y^*_{j,t} = u - \overline{v_i} + (\lambda + 2\beta)\Delta_{i,t}$, $i = 1, \ldots, N$, whose coefficient matrix, $(\lambda+\beta)I + \beta \mathbf{1}\mathbf{1}'$, has eigenvalues $\lambda+\beta$ and $\lambda+(N+1)\beta$, both strictly positive, and is therefore nonsingular. Solving these equations for N sellers gives unique solutions for $q^*_{i,t}(\overline{v_i}, \mathbb{E}Q_{-i,t}) = y^*_{i,t} - \Delta_{i,t}$ where

$$y^*_{i,t} = \frac{u - \overline{v}_i + (\lambda + 2\beta)\Delta_{i,t}}{\lambda + \beta} - \frac{\beta}{\lambda + \beta} \frac{Nu - \sum_{j=1}^N \overline{v}_j + (\lambda + 2\beta) \sum_{j=1}^N \Delta_{j,t}}{\lambda + (N+1)\beta}. \tag{29}$$

■

Implication for a Unique Equilibrium in Period t. Proposition A.1 shows the existence of a unique equilibrium for a given set of beliefs about sellers' previous $(t-1)$ period types and given next period values. Uniqueness of a period t separating equilibrium, given next period values, comes from applying the proof for *all* possible sets of equilibrium-path beliefs about previous period types.

Implication for a Unique Equilibrium in the Finite-Horizon Game. Uniqueness of an equilibrium in a finite-horizon game comes from applying the same result recursively. We know that the period T strategies are unique, given beliefs about types in $T-1$, which implies unique beginning-of-period-T values. We can then apply Proposition A.1, given on-path beliefs, to show uniqueness of the equilibrium in $T-1$. This will imply uniqueness of beginning of period $T-1$ values. We can then apply the same logic for $T-2$. Note that, at each stage, we are requiring that Conditions 1-4 hold for each seller, for each set of prior beliefs, and that all sellers want to produce positive amounts for all cost realizations. If a positive-measure set of types would choose zero, the fully separating equilibrium characterized here will fail because those types pool at the output constraint.

A.4 Belief Monotonicity (Condition 3) Given $t+1$ Strategies

We show that if every seller's period-$(t+1)$ strategy is the separating solution of differential equation (5) and boundary condition (6), with strictly decreasing quantity schedules, and all sellers produce strictly positive outputs in periods t and $t+1$, then seller i's period-t signaling payoff function satisfies Condition 3 whenever rivals' expected period-$(t+1)$ output is increasing in the type attributed to i. This property can be shown analytically in period $t = T-1$.

Fix the conditioning states $\boldsymbol{v}_{-i,t-1}$ and $\widehat{v}^{-i}_{i,t-1}$, and write $b \equiv \widehat{v}^{-i}_{i,t}$ for the period-t type that rivals attribute to i. Recall from the definition of the signaling payoff function that

$$\Pi^{i,t}\big(v_{i,t}, b, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v}^{-i}_{i,t-1}\big) = \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \pi\big(v_{i,t}, q_{i,t}, \boldsymbol{q}_{-i,t}(\boldsymbol{v}_{-i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v}^{-i}_{i,t-1})\big)\, g_{-i}(\boldsymbol{v}_{-i,t} \mid \boldsymbol{v}_{-i,t-1})\, d\boldsymbol{v}_{-i,t} + \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \delta\, V_{i,t+1}\big(v_{i,t}, b, \boldsymbol{v}_{-i,t}\big)\, g_{-i}(\boldsymbol{v}_{-i,t} \mid \boldsymbol{v}_{-i,t-1})\, d\boldsymbol{v}_{-i,t}. \tag{30}$$

The first integral is not affected by b. Hence b enters $\Pi^{i,t}$ only through the continuation value, and

$$\Pi_2^{i,t} \;=\; \delta \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} \frac{\partial V_{i,t+1}\big(v_{i,t}, b, \boldsymbol{v}_{-i,t}\big)}{\partial b}\, g_{-i}(\boldsymbol{v}_{-i,t} \mid \boldsymbol{v}_{-i,t-1})\, d\boldsymbol{v}_{-i,t}, \tag{31}$$

which does not depend on $q_{i,t}$.

Throughout, fix a period-$(t+1)$ public state $\big(b, \boldsymbol{v}_{-i,t}\big)$. By Lemma A.2 and Proposition A.1, seller j's period-$(t+1)$ equilibrium schedule can be written

$$q_{j,t+1}\big(v_{j,t+1}, b, \boldsymbol{v}_{-i,t}\big) \;=\; B_j\big(b, \boldsymbol{v}_{-i,t}\big) + h_{j,t+1}\big(v_{j,t+1}, b, \boldsymbol{v}_{-i,t}\big), \tag{32}$$

where B_j is the boundary quantity chosen by the highest-cost type $\overline{v}_j$, $h_{j,t+1}(\overline{v}_j, b, \boldsymbol{v}_{-i,t}) = 0$, and the increment $h_{j,t+1}$ does not depend on the period-$(t+1)$ strategies of other sellers (although it can depend on beliefs about lagged types). We define three objects:

$$R\big(b, \boldsymbol{v}_{-i,t}\big) \equiv \mathbb{E}\big[\, Q_{-i,t+1} \,\big|\, b, \boldsymbol{v}_{-i,t} \,\big], \qquad R_b \equiv \frac{\partial R}{\partial b}, \tag{33}$$

the expected combined period-$(t+1)$ output of i's rivals in the unique separating period-$(t+1)$

equilibrium, and its total sensitivity to the attributed type;

$$q^{BR}_{i,t+1}\big(v; b, \boldsymbol{v}_{-i,t}\big) \;\equiv\; \frac{u - \beta R\big(b, \boldsymbol{v}_{-i,t}\big) - v}{\lambda + 2\beta}, \tag{34}$$

i's static best-response quantity to R when its realized cost intercept is v; and, for every seller j,

$$\Delta_j\big(b, \boldsymbol{v}_{-i,t}\big) \;\equiv\; \mathbb{E}\big[\, h_{j,t+1}(v_{j,t+1}, b, \boldsymbol{v}_{-i,t}) \,\big|\, v_{j,t} \,\big], \tag{35}$$

j's expected increment, where the conditioning type is b for $j = i$ and j's true period-t type for $j \neq i$.

We make three observations that help to understand the proof.

(a) The channels through which b affects outcomes. In the separating equilibrium, i's period-$(t+1)$ quantity choice reveals its true type, so the period-$(t+2)$ continuation entering $V_{i,t+1}$, namely $\mathbb{E}\big[V_{i,t+2}(v_{i,t+1}, v_{i,t+1}, \boldsymbol{v}_{-i,t+1})\big]$, does not involve b: because strategies are Markov, rivals' period-$(t+2)$ strategies depend only on period-$(t+1)$ types and beliefs about them, all of which are revealed by period-$(t+1)$ play. Moreover, i's own increment $h_{i,t+1}$ does not depend on b: by Lemma A.2 it solves an initial value problem whose numerator integrates i's period-$(t+2)$ continuation derivative over rivals' transitions from their *true* period-t types $\boldsymbol{v}_{-i,t}$, and whose denominator, after substituting the boundary condition, equals $(\overline{v}_i - v_{i,t+1}) - (\lambda + 2\beta)\, h_{i,t+1}$; the attributed type b appears in neither, because b is payoff-irrelevant to i's own future once the period-$(t+1)$ signal supersedes it. However, this is not necessarily true of rivals' increments when $t+1 < T$; see the proof of part (ii) of the Lemma below. Therefore b enters $V_{i,t+1}$ only through period-$(t+1)$ flow profits, and, for these profits, only through two equilibrium objects: the level of expected rival

output, $R(b, \boldsymbol{v}_{-i,t})$, and the intercept of i's own schedule, $B_i(b, \boldsymbol{v}_{-i,t})$.[34]

(b) The boundary condition ties the schedule shift to R_b. Boundary condition (6), applied in period $t+1$, states that the highest-cost type plays its static best response: $u - \beta R - (\lambda + 2\beta)B_i - \overline{v}_i = 0$, i.e., $B_i = q^{BR}_{i,t+1}(\overline{v}_i)$. Differentiating with respect to b,

$$\frac{\partial B_i}{\partial b} \;=\; -\,\frac{\beta}{\lambda + 2\beta}\, R_b\big(b, \boldsymbol{v}_{-i,t}\big). \tag{36}$$

(c) Marginal flow profit measures over-production. With linear demand[35],

$$\Pi_3^{i,t+1}\big(v, q\big) \;=\; u - \beta R - (\lambda + 2\beta)\, q - v \;=\; (\lambda + 2\beta)\big(\, q^{BR}_{i,t+1}(v) - q\,\big). \tag{37}$$

Lemma A.3 *Suppose every seller's period-$(t+1)$ strategy is the separating solution of (5) and (6), with strictly decreasing quantity schedules, that the increments $h_{k,t+1}$, $k \neq i$, are differentiable in b, and that all sellers produce strictly positive outputs in periods t and $t+1$. Then, for every $\big(v_{i,t}, b, \boldsymbol{v}_{-i,t}\big)$:*

(i) $\dfrac{\partial V_{i,t+1}\big(v_{i,t}, b, \boldsymbol{v}_{-i,t}\big)}{\partial b} = -\,\beta\, R_b\big(b, \boldsymbol{v}_{-i,t}\big)\; \mathbb{E}\Big[\, q^{BR}_{i,t+1}\big(v_{i,t+1}; b, \boldsymbol{v}_{-i,t}\big) \;\Big|\; v_{i,t}\Big]$;

(ii) R_b decomposes as

$$R_b \;=\; \underbrace{-\,\kappa\,\frac{\partial \Delta_i}{\partial b}}_{\text{replacement}} \;+\; \underbrace{\frac{(\lambda + 2\beta)^2}{(\lambda + \beta)\big(\lambda + (N+1)\beta\big)} \sum_{k \neq i} \frac{\partial \Delta_k}{\partial b}}_{\text{amplification}}, \qquad \kappa \equiv \frac{(N-1)\,\beta\,(\lambda + 2\beta)}{(\lambda + \beta)\big(\lambda + (N+1)\beta\big)} \;>\; 0,$$

where the replacement term is strictly positive for all periods, and the amplification term is identically zero when $t+1 = T$, so that $R_b = -\kappa\, \partial \Delta_i / \partial b > 0$ when $t+1 = T$;

[34] When $t < T-1$, the second channel cannot be eliminated by application of the envelope theorem, as types below $\overline{v}_i$ produce more than a static best response, so the static first-order condition does not hold at the chosen quantity. The change in b affects rivals' inference, not just i's quantity choice.

[35] Recall that the third argument of $\Pi^{i,t+1}$ with full notation was q.

(iii) $\mathbb{E}\big[\, q^{BR}_{i,t+1} \,\big|\, v_{i,t}\big] \;\geq\; B_i \;>\; 0$. *Hence* $\operatorname{sign}\big(\partial V_{i,t+1}/\partial b\big) = -\operatorname{sign}(R_b)$*; in particular,* $\partial V_{i,t+1}/\partial b < 0$ *whenever* $R_b > 0$*, which holds for sure when* $t+1 = T$.

Proof. **(i)** Consider i's expected period-$(t+1)$ flow profit at a realized type $v_{i,t+1}$, given its equilibrium quantity $q = B_i(b) + h_{i,t+1}(v_{i,t+1})$. Because demand is linear, rivals' realized quantities enter this expectation only through their expected combined output R:

$$\mathbb{E}\pi\big(v_{i,t+1}; b\big) = \big(u - \beta q - \beta R(b) - v_{i,t+1}\big)\, q - \tfrac{\lambda}{2} q^2. \tag{38}$$

By observation (a), differentiating with respect to b picks up exactly two terms, the shift of i's own schedule and the shift of expected rival output:

$$\frac{\partial\, \mathbb{E}\pi}{\partial b} = \Pi_3^{i,t+1}\big(v_{i,t+1}, q\big)\, \frac{\partial B_i}{\partial b} \;-\; \beta\, q\, R_b. \tag{39}$$

Substituting (36) and (37),

$$\frac{\partial\, \mathbb{E}\pi}{\partial b} = -\,\beta R_b\big(\, q^{BR}_{i,t+1}(v_{i,t+1}) - q\,\big) \;-\; \beta\, q\, R_b = -\,\beta\, R_b\; q^{BR}_{i,t+1}\big(v_{i,t+1}\big), \tag{40}$$

so the equilibrium quantity q cancels pointwise. Intuitively, relative to choosing a static best-response, the signaler bears an extra cost from rival expansion equal to βR_b per over-produced unit (the price falls on those units too), but enjoys a benefit from the induced downward shift of its own schedule. The boundary condition (36) sizes that shift at exactly the best-response rate $-\beta R_b/(\lambda + 2\beta)$, and cutting output raises flow profit at rate $-\Pi_3^{i,t+1} = (\lambda + 2\beta)(q - q^{BR})$ — and the symmetry of the cross-partials ($\pi_{qR} = \pi_{Rq} = -\beta$) makes these two gap-proportional terms cancel unit-for-unit. A signaler therefore values changing rivals' beliefs exactly as if it would make a static best response, even though it will not. As the period-$(t+2)$ continuation is unaffected by

b, integrating over $v_{i,t+1} \mid v_{i,t}$ (using i's *true* lagged type) gives (i).

(ii) By observation (a), b affects the period-$(t+1)$ system of boundary intercepts in Proposition A.1 only through the expected increments Δ_j. From equation (29),

$$\frac{\partial\,\mathbb{E}[q_{j,t+1}]}{\partial\Delta_j} = \frac{(\lambda+2\beta)(\lambda+N\beta)}{(\lambda+\beta)\big(\lambda+(N+1)\beta\big)}, \qquad \frac{\partial\,\mathbb{E}[q_{j,t+1}]}{\partial\Delta_k} = -\,\frac{\beta(\lambda+2\beta)}{(\lambda+\beta)\big(\lambda+(N+1)\beta\big)} \;\; (k\neq j), \tag{41}$$

so summing over $j \neq i$ gives $\partial R/\partial\Delta_i = -\kappa$ and, for each $k \neq i$, $\partial R/\partial\Delta_k = (\lambda+2\beta)^2/(\lambda+\beta)\big(\lambda+(N+1)\beta\big)$; the chain rule then delivers the stated decomposition.

For the replacement term, suppressing the $\boldsymbol{v}_{-i,t}$ argument, recall that

$$\Delta_i(b) = \int_{\underline{v}_i}^{\overline{v}_i} h_{i,t+1}(v)\, g_i(v \mid b)\, dv, \tag{42}$$

where, by observation (a), $h_{i,t+1}$ does not depend on b. Differentiating with respect to b,

$$\frac{\partial\Delta_i(b)}{\partial b} = \int_{\underline{v}_i}^{\overline{v}_i} h_{i,t+1}(v)\frac{\partial g_i(v \mid b)}{\partial b}\, dv. \tag{43}$$

Because $g_i(\cdot \mid b)$ is a density for every b,

$$\int_{\underline{v}_i}^{\overline{v}_i} \frac{\partial g_i(v \mid b)}{\partial b}\, dv = 0. \tag{44}$$

Therefore, letting $v'(b)$ denote the crossing point in Assumption 1,

$$\frac{\partial\Delta_i(b)}{\partial b} = \int_{\underline{v}_i}^{\overline{v}_i} \big[h_{i,t+1}(v) - h_{i,t+1}\big(v'(b)\big)\big]\, \frac{\partial g_i(v \mid b)}{\partial b}\, dv. \tag{45}$$

The quantity function is strictly decreasing, so $h_{i,t+1}(v)$ is strictly decreasing in v. Hence, for

$v < v'(b)$,

$$h_{i,t+1}(v) - h_{i,t+1}\big(v'(b)\big) > 0 \quad \text{and} \quad \frac{\partial g_i(v \mid b)}{\partial b} < 0, \tag{46}$$

while for $v > v'(b)$,

$$h_{i,t+1}(v) - h_{i,t+1}\big(v'(b)\big) < 0 \quad \text{and} \quad \frac{\partial g_i(v \mid b)}{\partial b} > 0. \tag{47}$$

Thus the integrand is strictly negative except at $v = v'(b)$, and therefore

$$\frac{\partial \Delta_i(b)}{\partial b} < 0. \tag{48}$$

It follows that the replacement term

$$-\kappa \frac{\partial \Delta_i(b)}{\partial b} \tag{49}$$

is strictly positive.

For the amplification term: when $t+1 = T$, $h_{k,T}(v) = (\overline{v}_k - v)/(\lambda + 2\beta)$, which is independent of the state, so $\partial \Delta_k / \partial b = 0$ for all $k \neq i$ and the term vanishes. In this case $\Delta_i(b) = (\overline{v}_i - \mathbb{E}[v_{i,T} \mid b])/(\lambda + 2\beta)$, so that $R_b = \frac{(N-1)\beta}{(\lambda+\beta)\big(\lambda+(N+1)\beta\big)} \frac{\partial \mathbb{E}[v_{i,T}|b]}{\partial b}$, where the coefficient $\frac{(N-1)\beta}{(\lambda+\beta)\big(\lambda+(N+1)\beta\big)}$ is the derivative reported in equation (8) in the text. When $t+1 < T$, rivals' increments do depend on b: seller k's period-$(t+1)$ differential equation has a numerator that integrates k's period-$(t+2)$ continuation derivative over the transition of $v_{i,t+1}$ from the type that k attributes to i, namely b, so $h_{k,t+1}$, and therefore Δ_k, varies with b.

(iii) $q^{BR}_{i,t+1}(v)$ is strictly decreasing in v, so for every $v \in [\underline{v}_i, \overline{v}_i]$, $q^{BR}_{i,t+1}(v) \geq q^{BR}_{i,t+1}(\overline{v}_i) = B_i$, the boundary quantity, which is strictly positive under the maintained assumption that all types produce strictly positive outputs. Hence $\mathbb{E}[q^{BR}_{i,t+1} \mid v_{i,t}] \geq B_i > 0$; combining with (i) and $\beta > 0$ gives the sign. ■

Proposition A.2 (Belief Monotonicity) *Given $\boldsymbol{v}_{-i,t-1}$ and $\widehat{v}^{-i}_{i,t-1}$, and provided $R_b\big(\widehat{v}^{-i}_{i,t}, \boldsymbol{v}_{-i,t}\big) >$*

0 *for all* $\widehat{v}_{i,t}^{-i}$ *and* $\boldsymbol{v}_{-i,t}$*, i.e., rivals produce more when they believe* i *has a higher cost,*

$$\Pi_2^{i,t}\big(v_{i,t}, \widehat{v}_{i,t}^{-i}, q_{i,t}, \boldsymbol{v}_{-i,t-1}, \widehat{v}_{i,t-1}^{-i}\big) \;<\; 0 \tag{50}$$

for all $\big(v_{i,t}, \widehat{v}_{i,t}^{-i}\big)$*. When* $t = T-1$*, we know the proviso holds, so (50) holds unconditionally.*

Proof. Substituting Lemma A.3(i) into equation (31),

$$\Pi_2^{i,t} = -\,\delta\beta \int_{\underline{\boldsymbol{v}}_{-i}}^{\overline{\boldsymbol{v}}_{-i}} R_b\big(b, \boldsymbol{v}_{-i,t}\big)\, \mathbb{E}\Big[\, q_{i,t+1}^{BR}\big(v_{i,t+1}; b, \boldsymbol{v}_{-i,t}\big) \,\Big|\, v_{i,t}\Big]\, g_{-i}(\boldsymbol{v}_{-i,t} \mid \boldsymbol{v}_{-i,t-1})\, d\boldsymbol{v}_{-i,t}. \tag{51}$$

By Lemma A.3(iii) and the proviso, the integrand is strictly positive for every $\boldsymbol{v}_{-i,t}$, and $\delta, \beta > 0$. Hence $\Pi_2^{i,t} < 0$. Note that, when $t+1 < T$, both R_b and $\mathbb{E}[q_{i,t+1}^{BR}]$ depend on $\boldsymbol{v}_{-i,t}$ (through the increments and R respectively), so they must remain inside the integral; the sign argument applies pointwise. ■

We make a few remarks about what has been shown.

Remark A.1 (Verifying $R_b > 0$ when $t < T-1$) *For* $t < T-1$*, the proviso (i.e., rivals produce more when they believe* i *has a higher cost) concerns the amplification term in Lemma A.3(ii). Intuitively its sign is positive and reinforcing: a rival* k *who attributes a higher cost type to* i *anticipates larger static best-response quantities for itself in period* $t+2$*, and hence has stronger incentives to signal, which raises its increment* $h_{k,t+1}$ *and therefore* Δ_k*. However, proving that* R_b *is always positive would appear to require a complex multi-period analysis. Instead, we note that, by Lemma A.3, pointwise monotonicity of* $V_{i,t+1}$ *in* b *is equivalent to* $R_b > 0$*. Our computational procedure verifies, in every period, that the value functions, evaluated at a set of grid points, are consistent with belief monotonicity.*

Remark A.2 (Interpretation of numerical failures) *Lemma A.3(i) implies that the marginal*

value of a lower attributed type is proportional to i's expected period-$(t+1)$ static best-response quantity. Holding R_b away from zero, the marginal incentive to signal therefore becomes small as $\mathbb{E}[q_{i,t+1}^{BR}] \to 0$, i.e., as a seller approaches the point of preferring zero output against rivals' (over-produced) supply. This is consistent with our computational experience that Conditions 3 and 4 fail only in specifications where some sellers' outputs are zero or close to zero.

Remark A.3 (Extension to the specification with non-signaling sellers) *Our specification in Section 4 introduces non-signaling sellers, some of whom may have zero output for some cost realizations. The strict-positivity requirement above therefore applies to the signaling sellers, whose separating strategies we characterize. For belief monotonicity, the relevant object is total rival output: if expected total rival output increases when rivals attribute a higher cost to signaling seller i, then the same logic implies that i's continuation value decreases in its attributed type. We verify this condition numerically in the Section 4 specifications.*

B Computational Algorithms

This appendix describes the computational algorithms used to solve the dynamic signaling model. We focus on the baseline model used in the oil industry application, where a finite-horizon dynamic signaling model, with asymmetric sellers, is combined with a "fringe" of non-signaling sellers, before explaining what changes in other cases (e.g., when we assume an infinite-horizon game). We also describe how we simulate output, welfare and cost inefficiency outcomes.

B.1 Algorithm for Solving the Dynamic Signaling Game for the Application

There are N_S signaling quantity-setting sellers, indexed by $i = 1, \ldots, N_S$, with serially correlated private information vs, and a set of non-signaling sellers with uncorrelated private information vs. These other sellers will use static quantity-setting strategies as their current privately observed vs are uninformative about values in the next period.

B.1.1 Baseline Parameters

Demand and Common Shocks. AZ specify that inverse demand is $P_t = u_t - \beta Q_t$, where Q_t is period-t total output. We use AZ's estimate of $\beta = 0.03329$. AZ assume that each seller i has a cost function, $C_{i,t} = (w_t + v_{i,t})q_{i,t} + \frac{\lambda}{2}q_{i,t}^2$ where $q_{i,t}$ is i's output. Our baseline adopts AZ's estimate of $\lambda = 0.0007$.

AZ's static model assumes that u_t and w_t are random variables drawn independently across periods. These are assumed to be observed by all sellers. While we can include such shocks in our model (and have done so in examples), the fact that, in AZ's model, the us and ws are uncorrelated across periods means that the current values of these variables do not affect sellers' expectations of future profits or the future payoff from signaling private information. Allowing them to vary across periods increases computational costs. We therefore treat u and w as fixed, and set $u = 100$ and

$w = 0$.[36] We will not include w in our notation below.

$v_{i,t}$ Assumptions. AZ assume that $v_{i,t} \sim \overline{w}(\text{Beta}(a_i, b_i) + 1)$ where the parameters a and b vary across six groups of sellers. Draws are independent across periods. The estimated supports are $[28.2, 56.4]$. We choose narrower supports, reflecting the limited month-to-month variation in the calculated mean posteriors (Figure C.2) and the requirement that signaling sellers have strictly positive output in our signaling equilibrium.

For both signaling and non-signaling sellers, we will assume that the widths of the $v_{i,t}$ supports are 1.5.[37] While the support widths are the same, other assumptions on the $v_{i,t}$s differ between the signaling and non-signaling sellers.

Non-Signaling Sellers. We assume that their $v_{i,t}$s are drawn independently each period from symmetric truncated normal distributions on the supports, where the standard deviation of the untruncated distribution equals 0.6. The midpoint of the support for seller i, which will equal the expected value of $v_{i,t}$, is chosen so that if the seller has this cost intercept and uses a static BNE strategy, and every other seller is expected to produce its average output in the detrended data (Table C.1), then seller i's output will also equal its own average output in this data. The midpoints are therefore seller-specific.

Signaling Sellers. We assume that $v_{i,t} = \rho v_{i,t-1} + (1-\rho)\frac{\overline{v_i}+\underline{v_i}}{2} + \varepsilon_{i,t}$ where $\varepsilon_{i,t}$ is a draw from a normal distribution, which is truncated so that $v_{i,t}$ will remain on the support. Based on the estimated AR(1) coefficients for the signaling sellers, using the mean posterior vs implied by AZ's estimates, which we report in Table C.1, we set $\rho = 0.96$ and set the standard deviation of the untruncated (normal) distribution of the innovations equal to 0.5. This implies a standard deviation

[36] As explained in the text and Appendix C, we can calculate the value of u_t implied by the detrended data and a posterior expectation of w_t given the data and AZ's parameters. The implied u_ts and w_ts are fairly constant during the early part of the sample. The $u_t - w_t$ difference, which is what matters for output and profitability, varies less over the whole sample period.

[37] For example, the country-year average within-calendar year difference between the maximum and minimum $v_{i,t}$s implied by AZ's model is \$1.06.

of truncated innovations of 0.40, which is less than the standard deviation of the estimated residuals from the AR(1) regression for Saudi Arabia and Iraq, but more than the standard deviations of the estimated residuals for the other signaling sellers.[38]

As noted above, $\overline{v_i} - \underline{v_i} = 1.5$. To set the location of $\underline{v_i}$ we use two approaches. The first applies the same method used for the non-signaling countries, so that, when we assume static BNE behavior, the expected seller outputs for all countries match those in the detrended data. However, when we allow for signaling, the signaling countries will tend to produce more than the data implies when the supports are set in this way. We therefore also present results using a second approach where the $\underline{v_i}$s are calibrated so that the average output implied by the signaling equilibrium matches the data average.

B.1.2 Non-Signaling Supply

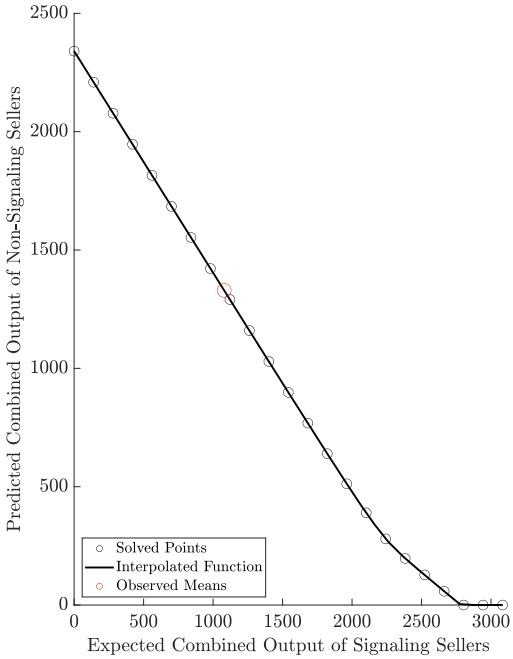
The non-signaling sellers will use static Bayesian Nash equilibrium strategies given an expected output of the signaling sellers. Before solving the signaling game, we pre-solve for a function that maps from the expected output of the signaling sellers to the mean combined output of the non-signaling sellers.

To do so, we specify a grid of 23 values for the expected output of the signaling sellers, covering a range from zero to levels where every non-signaling seller optimally produces zero given any cost realization. At each grid value, we solve for the expected combined non-signaling seller output. We find equilibrium strategies, which may involve zero output for some $v_{i,t}$ draws, by fixed point iteration, where, given the current guess of expected output of each non-signaling seller, we compute the threshold v for which each seller would want to produce zero output, and the seller's expected

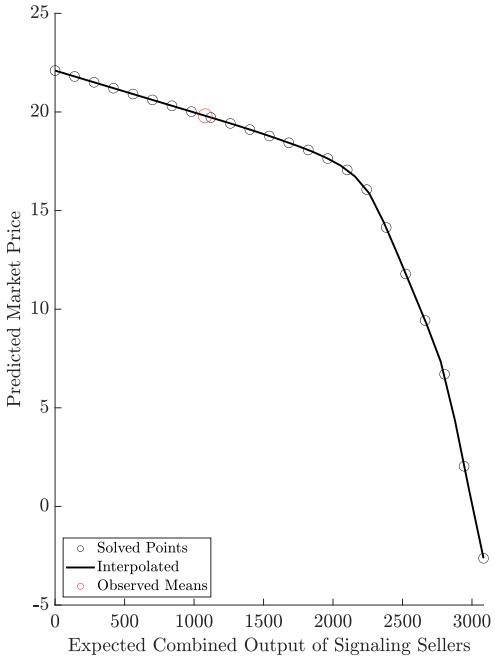
[38] Due to the truncation, the steady-state distribution of the $v_{i,t}$s will not correspond to a standard distribution. Therefore, as described below, when we simulate data, we will initially take draws for the signaling seller $v_{i,t}$s from uniform distributions but simulate their evolution for a number of periods so that the distributions of the vs can change, before calculating outcomes. Note that, as the non-signaling seller draws are assumed to be drawn, each period, from known truncated normal distributions, we use these distributions in the first period.

Figure B.1: Expected Combined Output of the Non-Signaling Sellers and the Expected Market Price Under the Baseline Specification. The red circle in (a) reflects the average outputs of each group observed in the detrended data. The red circle in (b) marks the predicted market price given these average outputs and the assumed demand function.

(a) Predicted Non-Signaling Output



(b) Predicted Market Price



output given its v distribution and this threshold. We form a function, $Q^{NS}(Q^S) = \sum_{i=1}^{N_{NS}} q_i^{NS}(Q^S)$, using piecewise cubic interpolation between the grid points.

Appendix Figure B.1(a) shows this function, with the black circles marking the calculated outcomes at the grid points. For most of the range, an additional unit of signaling output reduces non-signaling output by roughly 0.94 units. The relationship is approximately linear until prices (panel (b)) are low enough that several non-signaling sellers choose not to produce.

For use below, we create two MATLAB `griddedInterpolant` objects from this function.

- `eqf` is the $Q^{NS}(Q^S)$ function except, as it is useful to have a strictly decreasing function, we make a slight change by assuming that the combined non-signaling output takes on tiny, decreasing values, of the order of 10^{-12}, in the range where the model implies that the

combined output would be exactly zero. In the examples considered in the paper, these values are not relevant in equilibrium.

- `inveqf` is the inverse of `eqf`.

B.1.3 Timing and Definition of the State Space

We will define value functions, and expected profits, for signaling sellers at the start of each period, before a seller's $v_{i,t}$ has been revealed. Therefore, the value function for seller i will be defined on a state space $\left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ where $\widehat{v_{i,t-1}^{-i}}$ is the $v_{i,t-1}$ that other sellers will have inferred from i's action in the last period, and $\boldsymbol{v}_{-i,t-1}$ is the vector of the last period vs of the other signaling sellers.

Quantity choices are made when seller i knows $v_{i,t}$. Conditional on $v_{i,t}$ and the signal that i will send in period t, and given the first-order Markovian process for $v_{i,t}$ and the assumption of Markovian strategies, $v_{i,t-1}$ has no effect on i's current or future profits. Therefore, i's strategies will be defined, as functions of $v_{i,t}$, given states $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$.

B.1.4 Storage Grids

While strategies and states are continuous, we store signaling seller values and strategies at a finite number of points, and then interpolate where necessary. For these purposes, in our oil application, we will make use of vectors of v with three different numbers of points:

- a "coarse" vector with four equally spaced points from $\underline{v_i} + \iota_v$ to $\overline{v_i} - \iota_v$ where $\iota_v = 0.01$.
- a "standard" vector of vs with seven equally spaced points from $\underline{v_i} + \iota_v$ to $\overline{v_i} - \iota_v$ where $\iota_v = 0.01$.
- a "detailed" vector of vs with forty equally spaced points from $\underline{v_i} + \iota_{dv}$ to $\overline{v_i} - \iota_{dv}$ where $\iota_{dv} = 10^{-4}$.

Note that the storage grids exclude the endpoints of the support, but, when we solve the differential equation, we do initialize the solver at the highest value $\overline{v_i}$.

Note that in our Section 3 computations, the coarse and standard vectors are the same. When we solve the two-seller games, these vectors have 21 points, and the detailed vector has 400 points. When we solve games with variable numbers of sellers, the coarse/standard vectors have 4 points and the detailed vector has 25 points.

B.1.5 Integration

Our calculations require integration over the evolution of own and rival vs to calculate expected outputs, profits and value functions. All of the integrations performed will use trapezoidal integration, where we matrix multiply appropriate pre-computed vectors or matrices, with vectors of outputs or values (defined using the coarse, standard or detailed v vectors). We will use the detailed vectors when the underlying functions are non-linear (e.g., the signaling quantity functions). We have verified that increasing the number of grid points, which will make trapezoidal integration more accurate, or varying the values of ι_v or ι_{dv}, has negligible effects on our comparisons of output or welfare.

We solve the finite-horizon game backwards. We describe the method used to solve the game in the final period, when sellers will use static strategies, before describing the method used for any earlier period, t.

B.1.6 Period T Solution (Static BNE Strategies)

For each $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$, defined by the coarse vectors, we do the following steps 1–3:

Step 1. Calculate "expected incremental quantities" for each seller ($\Delta_{i,T}$). We define the expected incremental quantity, $\Delta_{i,T}$, as the quantity that i is expected (by other sellers) to

produce in excess of the quantity that it would produce if $v_{i,T} = \overline{v_i}$ given its equilibrium strategy. This can be calculated without knowing the equilibrium output of a $\overline{v_i}$ type. Recall,

$$q_{i,T}(v_{i,T}) = \frac{u - \beta \mathbb{E}Q_{-i,T} - v_{i,T}}{2\beta + \lambda} \tag{52}$$

where $\mathbb{E}Q_{-i,T}$ is the expected combined output of all rivals (including the non-signaling sellers), implying that

$$q_{i,T}(v_{i,T}) = q_{i,T}(\overline{v_i}) + \frac{\overline{v_i} - v_{i,T}}{2\beta + \lambda}. \tag{53}$$

Then,

$$\Delta_{i,T} = \frac{\overline{v_i} - \mathbb{E}(v_{i,T} | \widehat{v_{i,T-1}^{-i}})}{2\beta + \lambda}. \tag{54}$$

This calculation assumes that each signaling seller will produce a positive quantity for $v_{i,T} = \overline{v_i}$, which we verify.

Step 2. Solve for the equilibrium quantity intercepts and the expected non-signaling output. In this step we solve for the equilibrium output of each signaling seller when it has the highest-cost type, $v_{i,T} = \overline{v_i}$, and the expected output of the non-signaling fringe. This involves solving $N_S + 1$ equations. Denote $b_{i,T} = q^*_{i,T}(\overline{v_i})$. The first equation is

$$Q_T^S = \sum_{j=1}^{N_S} \left(b_{j,T} + \Delta_{j,T} \right), \tag{55}$$

as the expected output of the signaling sellers in period T. The expected output of the non-signaling sellers, using the `eqf` interpolant function, is therefore $\texttt{eqf}(Q_T^S)$. There are N_S equations of the form

$$b_{i,T} = \frac{u - \beta \sum_{j \neq i} \left(b_{j,T} + \Delta_{j,T} \right) - \beta \texttt{eqf}(Q_T^S) - \bar{v}_i}{2\beta + \lambda}, \qquad i = 1, \ldots, N_S. \tag{56}$$

For future use, we also calculate the expected combined output of all *other* sellers (regardless of whether they signal), i.e.,

$$\mathbb{E}Q_{-i,T} = \sum_{j \neq i} \left(b^*_{j,T} + \Delta_{j,T} \right) + \texttt{eqf}(Q_T^{S*}). \tag{57}$$

where the b^*s are the solved values, and Q_T^{S*} is the value implied by equation (55) given the b^*s.

Step 3. Storage of the quantity functions. We store each signaling seller's quantity functions and the expected combined output of all other sellers, and the expected combined output of the non-signaling sellers into storage grids defined by the coarse vectors for $(\boldsymbol{v}_{-i,T-1}, \widehat{v^{-i}_{i,T-1}})$ and by the detailed vector for $v_{i,T}$.

Having completed steps 1 to 3 for all sets of previous period types defined by the coarse vectors, we do step 4:

Step 4. Interpolate to standard vector storage grids. We use linear interpolation to interpolate from these coarse storage grids to storage grids based on the standard vectors for past types[39], and create interpolation functions (using MATLAB's `griddedInterpolant` command) that can be used to calculate quantities for any combination of $T-1$ and realized T types. The near-linearity of the quantity functions with respect to other sellers' lagged types makes these interpolations accurate.

Step 5. Integrate to calculate start-of-period T expected profits for each seller i on the state space defined by $\left(v_{i,T-1}, \boldsymbol{v}_{-i,T-1}, \widehat{v^{-i}_{i,T-1}} \right)$. Note that this state space distinguishes between seller i's actual previous type and its perceived previous type, so that we can quantify the incentives to signal alternative types in $T-1$. A seller i's expected profits, given a realization of

[39] We continue to use the detailed vectors for current types.

$v_{i,T}$ equals

$$\left(u - \beta \mathbb{E} Q_{-i,T}\left(\boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right) - \beta q_{i,T}\left(v_{i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right) - v_{i,T}\right) \tag{58}$$

$$\times q_{i,T}\left(v_{i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right) - \frac{\lambda}{2}\left\{q_{i,T}\left(v_{i,T}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right)\right\}^2$$

We perform a one-dimensional integration, over the evolution of $v_{i,T}$, defined by the detailed vector, given the assumed $v_{i,T-1}$ value. This gives us expected profits defined on the standard vectors. The period T value function equals expected period T profits.

Step 6. Verify the belief monotonicity property. If seller i's period-T value function satisfies belief monotonicity for all $v_{i,T-1}$ and all $\boldsymbol{v}_{-i,T-1}$, i.e., i's value increases when $\widehat{v_{i,T-1}^{-i}}$ decreases, then the signaling payoff function used in period $T-1$ will also satisfy its belief monotonicity condition. For each seller, we therefore verify that, for all values on the grids defined by the standard vectors,

$$V_{i,T}\left(v_{i,T-1}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{-i}}\right) > V_{i,T}\left(v_{i,T-1}, \boldsymbol{v}_{-i,T-1}, \widehat{v_{i,T-1}^{\prime -i}}\right) \text{ for all } \widehat{v_{i,T-1}^{\prime -i}} > \widehat{v_{i,T-1}^{-i}}. \tag{59}$$

This test passes for all of the specifications presented in the paper. This completes the analysis of period T.

B.1.7 Period $t \leq T-1$ Solution

We now describe the process followed in a generic period $t \leq T-1$ (i.e., where strategies will involve signaling). Each period-t solution consists of two nested problems: first, computing the incremental quantity functions $h_{i,t}$, and their expectations $\Delta_{i,t}$; and second, solving simultaneously for the boundary quantities. The incremental quantity functions and the boundary quantities are combined to give the equilibrium quantity functions.

For each combination of $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ defined by the coarse vectors, we do the following steps 1–3 to solve for the equilibrium quantity functions for each seller in turn.

Step 1. Calculate the "expected incremental quantities" ($\Delta_{i,t}$) for each signaling seller by solving the differential equation given an arbitrary assumption about the expected output of other sellers.

As discussed in Appendix A and Lemma A.2, the signaling quantity function of seller i can be written as $q_{i,t}^*(\overline{v_i}, \mathbb{E}Q_{-i,t}) + h_{i,t}(v_{i,t})$ where the first term is the static best response of i when it has the highest-cost type $\overline{v_i}$ and the expected output of the other sellers is $\mathbb{E}Q_{-i,t}$, and $h_{i,t}(v_{i,t})$ is a function that does not depend on the period t strategies of other sellers. Therefore we can compute $h_{i,t}(v_{i,t})$, and its expectation using an arbitrary guess of the quantity produced by rivals (we will use the static BNE quantities, i.e., those that were calculated for period T).

To solve the differential equations, we proceed through steps 1a-1g:

Step 1a. calculate the expected continuation values entering the signaling payoff function. Given the assumed $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$, we calculate a matrix, based on the standard vectors, that contains the discounted expected continuation values of seller i as a function of the realized $v_{i,t}$s and $\widehat{v_{i,t}^{-i}}$ by integrating the period $t+1$ value function over transitions from period $t-1$ to period t of the vs of the other signaling sellers.

Step 1b. construct an interpolation function. We create a cubic spline MATLAB `griddedInterpolant` object that can calculate the value in this matrix for any possible $\left(v_{i,t}, \widehat{v_{i,t}^{-i}}\right)$ combination.

Step 1c. calculate derivatives with respect to $\widehat{v_{i,t}^{-i}}$. We use the interpolation function to calculate numerical derivatives, at the points on the detailed vectors, using a step size of 10^{-4}. We store this derivative matrix and take its diagonal where the derivatives are evaluated at $v_{i,t} = \widehat{v_{i,t}^{-i}}$.

Step 1d. calculate i's static best response quantity, when it has type $v_{i,t} = \overline{v_i}$, to the (arbitrarily assumed) static BNE expected quantity of all rivals. Call this best response

q_i'.

Step 1e. solve the differential equation, for the function $q_{i,t}(v_{i,t}, \mathbb{E}Q_{-i,t})$, backwards from the boundary condition. To avoid the value of the differential equation being undefined, we begin from $(\overline{v_i}, q_i' + 10^{-5})$, and use MATLAB's `ode113` solver (with `InitialStep` equal to 0.00005 and `MaxStep` equal to 0.1) to solve the differential equation (5) to give $h_{i,t}(v_{i,t})$. The ODE solver interpolates the vector from step 1c, using 1-d spline interpolation, to calculate the numerator, and then calculates the denominator analytically using $u - \beta\mathbb{E}Q_{-i,t} - (2\beta + \lambda)q_{i,t} - v_{i,t}$, where $\mathbb{E}Q_{-i,t}$ is the guessed quantity of rivals.

Step 1f. check for violations of the single-crossing property. Recall that single-crossing requires that $\frac{\Pi_3^{i,t}(v_{i,t},\widehat{v_{i,t}^{-i}},q_{i,t},\boldsymbol{v}_{-i,t-1},\widehat{v_{i,t-1}^{-i}})}{\Pi_2^{i,t}(v_{i,t},\widehat{v_{i,t}^{-i}},q_{i,t},\boldsymbol{v}_{-i,t-1},\widehat{v_{i,t-1}^{-i}})}$ is a monotone function of $v_{i,t}$, holding $q_{i,t}$ and $\widehat{v_{i,t}^{-i}}$ fixed, at every point $(v_{i,t}, q_{i,t})$ in the graph of the quantity function. Application of the chain rule implies that *sufficient* conditions for single-crossing to hold are that (i) $\Pi_2^{i,t} < 0$ (i.e., belief monotonicity), (ii) $\Pi_3^{i,t} < 0$ for all $(v_{i,t}, q_{i,t})$ in the graph of the quantity function (i.e., outputs are more than the static best responses); and (iii) $\frac{\partial \Pi_2^{i,t}}{\partial v_{i,t}} > 0$ (i.e., it is less costly to be perceived as high v when v is actually higher). (i) has already been checked, so we verify that properties (ii) and (iii) hold numerically.[40]

Step 1g. calculate the expected incremental quantities.

$$\Delta_{i,t} = \mathbb{E}\left[h_{i,t}|\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)\right] \tag{60}$$

We use the $(v_{i,t}, q_{i,t})$ pairs returned by the ODE solver, and 1-dimensional spline interpolation, to calculate the quantities for the $v_{i,t}$ points on the detailed vector. The incremental quantity function

[40] To allow for numerical error, we treat a derivative as satisfying the positive (negative) sign test if it exceeds -10^{-3} (is below 10^{-3}). We have encountered cases where numerical issues can lead us to find a violation of the sufficient conditions if we do a simple sign test, when more detailed analysis using different grids shows that the sufficient conditions actually hold. Note that we are testing sufficient conditions for single-crossing: if these conditions fail, it is still possible that single-crossing holds. However, for the examples we consider, the sufficient conditions are satisfied (up to the tolerances listed above).

is $h_{i,t}(v_{i,t}) = q_{i,t}(v_{i,t}, \mathbb{E}Q_{-i,t}) - q_i'$. We calculate $\Delta_{i,t}$ by integrating $h_{i,t}$ over the period $t-1$-to-t evolution of $v_{i,t}$ from $\widehat{v_{i,t-1}^{-i}}$ (i.e., from i's perceived $t-1$ type).

Having found the h functions and values of Δ for each seller, we can then proceed as in period T. In the following description we only elaborate where we make changes from the period T procedures described above.

Step 2. Solve for the equilibrium quantity intercepts and the expected non-signaling output. Given the expected incremental quantities $\Delta_{i,t}$, we solve jointly for the equilibrium boundary quantities, and the expected output of the non-signaling sellers, using the period-T method (replacing the $\Delta_{i,T}$s with the $\Delta_{i,t}$s).[41]

For future use, we also calculate the expected quantity produced by all other sellers (regardless of whether they signal), i.e.,

$$\mathbb{E}Q_{-i,t}^* = \sum_{j \neq i} \left(b_{j,t}^* + \Delta_{j,t}\right) + \texttt{eqf}(Q_t^S). \tag{61}$$

Step 3. Storage of the quantity functions. As in period T.

Having completed steps 1 to 3 for all $\left(\boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$ defined by the coarse vectors for each signaling seller, we do step 4:

Step 4. Interpolate to standard vector storage grids. As in period T.

We calculate each seller's expected period t profits, in each $\left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}\right)$, defined by the standard vectors, in step 5a, and then each seller's values in step 5b.

Step 5a. Integrate to calculate start-of-period t expected period t profits, given equilibrium strategies for each state. When states are stacked, this gives a vector $\pi_{i,t}$ for seller i.

[41] When there are no non-signaling sellers, the linear equations in Proposition A.1 provide the solution.

Step 5b. Calculate the start-of-period t value function by adding the discounted expectation of period $t+1$ values to the expected profits. This is done by calculating $V_{i,t} = \pi_{i,t} + \delta \mathbb{M}_i * V_{i,t+1}$, where $\mathbb{M}_i$ is a discretized matrix that performs trapezoidal integration over the period t-to-$t+1$ v transitions of all sellers, assuming that all sellers, including i, use truthful strategies in period t (i.e., it assumes that the relevant $t+1$ states are those in which $\widehat{v_{i,t}^{-i}} = v_{i,t}$).

Step 6. Verify belief monotonicity. We use the same process that was described above for period T values.

We then repeat the process moving back one more period.

B.2 Changes to the Algorithm for Alternative Specifications

The models presented in Section 3 do not have non-signaling sellers. The algorithm proposed above can be applied if the quantities of the non-signaling fringe are set equal to zero. Here we discuss two other changes.

B.2.1 Symmetric Sellers.

When sellers are ex-ante symmetric we assume that they will use symmetric separating equilibrium strategies where each seller only cares about the types of its rivals, not their identities. We exploit symmetry to minimize the number of states for which we have to solve the differential equations. We only need to keep track of the expected profits and values of a single representative seller.

B.2.2 Infinite-Horizon Game.

The algorithm above is for a finite-horizon game, where given the (equilibrium) value function in period $t+1$ we can sequentially solve for equilibrium strategies, expected profits and values in period t. However, when we look for a stationary equilibrium in an infinite-horizon game, we are solving simultaneously for strategies and values.

The changes to the algorithm, however, are relatively small. Here we outline the changes to the algorithm presented above. We use policy function iteration, following Judd (1998).

Algorithm. An iteration counter is set equal to one. We set $\pi_{i,t}^{iter=1}(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}})$ equal to the expected (start-of-period) profits in a static BNE game in each discretized state based on the standard vectors. We then use a policy function iteration step to calculate seller values if these were the profits in every period. For the first iteration,

$$V_i^{iter=1}(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}}) = [\mathbb{I} - \delta \mathbb{M}_i]^{-1} \pi_{i,t}^{iter=1} \left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}} \right), \tag{62}$$

where $\mathbb{I}$ is the identity matrix, and $\mathbb{M}_i$ is the integration matrix mentioned above. As the number of sellers increases, and the number of discretized points in the state space expands, performing this inversion becomes prohibitive.

In subsequent iterations, we follow the finite-horizon period t steps 1-5a to calculate iteration-k strategies and expected profits. To update values, we use a policy iteration step with damping,

$$\begin{aligned} V_i^{iter=k} \left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}} \right) = \tau \left\{ [\mathbb{I} - \delta \mathbb{M}_i]^{-1} \pi_{i,t}^{iter=k} \left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}} \right) \right\} \\ + (1-\tau) V_i^{iter=k-1} \left(v_{i,t-1}, \boldsymbol{v}_{-i,t-1}, \widehat{v_{i,t-1}^{-i}} \right), \end{aligned} \tag{63}$$

where $\tau = 1 - \min \left\{ 0.9, \max \left(0.1, \frac{0.9 \times (iter-5)}{100} \right) \right\}$.

We determine that strategies have converged if the maximum change in quantity strategies (evaluated at any point on the grid defined by the coarse vectors) is less than 0.001%.

B.3 Simulating Outputs, Welfare and Efficiency Under Different Equilibrium Concepts

The main objective of the paper is to compare market outcomes under different equilibrium concepts. We now describe how we simulate the outcomes reported in the paper given equilibrium strategies.

Draws. When we compare outcomes under different information assumptions or equilibrium concepts, we do so using the same sets of draws for the $v_{i,t}$s. For signaling sellers, we allow the $v_{i,t}$s to evolve over 20 periods (Table 1) or 30 periods (oil application) before we use them to calculate outcomes, so that the distribution of the $v_{i,t}$s should be close to stationary.

Static Complete Information Nash Equilibrium. In this equilibrium, all sellers are assumed to know each other's $v_{i,t}$s. Given our assumptions, it is possible that, given their draws, some sellers will produce zero in equilibrium. We therefore use an iterative algorithm that starts by solving, analytically, for the market price and seller quantities assuming that no sellers are at this corner solution. If the result has strictly positive outputs for all sellers, this is the solution. If not, then sellers predicted to have negative quantities are set to be inactive, and the algorithm is repeated.

With the quantities we compute total output, the market price, total costs, consumer surplus and seller profits in each period-simulation. Total surplus is the sum of consumer surplus and seller profits.

The cost "inefficiency index" compares the equilibrium total costs with the total costs of producing the same total output with a cost-minimizing allocation across sellers, where the allocation is done by a planner with complete information. The efficient allocation of output often has sellers producing zero, so we use an iterative algorithm, where we calculate the optimal shadow cost of production for a given assumed set of active sellers and the implied output of each seller, to find the

allocation. The index is calculated as $100 \times \left(\frac{\text{Total Costs in Equilibrium}}{\text{Total Costs in Efficient Allocation}} - 1 \right)$, i.e., the percentage increase in total costs over the efficient solution. When sellers are symmetric, the values of the inefficiency index naturally tend to be closer to zero (i.e., the types of Cournot inefficiency that arise when sellers are consistently asymmetric will not be present).

Static Bayesian Nash Equilibrium. The static Bayesian Nash equilibrium outcomes that we report assume that sellers know all sellers' $v_{i,t-1}$s (i.e., the asymmetric information comes from other sellers not knowing how the cost intercept evolves from period $t-1$ to period t). This provides the most appropriate comparison to signaling outcomes, i.e., the differences will come from sellers' signaling strategies, not their information.

We find static BNE strategies using the period T algorithm steps 1-5 described for the finite-horizon game. This produces strategies defined on a grid of lagged vs, defined by the standard vectors, and each seller's current v (where we use the detailed vector). For the purposes of simulation, we create MATLAB `griddedInterpolant` objects for the strategies of each seller. If there are non-signaling sellers, we also create a `griddedInterpolant` object (as a function of lagged vs) for the equilibrium expected combined output of these sellers.

For each simulation, we then use each signaling seller's interpolation function, its $v_{i,t}$ and the previous period vs to calculate its realized output. We also calculate the expected combined output of the signaling sellers, given the interpolated combined output of the non-signaling sellers, using the `inveqf` function. Denote this quantity Q'^S. Where appropriate, we calculate the output of individual non-signaling sellers using their best response quantity functions, i.e.,

$$q_i^*(v_{i,t}) = \max\left\{0, \frac{u - \beta \mathbb{E}Q_{-i,t}(Q'^S) - v_{i,t}}{2\beta + \lambda}\right\} \tag{64}$$

where $\mathbb{E}Q_{-i,t}(Q'^S)$ is the function for the expected output of all sellers other than i, given an

expected output of the signaling sellers.

With the quantities in hand we then proceed to compute outcomes of interest in the same way as the static complete information case. Note that the inefficiency index is computed assuming that the cost-minimizing allocation is chosen with knowledge of the $v_{i,t}$s of all sellers.

Dynamic Signaling Equilibrium (with possible non-signaling sellers). We use the same approach as for the static Bayesian Nash equilibrium to compute outputs and outcomes. When we use period-specific strategies in a finite-horizon game, we create `griddedInterpolant` objects for strategies and the expected combined output of the non-signaling sellers, for each period $T-1$, $T-2$,..., so that we can simulate how strategies vary across periods.

B.4 Leaders-Followers Model

We present the "leaders-followers" model as an alternative type of model where a subset of sellers, who move simultaneously, are able to affect the quantities chosen by a different subset of sellers. Here we describe the model in slightly more detail and discuss how it is solved.

Assumptions. There is a subset of leader sellers, $l = 1, \ldots, L$, and a different subset of follower sellers $f = 1, \ldots, F$. The timing of the game each period t is as follows:

1. the period t $v_{l,t}$s of the leaders are publicly revealed. They do not know the period t $v_{f,t}$ draws of the followers.

2. the leaders simultaneously choose their quantities, $q_{l,t}$. These commitments are observed by the followers.

3. the followers simultaneously choose their quantities, $q_{f,t}$, using static Bayesian Nash strategies.

The game therefore resembles a Stackelberg game, although here there are several sellers moving first, and they do not know the exact cost draws of the followers. There are no signaling incentives,

as leader vs are publicly revealed each period.

Solution. We solve for the expected quantity produced by the followers as a function of the output produced by the leaders as in the dynamic signaling game (the only difference is that the quantity produced by the leaders will be known, but this does not affect the function).

Then for any realization of the draws we solve for the leaders' quantities and the combined output of the followers (determined by the function), where the leaders take into account how changes in their output will affect the output of the followers, i.e., there is an additional term in the leaders' first-order conditions, which take the form:

$$u - v_{l,t} - \beta \sum_{m \neq l} q_{m,t} - \beta Q_F(Q_L) - (2\beta + \lambda) q_{l,t} - \beta q_{l,t} Q'_F(Q_L) = 0 \tag{65}$$

where Q_F is the function that determines expected follower output and $Q_L = \sum_{l=1}^{L} q_{l,t}$. Corner solutions are handled in the same way as for the static complete-information equilibrium.[42]

B.5 Additional Results for the $N = 2$ Illustrative Specification

Table 1 in the text shows outcomes in the finite-horizon dynamic signaling model in period $T - 20$. Appendix Table B.1 additionally shows outcomes in $T - 1$ and when we use stationary strategies found by solving an infinite-horizon version of the model. Two results stand out, which are consistent with our analysis of $N = 2$ strategies. First, outcomes in period $T - 1$ differ quite significantly from those predicted by either of the static equilibria. Second, infinite-horizon predicted outcomes are very similar to those predicted by the $T - 20$ finite-horizon model. This is true even in column (5) where strategies converge more slowly than in the baseline specification.

[42] As mentioned above, the calculated expected output function assumes that the non-signaling followers would produce a negligible output rather than zero.

Table B.1: Welfare and Output in $N = 2$ Games with Different Parameters. Inverse Demand is $P = 200 - 0.05 \sum_j q_j$ and each seller has a quadratic cost function, $C_{i,t} = v_{i,t} q_{i,t} + \frac{\lambda}{2} q_{i,t}^2$ where $\lambda = 0.01$, and $v_{i,t}$ lies on an interval $[\underline{v}, \overline{v}]$.

Parameters	(1)	(2)	(3)	(4)	(5)
$\underline{v}$	20	15	20	20	15
$\overline{v}$	30	35	30	30	35
AR(1) ρ	0.8	0.8	0.9	0.8	0.9
σ_v	4.0	4.0	4.0	2.5	2.5
Static Complete Information Nash					
Mean Total Output	2,187.5	2,187.5	2,187.5	2,187.5	2,187.5
Std. Dev. Total Output	23.6	41.9	23.8	22.0	40.0
Mean Consumer Surplus (/1000)	119.6	119.7	119.6	119.6	119.7
Mean Profits (/1000)	131.7	132.0	131.7	131.7	131.9
Mean Welfare (/1000)	251.4	251.7	251.4	251.3	251.6
Inefficiency Index	0.38%	1.23%	0.38%	0.33%	1.12%
Static Incomplete Information Nash					
Mean Total Output	2,187.5	2,187.5	2,187.5	2,187.5	2,187.5
Std. Dev. Total Output	33.4	53.5	33.4	29.2	45.8
Mean Consumer Surplus (/1000)	119.7	119.7	119.7	119.7	119.7
Mean Profits (/1000)	131.7	131.9	131.7	131.7	131.9
Mean Welfare (/1000)	251.3	251.6	251.3	251.3	251.6
Inefficiency Index	0.44%	1.36%	0.45%	0.37%	1.18%
Finite-Horizon Dynamic Incomplete Information MPBE, Period $T-1$					
Mean Total Output	2,267.8	2,344.8	2,273.5	2,301.0	2,373.7
Std. Dev. Total Output	53.4	89.4	55.7	58.5	83.3
Mean Consumer Surplus (/1000)	128.6	137.6	129.3	132.4	141.0
Mean Profits (/1000)	126.9	122.0	126.5	124.7	120.0
Mean Welfare (/1000)	255.5	259.6	255.8	257.2	261.0
Inefficiency Index	0.37%	1.04%	0.37%	0.27%	0.79%
Finite-Horizon Dynamic Incomplete Information MPBE, Period $T-20$					
Mean Total Output	2,315.9	2,460.6	2,331.2	2,418.1	2,598.0
Std. Dev. Total Output	66.9	110.4	72.1	86.6	121.8
Mean Consumer Surplus (/1000)	134.2	151.7	136.0	146.4	169.1
Mean Profits (/1000)	123.7	113.0	122.7	116.1	100.6
Mean Welfare (/1000)	257.9	264.6	258.6	262.5	269.7
Inefficiency Index	0.33%	0.87%	0.32%	0.19%	0.43%
Infinite-Horizon Dynamic Incomplete Information MPBE					
Mean Total Output	2,315.7	2,460.4	2,331.0	2,417.9	2,598.2
Std. Dev. Total Output	66.8	110.2	71.9	86.5	121.0
Mean Consumer Surplus (/1000)	134.2	151.6	136.0	146.3	169.1
Mean Profits (/1000)	123.7	113.0	122.7	116.1	100.6
Mean Welfare (/1000)	257.9	264.6	258.6	262.5	269.7
Inefficiency Index	0.33%	0.87%	0.32%	0.19%	0.43%

Table B.2: Comparison of Simulated Outcomes in the Finite-Horizon Dynamic Signaling Model with 2, 4, or 6 Sellers When Different Numbers of Grid Points are Used in the Calculation of Equilibrium Strategies

Number of Sellers	2	4	6	2	4	6	2	4	6
Grid Points	4	4	4	6	6	6	8	8	8
$\underline{v}$	20	20	20	20	20	20	20	20	20
$\overline{v}$	30	30	30	30	30	30	30	30	30
AR(1) ρ	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8	0.8
σ_v	4	4	4	4	4	4	4	4	4
Finite-Horizon Dynamic Incomplete Information MPBE, Period $T-1$									
Mean Total Output	2267.8	2800.1	3022.9	2267.8	2800.0	3022.7	2267.8	2800.0	3022.7
Std. Dev. Total Output	53.6	78.1	91.2	53.6	78.0	91.1	53.6	78.0	91.1
Mean Consumer Surplus (/1000)	128.6	196.2	228.7	128.6	196.2	228.6	128.6	196.2	228.6
Mean Profits (/1000)	126.9	88.3	64.7	126.9	88.3	64.7	126.9	88.3	64.7
Mean Welfare (/1000)	255.5	284.5	293.4	255.5	284.5	293.4	255.5	284.5	293.4
Inefficiency Index	0.37%	0.96%	1.57%	0.37%	0.96%	1.57%	0.37%	0.96%	1.57%
Finite-Horizon Dynamic Incomplete Information MPBE, Period $T-15$									
Mean Total Output	2316.6	2863.0	3071.1	2316.4	2862.5	3070.6	2316.3	2862.4	3070.5
Std. Dev. Total Output	67.3	99.0	108.5	67.2	98.9	108.3	67.2	98.8	108.3
Mean Consumer Surplus (/1000)	134.3	205.2	236.1	134.3	205.1	236.0	134.2	205.1	236.0
Mean Profits (/1000)	123.7	81.0	58.1	123.7	81.0	58.2	123.7	81.0	58.2
Mean Welfare (/1000)	258.0	286.1	294.2	257.9	286.1	294.2	257.9	286.1	294.2
Inefficiency Index	0.33%	0.85%	1.43%	0.33%	0.85%	1.43%	0.33%	0.85%	1.44%

B.6 Sensitivity of Simulated Outcomes to the Number of State-Space Grid Points

For tractability we have to keep the number of points in the state-space (defined by beliefs about previous types) relatively low. The most extreme assumption is made in Section 3 when we solve a game with up to 13 sellers, representing each seller's possible types by only 4 points. While we cannot (for reasons of memory, as well as computation time) relax this assumption when there are 13 sellers, we can do so with fewer sellers. Table B.2 therefore shows simulated outcomes based on equilibrium dynamic signaling strategies calculated using 4, 6 or 8 grid points for each seller, for 2, 4 or 6 sellers. While the table reveals some small differences when precision varies, these are orders of magnitude smaller than the differences between equilibria discussed in the text.

C Common and Private Shocks Implied Using the Aryal and Zincenko (2024) Estimates

As explained in Section 4 of the text, AZ estimate the demand and cost parameters of a static incomplete information game where they allow for common shocks in both demand and countries' marginal cost intercepts, in addition to private information seller-specific shocks to each country's intercept. Here we explore the value of these shocks implied by the data, when we apply AZ's model and their estimated parameters period-by-period. We will particularly focus on how the implied seller-specific shocks display a high degree of persistence across periods.

Table C.1: Country-Level Summary Statistics and Serial Correlation from the Detrended Oil Data. The mean posterior $v_{i,t}$ is calculated for each month as described below, and the table reports the cross-month average for the country. The AR(1) specifications are: $v'_{i,t} = \rho_i v'_{i,t-1} + \kappa_i + \zeta_{i,t}$ where $v'_{i,t}$ is the mean posterior of $v_{i,t}$ given the data and AZ's estimated parameters. The table reports the estimates of κ_i, ρ_i, the R^2 from the regression and the standard deviation of the estimated residuals $\zeta_{i,t}$. MM = millions.

	Average	Average	AR(1) Regression Results					
	Output	Mean Posterior	Constant (κ_i)	ρ_i	Constant (κ_i)	ρ_i		Innovations
Country	(MM bbl/month)	$v_{i,t}$	Coefficient	Coefficient	Std. Err.	Std. Err.	R^2	Std. Dev.
Algeria	57.51	49.51	0.0528	0.9992	0.4176	0.0084	0.9770	0.1813
Angola	58.04	49.47	0.2977	0.9943	0.4844	0.0098	0.9689	0.1933
Canada	134.05	45.43	1.1831	0.9740	0.5495	0.0121	0.9515	0.3168
China	138.32	45.15	0.6743	0.9853	0.4783	0.0106	0.9632	0.2123
Ecuador	16.93	52.17	0.2602	0.9952	0.4489	0.0086	0.9759	0.1824
Egypt	23.21	51.82	0.1385	0.9976	0.3782	0.0073	0.9826	0.1807
Iran	121.84	46.74	-0.0525	1.0015	0.4516	0.0097	0.9702	0.2544
Iraq	116.88	47.07	1.9427	0.9589	0.7171	0.0152	0.9231	0.5306
Kuwait	90.17	48.87	0.3843	0.9922	0.4457	0.0091	0.9728	0.2030
Libya	40.88	50.63	0.4389	0.9916	0.4433	0.0088	0.9749	0.2578
Mexico	93.05	47.05	-0.0191	1.0007	0.3601	0.0077	0.9810	0.2686
Nigeria	70.72	48.62	0.4011	0.9920	0.5101	0.0105	0.9644	0.2394
Norway	77.37	49.25	0.4580	0.9909	0.4414	0.0090	0.9737	0.3630
Qatar	61.50	50.79	0.4878	0.9906	0.4658	0.0092	0.9725	0.1812
Russia	339.84	39.64	1.6960	0.9579	0.2943	0.0074	0.9806	0.2375
Saudi Arabia	362.79	38.10	1.9062	0.9504	0.6585	0.0173	0.9015	0.4391
U.A.E.	106.57	47.76	0.9333	0.9807	0.4838	0.0101	0.9659	0.2128
U.K.	55.02	50.75	0.1563	0.9972	0.3187	0.0063	0.9871	0.2941
U.S.	363.49	41.10	-0.1764	1.0037	0.2076	0.0050	0.9918	0.4699
Venezuela	80.96	47.86	0.0461	0.9995	0.4530	0.0095	0.9712	0.3524

AZ's Data. We use AZ's raw data. Their output data cover the twenty countries listed in Table C.1 over the period from January 1992 to December 2019. We use AZ's procedure to detrend the output data and deflate prices (deflation is done to Q4 2019).[43] Table C.1 reports mean detrended outputs, in millions of barrels per month.

AZ's Empirical Model. As a reminder, AZ assume that inverse demand in period t is $P_t = u_t - \beta \sum_i q_{i,t}$ and each seller i has quadratic costs $C(q_{i,t}, v_{i,t}, w_t) = (v_{i,t} + w_t)q_{i,t} + \frac{\lambda}{2}q_{i,t}^2$, where u_t and w_t are common demand and cost shocks that are observed by all sellers in period t. The $v_{i,t}$s are private information. AZ make additional distributional assumptions

$$u_t \sim TrN(\mu_u, \sigma_u^2) \text{ on } [\underline{u}, \infty) \text{ with } \underline{u} > 0 \tag{66}$$

where $TrN(\mu_u, \sigma_u^2)$ denotes a truncated normal distribution where the untruncated distribution has mean μ_u and variance σ_u^2.

$$w_t|_{u_t=u} \sim \overline{w} \times (2\text{Beta}(a_w(u), a_w(u)) - 1) \text{ where } a_w(u) = \exp(\tilde{a_1} + \tilde{a_2}u) \text{ and } \tilde{a_1}, \tilde{a_2}, \overline{w} > 0, \tag{67}$$

which should be read as follows: conditional on $u_t = u$, w_t has the same distribution as a Beta random variable, with parameters $a_w(u)$ and $a_w(u)$, that is multiplied by $2\overline{w}$ and then has $\overline{w}$ subtracted. w_t therefore has a support $[-\overline{w}, \overline{w}]$. Using the same reading,

$$v_{i,t} \sim \overline{w} \times (\text{Beta}(a_i, b_i) + 1) \text{ where } a_i, b_i, \overline{w} > 0. \tag{68}$$

[43] AZ's replication package can be downloaded from `https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/KELBWK`. Their detrending procedure uses a genetic algorithm, and we set the random-number seed to `rng(150000031,'twister')` so that we can reproduce our results.

AZ allow a_i and b_i to vary across six country groups, but the supports of the vs, $[\overline{w}, 2\overline{w}]$, are the same for all groups. AZ estimate the parameters ($\underline{u}$, μ_u, σ_u^2, $\tilde{a_1}$, $\tilde{a_2}$, $\overline{w}$ and the a_is and b_is) using maximum likelihood assuming that sellers play a static Bayesian Nash equilibrium with linear quantity strategies. Their estimates imply that the common shocks and the country-level shocks have wide supports (e.g., the vs have supports from 28.206 to 56.412).

Implied Shocks. AZ's static specification treats all shocks as independent across months, which is natural for their one-shot analysis. For our model, it is important to know if the seller-specific cost shocks are persistent. Country-level detrended outputs are highly persistent: the mean estimated AR(1) coefficient from country-level regressions is 0.969 (consistent with AZ's Figures 1 and 4), with coefficients that are broadly similar for both OPEC and non-OPEC members (i.e., persistence is not limited to countries that may be colluding). We can also analyze the serial correlation in the $v_{i,t}$s, and the values of the common shocks, by applying AZ's estimated static model to the data period-by-period. To do so, we use AZ's parameter estimates (AZ Table 3), which we will denote $\widehat{\Theta}$, and the detrended output series ($\mathbf{q}_t$) and the deflated price. The implied value of u_t, which we will denote $\widehat{u_t}$, can be calculated from the inverse demand equation given AZ's estimate of β. Appendix Figure C.1(a) shows that $\widehat{u_t}$ is quite stable through 2004, before increasing and becoming more volatile.

Then, using the estimated distributions as priors, we compute posterior densities for w_t and the $v_{i,t}$s.[44] For any w_t,

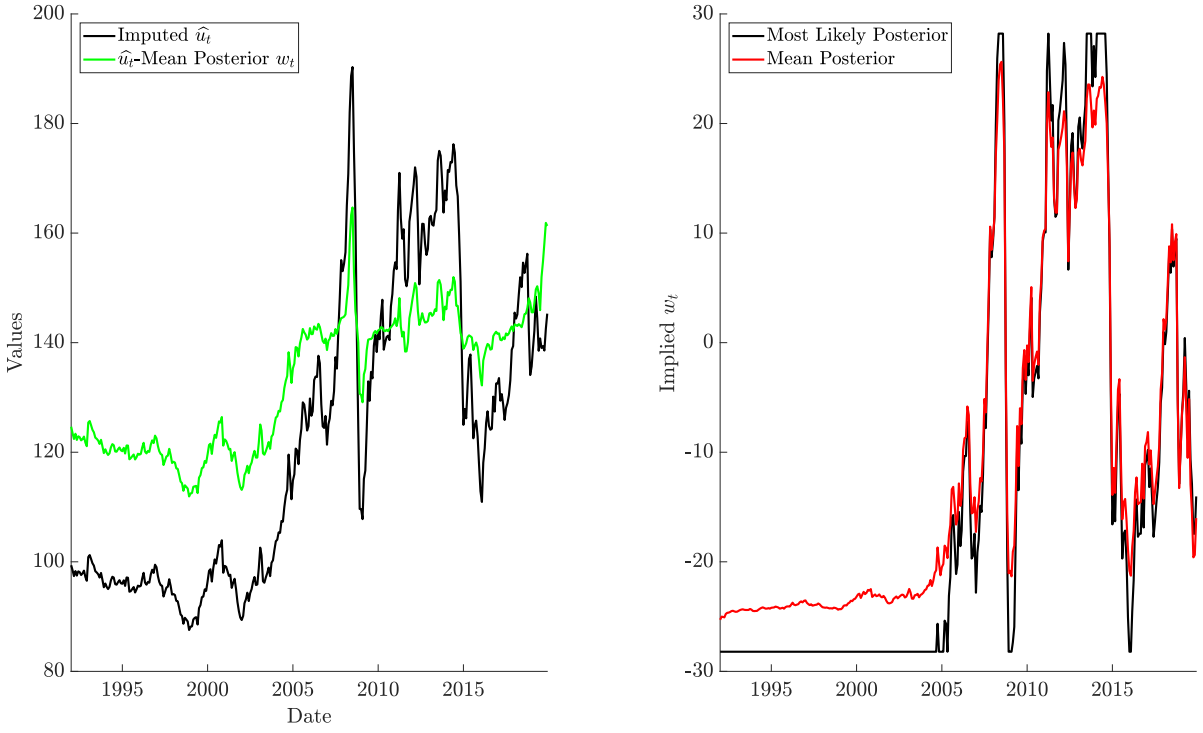
$$g\left(w_t|\widehat{\Theta}, \mathbf{q_t}, \widehat{u_t}\right) = \frac{f_w(w_t|\widehat{u_t}, \widehat{\Theta}) \prod_{i=1}^{N} f_{v(i)}\left(v_{i,t}(\mathbf{q}_t, w_t, \widehat{u_t}, \widehat{\Theta})\right)}{\int_{-\overline{w}}^{\overline{w}} f_w(w_t'|\widehat{u_t}, \widehat{\Theta}) \prod_{i=1}^{N} f_{v(i)}\left(v_{i,t}(\mathbf{q}_t, w_t', \widehat{u_t}, \widehat{\Theta})\right) dw_t'} \tag{69}$$

where f_w and $f_{v(i)}$ are the densities implied by (67) and (68) given $\widehat{u_t}$, the data and parameters,

[44] Note that the priors are formed using $\widehat{\Theta}$, so they do not account for the uncertainty in the estimated parameter values.

Figure C.1: Implied Common Demand Shocks (u_t) and Measures of the Common Cost Shocks (w_t) Given AZ's Estimated Parameters and the Detrended Data. The red line in panel (b) shows the mean posterior value of w_t, and the black line shows the value of w_t for which the posterior density is highest. AZ estimate a support of w_t of [-28.206,28.206] and a support of u_t of [84.443, ∞).

(a) Common Demand Shock (u_t) (b) Posterior Common Cost Shock (w_t)



and $N = 20$ is the number of sellers. Given w_t,

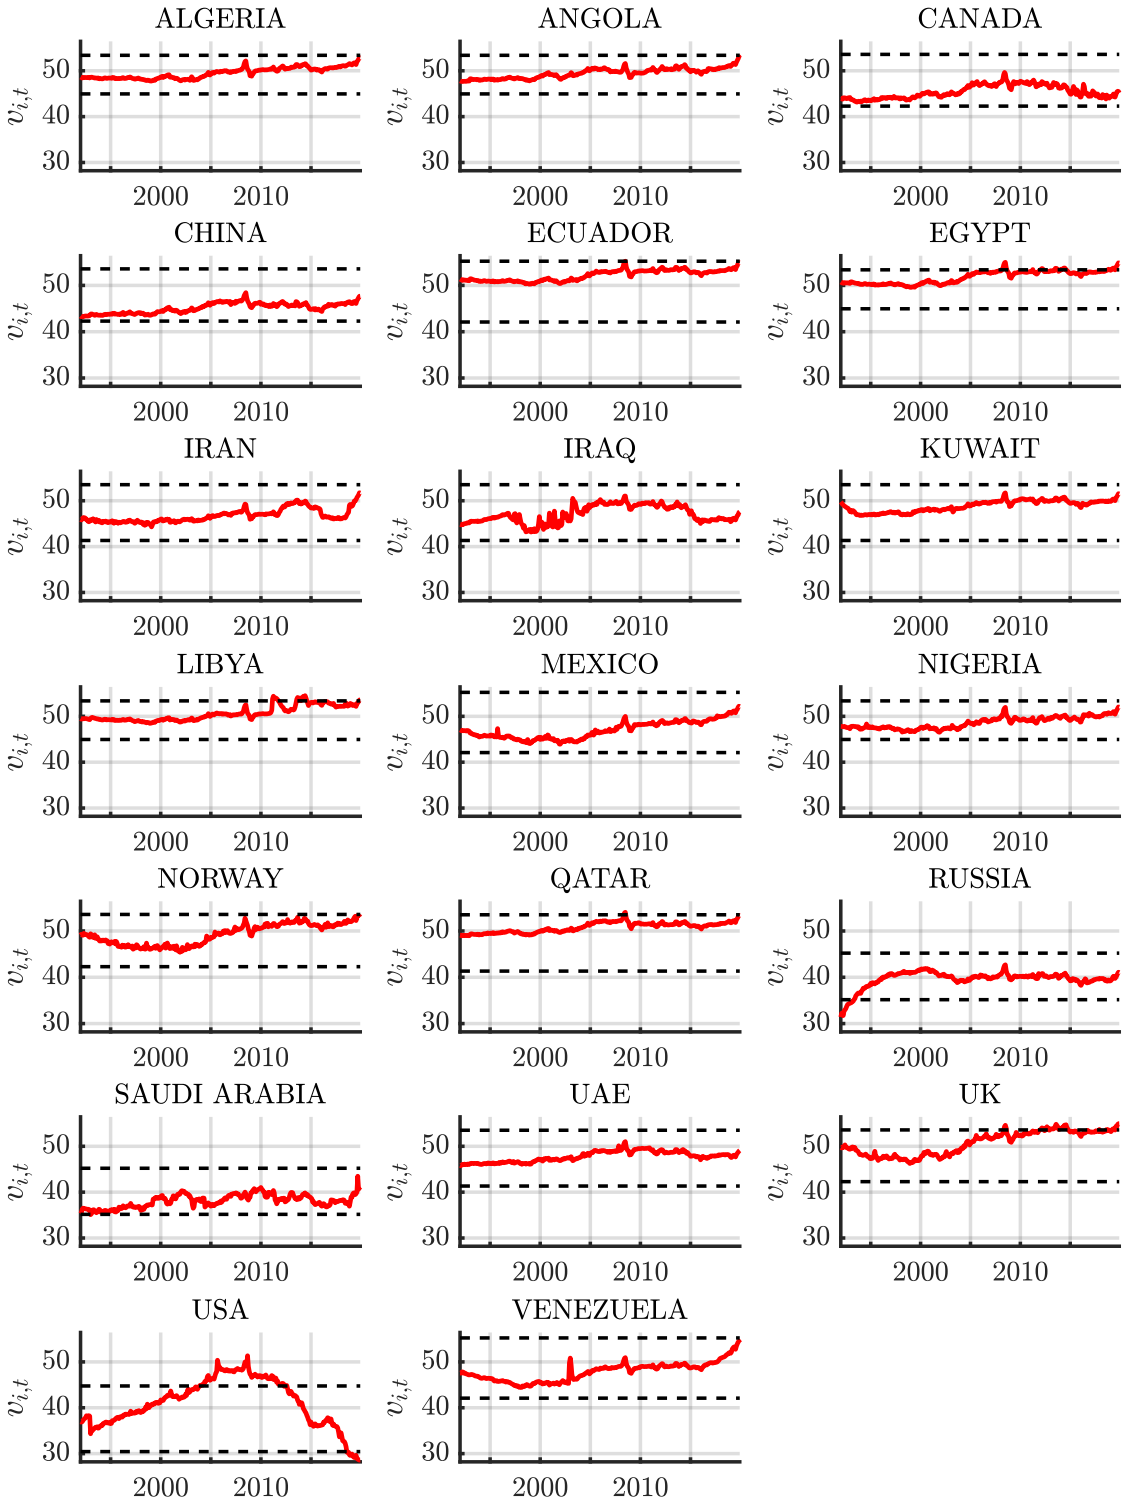
$$v_{i,t}\left(\mathbf{q}_t, w_t, \widehat{u}_t, \widehat{\Theta}\right) = \frac{\widehat{\lambda}+2\widehat{\beta}}{\widehat{\lambda}+(N+1)\widehat{\beta}}\left\{\widehat{u}_t - w_t - \frac{1}{\widehat{\lambda}+\widehat{\beta}}\left[(\widehat{\lambda}+N\widehat{\beta})\widehat{\mu_{v_i}} - \widehat{\beta}\sum_{j\neq i}\widehat{\mu_{v_j}}\right]\right\} + \widehat{\mu_{v_i}} - (\widehat{\lambda}+2\widehat{\beta})q_{i,t}, \tag{70}$$

where $q_{i,t}$ is seller i's detrended output and $\widehat{\mu_{v_i}}$ is the mean of the v_i distribution implied by $\widehat{\Theta}$.

Appendix Figure C.1(b) shows the value of w_t which maximizes g (the calculation is based on 200 equally spaced values of w_t in each period), and the mean of each period's posterior, i.e., $\int_{-\overline{w}}^{\overline{w}} w' g\left(w'|\widehat{\Theta}, \mathbf{q}_t, \widehat{u}_t\right) dw'$. These series are also stable before 2004.

In our application, we will assume that u_t and w_t are fixed. In the static model, with linear

Figure C.2: Mean Posterior Values of $v_{i,t}$ Given AZ's Estimated Parameters and the Detrended Data. The dashed lines show the 10th and 90th percentiles of the unconditional distribution of the $v_{i,t}$s implied by AZ's estimates (i.e., no posteriors). AZ estimate a support of $v_{i,t}$ of [28.206,56.412] for all countries.



strategies, it is the difference between the demand shock u_t and the cost shock w_t that matters for output choices. Appendix Figure C.1(a) shows that because the $\widehat{u_t}$ and mean posterior w_t series are highly (and positively) correlated, the difference between them is also fairly stable over a longer period.[45] The difference between u_t and w_t will also determine the level of profits in the dynamic signaling model. Because the equilibrium incentives to signal depend on how signals change rivals' outputs in the next period, our assumption that there is no variation in $u_t - w_t$ is only likely to matter if, contrary to our assumption, the difference varies significantly from month-to-month, which does not appear to be the case. As discussed in the text, if we assumed different fixed values for u_t and w_t our main results would also not change because the imputed mean level of different sellers' $v_{i,t}$s would adjust to rationalize their average output levels.

The posterior density for w_t can be used to calculate a posterior density for each $v_{i,t}$ using equation (70). The time-series of mean posteriors are shown by the red lines in Appendix Figure C.2 (the time-series of the vs that maximize the posteriors are almost identical and are not shown). For comparison purposes only, the dashed black lines show the 10th and 90th percentiles of the distributions estimated by AZ for the country's $v_{i,t}$s.

These figures are consistent with the $v_{i,t}$s being very persistent across months, and the mean posteriors for individual countries change relatively little over periods of several years or more, with some exceptions associated with periods of conflict (e.g., Iraq in the early 2000s), and the rapid growth of fracking in the U.S. after 2010.

We use the computed mean posteriors ($v'_{i,t}$) to estimate an AR(1) regression ($v'_{i,t} = \rho_i v'_{i,t-1} + \kappa_i + \zeta_{i,t}$) for each country. Table C.1 reports the estimated coefficients from these regressions and the standard deviation of the estimated ζ innovations, as well as the R^2s, the average $v'_{i,t}$ and the country's average detrended output. Consistent with Figure C.2, the estimated ρs are all very close

[45] The correlation between the imputed $\widehat{u_t}$ and the mean posterior w_t is 0.960. This correlation reflects the data, not the parameter estimates, which imply that changes in u_t only affect the density of w_t slightly (see AZ Figure 5 for an illustration).

to, and in many cases not significantly different from, 1. The standard deviations of the implied innovations are small (less than 25 cents for one-half of the countries). The typical changes in the intercepts are also small. For example, for Saudi Arabia, 87.2% of the absolute changes in the mean posterior $v_{i,t}$s are less than 50 cents, and for most countries the proportion of absolute changes less than 50 cents is over 95%. In comparison, AZ's estimates imply that the standard deviation of a $v_{i,t}$ draw is at least \$4 for most countries, with their independence assumption implying that, for most countries, over 90% of absolute changes are more than 50 cents (with most changes more than \$4). The results support our assumption that sellers' marginal costs are persistent, so that, if the $v_{i,t}$s are private information, there should be incentives to signal, although we recognize that other possible causes of output persistence, such as short-run capacity constraints, have not been controlled for.