

Dynamic Oligopoly with Incomplete Information*

Alessandro Bonatti

Gonzalo Cisternas

Juuso Toikka

July 9, 2015

Abstract

We consider signaling and learning dynamics in a Cournot oligopoly where firms have private information about their production costs and only observe the market price, which is subject to unobservable demand shocks. An equilibrium is *Markov* if it depends on the history of play only through the firms' beliefs about costs and calendar time. We characterize symmetric linear Markov equilibria as solutions to a boundary value problem. In every such equilibrium, given a long enough horizon, play converges to the static complete information outcome for the realized costs, but each firm only learns its competitors' average cost. The weights assigned to private and public information under the equilibrium strategies are non-monotone over time. We explain this by decomposing incentives into signaling and learning, and discuss implications for prices, quantities, and profits.

1 Introduction

Despite the role of asymmetric information in explaining specific anti-competitive practices such as limit pricing (Milgrom and Roberts, 1982b) or price rigidity in cartels (Athey and Bagwell, 2008), the question of learning and information aggregation in product markets under incomplete information has received little theoretical attention. Motivated by the recent interest in the industrial organization of such markets (see, *e.g.*, Fershtman and Pakes, 2012; Doraszelski, Lewis, and Pakes, 2014), we provide a theoretical analysis of “Markov perfect oligopolistic competition” in a new market where firms have incomplete information about their rivals' production costs.

*We thank Johannes Hörner, Greg Pavlov, Andy Skrzypacz, Philipp Strack, and Alexander Wolitzky for useful discussions, and the seminar audiences at Northwestern, Stanford, Cowles, and Warwick for comments. Bonatti: Sloan School of Management, MIT, 100 Main Street, Cambridge MA 02142, bonatti@mit.edu. Cisternas: Sloan School of Management, MIT, 100 Main Street, Cambridge MA 02142, gcistern@mit.edu. Toikka: Department of Economics, MIT, 77 Massachusetts Avenue, Cambridge MA 02139, toikka@mit.edu.

More specifically, we study a dynamic Cournot oligopoly where each firm has private information about its cost and only observes the market price, which is subject to unobservable demand shocks. The presence of asymmetric information introduces the possibility of signaling with the resulting dynamics capturing the jockeying for position among oligopolists before the market reaches its long-run equilibrium. We are interested in the allocative and welfare implications of the firms' attempts to signal their costs. For example, how does signaling impact the distribution of market shares, the total industry profits and the consumer surplus over time?

The main challenge in analyzing such settings is keeping track of the firms' beliefs. We therefore consider a linear Gaussian environment, casting the analysis in continuous time. That is, the market demand function and the firms' cost functions are assumed linear in quantities, the constant marginal costs are drawn once and for all from a symmetric normal prior, and the noise in the market price is given by the increments of a Brownian motion. Restricting attention to equilibria in strategies that are linear in the history, we can then derive the firms' beliefs using the Kalman filter.

When costs are private information, a natural way to impose a Markov restriction on behavior is to allow current outputs to depend on the history only through the firms' beliefs about the costs. But when outputs are unobservable, these beliefs are also private information: not observing its output, a firm's rivals cannot tell what inference the firm made from the price. Thus, if the firm plays as a function of its belief—that is, if the belief is part of its “state”—then its rivals have to entertain (private) beliefs about this belief, and so on, making the problem seemingly intractable.¹ However, building on Foster and Viswanathan's (1996) analysis of a multi-agent version of Kyle's (1985) insider trading model, we show that under symmetric linear strategies each firm's belief can be written as the weighted sum of its own cost and the public posterior expectation about the average industry cost conditional on past prices. In particular, own cost and the public belief are sufficient statistics for the firm's beliefs despite there being a non-trivial hierarchy of beliefs. The same is true even if the firm unilaterally deviates from the symmetric linear strategy profile once we appropriately augment these statistics to account for the bias in the public belief caused by the deviation.

The representation of beliefs allows representing all symmetric linear Markov strategies as affine, time-dependent functions of the firm's own cost and the public belief. We consider equilibria in such strategies, and show that they are characterized by solutions to a boundary value problem, which is the key to our analysis.

The boundary value problem characterizing Markov equilibria consists of a system of differential equations for the coefficients of the equilibrium strategy and the posterior variance of the public belief. As is well known, there is no general existence theory for

¹This is the “forecasting the forecasts of others problem” of Townsend (1983).

such problems. Indeed, the biggest technical challenge in our analysis is establishing the existence of a solution to the boundary value problem, or, equivalently, the existence of a symmetric linear Markov equilibrium. We provide a sufficient condition for existence in terms of the model's parameters, which amounts to requiring that the incentive to signal not be too strong. The condition is not tight but not redundant either: a Markov equilibrium may fail to exist if the signaling incentive is sufficiently strong. On the other hand, it turns out that surprisingly much can be said analytically about the properties of Markov equilibria.

Regarding learning and long-run behavior, we show that in every symmetric linear Markov equilibrium, play converges to the static complete information outcome for the realized costs provided that the horizon is taken to be long enough. However, the firms only learn the average cost of their rivals because of the identification problem caused by the *ex ante* symmetry of firms and the one-dimensional market price.

The equilibrium strategy assigns weights to the firm's own cost and the public information that are non-monotone in time. We show that after an appropriate regrouping of terms, this can be understood as arising from learning and signaling. Roughly, equilibrium coefficients under myopic play, which only reflects learning, are monotone over time. Similarly, the coefficients that correspond to the signaling component of a forward-looking player's best response are monotone over time, but in the opposite direction than under myopic play. It is these two opposite monotone effects that sum up to the non-monotone coefficients of the equilibrium strategy.

Signaling results in the expected total quantity being above the corresponding static complete information level, which in turn implies that the expected market price is depressed below its static complete information level. Moreover, at any point in time, consumers only observing historical prices expect prices to increase in the future. This suggests industry profits improve over time. Because firms assign non-monotone weights to their own costs, however, the difference in any two firms' outputs is non-monotone over time conditional on their realized costs. This has implications for productive efficiency. In particular, this may lead the expected profitability of the market to be the highest in the medium run.

Related Literature The early literature on the effects of asymmetric information on dynamic oligopolistic competition mostly considers two period models, often with one-sided private information, focusing on issues such as limit pricing and predation that are complementary to our analysis of learning and signaling over a long horizon. See, *e.g.*, Milgrom and Roberts (1982a), Riordan (1985), Fudenberg and Tirole (1986), or the dynamic analysis of reputation by Milgrom and Roberts (1982b). More related are Mailath (1989) and Mester (1992), who construct separating equilibria in two and

three-period oligopoly games with private costs and observable actions, and Mirman, Samuelson, and Urbano (1993), who provide a theory of signal-jamming in a duopoly with common demand uncertainty. Our model has elements of both as noisy signaling eventually leads to the play converging to the complete information outcome, but in the meantime the equilibrium strategies exhibit signal-jamming.

More recently, Athey and Bagwell (2008) study collusion among patient firms in a repeated Bertrand oligopoly with fully persistent private costs as in our model. Their main interest is in identifying conditions under which the best collusive equilibrium features rigid pricing at the cost of productive efficiency. In such equilibria, all cost types pool at the same price and there is no learning. In contrast, we consider a Cournot oligopoly with a fixed discount rate, and focus on Markov equilibria where firms actively signal their costs and learning is central to the analysis.²

Fershtman and Pakes (2012) consider the steady state of a learning-inspired adjustment process in a dynamic oligopoly with privately observed states in an effort to incorporate incomplete information into the analysis of Markov-perfect industry dynamics (Ericson and Pakes, 1995, Doraszelski and Satterthwaite, 2010, Weintraub, Benkard, and Van Roy, 2008). Their approach is entirely computation-oriented, whereas we apply a standard solution concept and focus on analytical results.³

Our analysis is also related to the literature on information sharing in oligopoly, beginning with Vives (1984) and Gal-Or (1985), and generalized by Raith (1996). More recent contributions to static oligopolistic competition with asymmetric information include, among others, the analysis of supply-function equilibria by Vives (2011) and Bernhardt and Taub (2015).

A large literature studies strategic use of information and its aggregation through prices in financial markets following the seminal analysis by Kyle (1985). Most closely related to our work is the multiagent version of Kyle’s model developed by Foster and Viswanathan (1996) mentioned above, and its continuous-time analog studied by Back, Cao, and Willard (2000). We share their focus on linear equilibria in a Gaussian environment. However, strategic trading in a financial market with common values differs starkly from product market competition under strategic substitutes and private values. In the former, the players limit their trades in order to retain their informational advantage, whereas in the latter, they engage in excess production in

²The focus on Markov equilibria and the use of continuous time methods also distinguishes our work from the literature on repeated Bayesian games with fully or partially persistent types, which has almost exclusively restricted attention to patient players, typically focusing on cooperative equilibria. See, *e.g.*, Aumann and Maschler (1995), Hörner and Lovo (2009), Escobar and Toikka (2013), Peski (2014), or Hörner, Takahashi, and Vieille (forthcoming). There is also a literature on learning in repeated games of incomplete information under myopic play, see, *e.g.*, Nyarko (1997).

³Nevertheless, finding a solution to our system numerically is trivial. Thus, in contrast to the discrete dynamic oligopoly models our problem is computationally much simpler.

an effort to signal their costs and to discourage their rivals, leading to qualitatively different equilibrium behavior. The differences between the games also result in the analysis being technically substantially different.

Finally, Cisternas (2015) develops methods for continuous-time games where learning is common in equilibrium but private beliefs arise after deviations. In contrast, our firms start with private costs and have private beliefs even on the equilibrium path.

Outline The rest of the paper is organized as follows. We setup the model in the next section, and consider the firms' beliefs under linear strategies in Section 3. We then turn to Markov strategies and Markov equilibria in Section 4, and discuss properties of such equilibria in Section 5. We conclude in Section 6 with a discussion of the modeling assumptions and possible extensions. All proofs are in the Appendix.

2 Model

We consider a Cournot game with privately known costs and imperfect monitoring, played in continuous time over the compact interval $[0, T]$. There are $n \geq 2$ firms, each with a privately known (marginal) cost C^i ($i = 1, \dots, n$). The firms' common prior is that the costs are i.i.d. normal random variables with mean π_0 and variance g_0 .⁴

At each time $t \in [0, T]$, each firm i supplies a quantity $Q_t^i \in \mathbb{R}$. The firms do not observe each others' quantities, but observe the (cumulative) price process

$$dY_t = (\bar{p} - \sum_i Q_t^i)dt + \sigma dZ_t, \quad (1)$$

where $\bar{p} > 0$ is the demand intercept, $\sigma^2 > 0$ is the variance, and Z is a standard Brownian motion that is independent of the firms' costs. The current market price is given by the increment of the controlled process Y . The resulting flow payoff to firm i is $Q_t^i(dY_t - C^i dt)$, which the firm discounts at rate $r \geq 0$ common to all firms.

A pure strategy for a firm determines current output as a function of the firm's cost, past prices, and own past outputs. However, because of the noise in the market price, no firm can ever observe that another firm has deviated from a given strategy.⁵ For the analysis of equilibrium outcomes it therefore suffices to know the quantities each firm's strategy specifies at histories that are consistent with the strategy being followed, i.e., on the path play. Thus, abusing terminology, we define a strategy to be only a function of the firm's cost and prices, leaving off-path behavior unspecified.

⁴We discuss the role of assumptions such as finite horizon, symmetry, and independence in Section 6.

⁵As the firms' quantities only affect the drift of Y , the monitoring structure has full support in the sense that any two (admissible) strategy profiles induce equivalent measures over the space of sample paths of Y .

This notion of strategy can be viewed as extending public strategies studied in repeated games with imperfect public monitoring to a setting with private costs.

Formally, a (*pure*) *strategy* for firm i is a process Q^i that is progressively measurable with respect to the filtration generated by (C^i, Y) .⁶ A strategy profile (Q^1, \dots, Q^n) is *admissible* if (i) for each i , $\mathbb{E}[\int_0^T (Q_t^i)^2 dt] < \infty$, in which case we write $Q^i \in L_2[0, T]$, and (ii) equation (1) has a unique solution $Y \in L_2[0, T]$. The expected payoff of firm i under an admissible strategy profile is well defined and given by

$$\mathbb{E} \left[\int_0^T e^{-rt} Q_t^i dY_t - C^i \int_0^T e^{-rt} Q_t^i dt \right] = \mathbb{E} \left[\int_0^T e^{-rt} (\bar{p} - \sum_j Q_t^j - C^i) Q_t^i dt \right]. \quad (2)$$

Payoff from all other strategy profiles is set to $-\infty$. In what follows, a strategy profile is always understood to mean an admissible one unless noted otherwise.

A *Nash equilibrium* is a strategy profile (Q^1, \dots, Q^n) from which no firm has a profitable deviation.⁷ We focus on equilibria in strategies that are linear in histories to facilitate tractable updating of beliefs, but we allow firms to contemplate deviations to arbitrary strategies. Formally, firm i 's strategy Q^i is *linear* if there exist (Borel measurable) functions $\alpha, \delta : [0, T] \rightarrow \mathbb{R}$ and $f : [0, T]^2 \rightarrow \mathbb{R}$ such that⁸

$$Q_t^i = \alpha_t C^i + \int_0^t f_s^i dY_s + \delta_t, \quad t \in [0, T]. \quad (3)$$

A profile of linear strategies is *symmetric* if the functions (α, f, δ) are the same for all firms. Our interest is in Nash equilibria in symmetric linear strategies that condition on the history only through its effect on the firms' beliefs about the cost vector (C^1, \dots, C^n) and calendar time. Such equilibria, defined formally below, are a natural extension of Markov perfect equilibrium to our model.

⁶More precisely, let \mathcal{B}^i be the sigma-algebra on \mathbb{R} generated by C^i , and let $\mathcal{F} = \{\mathcal{F}_t\}$ be the filtration on $C[0, T]$, the space of continuous functions on $[0, T]$, where each \mathcal{F}_t is generated by sets $\{f \in C[0, T] : f_s \in \Gamma\}$, where $s \leq t$ and Γ is a Borel set in \mathbb{R} . (Heuristically, \mathcal{F} corresponds to observing the past of the process Y .) A strategy is a process Q^i that is progressively measurable with respect to $\bar{\mathcal{F}}^i := \{\bar{\mathcal{F}}_t^i\}$, where $\bar{\mathcal{F}}_t^i := \mathcal{B}^i \otimes \mathcal{F}_t$.

⁷The best-response problem against a profile Q^{-i} of other players' strategies can be viewed as a stochastic control problem with a partially observable state (see, *e.g.*, Davis and Varaiya, 1973, or the general formulation in Davis, 1979). In particular, any admissible strategy profile (Q^i, Q^{-i}) induces a probability measure $\mathbb{P}^{(Q^i, Q^{-i})}$ such that (1) holds. A deviation to any strategy \hat{Q}^i such that (\hat{Q}^i, Q^{-i}) is admissible amounts to changing the measure to $\mathbb{P}^{(\hat{Q}^i, Q^{-i})}$ (*cf.* Sannikov, 2007); payoff from other deviations is $-\infty$.

⁸A necessary and sufficient condition for a linear strategy profile to be admissible is that all the functions α , δ , and f be square-integrable over their respective domains (Kallianpur, 1980, Theorem 9.4.2). Note that in discrete time, any affine function of own cost and past prices takes the form $q_t^i = \alpha_t c_i + \sum_{s < t} f_s^i (y_s - y_{s-1}) + \delta_t$. Equation (3) can be viewed as a limit of such strategies.

3 Beliefs under Linear Strategies

As a step towards Markov equilibria, we derive sufficient statistics for the firms' beliefs about costs under symmetric linear strategies and unilateral deviations from them.

Fix firm i , and suppose the other firms are playing symmetric linear strategies so that $Q_t^j = \alpha_t C^j + B_t(Y^t)$ for $j \neq i$, where $B_t(Y^t) := \int_0^t f_s^t dY_s + \delta_t$. Regardless of its own strategy, firm i can always subtract the effect of its own quantity and that of the public component $B_t(Y^t)$ of the other firms' quantities on the price, and hence the relevant signal for firm i about C^j , $j \neq i$, is

$$dY_t^i := -\alpha_t \sum_{j \neq i} C^j dt + \sigma dZ_t = dY_t - (\bar{p} - Q_t^i - (n-1)B_t(Y^t))dt. \quad (4)$$

Therefore, firm i 's belief about the other firms' costs can be derived by applying the Kalman filter with Y^i as the signal and C^{-i} as the unknown vector. Moreover, since the other firms are *ex ante* symmetric and play symmetric strategies, firm i can only ever hope to filter the sum of their costs. The following lemma formalizes these observations.

Lemma 1. *Under any symmetric linear strategy profile and any strategy of firm i , firm i 's posterior at $t \in [0, T]$ is that C^j , $j \neq i$, are jointly normal, each with mean $M_t^i := \frac{1}{n-1} \mathbb{E}[\sum_{j \neq i} C^j | \mathcal{F}_t^{Y^i}]$, and with a symmetric covariance matrix $\Gamma_t = \Gamma(\gamma_t^M)$, where the function $\Gamma : \mathbb{R} \rightarrow \mathbb{R}^{2(n-1)}$ is independent of t , and*

$$\gamma_t^M := \mathbb{E} \left[\left(\sum_{j \neq i} C^j - (n-1)M_t^i \right)^2 \middle| \mathcal{F}_t^{Y^i} \right] = \frac{(n-1)g_0}{1 + (n-1)g_0 \int_0^t (\frac{\alpha_s}{\sigma})^2 ds}$$

is a deterministic nonincreasing function of t .

The upshot of Lemma 1 is that firm i 's belief is summarized by the pair (M_t^i, γ_t^M) . The expectation about the other firms' average cost, M_t^i , is firm i 's private information as the other firms do not observe i 's quantity and hence do not know what inference it made. (Formally, Q^i enters Y^i .) The posterior variance γ_t^M is a deterministic function of time because the function α in the other firms' strategy is taken as given.

By Lemma 1, asking symmetric linear strategies to condition on history only through beliefs amounts to requiring each firm i 's output at time t to only depend on C^i , M_t^i , and t . From the perspective of the normal form of the game, this is simply a measurability requirement on the firms' strategies, and causes no immediate problems. However, showing the existence of a Nash equilibrium in strategies of this form requires verifying the optimality of the strategies to each firm, and for this it is essentially necessary to use dynamic optimization. But formulating firm i 's best-response problem as a dynamic optimization problem, we then have M^j , $j \neq i$, appearing as unobservable states in firm i 's problem, and we thus need to consider i 's second-order

beliefs about them. Indeed, it could even be the case that firm i 's best response then has to explicitly condition on these second-order beliefs, requiring them to be added to the state, and so on, leading to an infinite regress problem.

It turns out, however, that for linear Gaussian models there is an elegant solution, first applied to a strategic setting by Foster and Viswanathan (1996). The key observation is that each firm's private belief can be expressed as a weighted sum of its own cost and the public belief about the average cost conditional on past prices. Thus, even when the other firms' behavior conditions on their beliefs, firm i only needs to have a belief about their costs as the public belief is public information. Firm i 's belief in turn is just a function of its cost and the public belief, closing the loop.

More specifically, consider the posterior expectation about the average firm cost conditional on the price process Y under a symmetric linear strategy profile. This *public belief* is defined as $\Pi_t := \frac{1}{n} \mathbb{E}[\sum_j C^j | \mathcal{F}_t^Y]$, with corresponding posterior variance $\gamma_t := \mathbb{E}[(\sum_j C^j - n\Pi_t)^2 | \mathcal{F}_t^Y]$.⁹ It can be computed using the Kalman filter with Y as the signal and the sum $\sum_j C^j$ as the unknown parameter (see Lemma A.1 in the Appendix), and it corresponds to the belief of an outsider who knows the strategy, but only observes the prices (*cf.* market makers in Foster and Viswanathan, 1996). We note for future reference that the posterior variance of the public belief is a deterministic function of time given by

$$\gamma_t = \frac{ng_0}{1 + ng_0 \int_0^t (\frac{\alpha_s}{\sigma})^2 ds}, \quad t \in [0, T]. \quad (5)$$

The public belief can be used to express private beliefs as follows.

Lemma 2. *Under any symmetric linear strategy profile, for each firm i ,*

$$M_t^i = z_t \Pi_t + (1 - z_t) C^i, \quad t \in [0, T],$$

where

$$z_t := \frac{n}{n-1} \frac{\gamma_t^M}{\gamma_t} = \frac{n^2 g_0}{n(n-1)g_0 + \gamma_t} \in \left[1, \frac{n}{n-1}\right] \quad (6)$$

is a deterministic nondecreasing function of t .

That is, on the path of play of a symmetric linear strategy profile, each firm's private belief M_t^i is a weighted average of the public belief Π_t and its cost C^i , with the weight z_t a deterministic function of time. Heuristically, C^i captures the firm's private information about both its cost and its past outputs (whose private part equals $\alpha_s C^i$ at time s), and hence it is the only additional information the firm has compared to an outsider observing prices. The functional form comes from the properties of normal

⁹We use the posterior variance of $n\Pi_t$ for notational convenience.

distributions, since under linear strategies the system is Gaussian. Moreover, since γ_t^M is also only a function of time by Lemma 1, the tuple (C^i, Π_t, t) is a sufficient statistic for firm i 's posterior belief at time t .¹⁰

If firm i unilaterally deviates, then the formula in Lemma 2 does not apply to its belief because the public belief Π_t assumes that all firms play the same linear strategy. (The formula still holds for the other firms, because they do not observe the deviation.) At such off path histories, it is convenient to represent firm i 's belief in terms of a counterfactual public belief, which corrects for the difference in firm i 's quantities, and which coincides with Π_t if i has not deviated.

Lemma 3. *Under any symmetric linear strategy profile and any strategy of firm i ,*

$$M_t^i = z_t \hat{\Pi}_t^i + (1 - z_t)C^i, \quad t \in [0, T],$$

where z_t is as in Lemma 2, and the process $\hat{\Pi}^i$ is defined by

$$d\hat{\Pi}_t^i = \lambda_t \alpha_t (1 + (n - 1)(1 - z_t)) (\hat{\Pi}_t^i - C^i) dt + \lambda_t \sigma dZ_t^i, \quad \hat{\Pi}_0^i = \pi_0,$$

where

$$\lambda_t := -\frac{\alpha_t \gamma_t}{n\sigma^2}, \quad \text{and} \quad dZ_t^i := \frac{dY_t^i + (n - 1)\alpha_t (z_t \hat{\Pi}_t^i + (1 - z_t)C^i) dt}{\sigma}$$

is a standard Brownian motion (with respect to \mathcal{F}^{Y^i}) called firm i 's innovation process. Moreover, if firm i plays on $[0, t)$ the same strategy as the other firms, then $\hat{\Pi}_t^i = \Pi_t$.

The counterfactual public belief $\hat{\Pi}^i$ evolves independently of firm i 's strategy by construction. However, it is defined in terms of the process Y^i defined in (4), and hence its computation requires knowledge of firm i 's past quantities. Thus $\hat{\Pi}_t^i$ is in general firm i 's private information. Nevertheless, if firm i plays the same strategy as the other firms, then the counterfactual and actual public beliefs coincide (*i.e.*, $\hat{\Pi}_t^i = \Pi_t$) and we obtain Lemma 2 as a special case. In general, however, firm i 's posterior at time t is captured by $(C^i, \hat{\Pi}^i, t)$.¹¹

¹⁰In fact, each firm i 's entire time- t hierarchy of beliefs is captured by (C^i, Π_t, t) . For example, firm i 's first-order belief about firm j 's cost C^j is normal with mean $z_t \Pi_t + (1 - z_t)C^i$ and variance a function of γ_t^M , where z_t and γ_t^M are only functions of t . Thus to find, say, firm k 's second-order belief about firm i 's first-order belief about C^j , we only need k 's first-order belief about C^i because (Π_t, t) are public. But k simply believes that C^i is normal with mean $z_t \Pi_t + (1 - z_t)C^k$ and variance a function of γ_t^M . And so on.

¹¹If firm i has deviated from the symmetric linear strategy profile, then its time- t hierarchy of beliefs is captured by $(C^i, \Pi_t, \hat{\Pi}_t^i, t)$: its first-order belief uses $\hat{\Pi}_t^i$ instead of Π_t , but since each firm $j \neq i$ still forms its (now biased) beliefs using (C^j, Π_t, t) , Π_t is needed for the computation of higher order beliefs.

4 Markov Equilibria

In games of complete information, a Markov (perfect) equilibrium requires behavior to only depend on the payoff-relevant part of history. In our model, only the costs and calendar time are directly payoff relevant, but because the firms do not know each others' costs, it is in general necessary to let behavior to depend on the history through its effect on the firms' beliefs about costs. Our Markov restriction is to not allow any more history dependence than that.

With this motivation, we say that a strategy profile is *Markov* if each firm's strategy depends on the history only through calendar time and the firm's belief about the cost vector (C^1, \dots, C^n) . Based on our analysis in Section 3, we have the following characterization of symmetric linear Markov strategies.

Lemma 4. *A symmetric linear strategy profile is Markov if and only if there exist functions $\alpha, \beta, \delta : [0, T] \rightarrow \mathbb{R}$, called the coefficients, such that for each firm i ,*

$$Q_t^i = \alpha_t C^i + \beta_t \Pi_t + \delta_t, \quad t \in [0, T].$$

That a strategy of this form only conditions on calendar time and firm i 's belief about costs (including its own) is immediate from the fact that i 's belief is summarized by (C^i, Π_t, t) . The other direction combines this representation of beliefs with the observation that Π_t is itself a linear function of history, and hence for a strategy conditioning on it to be linear in the sense of (3), it has to take the above form.¹²

We then define our notion of Markov equilibrium as follows.

Definition 1. *A symmetric linear Markov equilibrium is a Nash equilibrium in symmetric linear strategies such that (i) the strategy profile is Markov, and (ii) the coefficients (α, β, δ) of the equilibrium strategy are continuously differentiable.*

We identify a symmetric linear Markov equilibrium with the coefficients (α, β, δ) of the equilibrium strategy. Their differentiability is included in the above definition to avoid having to keep repeating it as a qualifier in what follows.

We do not require perfection in the definition of Markov equilibria, since given the full support of the price process Y , the only off-path histories at which a firm can find itself are those that follow its own deviations. Thus, requiring perfection would not restrict the set of equilibrium outcomes. Nevertheless, we obtain a partial description of optimal off-path behavior in our best-response analysis, to which we turn next.

¹²Since our strategies only prescribe behavior on the path of play, the observation in footnote 10 implies that Lemma 4 continues to hold verbatim if "firm's belief" is replaced with "firm's hierarchy of beliefs" in the definition of a Markov strategy profile.

4.1 Best-Response Problem

In order to characterize existence and properties of Markov equilibria, we now explicitly formulate firm i 's best-response problem to a symmetric linear Markov strategy profile as a dynamic stochastic optimization problem.

To this end, fix firm i , and suppose the other firms play a symmetric linear Markov strategy profile (α, β, δ) with differentiable coefficients. We observe first that the tuple $(C^i, \Pi_t, \hat{\Pi}_t^i, t)$ is the relevant *state* for firm i 's problem. To see this, note that the integrand in the expected payoff (2) is linear in the other firms' outputs, and hence firm i 's flow payoff at time t depends only on the other firms' expected output conditional on i 's information. By Lemmas 1 and 4, this is given by $(n-1)(\alpha_t M_t^i + \beta_t \Pi_t + \delta_t)$, where the private belief satisfies $M_t^i = z_t \hat{\Pi}_t^i + (1-z_t)C^i$ by Lemma 3. Furthermore, the coefficients (α, β, δ) and the weight z are deterministic functions of time (as are γ and λ that appear in the laws of motion for Π and $\hat{\Pi}^i$). Thus $(C^i, \Pi_t, \hat{\Pi}_t^i, t)$ fully summarizes the state of the system.

Using the state $(C^i, \Pi_t, \hat{\Pi}_t^i, t)$, the normal form of firm i 's best-response problem can be written as

$$\sup_{Q^i \in L_2[0, T]} \mathbb{E} \left[\int_0^T e^{-rt} [\bar{p} - Q_t^i - (n-1)(\alpha_t M_t^i + \beta_t \Pi_t + \delta_t) - C^i] Q_t^i dt \right]$$

subject to

$$\begin{aligned} d\Pi_t &= \lambda_t [(\alpha_t + \beta_t)\Pi_t + \delta_t - Q_t^i + (n-1)\alpha_t(\Pi_t - M_t^i)]dt + \lambda_t \sigma dZ_t^i, & \Pi_0 &= \pi_0, \\ d\hat{\Pi}_t^i &= \lambda_t [\alpha_t(\hat{\Pi}_t^i - C^i) + (n-1)\alpha_t(\hat{\Pi}_t^i - M_t^i)]dt + \lambda_t \sigma dZ_t^i, & \hat{\Pi}_0^i &= \pi_0, \\ M_t^i &= z_t \hat{\Pi}_t^i + (1-z_t)C^i. \end{aligned}$$

The only sources of randomness in the problem are the initial draw of C^i and firm i 's innovation process Z^i defined in Lemma 3, which is a standard Brownian motion.

The law of motion of the public belief Π is simply the dynamic from Lemma A.1 written from firm i 's perspective.¹³ Conditional on prices, Π is a martingale, but from i 's perspective it has a drift, which consist of two components. The first component, $(\alpha_t + \beta_t)\Pi_t + \delta_t - Q_t^i$, captures the difference between the public expectation of firm i 's output and firm i 's actual output. The second, $(n-1)\alpha_t(\Pi_t - M_t^i)$, captures the

¹³Noting that under Markov strategies, $B_t(Y^t) = \beta_t \Pi_t + \delta_t$, we have by Lemma A.1 and equation (4),

$$\begin{aligned} d\Pi_t &= \lambda_t [dY_t - (\bar{p} - \alpha_t n \Pi_t - n B_t(Y^t))]dt \\ &= \lambda_t [dY_t^i + (\alpha_t n \Pi_t + \beta_t \Pi_t + \delta_t - Q_t^i)dt] \\ &= \lambda_t [\sigma dZ_t^i + (\alpha_t n \Pi_t + \beta_t \Pi_t + \delta_t - Q_t^i - \alpha_t(n-1)M_t^i)dt], \end{aligned}$$

where the last step is by definition of the innovation process $dZ^i := \sigma^{-1}[dY^i + (n-1)\alpha_t M_t^i dt]$ in Lemma 3.

difference between the public's and firm i 's expectations about the other firms' outputs due to firm i 's superior information about their costs. Since Q^i enters the drift, firm i controls the public belief Π . This allows the firm to (noisily) signal its cost and makes the problem dynamic.

The other stochastically evolving state variable, the counterfactual public belief $\hat{\Pi}^i$, evolves exogenously. (Its law of motion is given in Lemma 3.) The interpretation of its drift is the same as that of Π , except that $\hat{\Pi}^i$ is calculated assuming that firm i plays the strategy (α, β, δ) and hence the difference in its expected and realized quantity is just $\alpha_t(\hat{\Pi}_t^i - C^i)$. Note that $d(\Pi_t - \hat{\Pi}_t^i) = \lambda_t[\alpha_t n(\Pi_t - \hat{\Pi}_t^i) + \alpha_t C^i + \beta_t \Pi_t + \delta_t - Q_t^i]dt$, from which it is immediate that $\Pi_t = \hat{\Pi}_t^i$ if firm i has indeed played according to (α, β, δ) in the past.

Firm i 's best-response problem can be formulated recursively as the Hamilton-Jacobi-Bellman (HJB) equation

$$rV(c, \pi, \hat{\pi}, t) = \sup_{q \in \mathbb{R}} \left\{ [\bar{p} - q - (n-1)(\alpha_t(z_t \hat{\pi} + (1-z_t)c) + \beta_t \pi + \delta_t) - c]q + \mu_t(q) \frac{\partial V}{\partial \pi} + \hat{\mu}_t \frac{\partial V}{\partial \hat{\pi}} + \frac{\partial V}{\partial t} + \frac{\lambda_t^2 \sigma^2}{2} \left(\frac{\partial^2 V}{\partial \pi^2} + 2 \frac{\partial^2 V}{\partial \pi \partial \hat{\pi}} + \frac{\partial^2 V}{\partial \hat{\pi}^2} \right) \right\}, \quad (7)$$

where the drifts of Π and $\hat{\Pi}^i$ are, as above,

$$\begin{aligned} \mu_t(q) &:= \lambda_t [(\alpha_t + \beta_t)\pi + \delta_t - q + (n-1)\alpha_t(\pi - (z_t \hat{\pi} + (1-z_t)c))], \\ \hat{\mu}_t &:= \lambda_t \alpha_t [1 + (n-1)(1-z_t)](\hat{\pi} - c), \end{aligned}$$

written here using Lemma 3 to express firm i 's belief as $z_t \hat{\pi} + (1-z_t)c$.

The objective function in the maximization problem on the right-hand side of (7) is linear-quadratic in q with $-q^2$ the only quadratic term, and thus it is strictly concave. Therefore, there is a unique maximizer $q^*(c, \pi, \hat{\pi}, t)$ given by the first-order condition

$$q^*(c, \pi, \hat{\pi}, t) = \frac{\bar{p} - (n-1)[\alpha_t(z_t \hat{\pi} + (1-z_t)c) + \beta_t \pi + \delta_t] - c}{2} - \frac{\lambda_t}{2} \frac{\partial V}{\partial \pi}, \quad (8)$$

where the first term is the myopic best response, and the second term captures the dynamic incentive to signal.

It is worth noting that here continuous time greatly simplifies the analysis. The same arguments can be used in discrete time to derive a Bellman equation analogous to (7). The public belief still enters the flow payoff linearly, so the value function is convex in π . However, the quantity q then affects the *level* of π linearly, which means that the optimization problem in the Bellman equation has a convex term in q . Moreover, this term involves the value function—an endogenous object—which makes

it hard to establish the existence and uniqueness of an optimal quantity. In contrast, in continuous time q only affects the drift of π , which in turn affects the value of the problem linearly. This renders the HJB equation strictly concave in q by inspection.

4.2 Characterization

We can view any symmetric linear Markov equilibrium as a solution to the HJB equation (7) satisfying the fixed point requirement that the optimal policy coincide with the strategy to which the firm is best responding. This leads to a boundary value problem characterization of such equilibria, which is the key to our analysis.

More specifically, we proceed as follows. We show first that if (α, β, δ) is a symmetric linear Markov equilibrium, then the solution to the HJB equation (7) is a (continuation) value function of the form

$$V(c, \pi, \hat{\pi}, t) = v_0(t) + v_1(t)\pi + v_2(t)\hat{\pi} + v_3(t)c + v_4(t)\pi\hat{\pi} + v_5(t)\pi c + v_6(t)\hat{\pi}c + v_7(t)c^2 + v_8(t)\pi^2 + v_9(t)\hat{\pi}^2 \quad (9)$$

for some differentiable $v_k : \mathbb{R} \rightarrow \mathbb{R}$, $k = 0, \dots, 9$, and a linear optimal policy exists on and off the path of play.¹⁴ Substituting for $\partial V/\partial \pi$ in the first-order condition (8) using (9), we see that the best response to the equilibrium strategy can be written as

$$q^*(c, \pi, \hat{\pi}, t) = \alpha_t^* c + \beta_t^* \pi + \delta_t^* + \xi_t^* (\hat{\pi} - \pi).$$

The fixed point requirement is thus simply that $(\alpha^*, \beta^*, \delta^*) = (\alpha, \beta, \delta)$.

The off-path coefficient ξ^* is a free variable given our focus on Nash equilibria, but this argument shows that optimal off-path behavior is necessarily linear, and that a best response exists on and off the path of play.

After imposing the fixed point, the HJB equation (7) reduces to a system of ordinary differential equations (ODEs) for the coefficients v_k of the value function V and the posterior variance γ . However, it turns out to be more convenient to consider an equivalent system of ODEs for γ and the coefficients $(\alpha, \beta, \delta, \xi)$ of the optimal policy along with the relevant boundary conditions. This identifies symmetric linear Markov equilibria with solutions to a boundary value problem. A verification argument establishes the converse.

¹⁴The proof uses the fact that the best-response problem is a stochastic linear-quadratic regulator (see, *e.g.*, Yong and Zhou, 1999, Chapter 6). Note that the posterior variance γ_t depends non-linearly on the coefficient α , and so do the weight z_t and the sensitivity of the public belief to the price, $\lambda_t = -\alpha_t \gamma_t / (n\sigma^2)$. Hence, even though the best-response problem is linear-quadratic because it takes α as given, our game is not a linear-quadratic game in the sense of the literature on differential games (see, *e.g.*, Friedman, 1971).

For a formal statement, define the functions $\alpha^m, \beta^m, \delta^m, \xi^m : \mathbb{R} \rightarrow \mathbb{R}$ by

$$\begin{aligned}\alpha^m(x) &:= -\frac{(n-1)ng_0 + x}{(n-1)ng_0 + (n+1)x}, & \delta^m(x) &:= \frac{\bar{p}}{n+1}, \\ \beta^m(x) &:= \frac{(n-1)n^2g_0}{(n+1)[(n-1)ng_0 + (n+1)x]}, & \xi^m(x) &:= \frac{(n-1)n^2g_0}{2[(n-1)ng_0 + (n+1)x]}.\end{aligned}\tag{10}$$

In the proof of the following result, we show that these are the myopic equilibrium coefficients given posterior variance x . In particular, firm i 's time- T output under the equilibrium best-response policy is $Q_T^i = \alpha^m(\gamma_T)C^i + \beta^m(\gamma_T)\Pi_T + \delta^m(\gamma_T) + \xi^m(\gamma_T)(\hat{\Pi}_T^i - \Pi_T)$.

Recalling from (6) that z_t is only a function of the current γ_t , we have the following characterization of equilibria.

Theorem 1. *(α, β, δ) is a symmetric linear Markov equilibrium with posterior variance γ if and only if $\delta = -\bar{p}(\alpha + \beta)$ and there exists ξ such that $(\alpha, \beta, \xi, \gamma)$ is a solution to*

$$\dot{\alpha}_t = r(\alpha_t - \alpha^m(\gamma_t))\frac{\alpha_t}{\alpha^m(\gamma_t)} - \frac{\alpha_t^2\beta_t\gamma_t[(n-1)n\alpha_t(z_t-1) + 1]}{n\sigma^2},\tag{11}$$

$$\begin{aligned}\dot{\beta}_t &= r(\beta_t - \beta^m(\gamma_t))\frac{\alpha_t}{\alpha^m(\gamma_t)} \\ &+ \frac{\alpha_t\beta_t\gamma_t[n\alpha_t(n+1 - (n-1)z_t - (n^2-1)\beta_t(z_t-1)) + (n-1)\beta_t]}{n(n+1)\sigma^2},\end{aligned}\tag{12}$$

$$\begin{aligned}\dot{\xi}_t &= r(\xi_t - \xi^m(\gamma_t))\frac{\alpha_t}{\alpha^m(\gamma_t)} \\ &+ \frac{\alpha_t\gamma_t\xi_t}{n\sigma^2}[\xi_t - (n\alpha_t((n-1)\beta_t(z_t-1) - 1) + \beta_t)] - \frac{(n-1)\alpha_t^2\beta_t\gamma_tz_t}{2\sigma^2},\end{aligned}\tag{13}$$

$$\dot{\gamma}_t = -\frac{\alpha_t^2\gamma_t^2}{\sigma^2},\tag{14}$$

with boundary conditions $\alpha_T = \alpha^m(\gamma_T)$, $\beta_T = \beta^m(\gamma_T)$, $\xi_T = \xi^m(\gamma_T)$, and $\gamma_0 = ng_0$.

In particular, such an equilibrium exists if and only if the above boundary value problem has a solution. A sufficient condition for existence is

$$\frac{g_0}{\sigma^2} < \max\left\{\frac{r}{\kappa(n)}, \frac{1}{3nT}\right\},\tag{15}$$

where the function $\kappa : \mathbb{N} \rightarrow \mathbb{R}_{++}$ defined in (A.10) satisfies $\kappa(n) \leq n - 2 + \frac{1}{n}$ for all n .

The derivation of the boundary value problem for $(\alpha, \beta, \xi, \gamma)$ proceeds along the lines sketched above. This is the standard argument for characterizing solutions to HJB equations, save for the facts that (i) here we are simultaneously looking for a fixed point, and hence also the flow payoff is determined endogenously as it depends on the strategy played by the other firms, and (ii) we derive a system of differential equations for the optimal policy rather than for the value function.

The identity $\delta = -\bar{p}(\alpha + \beta)$ provides a surprising, but very welcome, simplification

for equilibrium analysis, and allows us to eliminate δ from the boundary value problem. A similar relationship holds in a static Cournot oligopoly with complete information and asymmetric costs.¹⁵ We establish the result by direct substitution into the ODE for δ . Since it is an equilibrium relationship, it doesn't seem possible to establish it by only considering the best-response problem even in a static model.

The hard part in the proof of Theorem 1 is establishing existence. This requires showing the existence of a solution to the nonlinear boundary value problem defined by equations (11)–(14) and the relevant boundary conditions. As is well known, there is no general existence theory for such problems. We thus have to use *ad hoc* arguments, which require detailed study of the system's behavior. On the upside, we obtain as a by-product a relatively complete description of equilibrium behavior, which we discuss in the next section. However, due to the complexity of the system, we have not been able to prove or disprove uniqueness, even though numerical analysis strongly suggests that a symmetric linear Markov equilibrium is unique whenever it exists. (All the results to follow apply to every such equilibrium.)

Our existence proof can be sketched as follows. As ξ only enters equation (13), it is convenient to first omit it from the system and establish existence for the other three equations. For this we use the so-called shooting method. That is, we choose a time- T value for γ , denoted γ_F (mnemonic for final). This determines the time- T values of α and β by $\alpha_T = \alpha^m(\gamma_F)$ and $\beta_T = \beta^m(\gamma_F)$. We then follow equations (11), (12), and (14) backwards in time from T to 0. This gives some γ_0 , provided that none of the three equations diverges before time 0. Thus we need to show that γ_F can be chosen such that there exists a global solution to (11), (12), and (14) on $[0, T]$, and the resulting γ_0 satisfies $\gamma_0 = n\gamma_0$. For the latter, note that we have $\gamma_0 \geq \gamma_F$ since $\dot{\gamma} \leq 0$. Furthermore, setting $\gamma_F = 0$ yields $\gamma_0 = 0$. As the system is continuous in the terminal value γ_F , this implies that the boundary condition for γ_0 is met for some $\gamma_F \in (0, n\gamma_0]$. The sufficient condition given in the theorem ensures that α and β remain bounded as we vary γ_F in this range.

The proof is completed by showing that there exists a solution on $[0, T]$ to equation (13), viewed as a quadratic first-order ODE in ξ with time-varying coefficients given by the solution (α, β, γ) to the other three equations. We use a novel approach where we first establish the existence of ξ , and hence of equilibria, for g_0 sufficiently small, in which case the system resembles the complete information case. We then observe that if ξ is the first to diverge as g_0 approaches some \bar{g}_0 from below, then some of the coefficients of the equilibrium value function V in (9) diverge. This allows us to construct a non-local deviation that is profitable for g_0 close enough to \bar{g}_0 and hence

¹⁵For example, given $n = 2$ and demand $p = \bar{p} - q_1 - q_2$, if we define $\pi = (c_1 + c_2)/2$, then the equilibrium quantities are $q_i = ac_i + b\pi + d$ ($i = 1, 2$), where $a = -1$, $b = 2/3$, and $d = \bar{p}/3$, and hence $d = -\bar{p}(a + b)$.

contradicts the existence of an equilibrium for all $g_0 < \bar{g}_0$.

The sufficient condition (15) for existence in Theorem 1 is not tight. Numerical analysis suggests that equilibria exist for parameters in a somewhat larger range. However, the condition is not redundant either. For example, it is possible to prove that, given any values for the other parameters, if $r = 0$, then there exists a sufficiently large but finite \bar{T} such that a symmetric linear Markov equilibrium fails to exist for $T > \bar{T}$. In terms of the decomposition of the firms' equilibrium incentives provided in the next section, lack of existence appears to be due to the signaling incentive becoming too strong. Consistent with this interpretation, the condition given in Theorem 1 becomes harder to satisfy if r or σ^2 decreases, or if g_0 or T increases.

5 Equilibrium Properties

We then turn to the properties of equilibrium strategies and implications for prices, quantities, and profits. We first summarize properties of the equilibrium coefficients.

Proposition 1. *Let (α, β, δ) be a symmetric linear Markov equilibrium. Then*

1. $(-\alpha_t, \beta_t, \delta_t) \geq (-\alpha^m(\gamma_t), \beta^m(\gamma_t), \delta^m(\gamma_t)) > 0$ for all t .
2. α is initially decreasing and if T is sufficiently large, it is eventually increasing.¹⁶
3. β is initially increasing and if T is sufficiently large, it is eventually decreasing.
4. δ is eventually decreasing.
5. If $r = 0$, then α is quasiconvex, β is quasiconcave, and δ is decreasing.

The first part of Proposition 1 shows that the equilibrium coefficients are everywhere larger in absolute value than the myopic equilibrium coefficients (for the current beliefs) defined in (10). As the latter are signed and bounded away from zero, so are the former. In particular, each firm's output is decreasing in cost and increasing in the public belief.

The second and third part of the proposition imply that the equilibrium coefficients on private information, α , and on public information, β , are necessarily nonmonotone for T sufficiently large. As we discuss below, this seemingly surprising pattern is a natural consequence of learning and signaling. In contrast, the myopic equilibrium coefficients, which only reflect learning, are monotone: $\alpha^m(\gamma_t)$ is decreasing, $\beta^m(\gamma_t)$ is increasing, and $\delta^m(\gamma_t)$ is constant in t by inspection of (10).

The last part of Proposition 1 completes the qualitative description of equilibrium coefficients for $r = 0$, in which case $-\alpha$ and β are single peaked and δ is decreasing.

¹⁶A function $[0, T] \rightarrow \mathbb{R}$ satisfies a property *initially* if it satisfies it in an open neighborhood of 0. Similarly, the function satisfies a property *eventually* if it satisfies it in an open neighborhood of T .




Figure 1: Equilibrium Coefficients, $(r, \sigma, n, \bar{p}, T, g_0,) = (0.1, 1, 2, 5, 5, 2)$.

In fact, numerical analysis suggests that these properties always hold even for $r > 0$, but we are not aware of a proof. Figure 1 illustrates a typical equilibrium.

As an immediate corollary to Proposition 1, we obtain a characterization of long-run behavior. To see this, note that α is bounded away from zero, since $\alpha_t \leq \alpha^m(\gamma_t) \leq -1/2$ for all t , where the second inequality is by definition of α^m in (10). By inspection of (14), this implies that learning will never stop. Moreover, since the bound on α is independent of the length of the horizon, the rate of convergence is uniform across T , in the following sense.

Corollary 1. *For all $\varepsilon > 0$, there exists $t_\varepsilon < \infty$ such that for all $T \geq t \geq t_\varepsilon$, every symmetric linear Markov equilibrium of the T -horizon game satisfies $\gamma_t < \varepsilon$.*

This implies that the public belief converges to the true average cost, and hence each firm learns its rivals' average cost, asymptotically as we send the horizon T to infinity. Because of the identification problem arising from a one-dimensional signal and symmetric strategies, the firms cannot learn the cost of any given rival when there are more than two firms. However, with linear demand and constant marginal costs, knowing the average is sufficient for the firms to play their complete information best responses even in this case. Thus, under Markov strategies, play converges asymptotically to the static complete information Nash equilibrium for the realized costs.

Formally, let $Q_t := (Q_t^1, \dots, Q_t^n)$, and let $q^N : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the Nash equilibrium map of costs to quantities in the static, complete information version of our model.

Corollary 2. *Suppose $r\sigma^2 > g_0\kappa(n)$. Then for all $\varepsilon > 0$, there exists $t_\varepsilon < \infty$ such that for all $T \geq t \geq t_\varepsilon$, every symmetric linear Markov equilibrium of the T -horizon game*

satisfies $\mathbb{P}[\|Q_t - q^N(C)\| < \varepsilon] > 1 - \varepsilon$.¹⁷

The key to the proof is the fact that under the sufficient condition for existence, the equilibrium coefficients can be shown to converge over time to the static complete information values at a rate bounded from below uniformly in T . Corollary 2 then follows by noting that the public belief converges to the true average cost in distribution at a similarly uniform rate by Corollary 1. In particular, t_ε being independent of T suggests that it is the Markov restriction rather than the finite horizon that is driving the convergence to the static complete information Nash outcome, and, indeed, our other results. We confirm this in Section 6 by showing that as $T \rightarrow \infty$, our equilibria converge to an equilibrium of the infinite horizon version of the model.

5.1 Signaling and Learning

In order to explain the qualitative properties of equilibrium strategies, we consider here how signaling and learning affect the firms' incentives. For the deterministic part of the equilibrium strategy, δ , the intuition is well understood in terms of signal-jamming in a game with strategic substitutes.¹⁸ Indeed, compared to the myopic equilibrium where δ^m is constant, the equilibrium δ results in higher output with the difference (eventually) decreasing over time.

For the weights on the own cost and the public belief, *i.e.*, α and β , the intuition seems less clear at first. From firm i 's perspective, the public belief is not just the average cost of its rivals, but also includes its own cost. Furthermore, conditioning on C^i serves two purposes: it accounts both for firm i 's cost of production as well as its belief about the other firms' average cost as $M_t^i = z_t \Pi_t + (1 - z_t)C^i$.

To separate these effects, we proceed as follows. Rewrite firm i 's strategy as conditioning explicitly on its cost C^i and its belief M_t^i . That is, fix a symmetric linear Markov equilibrium (α, β, δ) , and define $\hat{\alpha}_t := \alpha_t - \beta_t(1 - z_t)/z_t$ and $\hat{\beta}_t := \beta_t/z_t$. Then, by Lemma 2, firm i 's equilibrium quantity on the path of play is given by

$$Q_t^i = \alpha_t C^i + \beta_t \Pi_t + \delta_t = \hat{\alpha}_t C^i + \hat{\beta}_t M_t^i + \delta_t, \quad t \in [0, T].$$

By inspection of the first-order condition (8), there are two drivers of firm i 's output: myopic flow profits, and the value of signaling. The myopic time- t best response to the equilibrium strategy is found by setting the second term $\partial V/\partial \pi \equiv 0$ in (8). Expressed

¹⁷Here, \mathbb{P} denotes the joint law of (C, Q_t) under the equilibrium strategies in the game with horizon T . Even without the condition on r , the result continues to hold for all t sufficiently close to T .

¹⁸See, *e.g.*, Riordan (1985), Fudenberg and Tirole (1986), or Mirman, Samuelson, and Urbano (1993).

in terms of C^i and M_t^i as above, this gives $Q_t^{br} = \hat{\alpha}_t^{br} C^i + \hat{\beta}_t^{br} M_t^i + \delta_t^{br}$, where

$$\hat{\alpha}_t^{br} = -\frac{(n-1)\beta_t(z_t-1)}{2z_t} - \frac{1}{2}, \quad \hat{\beta}_t^{br} = -\frac{(n-1)(\beta_t + \alpha_t z_t)}{2z_t}, \quad \delta_t^{br} = \frac{\bar{p} - (n-1)\delta_t}{2}.$$

The difference between the equilibrium strategy and the myopic best response, or

$$Q_t^i - Q_t^{br} = (\hat{\alpha}_t - \hat{\alpha}_t^{br})C^i + (\hat{\beta}_t - \hat{\beta}_t^{br})M_t^i + (\delta_t - \delta_t^{br}), \quad (16)$$

which corresponds to the second term in (8), is then by construction only due to signaling. Accordingly, we refer to the coefficients on the right as *signaling components*.

Proposition 2. *In every symmetric linear Markov equilibrium, the signaling components satisfy*

1. $\hat{\alpha}_t - \hat{\alpha}_t^{br} < 0$, $\hat{\beta}_t - \hat{\beta}_t^{br} > 0$, and $\delta_t - \delta_t^{br} > 0$ for all $0 \leq t < T$, and we have $\hat{\alpha}_T - \hat{\alpha}_T^{br} = \hat{\beta}_T - \hat{\beta}_T^{br} = \delta_T - \delta_T^{br} = 0$.
2. If $r = 0$, then $|\hat{\alpha}_t - \hat{\alpha}_t^{br}|$, $|\hat{\beta}_t - \hat{\beta}_t^{br}|$, and $|\delta_t - \delta_t^{br}|$ are decreasing in t .¹⁹

Armed with Proposition 2, we are now in a position to explain equilibrium signaling and the nonmonotonicity of the equilibrium coefficients. Note first that the *ex ante* expected signaling quantity is given by

$$\mathbb{E}[Q_t^i - Q_t^{br}] = (\hat{\alpha}_t - \hat{\alpha}_t^{br})\pi_0 + (\hat{\beta}_t - \hat{\beta}_t^{br})\pi_0 + (\delta_t - \delta_t^{br}) = (\delta_t - \delta_t^{br})\left(1 - \frac{\pi_0}{\bar{p}}\right),$$

where we have used $\delta_t = -\bar{p}(\hat{\alpha}_t + \hat{\beta}_t)$ and $\delta_t^{br} = -\bar{p}(\hat{\alpha}_t^{br} + \hat{\beta}_t^{br})$. Thus in the relevant range where $\pi_0 < \bar{p}$, the expected signaling quantity is positive as the firms are engaging in excess production in an effort to convince their rivals to scale back production. Moreover, when $r = 0$, the expected excess production is monotonically decreasing over time, reflecting the shorter time left to benefit from any induced reduction in the rivals' output, and the fact that beliefs are less sensitive to output when the firms already have a fairly precise estimate of their rivals' average cost.

The costs and benefits of signaling depend on firm i 's own cost and its belief about the other firms' average cost. In particular, a lower cost first makes it cheaper to produce additional output and then results in higher additional profits from the expansion of market share when other firms scale back their outputs in response. This is captured by the signaling component $\hat{\alpha}_t - \hat{\alpha}_t^{br}$ multiplying C^i in (16) being negative. If $r = 0$, it is decreasing in absolute value over time for the same reasons why the expected signaling quantity discussed above is decreasing and vanishing at the end.

¹⁹As with some of our other results for $r = 0$, numerical analysis strongly suggests that this result holds for all $r > 0$, but proving it appears difficult without the tractability gained by assuming $r = 0$.

The existence of the strictly positive signaling component $\hat{\beta}_t - \hat{\beta}_t^{br}$ multiplying firm i 's belief M_t^i in (16) is due to the belief being private. That is, firm i produces more when it believes that its rivals' costs are high both because it expects them to not produce much today (captured by $\hat{\beta}_t^{br} > 0$), and because by producing more, it signals to its rivals that it thinks that their costs are high and that it will hence be producing aggressively in the future. Again, this signaling component is monotone decreasing over time when $r = 0$.




Figure 2: Learning and Signaling Incentives, $(r, \sigma, n, \bar{p}, T, g_0) = (0, 1, 2, 1, 4.1, 2)$.

Turning to the non-monotonicity of the equilibrium coefficients, consider Figure 2, which illustrates the equilibrium coefficients $\hat{\alpha} < 0$ and $\hat{\beta} > 0$, the coefficients $\hat{\alpha}^{br} < 0$ and $\hat{\beta}^{br} > 0$ of the myopic best response to the equilibrium strategy, the signaling components $\hat{\alpha} - \hat{\alpha}^{br} \leq 0$ and $\hat{\beta} - \hat{\beta}^{br} \geq 0$, and the implied coefficients on C^i and M_t^i under the myopic equilibrium coefficients in (10) (dashed).

The myopic equilibrium (dashed) reflects only the effect of learning. There, the weights on own cost and belief are increasing in absolute value. This is best understood by analogy with a static Cournot game of incomplete information, where each of two firms privately observes an unbiased signal about its opponent's cost.²⁰ In this setting, a higher-cost firm knows that its rival will observe, on average, a higher signal. As the private signals become more precise, firms assign greater weight to their beliefs about their rival's cost. In a setting with strategic substitutes, each firm then consequently also assigns greater weight to its own cost in response, *i.e.*, a high-cost firm scales back production further when signals are more precise as it expects its rival to be more aggressive. This explains why also in the myopic equilibrium of our game, the weights

²⁰A similar game is studied in the literature on *ex ante* information sharing in oligopoly, see Raith (1996).

on C^i and M_t^i are increasing in absolute value over time as the firms' information becomes more precise (*i.e.*, as γ decreases).

The myopic best reply to the equilibrium strategy reflects these forces, but it is also affected by the shape of the equilibrium coefficients. As the equilibrium β (and $\hat{\beta}$) is initially much larger than the corresponding weight in the myopic equilibrium, the myopic best reply initially places a correspondingly higher weight $\hat{\alpha}^{br}$ on the firm's own cost, and hence lies below the lower of the dashed curves. Proposition 1 shows that β is eventually decreasing (for T large enough), which explains why $\hat{\alpha}^{br}$ is eventually slightly increasing in Figure 2. Similarly, as the equilibrium α (and $\hat{\alpha}$) is larger than the weight on the cost in the myopic equilibrium, the price is a more informative signal, and hence $\hat{\beta}^{br}$ lies above the corresponding dashed curve. As the equilibrium α is eventually increasing by Proposition 1, the opponents' output becomes eventually less sensitive to their cost, and the myopic best response then places a smaller weight on the belief about their cost. This is why $\hat{\beta}^{br}$ is eventually slightly decreasing in Figure 2.

Finally, the difference between the equilibrium coefficients and those of the myopic best reply is given by the signaling components $\hat{\alpha} - \hat{\alpha}^{br}$ and $\hat{\beta} - \hat{\beta}^{br}$, which are decreasing in absolute value by Proposition 2.

Therefore, we see that $\hat{\alpha}$ and $\hat{\beta}$ is the sum of a monotone signaling component, and of an almost monotone myopic component reflecting learning. Since $\hat{\alpha}$ and $\hat{\beta}$ simply amount to a decomposition of the equilibrium coefficients α and β , these two effects are responsible for the non-monotonicity of the latter as well.

The properties of the equilibrium coefficients have immediate implications for several outcome variables, which we turn to next.

5.2 Prices and Quantities

The relationship $\delta = -\bar{p}(\alpha + \beta)$ between the coefficients of the equilibrium strategy from Theorem 1 yields a simple expression for the expected total quantity in the market conditional on past prices: for any t and $s \geq t$, we have

$$\mathbb{E}\left[\sum_i Q_s^i \mid \mathcal{F}_t^Y\right] = n(\alpha_s \Pi_t + \beta_s \Pi_t + \delta_s) = n\delta_s \left(1 - \frac{\Pi_t}{\bar{p}}\right).$$

Thus the total expected output inherits the properties of the coefficient δ when $\Pi_t \leq \bar{p}$. (For $t = 0$ the condition can be satisfied simply by assuming that $\pi_0 \leq \bar{p}$; for $t > 0$ it can be made to hold with arbitrarily high probability by a judicious choice of parameters.) Proposition 1 then implies that the total expected output is eventually decreasing in s , and lies everywhere above its terminal value $(\bar{p} - \Pi_t)n/(n + 1)$, which is the complete information Nash total output for an industry with average cost Π_t . That is, if $\Pi_t \leq \bar{p}$ (respectively, $\Pi_t > \bar{p}$), then the expected current market supply conditional on public

information is higher (lower) than the market supply in a complete information Cournot market with average cost Π_t .

In order to describe the behavior of prices, we average out the demand shocks by defining for any t and $s \geq t$ the *expected price*

$$\mathbb{E}_t[P_s] := \bar{p} - \mathbb{E}\left[\sum_i Q_s^i \mid \mathcal{F}_t^Y\right] = \bar{p} - n\delta_s\left(1 - \frac{\Pi_t}{\bar{p}}\right),$$

which is just the expected time- s drift of the process Y conditional on its past up to time t . The above properties of the expected total market quantity then carry over to the expected price with obvious sign reversals. We record these in the following proposition, which summarizes some properties of equilibrium outputs and prices.

Proposition 3. *The following hold for every symmetric linear Markov equilibrium:*

1. *If $\Pi_t \leq \bar{p}$ (respectively, $\Pi_t > \bar{p}$), then for all $s \geq t$, the expected price $\mathbb{E}_t[P_s]$ is lower (respectively, higher) than the complete information equilibrium price in a Cournot market with average cost Π_t . As $s \rightarrow T$, the expected price converges to the complete information equilibrium price given average cost Π_t . If $r = 0$, then convergence is monotone. If in addition $\Pi_t < \bar{p}$, then $\mathbb{E}_t[P_s]$ is increasing in s .*
2. *The difference between any two firms' output levels conditional on their costs, $Q_t^i - Q_t^j = \alpha_t(C^i - C^j)$, is deterministic and, for T sufficiently large, nonmonotone.*

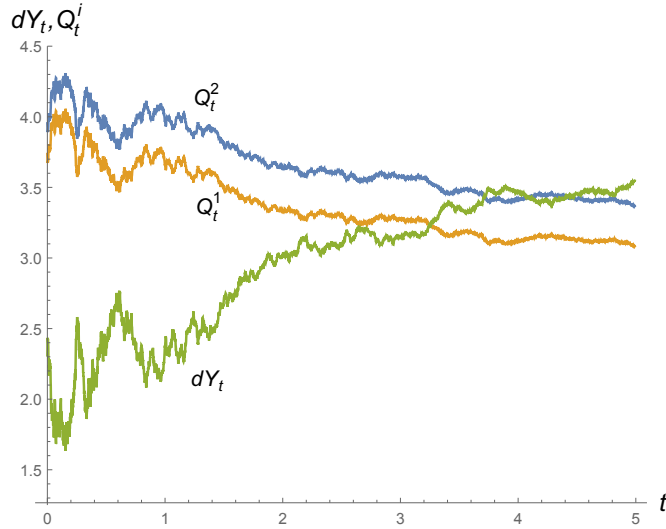


Figure 3: Price and Output Paths, $(r, \sigma, n, \bar{p}, T, g_0, \pi_0) = (0.75, 0.75, 2, 10, 5, 2, 0)$.

The first part of Proposition 3 implies that as long as the public belief about the average cost lies below the demand intercept, then conditional on past prices, future prices are expected to increase, monotonically so if $r = 0$. In particular, this is true of

the time-0 expectation as long as $\pi_0 \leq \bar{p}$. The finding is illustrated in Figure 3, which shows simulated price and output paths for two firms with costs $(C^1, C^2) = (1/2, 1/5)$.

The second part follows simply by definition of Markov strategies and the non-monotonicity of α for T large. As we discuss further below, it has implications for productive efficiency and hence for market profitability.

5.3 Expected Profits

We conclude our discussion of equilibrium properties by considering the implications of learning and signaling on the firms' profits and on consumer surplus. In particular, we are interested in the *ex ante* profits accruing to each firm over time, and in their magnitude relative to the expectation of complete-information profits.

Using the symmetry of the cost distribution, we define each firm's ex-ante expected time- t profit level and the corresponding expected consumer surplus as

$$W_t := \mathbb{E}[(\bar{p} - C^i - \sum_j Q_t^j)Q_t^i], \quad \text{and} \quad CS_t := \mathbb{E}[\frac{1}{2}(\sum_j Q_t^j)^2].$$

We compare these *ex ante* flows to the expected flow profit and consumer surplus under complete information, which are given by

$$W^{co} := \frac{(\bar{p} - \pi_0)^2 + g_0(n^2 + n - 1)}{(n + 1)^2}, \quad \text{and} \quad CS^{co} := \frac{n^2(\bar{p} - \pi_0)^2 + ng_0}{2(n + 1)^2}.$$

Proposition 4. *In every symmetric linear Markov equilibrium,*

1. $W_t < W^{co}$ for $t = 0$ and $t = T$.
2. Assume $g_0^{-1}r\sigma^2 \leq (n + 1)^{-1}(n - 1)^2$ and let π_0 be sufficiently close to \bar{p} . Then, for T sufficiently large, there exists $t < T$ such that $W_t > W^{co}$.
3. $CS_t \geq CS^{co}$ for all $t \in [0, T]$.

Figure 4 compares the expected profit levels under complete and incomplete information. The left and right panel contrast markets with a low and high mean of the cost distribution.

To obtain some intuition, we note the two main forces at play. On one hand, as seen above, signal-jamming adds to the expected total output over myopic players, which drives down profits. This wasteful spending is (eventually) declining over time by Proposition 2. On the other hand, learning about costs improves productive efficiency, and the firms' active signaling (*i.e.*, α being above its myopic value) increases the speed at which this happens.

At first sight, it seems that both of the above forces would simply result in the expected flow profit being increasing over time as in the left panel of Figure 4. But




Figure 4: $\pi_0 = 0$ (left), $\pi_0 = \bar{p}$ (right), and $(r, \sigma, n, \bar{p}, T, g_0) = (0.2, 1, 2, 5, 15.44, 2)$.

recall from Proposition 3 that the difference in output between any two firms i and j conditional on their costs is given by $\alpha_t(C^i - C^j)$, which is nonmonotone for T sufficiently large, because the sensitivity of output to cost, α , is then nonmonotone. This effect, driven by signaling, enhances productive efficiency in the medium run and can lead the expected flow profit W_t to surpass the expected complete information profit W^{co} at some interior t as in the right panel of Figure 4.

From an *ex ante* perspective, the enhanced productive efficiency corresponds to firms “taking turns” being the market leader, instead of the more even split of the market under complete information. The conditions in the second part of Proposition 4 ensure (i) that each firm is sufficiently patient, so that the value of signaling is sufficiently large,²¹ and (ii) that the “average profitability of the market” is not too high relative to the variance of output, so that this effect is important enough to outweigh the “price level” effect.

Thus, after the initial phase of high output levels, the overall effect of firms “jockeying for position” can improve industry profits.²² Finally, it is not hard to see that the expected consumer surplus always lies above the complete-information level. Indeed, expected prices are everywhere below the complete-information level and any variation in the price (conditional on the realized costs) is only beneficial to consumers who can adjust their demands. Therefore, the effects of signaling and signal-jamming incentives described in Proposition 2 can lead to an increase in total surplus as well.

²¹It can be verified that the upper bound on r in Proposition 4 is compatible with the lower bound on r in the first case of the sufficient condition for existence in Theorem 1 for some parameter values. For example, if $g_0/\sigma^2 = 1$, then this is the case for any $n = 2, \dots, 10$ (see the proof of Proposition 4).

²²In the static literature on *ex ante* information sharing in oligopoly (see, *e.g.*, Raith, 1996), output is most sensitive to costs under complete information, while the expected total quantity is constant in the precision of the information revealed. As a result, sharing full information about costs is beneficial in Cournot models. Instead, in our dynamic model with forward-looking firms, total expected quantity is decreasing, but output is most sensitive to cost for some intermediate time, leading to the richer picture outlined above.

6 Concluding Remarks

Our analysis makes use of a number of simplifying assumptions. The restriction to a Gaussian information structure and to equilibria in linear strategies—and hence to quadratic payoffs and unbounded, potentially negative outputs—is essential for the analysis as it allows us to use the Kalman filter to derive the firms’ posterior beliefs. Similarly, the representation of each firm’s private belief as a weighted sum of the public belief and the firm’s cost relies crucially on the properties of normal distributions.

Below we discuss the other main assumptions of our model.

6.1 Infinite Horizon

We have assumed a finite horizon throughout. However, Corollaries 1 and 2 show that as $t \rightarrow \infty$, beliefs and equilibrium behavior converge to the static complete information Nash outcome at a rate bounded uniformly in the horizon T . This suggests that our results are due to the Markov restriction rather than the finite horizon. To confirm this, we now show under a slight strengthening of the sufficient condition for existence (15) that symmetric linear Markov equilibria converge (along a subsequence) to an equilibrium of the infinite-horizon game as $T \rightarrow \infty$. This result is of independent interest as it provides a method of approximating equilibria of the infinite-horizon game by using our boundary value problem.

For the formal statement, it is convenient to use Theorem 1 to identify Markov equilibria of the T -horizon game with the tuple $(\alpha^T, \beta^T, \delta^T, \xi^T, \gamma^T)$. Moreover, we extend the functions to all of $[0, \infty)$ by setting $(\alpha_t^T, \beta_t^T, \delta_t^T, \xi_t^T, \gamma_t^T) = (\alpha_T^T, \beta_T^T, \delta_T^T, \xi_T^T, \gamma_T^T)$ for $t > T$. We then define a sequence of symmetric linear Markov equilibria to be any sequence of such tuples indexed by a strictly increasing, unbounded sequence of horizons. By the infinite-horizon game we mean the game obtained by setting $T = \infty$ in Section 2. (Note that the first time we use $T < \infty$ in the above analysis is when we impose boundary values on the equilibrium coefficients in Section 4.2.)

Proposition 5. *Suppose $g_0/\sigma^2 < 4r/(27n)$. Then any sequence of symmetric linear Markov equilibria contains a subsequence that converges uniformly to a symmetric linear Markov equilibrium $(\alpha^*, \beta^*, \delta^*, \xi^*, \gamma^*)$ of the infinite-horizon game.²³ Moreover, $\delta^* = -\bar{p}(\alpha^* + \beta^*)$ and $(\alpha^*, \beta^*, \xi^*, \gamma^*)$ is a solution to the system (11)–(13) on $[0, \infty)$ with $\lim_{t \rightarrow \infty} \alpha_t^* = \alpha^m(0)$, $\lim_{t \rightarrow \infty} \beta_t^* = \beta^m(0)$, $\lim_{t \rightarrow \infty} \xi_t^* = \xi^m(0)$, and $\gamma_0^* = ng_0$.*

The condition on g_0 , σ^2 , r , and n strengthens (15) to ensure that all the functions are bounded uniformly in T , facilitating the convergent argument. In particular, under

²³There is no reason to allow Markov strategies to condition on calendar time in the infinite-horizon game. However, allowing for it is innocuous because α^T is bounded away from zero uniformly in T , and hence the limit function γ^* is strictly decreasing, implying that conditioning on t is redundant.

(15), $(\alpha^T, \beta^T, \delta^T, \gamma^T)$ are uniformly bounded and have well-defined limits (along a subsequence), but the stronger condition allows us to also bound ξ^T and ultimately establish that the limit is an equilibrium of the infinite-horizon game.

As each δ^T lies everywhere above the complete information level, so does δ^* . This implies that our predictions for expected outputs and prices carry over to the infinite-horizon model. Furthermore, depending on the parameters, α^* is either non-monotone or everywhere decreasing, implying that the possibility of non-monotone market shares and expected profits carries over as well.

6.2 Asymmetric, Correlated, and Interdependent Costs

Symmetry of the prior distribution and of the equilibrium strategies is important for tractability. The asymmetric case presents no new conceptual issues, but the public belief Π and variance γ become vector-valued, and the analysis of the resulting boundary value problem seems a daunting task (see Lambert, Ostrovsky, and Panov (2014) for an extension of the static Kyle (1989) model to the asymmetric case).

In contrast, the assumption about independent costs can be easily relaxed. Correlated costs bring qualitatively no new insights, and the analysis under independence extends to this setting. In particular, every continuation equilibrium in our model corresponds to an equilibrium of a game played over the residual horizon with negatively correlated costs. Conversely, all our equilibria are continuation equilibria of longer games where the costs are positively correlated.

Similarly, we can introduce interdependent values, modeled as firm i 's cost being the sum $C^i + k \sum_{j \neq i} C^j$ for some $0 < k \leq 1$. Cost interdependence reduces the incentive to signal, since any given firm having a lower cost implies that the costs of the other firms are lower as well, and hence induces them to produce more. In the extreme case of pure common values ($k = 1$), the firms initially scale back production, with the burst of production toward the end resembling the aggressive behavior at the end of the horizon in models of insider trading in financial markets.

The special case of pure common values can alternatively be interpreted as a setting where costs are known and uncertainty is about the intercept \bar{p} of the demand function, of which each firm has received an initial private signal. A more general model would have both cost and demand uncertainty similar to Sadzik and Woolnough (2014) who generalize the model of Kyle (1985) by endowing the insider with private information about both the fundamental value and the amount of noise traders.

Finally, our model with fixed costs captures in a stylized way a new market where firms eventually converge to a static equilibrium. It is also of interest to consider settings where costs vary over time as in Fershtman and Pakes (2012) or the second model in Athey and Bagwell (2008). We pursue this in ongoing work.

Appendix

A.1 Preliminary Lemma

Under symmetric linear strategies, $dY_t = (\bar{p} - \alpha_t \sum_i C^i - nB_t(Y^t))dt + \sigma dZ_t$, with $B_t(Y^t) := \int_0^t f_s^t dY_s + \delta_t$. The following result is standard (Liptser and Shiryaev, 1977).

Lemma A.1. *Under any symmetric linear strategy profile, $\Pi_t := \frac{1}{n} \mathbb{E}[\sum_j C^j | \mathcal{F}_t^Y]$ and $\gamma_t := \mathbb{E}[(\sum_j C^j - n\Pi_t)^2 | \mathcal{F}_t^Y]$ are given by the unique solution to the system*

$$\begin{aligned} d\Pi_t &= -\frac{\alpha_t \gamma_t}{n\sigma^2} [dY_t - (\bar{p} - \alpha_t n\Pi_t - nB_t(Y^t))dt], \quad \Pi_0 = \pi_0, \\ \dot{\gamma}_t &= -\left(\frac{\alpha_t \gamma_t}{\sigma}\right)^2, \quad \gamma_0 = n g_0. \end{aligned}$$

In particular, the solution to the second equation is given by (5).

A.2 Proofs of Lemmas 1 to 4

Proof of Lemma 1. Let $e := (1, \dots, 1)' \in \mathbb{R}^{n-1}$ be a column vector of ones, and let I denote the $(n-1) \times (n-1)$ identity matrix. The argument in the text before the Lemma shows that firm i 's belief can be found by filtering the (column) vector $C^{-i} := (C^1, \dots, C^{i-1}, C^{i+1}, \dots, C^n)'$ from the one-dimensional process

$$dY^i = -\alpha_t e' C^{-i} dt + \sigma dZ_t.$$

By standard formulas for the Kalman filter (see, *e.g.*, Liptser and Shiryaev, 1977, Theorem 10.2), the posterior mean $M_t^{-i} := \mathbb{E}[C^{-i} | \mathcal{F}_t^{Y^i}]$ and the posterior covariance matrix $\Gamma_t := \mathbb{E}[(C^{-i} - M_t^{-i})(C^{-i} - M_t^{-i})' | \mathcal{F}_t^{Y^i}]$ are the unique solutions to the system

$$dM_t^{-i} = -\frac{\alpha_t}{\sigma} \Gamma_t e \frac{dY^{-i} - \alpha_t e' M_t^{-i} dt}{\sigma}, \quad M_0^{-i} = \pi_0 e, \quad (\text{A.1})$$

$$\dot{\Gamma}_t = -\frac{\alpha_t^2}{\sigma^2} \Gamma_t e e' \Gamma_t, \quad \Gamma_0 = g_0 I, \quad (\text{A.2})$$

where for Γ_t uniqueness is in the class of symmetric nonnegative definite matrices.

We first guess and verify the form of the solution for Γ_t . Let $A_t := \Gamma_t e e' \Gamma_t$. It is easy to see that its (i, j) -th component satisfies

$$A_t^{ij} = \sum_{k=1}^{n-1} \Gamma_t^{ik} \sum_{\ell=1}^{n-1} \Gamma_t^{\ell j}.$$

Thus we guess that the solution takes the form $\Gamma^{ii} = \gamma_t^1$, $\Gamma_t^{ij} = \gamma_t^2$, $i \neq j$, for some

functions γ^1 and γ^2 . The matrix equation (A.2) then reduces to the system

$$\begin{aligned}\dot{\gamma}_t^1 &= -\frac{\alpha_t^2}{\sigma^2}(\gamma_t^1 + (n-2)\gamma_t^2)^2, & \gamma_0^1 &= g_0, \\ \dot{\gamma}_t^2 &= -\frac{\alpha_t^2}{\sigma^2}(\gamma_t^1 + (n-2)\gamma_t^2)^2, & \gamma_0^2 &= 0.\end{aligned}$$

Consequently, $\gamma_t^M := (n-1)[\gamma_t^1 + (n-2)\gamma_t^2]$ satisfies

$$\dot{\gamma}_t^M = -\left(\frac{\alpha_t \gamma_t^M}{\sigma}\right)^2, \quad \gamma_0^M = (n-1)g_0,$$

whose solution is

$$\gamma_t^M = \frac{(n-1)g_0}{1 + (n-1)g_0 \int_0^t \frac{\alpha_s^2}{\sigma^2} ds}.$$

We can then solve for γ^1 and γ^2 by noting that $\dot{\gamma}_t^i = \dot{\gamma}_t^M / (n-1)^2$ for $i = 1, 2$, and hence integration yields

$$\Gamma_t^{ii} = \gamma_t^1 = \frac{\gamma_t^M}{(n-1)^2} + \frac{(n-2)g_0}{n-1} \quad \text{and} \quad \Gamma_t^{ij} = \gamma_t^2 = \frac{\gamma_t^M}{(n-1)^2} - \frac{g_0}{n-1}, \quad i \neq j.$$

It remains to verify that Γ_t so obtained is nonnegative definite. To this end, note that $\gamma_t^1 = \gamma_t^2 + g_0$, and hence $\Gamma_t = g_0 I + \gamma_t^2 E$, where E is a $(n-1) \times (n-1)$ matrix of ones. Therefore, for any nonzero (column) vector $x \in \mathbb{R}^{(n-1)}$ we have

$$x' \Gamma_t x = g_0 \|x\|_2^2 + \gamma_t^2 \|x\|_1^2 \geq \|x\|_1^2 \left(\frac{g_0}{n-1} + \gamma_t^2 \right) = \|x\|_1^2 \frac{\gamma_t^M}{(n-1)^2} > 0,$$

where the first inequality follows from $\sqrt{n-1} \|x\|_2 \geq \|x\|_1$ and the second inequality from $\gamma_t^M > 0$. We conclude that Γ_t is nonnegative definite, and hence it is indeed our covariance matrix. By inspection, it is of the form $\Gamma_t = \Gamma(\gamma_t^M)$ as desired.

In order to establish the form of the posterior mean, note that $(\Gamma_t e)^i = \gamma_t^M / (n-1)$. Thus (A.1) implies that $M_t^{-i} = M_t^i e$, where M_t^i evolves according to

$$dM_t^i = -\frac{\alpha_t}{\sigma} \frac{\gamma_t^M}{n-1} \frac{dY^i + \alpha_t (n-1) M_t^i dt}{\sigma}, \quad (\text{A.3})$$

and where

$$dZ_t^i := \frac{dY^i + (n-1)\alpha_t M_t^i dt}{\sigma}$$

is a standard Brownian motion (with respect to \mathcal{F}^{Y^i}) known as firm i 's *innovation process*. It is readily verified that $((n-1)M_t^i, \gamma_t^M)$ are the posterior mean and variance for the problem

$$dY_t^i = -\alpha_t \nu dt + \sigma dZ_t^i, \quad \nu \sim \mathcal{N}((n-1)\pi_0, (n-1)g_0),$$

which amounts to filtering the other firms' total cost. Thus M_t^i is the posterior expectation about the other firms' average cost as desired. \square

Proof of Lemma 2. The result is a special case of Lemma 3. (The formula for z_t follows by direct calculation from the formulas for γ_t^M and γ_t given in Lemma 1 and equation (5), respectively.) \square

Proof of Lemma 3. Fix a symmetric linear strategy profile, and let

$$\lambda_t := -\frac{\alpha_t \gamma_t}{n\sigma^2} \quad \text{and} \quad \lambda_t^M := -\frac{\alpha_t \gamma_t^M}{(n-1)\sigma^2}, \quad t \in [0, T].$$

Note that $z_t := n\gamma_t^M / [(n-1)\gamma_t] = \lambda_t^M / \lambda_t$. Recall the law of motion of the private belief M^i in (A.3), and define the process $\hat{\Pi}^i$ by

$$\begin{aligned} \hat{\Pi}_t^i &:= \exp\left(n \int_0^t \lambda_u \alpha_u du\right) \pi_0 \\ &\quad + \int_0^t \exp\left(n \int_s^t \lambda_u \alpha_u du\right) \lambda_s \left[-\alpha_s (C^i + (n-1)M_s^i) ds + \frac{dM_s^i}{\lambda_s^M}\right]. \end{aligned}$$

The process $\hat{\Pi}^i$ is in firm i 's information set because it is a function of its belief M^i and cost C^i . We prove the first part of the Lemma by showing that

$$M_t^i - C^i = z_t (\hat{\Pi}_t^i - C^i), \quad t \in [0, T]. \quad (\text{A.4})$$

To this end, note that the law of motion of $\hat{\Pi}^i$ is given by

$$d\hat{\Pi}_t^i = \lambda_t \alpha_t [\hat{\Pi}_t^i - C^i + (n-1)(\hat{\Pi}_t^i - M_t^i)] dt + \frac{\lambda_t}{\lambda_t^M} dM_t^i, \quad \hat{\Pi}_0^i = \pi_0. \quad (\text{A.5})$$

Let $W_t := z_t (\hat{\Pi}_t^i - C^i)$. Applying Ito's rule and using that $z_t \lambda_t = \lambda_t^M$ gives²⁴

$$\begin{aligned} dW_t &= \lambda_t^M \alpha_t [(n-1)z_t - n](\hat{\Pi}_t - C^i) dt + \lambda_t^M \alpha_t [\hat{\Pi}_t - C^i + (n-1)(\hat{\Pi}_t - M_t^i)] dt + dM_t^i \\ &= (n-1)\lambda_t^M \alpha_t [z_t (\hat{\Pi}_t - C^i) - (M_t^i - C^i)] dt + dM_t^i \\ &= (n-1)\lambda_t^M \alpha_t [W_t - (M_t^i - C^i)] dt + dM_t^i. \end{aligned}$$

Therefore, we have

$$d[W_t - (M_t^i - C^i)] = (n-1)\lambda_t^M \alpha_t [W_t - (M_t^i - C^i)] dt,$$

²⁴Observe that

$$\dot{z}_t = \frac{n}{n-1} \frac{\dot{\gamma}_t^M \gamma_t - \gamma_t^M \dot{\gamma}_t}{\gamma_t^2} = -\frac{n}{n-1} \frac{\alpha_t^2 (\gamma_t^M)^2}{\sigma^2 \gamma_t} + z_t \frac{\alpha_t^2 \gamma_t}{\sigma^2} = (n-1)\lambda_t^M \alpha_t z_t - n\lambda_t^M \alpha_t,$$

where we have used that $\dot{\gamma}_t = -(\alpha_t \gamma_t / \sigma)^2$ and $\dot{\gamma}_t^M = -(\alpha_t \gamma_t^M / \sigma)^2$.

which admits as its unique solution

$$W_t - (M_t^i - C^i) = [W_0 - (M_0^i - C^i)] \exp \left((n-1) \int_0^t \lambda_s^M \alpha_s ds \right).$$

But $W_0 - (M_0^i - C^i) = z_0(\hat{\Pi}_0^i - C^i) - (M_0^i - C^i) = 0$, since $z_0 = 1$ and $\hat{\Pi}_0^i = M_0^i = \pi_0$. Consequently, $W_t - (M_t^i - C^i) \equiv 0$, which establishes (A.4).

The law of motion for $\hat{\Pi}^i$ given in the Lemma now follows from (A.5) by using (A.4) to substitute for M_t^i , and by using (A.3) to substitute for dM_t^i .

It remains to show that $\hat{\Pi}_s^i = \Pi_s$ if firm i plays the same strategy on $[0, s)$ as the other firms. Note that then by (4), we have from the perspective of firm i

$$dY_t - (\bar{p} - nB_t(Y))dt = dY_t^i - \alpha_t C^i dt = -\alpha_t [C^i + (n-1)M_t^i]dt + \frac{dM_t^i}{\lambda_t^M}, \quad t \in [0, s),$$

where the second equality follows by (A.3). Therefore, the law of motion of Π in Lemma A.1 is from firm i 's perspective given on $[0, s)$ by

$$\begin{aligned} d\Pi_t &= -\frac{\alpha_t \gamma_t}{n\sigma^2} [dY_t - (\bar{p} - \alpha_t n\Pi_t - nB_t(Y))dt] \\ &= \lambda_t \alpha_t [n\Pi_t - C^i - (n-1)M_t^i]dt + \frac{\lambda_t}{\lambda_t^M} dM_t \\ &= \lambda_t \alpha_t [\Pi_t - C^i + (n-1)(\Pi_t - M_t)]dt + \frac{\lambda_t}{\lambda_t^M} dM_t, \end{aligned}$$

with initial condition $\Pi_0 = \pi_0$. By inspection of (A.5) we thus have $\Pi_t = \hat{\Pi}_t^i$ for all $t \leq s$. (This also shows that if firm i has ever unilaterally deviated from the symmetric linear strategy profile in the past, then $\hat{\Pi}_t^i$ equals the counterfactual value of the public belief that would have obtained had firm i not deviated.) \square

Proof of Lemma 4. Lemmas 1 and 2 imply that if all firms play a symmetric linear strategy profile, then there is a one-to-one correspondence between (C^i, Π_t, t) and firm i 's time- t belief about (C^1, \dots, C^n) and calendar time. Thus, if $Q_t^i = \alpha_t C^i + \beta_t \Pi_t + \delta_t$, $t \in [0, T]$, then firm i 's quantity is only a function of its belief and calendar time. Using the law of motion from Lemma A.1, it is straightforward to verify that the public belief is of the form $\Pi_t = \int_0^t k_s^t dY_s + \text{constant}_t$. Thus conditioning on it agrees with our definition of a linear strategy in (3).

Conversely, suppose that a symmetric linear strategy profile (α, f, δ) is only a function of beliefs and calendar time. Given the one-to-one correspondence noted above, we then have for each firm i and all t ,

$$Q_t^i = \psi_t(C^i, \Pi_t)$$

for some function $\psi_t : \mathbb{R}^2 \rightarrow \mathbb{R}$. Let $\text{supp}(\alpha)$ denote the essential support of α on $[0, T]$, and let $\tau := \min \text{supp}(\alpha)$. Then private and public beliefs about firms $j \neq i$ are simply given by the prior at all $0 \leq t \leq \tau$ (i.e., $\Pi_t = \pi_0$, $z_t = 1$, and thus $M_t^i = \pi_0$), and hence the strategy can only condition on firm i 's (belief about its) own cost and on calendar time on $[0, \tau]$. Thus, by linearity of the strategy, we have $\psi_t(C^i, \Pi_t) = \alpha_t C^i + \delta_t$ for $t \leq \tau$, which shows that the strategy takes the desired form on this (possibly empty) subinterval. Note then that for any $t > \tau$, we have

$$Q_t^i = \alpha_t C^i + \int_0^t f_s^t dY_s + \delta_t = \psi_t(C^i, \Pi_t) = \psi_t\left(C^i, \int_0^t k_s^t dY_s + \text{constant}_t\right),$$

where the argument of ψ_t can take on any value in \mathbb{R}^2 given the distribution of C^i and the form of the noise in the price process Y . Thus, for the equality to hold, ψ_t must be an affine function, i.e., $\psi_t(C^i, \Pi_t) = a_t C^i + b_t \Pi_t + d_t$ for some constants (a_t, b_t, d_t) , establishing the result. \square

A.3 Proof of Theorem 1

The proof proceeds as a series of lemmas.

Lemma A.2. *If (α, β, δ) is a symmetric linear Markov equilibrium with posterior variance γ , then (i) $(\alpha, \beta, \xi, \gamma)$ with ξ defined by (13) is a solution to the boundary value problem, and (ii) $\delta = -\bar{p}(\alpha + \beta)$.*

Proof. Fix such an equilibrium (α, β, δ) with variance γ , and fix some firm i . By inspection, the best-response problem in Section 4.1 is a stochastic linear-quadratic regulator (see, e.g., Yong and Zhou, 1999, Chapter 6). Moreover, (α, β, δ) is an optimal policy (a.s.) on the path of play, i.e., at states where $\Pi_t = \hat{\Pi}_t^i$.

We argue first that the value function takes the form given in (9). Along the way, we also establish the existence of an optimal policy at off-path states $(C^i, \Pi_t, \hat{\Pi}_t^i, t)$ where $\Pi_t \neq \hat{\Pi}_t^i$. Introducing the shorthand S_t for the state, we can follow Yong and Zhou (1999, Chapter 6.4) and write the best-response problem at any state S_t as an optimization problem in a Hilbert space where the choice variable is a square-integrable output process Q^i on $[t, T]$ and the objective function takes the form

$$\frac{1}{2} [\langle L_t^1 Q^i, Q^i \rangle + 2 \langle L_t^2(S_t), Q^i \rangle + L_t^3(S_t)]$$

for certain linear functionals L_t^i , $i = 1, 2, 3$.²⁵ Since an equilibrium exists, the value of the problem at S_t is finite, and hence $L_t^1 \leq 0$ by Theorem 4.2 of Yong and Zhou (1999, p. 308). By the same theorem, if $L_t^1 < 0$, then there exists a unique optimal

²⁵ Cf. display (4.17) on page 307 in Yong and Zhou (1999).

policy, which is linear in S_t , and the value function is of the form (9). This leaves the case $L_t^1 = 0$. In that case the objective function is linear in Q^i on the vector space of square-integrable processes, and hence its supremum is unbounded, which contradicts the existence of an equilibrium, unless $L_t^2 = 0$. But if $L_t^1 = L_t^2 = 0$, then any policy is optimal, which is clearly impossible as, for example, $Q^i \equiv 0$ yields a zero payoff whereas the expected payoff from myopically maximizing the flow payoff is strictly positive. We thus conclude that $L^1 < 0$ and there exists a unique optimal policy.

We note then that the value function V is continuously differentiable in t and twice continuously differentiable in $(c, \pi, \hat{\pi})$.²⁶ Thus it satisfies the HJB equation (7). This implies that the linear optimal policy $q = \alpha_t c + \beta_t \pi + \delta_t + \xi_t(\hat{\pi} - \pi)$, where $(\alpha_t, \beta_t, \delta_t)$ are the equilibrium coefficients, satisfies the first-order condition (8). This gives

$$\alpha_t c + \beta_t \pi + \delta_t + \xi_t(\hat{\pi} - \pi) = \frac{\bar{p} - (n-1)[\alpha_t(z_t \hat{\pi} + (1-z_t)c) + \beta_t \pi + \delta_t] - c}{2} - \lambda_t \frac{v_1(t) + v_4(t)\hat{\pi} + v_5(t)c + 2v_8(t)\pi}{2},$$

where we have written out $\partial V/\partial \pi$ using (9). As this equality holds for all $(c, \pi, \hat{\pi}) \in \mathbb{R}^3$, we can match the coefficients of $c, \pi, \hat{\pi}$, and constants on both sides to obtain the system

$$\begin{aligned} \alpha_t &= -\frac{(n-1)\alpha_t(1-z_t) + 1}{2} + \frac{\alpha_t \gamma_t}{2n\sigma^2} v_5(t), \\ \beta_t - \xi_t &= -\frac{(n-1)\beta_t}{2} + \frac{\alpha_t \gamma_t}{n\sigma^2} v_8(t), \\ \delta_t &= \frac{\bar{p} - (n-1)\delta_t}{2} + \frac{\alpha_t \gamma_t}{2n\sigma^2} v_1(t), \\ \xi_t &= -\frac{(n-1)\alpha_t z_t}{2} + \frac{\alpha_t \gamma_t}{2n\sigma^2} v_4(t), \end{aligned} \tag{A.6}$$

where we have used $\lambda_t = -\alpha_t \gamma_t / (n\sigma^2)$.

We can now show that $(\alpha, \beta, \xi, \gamma)$ satisfy the boundary conditions given in the theorem. Note that $v_k(T) = 0$, $k = 1, \dots, 9$. Thus we obtain $(\alpha_T, \beta_T, \delta_T, \xi_T)$ from (A.6) by solving the system with $(v_1(T), v_4(T), v_5(T), v_8(T)) = (0, \dots, 0)$. Recalling the expression for z_T in terms of γ_T from (6), a straightforward calculation yields $\alpha_T = \alpha^m(\gamma_T)$, $\beta_T = \beta^m(\gamma_T)$, $\delta_T = \delta^m(\gamma_T)$, and $\xi_T = \xi^m(\gamma_T)$, where the functions $(\alpha^m, \beta^m, \delta^m, \xi^m)$ are defined in (10). The condition $\gamma_0 = ng_0$ is immediate from (5).

As γ satisfies (14) by construction, it remains to show that $(\alpha, \beta, \xi, \gamma)$ satisfy equations (11)–(13) and that $\delta = -\bar{p}(\alpha + \beta)$. Applying the envelope theorem to the HJB

²⁶Differentiability of each v_i in (9) can be verified using the fact that V is the value under the unique optimal policy $(\alpha, \beta, \delta, \xi)$, where (α, β, δ) are differentiable by assumption and ξ , which only enters V through an integral, is continuous.

equation (7) we have

$$r \frac{\partial V}{\partial \pi} = -(n-1)\beta_t q^*(c, \pi, \hat{\pi}, t) + \mu_t \frac{\partial^2 V}{\partial \pi^2} + \frac{\partial \mu_t}{\partial \pi} \frac{\partial V}{\partial \pi} + \hat{\mu}_t \frac{\partial^2 V}{\partial \pi \partial \hat{\pi}} + \frac{\partial^2 V}{\partial \pi \partial t}, \quad (\text{A.7})$$

where we omit third-derivative terms as V is quadratic. By inspection of (9), the only coefficients of V that enter this equation are $v_1(t)$, $v_4(t)$, $v_5(t)$, and $v_8(t)$ as well as their derivatives $\dot{v}_1(t)$, $\dot{v}_4(t)$, $\dot{v}_5(t)$, and $\dot{v}_8(t)$. Therefore, we first solve (A.6) for $(v_1(t), v_4(t), v_5(t), v_8(t))$ in terms of $(\alpha_t, \beta_t, \delta_t, \xi_t, \gamma_t)$, and then differentiate the resulting expressions to obtain $(\dot{v}_1(t), \dot{v}_4(t), \dot{v}_5(t), \dot{v}_8(t))$ in terms of $(\alpha_t, \beta_t, \delta_t, \xi_t, \gamma_t)$ and $(\dot{\alpha}_t, \dot{\beta}_t, \dot{\delta}_t, \dot{\xi}_t, \dot{\gamma}_t)$. (Note that (A.6) holds for all t and (α, β, δ) are differentiable by assumption; differentiability of ξ follows by (A.6).) Substituting into (A.7) then yields an equation for $(\alpha_t, \beta_t, \delta_t, \xi_t, \gamma_t)$ and $(\dot{\alpha}_t, \dot{\beta}_t, \dot{\delta}_t, \dot{\xi}_t, \dot{\gamma}_t)$ in terms of $(c, \pi, \hat{\pi})$ and the parameters of the model. Moreover, as this equation holds for all $(c, \pi, \hat{\pi}) \in \mathbb{R}^3$, we can again match coefficients to obtain a system of four equations that are linear in $(\dot{\alpha}_t, \dot{\beta}_t, \dot{\delta}_t, \dot{\xi}_t)$. A very tedious but straightforward calculation shows that these equations, solved for $(\dot{\alpha}_t, \dot{\beta}_t, \dot{\delta}_t, \dot{\xi}_t)$, are equations (11)–(13) and

$$\dot{\delta}_t = r\alpha_t \frac{\delta_t - \delta^m(\gamma_t)}{\alpha^m(\gamma_t)} + \frac{(n-1)\alpha_t\beta_t\gamma_t}{n(n+1)\sigma^2} [\delta_t - n\alpha_t(z_t - 1)((n+1)\delta_t - \bar{p})]. \quad (\text{A.8})$$

The identity $\delta = -\bar{p}(\alpha + \beta)$ can be verified by substituting into (A.8) and using (11) and (12), and noting that the boundary conditions satisfy it by inspection of (10). \square

Lemma A.3. *If $(\alpha, \beta, \xi, \gamma)$ is a solution to the boundary value problem, then (α, β, δ) with $\delta = -\bar{p}(\alpha + \beta)$ is a symmetric linear Markov equilibrium with posterior variance γ .*

Proof. Suppose that $(\alpha, \beta, \xi, \gamma)$ is a solution to the boundary value problem and let $\delta = -\bar{p}(\alpha + \beta)$. Then (α, β, δ) are bounded functions on $[0, T]$, and hence they define an admissible symmetric linear Markov strategy (see footnote 8 on page 6). Moreover, (5) is the unique solution to (14) with $\gamma_0 = ng_0$, and hence γ is the corresponding posterior variance of the public belief.

To prove the claim, we assume that the other firms play according to (α, β, δ) , and we construct a solution V to firm i 's HJB equation (7) such that V takes the form (9) and the optimal policy is $q^*(c, \pi, \hat{\pi}, t) = \alpha_t c + \beta_t \pi + \delta_t + \xi_t(\hat{\pi} - \pi)$. We then use a verification theorem to conclude that this indeed constitutes a solution to firm i 's best response problem.

We construct V as follows. By Proposition 1, $(\alpha, \beta, \delta, \xi)$ are bounded away from 0, and so is γ because T is finite.²⁷ We can thus define (v_1, v_4, v_5, v_8) by (A.6). Then, by construction, $q^*(c, \pi, \hat{\pi}, t) = \alpha_t c + \beta_t \pi + \delta_t + \xi_t(\hat{\pi} - \pi)$ satisfies the first-order condition (8), which is sufficient for optimality by concavity of the objective function in (7).

²⁷For ξ , this follows from $\xi_t \geq \xi_t^m := \xi^m(\gamma_t) > 0$. The second inequality is by (10). To see the first, notice

The remaining functions $(v_0, v_2, v_3, v_6, v_7, v_9)$ can be obtained from (7) by substituting the optimal policy $q^*(c, \pi, \hat{\pi}, t)$ for q on the right-hand side and matching the coefficients of $(c, \hat{\pi}, c\hat{\pi}, c^2, \hat{\pi}^2)$ and the constants on both sides of the equation so obtained. This defines a system of six linear first-order ODEs (with time-varying coefficients) for $(v_0, v_2, v_3, v_6, v_7, v_9)$, stated here for future reference:

$$\begin{aligned}
\dot{v}_0(t) &= rv_0(t) - \delta_t(\bar{p} - n\delta_t) - \frac{\alpha_t^2 \gamma_t^2}{n^2 \sigma^2} v_9(t) - \frac{\alpha_t \gamma_t (n\beta_t + \beta_t + 2\xi_t) + 2(n-1)\alpha_t^2 \gamma_t z_t}{2n}, \\
\dot{v}_2(t) &= (n-1)\alpha_t z_t (\bar{p} - n\delta_t) + \frac{nr\sigma^2 + \alpha_t^2 \gamma_t (n(1-z_t) + z_t)}{n\sigma^2} v_2(t), \\
\dot{v}_3(t) &= (n-1)\alpha_t (z_t - 1)(n\delta_t - \bar{p}) + rv_3(t) + \delta_t + \frac{\alpha_t^2 \gamma_t ((n-1)z_t - n)}{n\sigma^2} v_2(t), \\
\dot{v}_6(t) &= \frac{nr\sigma^2 + \alpha_t^2 \gamma_t (n(1-z_t) + z_t)}{n\sigma^2} v_6(t) + \frac{2\alpha_t^2 \gamma_t ((n-1)z_t - n)}{n\sigma^2} v_9(t) \\
&\quad + \alpha_t (-2n\xi_t - (n-1)z_t(2n\alpha_t - 2\xi_t + 1) + 2(n-1)^2 \alpha_t z_t^2), \\
\dot{v}_7(t) &= rv_7(t) + \alpha_t (n-1)(z_t - 1) - \alpha_t^2 (n(1-z_t) + z_t)^2 + \frac{\alpha_t^2 \gamma_t ((n-1)z_t - n)}{n\sigma^2} v_6(t), \\
\dot{v}_9(t) &= \frac{nr\sigma^2 + 2\alpha_t^2 \gamma_t (n(1-z_t) + z_t)}{n\sigma^2} v_9(t) - ((n-1)\alpha_t z_t + \xi_t)^2.
\end{aligned} \tag{A.9}$$

By linearity, the system has a unique solution on $[0, T]$ that satisfies the boundary condition $(v_0(T), v_2(T), v_3(T), v_6(T), v_7(T), v_9(T)) = (0, \dots, 0)$. Defining V by (9) with the functions v_k , $k = 1, \dots, 9$, defined above then solves the HJB equation (7) by construction.

Finally, because V is linear-quadratic in $(c, \pi, \hat{\pi})$ and the functions v_k are uniformly bounded, V satisfies the quadratic growth condition in Theorem 3.5.2 of Pham (2009). Therefore, V is indeed firm i 's value function and $(\alpha, \beta, \delta, \xi)$ is an optimal policy. Moreover, on-path behavior is given by (α, β, δ) as desired. \square

We then turn to existence. As discussed in the text following the theorem, we use the shooting method, omitting first equation (13) from the system.

Define the *backward system* as the initial value problem defined by (11), (12), and (14) with $\gamma_T = \gamma_F$, $\alpha_T = \alpha^m(\gamma_F)$, and $\beta_T = \beta^m(\gamma_F)$ for some $\gamma_F \in \mathbb{R}_+$. By inspection, the backward system is locally Lipschitz continuous (note that $g_0 > 0$ by definition). For $\gamma_F = 0$, its unique solution on $[0, T]$ is given by $\alpha_t = \alpha^m(0)$, $\beta_t = \beta^m(0)$, and $\gamma_t = 0$ for all t . By continuity, it thus has a solution on $[0, T]$ for all γ_F in some interval $[0, \tilde{\gamma}_F)$ with $\tilde{\gamma}_F > 0$. Let $G := [0, \bar{\gamma}_F)$ be the maximal such interval with respect to set inclu-

that $\dot{\xi}_t$ is decreasing in β_t . Therefore, bounding β_t with β_t^m by Proposition 1, we obtain

$$(\beta_t, \xi_t) = (\beta_t^m, \xi_t^m) \Rightarrow \dot{\xi}_t - \dot{\xi}_t^m = -\frac{g_0^2 (n-1)^3 n^3 \alpha_t (2\alpha_t - 1) \gamma_t}{4(n+1)\sigma^2 (g_0(n-1)n + (n+1)\gamma_t)^2} \leq 0,$$

because $\alpha_t < \alpha_t^m \leq -1/2$ for all t by Proposition 1. This implies that ξ can only cross its myopic value from above, which occurs at time $t = T$.

sion. (I.e., $\bar{\gamma}_F = \sup \{\tilde{\gamma}_F \in \mathbb{R}_+ : \text{backward system has a solution for all } \gamma_F \in [0, \tilde{\gamma}_F)\}$.)
 Finally, define the function $\kappa : \mathbb{N} \rightarrow \mathbb{R}_{++}$ by

$$\kappa(n) := \inf_{a \in (-\infty, -1]} \left\{ -\frac{(n-1)2\sqrt{a^5(a+1)n(2an+n+1)(a(n-1)n-1)}}{(a+an+1)^2} + \frac{a^2(a(n(a-(3a+2)n)+1)+1)}{(a+an+1)^2} \right\}. \quad (\text{A.10})$$

Lemma A.4. *Suppose (15) holds, i.e., $g_0/\sigma^2 < \max\{r/\kappa(n), 1/(3nT)\}$. Then there exists $\gamma_F \in G$ such that the solution to the backward system satisfies $\gamma_0 = ng_0$.*

Proof. Suppose $g_0/\sigma^2 < \max\{r/\kappa(n), 1/(3nT)\}$. The backward system is continuous in γ_F , and $\gamma_F = 0$ results in $\gamma_0 = 0$. Thus it suffices to show that $\gamma_0 \geq ng_0$ for some $\gamma_F \in G$. Suppose, in negation, that the solution to the backward system has $\gamma_0 < ng_0$ for all $\gamma_F \in G$. Since γ is monotone by inspection of (14), we then have $\gamma_F = \gamma_T \leq \gamma_0 < ng_0$ for all $\gamma_F \in G$, and thus $\bar{\gamma}_F \leq ng_0 < \infty$. We will show that this implies that the solutions (α, β, γ) are bounded uniformly in γ_F on G , which contradicts the fact that, by definition of G , one of them diverges at some $t \in [0, T)$ when $\gamma_F = \bar{\gamma}_F$.

To this end, let $\gamma_F \in G$, and let (α, β, γ) be the solution to the backward system.

By monotonicity of γ , we have $0 \leq \gamma_t \leq \gamma_0 < ng_0$ for all t , and hence γ is bounded uniformly across γ_F in G as desired.

Note then that, by the arguments in the proof of the first part of Proposition 1 below, we have $(-\alpha, \beta, \delta) \geq 0$. The identity $-\bar{p}(\alpha + \beta) = \delta$ then implies $\alpha \leq -\beta \leq 0$. Therefore, to bound α and β , it suffices to bound α from below.

We first derive a lower bound for α when $\rho := ng_0/\sigma^2 < 1/(3T)$. Consider

$$\dot{x}_t = \rho x_t^4, \quad x_T = -1. \quad (\text{A.11})$$

By (10), we have $x_T \leq \alpha^m(\gamma_F) = \alpha_T$ for all $\gamma_F \geq 0$. Furthermore, recalling that $\gamma_t \leq ng_0$, $z_t \in [1, n/(n-1)]$, and $-\alpha_t \geq \beta_t \geq 0$ for all t , we can verify using equation (11) that $\rho\alpha_t^4 \geq \dot{\alpha}_t$ for all $\alpha_t \leq -1$. Working backwards from T , this implies $x_t \leq \alpha_t$ for all t at which x_t exists. Furthermore, the function x is by definition independent of γ_F , so to bound α it suffices to show that (A.11) has a solution on $[0, T]$. This follows, since the unique solution to (A.11) is

$$x_t = \frac{1}{\sqrt[3]{3\rho(T-t) - 1}}, \quad (\text{A.12})$$

which exists on all of $[0, T]$, because $3\rho(T-t) - 1 \leq 3\rho T - 1 < 0$ by assumption.

We then consider the case $g_0/\sigma^2 < r/\kappa(n)$. We show that there exists a constant $\bar{a} < -1$ such that $\alpha \geq \bar{a}$. In particular, denoting the right-hand side of $\dot{\alpha}$ in (11) by

$f(\alpha_t, \beta_t, \gamma_t)$, we show that there exists $\bar{a} < -1$ such that $f(\bar{a}, b, g) \leq 0$ for all $b \in [0, -\bar{a}]$ and $g \in [0, ng_0]$. Since $0 \leq \beta \leq -\alpha$ and $0 \leq \gamma \leq ng_0$, this implies that following (11) backwards from any $\alpha_T > -1$ yields a function bounded from below by \bar{a} on $[0, T]$.

For $a \leq -1$ and $r > 0$, let

$$D(a, r) := (\bar{a}^2 g_0 (n-1)(an+1) - r\sigma^2(a(1+n)+1))^2 - 4a^2(a+1)g_0(n-1)r\sigma^2(a(n-1)n-1).$$

We claim that there exists $\bar{a} \leq -1$ such that $D(\bar{a}, r) < 0$. Indeed, $D(a, r)$ is quadratic and convex in r . It is therefore negative if $r \in [r_1, r_2]$, where $r_1 = r_1(a)$ and $r_2 = r_2(a)$ are the two roots of $D(a, r) = 0$. One can verify that for any $a \leq -1$, $D(a, r) = 0$ admits two real roots $r_1 = r_1(a) \leq r_2 = r_2(a)$, with strict inequality if $a < -1$, which are both continuous functions of a that grow without bound as $a \rightarrow -\infty$. Thus, there exists \bar{a} such that $D(\bar{a}, r) < 0$ if $r > \inf_{a \in (-\infty, -1)} r_1(a)$. But, by definition, the objective function in the extremum problem in (A.10) is $(\sigma^2/g_0)r_1(a)$, and hence the existence of \bar{a} follows from $r > \kappa(n)g_0/\sigma^2$. We fix some such \bar{a} for the rest of the proof.

Consider any $g \in [0, ng_0]$. Let $z := n^2 g_0 / [n(n-1)g_0 + g]$. By inspection of (11), if $(n-1)n\bar{a}(z-1) + 1 \geq 0$, then $f(\bar{a}, b, g) \leq 0$ for all $b \in [0, -\bar{a}]$, since $\bar{a} \leq -1 \leq \alpha^m(g)$, which implies that the r -term is negative. On the other hand, if $(n-1)n\bar{a}(z-1) + 1 < 0$, then $f(\bar{a}, b, g) \leq f(\bar{a}, -\bar{a}, g)$ for all $b \in [0, -\bar{a}]$. Thus it suffices to show $f(\bar{a}, -\bar{a}, g) \leq 0$.

Note that

$$f(\bar{a}, -\bar{a}, g) = \frac{\bar{a} (g(n-1)ng_0\bar{a}^2(n\bar{a}+1) - g^2\bar{a}^2((n-1)n\bar{a}-1))}{n\sigma^2(g_0(n-1)n+g)} + \frac{r\sigma^2\bar{a}(g_0(-(n-1)n^2(\bar{a}+1) - gn((n+1)\bar{a}+1))}{n\sigma^2(g_0(n-1)n+g)}.$$

The numerator on the right-hand side is quadratic and concave in g , while the denominator is strictly positive. Thus, if there exists no real root g to the numerator, $f(\bar{a}, -\bar{a}, g)$ is negative. In particular, the equation $f(\bar{a}, -\bar{a}, g) = 0$ admits no real root g if the discriminant is negative. But this discriminant is exactly $D(\bar{a}, r)$, which is negative by definition of \bar{a} . \square

Lemma A.4 shows that there exists a solution (α, β, γ) to equations (11), (12), and (14) satisfying boundary conditions $\alpha_T = \alpha^m(\gamma_T)$, $\beta_T = \beta^m(\gamma_T)$, and $\gamma_0 = ng_0$ when (15) holds. Therefore, it only remains to establish the following:

Lemma A.5. *Suppose (15) holds, and let (α, β, γ) be a solution to equations (11), (12), and (14) with $\alpha_T = \alpha^m(\gamma_T)$, $\beta_T = \beta^m(\gamma_T)$, and $\gamma_0 = ng_0$. Then there exists a solution ξ to equation (13) on $[0, T]$ with $\xi_T = \xi^m(\gamma_T)$.*

Proof. Let $g_0 < \max\{r\sigma^2/\kappa(n), \sigma^2/(3nT)\}$ and let (α, β, γ) be as given in the lemma. We first establish the result for all $g_0 > 0$ sufficiently small.

Recall that for any $g_0 < \sigma^2/(3nT)$ we can bound α from below by x given in (A.12). In particular, for $g_0 \leq 7\sigma^2/(24nT)$, we have

$$0 \geq \alpha_t \geq x_t = \frac{1}{\sqrt[3]{3\frac{ng_0}{\sigma^2}(T-t) - 1}} \geq \frac{1}{\sqrt[3]{3\frac{ng_0T}{\sigma^2} - 1}} \geq -2.$$

Combining this with $0 \leq \gamma_t \leq \gamma_0 = ng_0$, we see that the coefficient on ξ_t^2 in (13), $\alpha_t\gamma_t/(n\sigma^2)$, is bounded in absolute value by $2g_0/\sigma^2$. Thus for any g_0 small enough, (13) is approximately linear in ξ_t and hence it has a solution on $[0, T]$.

Define now \bar{g}_0 as the supremum over \tilde{g}_0 such that a solution to the boundary value problem exists for all $g_0 \in (0, \tilde{g}_0)$. By the previous argument, $\bar{g}_0 > 0$. We complete the proof of the lemma by showing that $\bar{g}_0 \geq \max\{r\sigma^2/\kappa(n), \sigma^2/(3nT)\}$.

Suppose towards contradiction that $\bar{g}_0 < \max\{r\sigma^2/\kappa(n), \sigma^2/(3nT)\}$. Then for $g_0 = \bar{g}_0$ there exists a solution (α, β, γ) to (11), (12), and (14) satisfying the boundary conditions by Lemma A.4, but following equation (13) backwards from $\xi_T = \xi^m(\gamma_T)$ yields a function ξ that diverges to ∞ at some $\tau \in [0, T]$. We assume $\tau > 0$ without loss of generality, since if $\lim_{t \downarrow 0} \xi_t = \infty$, then ξ_t can be taken to be arbitrarily large for $t > 0$ small enough, which is all that is needed in what follows.

Since the boundary value problem has a solution for all $g_0 < \bar{g}_0$, a symmetric linear Markov equilibrium exists for all $g_0 < \bar{g}_0$. So fix any such g_0 and any firm i . The firm's equilibrium continuation payoff at time $s < \tau$ given state $(C^i, \Pi_s, \hat{\Pi}_s^i, s) = (0, 0, 0, s)$ is $V(0, 0, 0, s) = v_0(s)$. Total equilibrium profits are bounded from above by the (finite) profit of an omniscient planner operating all the firms, and hence $v_0(s)$ is bounded from above by the expectation of the planner's profit conditional on $C^i = \Pi_s = \hat{\Pi}_s^i = 0$ and γ_s . The expectation depends in general on g_0 (through γ_s and z_s), but we obtain a uniform bound by taking the supremum over $g_0 \leq \bar{g}_0$. Denote this bound by B .

Let $\Delta > 0$, and suppose firm i deviates and produces $Q_t^i = \beta_t \Pi_t + \delta_t - \Delta$ for all $t \in [s, \tau)$, and then reverts back to the equilibrium strategy at τ . Then $d(\Pi_t - \hat{\Pi}_t^i) = \lambda_t[\alpha_t n(\Pi_t - \hat{\Pi}_t^i) + \Delta]dt$ (see Section 4.1), and hence

$$\Pi_\tau - \hat{\Pi}_\tau^i = \Delta \int_s^\tau \exp\left(-\int_\tau^t \lambda_u \alpha_u n du\right) dt > 0. \quad (\text{A.13})$$

Since Π and Q^i still have linear dynamics on $[s, \tau)$, their expectation and variance are bounded, and hence so is firm i 's expected payoff from this interval. Moreover, since (α, β, γ) (and hence also $\delta = -\bar{p}(\alpha + \beta)$) exist and are continuous in g_0 at \bar{g}_0 , the supremum of this expected payoff over $g_0 \leq \bar{g}_0$ is then also finite.

Firm i 's continuation payoff from reverting back to the equilibrium best-response

policy $(\alpha, \beta, \delta, \xi)$ at time τ is given by

$$V(0, \pi, \hat{\pi}, \tau) = v_0(\tau) + v_1(\tau)\pi + v_2(\tau)\hat{\pi} + v_4(\tau)\pi\hat{\pi} + v_8(\tau)\pi^2 + v_9(\tau)\hat{\pi}^2 \geq 0,$$

where the inequality follows, since the firm can always guarantee zero profits by producing nothing. By inspection of (A.6) and (A.9), we observe that

- (i) $v_4(\tau) \propto -\xi_\tau$ and $v_8(\tau) \propto \xi_\tau$;
- (ii) $v_1(\tau)$ and $v_2(\tau)$ are independent of ξ ;
- (iii) $v_9(\tau)$ depends on ξ , but is either finite or tends to ∞ as ξ grows without bound;
- (iv) $v_0(\tau) = V(0, 0, 0, \tau) \geq 0$.

Therefore, letting $g_0 \rightarrow \bar{g}_0$ and hence $\xi_\tau \rightarrow \infty$, we have for all $\pi > 0 \geq \hat{\pi}$,

$$V(0, \pi, \hat{\pi}, \tau) \rightarrow \infty.$$

Moreover, such pairs $(\pi, \hat{\pi})$ have strictly positive probability under the deviation by (A.13), because $\hat{\Pi}^i$ is an exogenous Gaussian process. Together with the lower bound $V(0, \pi, \hat{\pi}, \tau) \geq 0$ for all $(\pi, \hat{\pi})$ this implies that the time- s expectation of the deviation payoff tends to infinity as $g_0 \rightarrow \bar{g}_0$, and hence it dominates B for g_0 close enough to \bar{g}_0 . But this contradicts the fact that a symmetric linear Markov equilibrium exist for all $g_0 < \bar{g}_0$. \square

A.4 Proofs for Section 5

We start with a lemma that is used in the proof of Corollary 2, and later in the proof of Proposition 5. Let $g_0/\sigma^2 < r/\kappa(n)$ so that a symmetric linear equilibrium exists for all T , and select for each T some such equilibrium $f^T := (\alpha^T, \beta^T, \delta^T, \gamma^T)$, where γ^T is the corresponding posterior variance. Extend each f^T to all of $[0, \infty)$ by setting $f^T(t) = f^T(T)$ for $t > T$. We continue to use f^T to denote the function so extended. Denote the sup-norm by $\|f^T\|_\infty := \sup_t \|f^T(t)\|$, where $\|f^T(t)\| := \max_i |f_i^T(t)|$.

Since $g_0/\sigma^2 < r/\kappa(n)$, each α^T is bounded in absolute value uniformly in T by some $\bar{a} < \infty$ (see the proof of Lemma A.4). Thus, $0 < \beta^T \leq -\alpha^T < \bar{a}$ and $0 < \delta^T = -\bar{p}(\alpha^T + \beta^T) < \bar{p}\bar{a}$ for all $T > 0$. This implies, in particular, that the “non- r term” on the right-hand side of \dot{f}_i^T is bounded in absolute value by $\gamma_t^T K$ for some $K < \infty$ independent of i and T .

Lemma A.6. *For all $\varepsilon > 0$, there exists $t_\varepsilon < \infty$ such that for all $T \geq t \geq t_\varepsilon$, $\|f^T(t) - (\alpha^m(0), \beta^m(0), \delta^m(0), 0)\| < \varepsilon$.*

Proof. For γ , the claim follows by Corollary 1. We prove the claim for α ; the same argument can be applied to β and δ . By Corollary 1, for any $\eta > 0$, there exists t_η

such that $0 \leq \gamma_t^T < \eta$ for all $T \geq t \geq t_\eta$. Furthermore, by taking t_η to be large enough, we also have $|\alpha^m(\gamma_t^T) + 1| < \eta$ for all $T \geq t \geq t_\eta$ by continuity of α^m . This implies, in particular, that $\alpha_t^T \leq \alpha^m(\gamma_t^T) < -1 + \eta$ for all $T \geq t \geq t_\eta$, establishing an upper bound on α^T uniformly in T .

To find a lower bound, fix $T > t_\eta$. Define $b : [t_\eta, T] \rightarrow \mathbb{R}$ as the unique solution to $\dot{b}_t = r(b_t + 1) + \eta K$ with $b_T = -1$, where K is the constant from the remark just before Lemma A.6. Then, by construction, $-1 - \eta K/r \leq b_t \leq -1$ for all t in $[t_\eta, T]$. Furthermore, we have $\alpha^T > b$ on $[t_\eta, T]$. To see this, note that $\alpha_T^T = \alpha^m(\gamma_T^T) > -1 = b_T$, and if for some t in $[t_\eta, T)$ we have $\alpha_t^T = b_t$, then

$$\begin{aligned} \dot{\alpha}_t^T &\leq r \frac{\alpha_t^T}{\alpha^m(\gamma_t^T)} (\alpha_t^T - \alpha^m(\gamma_t^T)) + \gamma_t^T K \\ &= r \frac{\alpha_t^T}{\alpha^m(\gamma_t^T)} (\alpha_t^T + 1) - r \frac{\alpha_t^T}{\alpha^m(\gamma_t^T)} (\alpha^m(\gamma_t^T) + 1) + \gamma_t^T K \\ &< r \frac{\alpha_t^T}{\alpha^m(\gamma_t^T)} (\alpha_t^T + 1) + \eta K \\ &\leq r(\alpha_t^T + 1) + \eta K = \dot{b}_t, \end{aligned}$$

where the first inequality is by definition of K , the second uses $\alpha^m(\gamma_t^T) \geq -1$ and $t \geq t_\eta$, and the third follows from $\alpha_t^T = b_t \leq -1 \leq \alpha^m(\gamma_t^T)$. Thus, at any point of intersection, α^T crosses b from above, and hence the existence of an intersection contradicts $\alpha_T^T > b_T$. We conclude that $\alpha_t^T > b_t \geq -1 - \eta K/r$ for all $T \geq t \geq t_\eta$. Note that even though b depends on T , the lower bound is uniform in T .

To conclude the proof, fix $\varepsilon > 0$, and put $\eta = \min\{\varepsilon, r\varepsilon/K\}$. Then, by the above arguments, there exists $t_\varepsilon = t_\eta$ such that $\alpha_t^T \in (-1 - \varepsilon, -1 + \varepsilon)$ for all $T \geq t \geq t_\varepsilon$. \square

Proof of Corollary 2. Corollary 1 and Lemma A.1 imply that for every $\eta > 0$, there exists $t_\eta < \infty$ such that for all $T > t_\eta$, every symmetric linear Markov equilibrium satisfies $\mathbb{P}[|\Pi_t - n^{-1} \sum_i C^i| < \eta] > 1 - \eta$ for all $t > t_\eta$. Furthermore, we have

$$|Q_t^i - q_i^N(C)| \leq |\alpha_t - \alpha^m(0)| |C^i| + |\beta_t - \beta^m(0)| |\Pi_t| + \beta^m(0) \left| \Pi_t - \frac{\sum_i C^i}{n} \right| + |\delta_t - \delta^m(0)|.$$

By the above observation about Π and Lemma A.6, each term on the right converges in distribution to zero as $t \rightarrow \infty$ (uniformly in T). Since zero is a constant, this implies that the entire right-hand side converges to zero in distribution. In particular, if we denote the right-hand side by X_t , then for any $\varepsilon > 0$, there exists t_ε such that for every $T \geq t \geq t_\varepsilon$, we have $\mathbb{P}[|X_t| < \varepsilon] \geq 1 - \varepsilon$. But $\{|X_t| < \varepsilon\} \subset \{|Q_t^i - q_i^N(C)| < \varepsilon\}$, and hence $\mathbb{P}[|Q_t^i - q_i^N(C)| < \varepsilon] > 1 - \varepsilon$. \square

Proof of Proposition 1. (1.) Consider a symmetric linear Markov equilibrium (α, β, δ) with posterior variance γ . Denote the myopic equilibrium values by

$$(\alpha_t^m, \beta_t^m, \delta_t^m) := (\alpha^m(\gamma_t), \beta^m(\gamma_t), \delta^m(\gamma_t)).$$

By Theorem 1, (α, β) are a part of a solution to the boundary value problem, and hence δ satisfies (A.8). The boundary conditions require that $\alpha_T = \alpha_T^m < 0$ and $\beta_T = \beta_T^m > 0$. We first show that $\alpha \leq 0$ for all t . This is immediate, since $\alpha_T < 0$ and $\dot{\alpha}_t = 0$ if $\alpha_t = 0$. Next, we show that δ_t lies everywhere above its (constant) myopic value δ_t^m . To establish this, notice that $\delta_T = \delta_T^m$, and $\dot{\delta}_T < 0$ by (A.8). Furthermore

$$\delta_t = \delta_t^m \Rightarrow \dot{\delta}_t - \dot{\delta}_t^m = \frac{(n-1)p\alpha_t\beta_t\gamma_t}{n(n+1)^2\sigma^2} \leq 0.$$

Now suppose towards a contradiction that β_t crosses β_t^m from below at some $t < T$. Then evaluate $\dot{\beta}_t$ at the crossing point and obtain

$$\beta_t = \beta_t^m \Rightarrow \dot{\beta}_t - \dot{\beta}_t^m = -\frac{g_0^2(n-1)^3n^3\alpha_t\gamma_t((n+1)\alpha_t-1)}{(n+1)^3\sigma^2(g_0(n-1)n+(n+1)\gamma_t)^2} < 0,$$

a contradiction. Therefore $\beta_t \geq \beta_t^m$.

The results shown above ($\alpha_t \leq 0$, $\delta_t/\bar{p} = -\alpha_t - \beta_t \geq 1/(n+1)$, and $\beta_t \geq \beta_t^m$) imply that, if for some t , $\alpha_t = \alpha_t^m$, then also $\beta_t = \beta_t^m$, since $-\alpha_t^m - \beta_t^m = 1/(n+1)$. Using this we evaluate $\dot{\alpha}_t$ at $\alpha_t = \alpha_t^m$ to obtain

$$(\alpha_t, \beta_t) = (\alpha_t^m, \beta_t^m) \Rightarrow \dot{\alpha}_t - \dot{\alpha}_t^m = \frac{g_0(n-1)^2n\gamma_t(g_0(n-1)n+\gamma_t)^3}{(n+1)\sigma^2(g_0(n-1)n+(n+1)\gamma_t)^4} > 0,$$

which establishes $\alpha_t \leq \alpha_t^m$ for all t .

(2.–3.) The boundary conditions imply $\gamma_0 = ng_0$. Substituting into $\dot{\alpha}_t$ gives

$$\dot{\alpha}_0 = -r(2\alpha_0+1)\alpha_0 - \frac{g_0\alpha_0^2\beta_0}{\sigma^2} < 0,$$

since both terms are negative as by part (1.), $-\alpha_0 \leq -\alpha_0^m = 1/2$. Similarly, we have

$$\dot{\beta}_0 = \frac{r\alpha_0(n-1-2(n+1)\beta_0)}{n+1} + \frac{g_0\alpha_0\beta_0(2n\alpha_0+(n-1)\beta_0)}{(n+1)\sigma^2} > 0,$$

since $n \geq 2$, $\alpha_t + \beta_t < 0$, and $\beta_t > \beta_t^m$. Boundary conditions $(\alpha_T, \beta_T) = (\alpha_T^m, \beta_T^m)$

imply

$$\begin{aligned}\dot{\alpha}_T &= \frac{(n-1)\gamma_T z_T ((n^2-1)z_T - n^2 - 1)}{n(n+1)\sigma^2 (n+1 - z_T(n-1))^4}, \\ \dot{\beta}_T &= \frac{(n-1)\gamma_T z_T ((n-1)^3 z_T^2 - (n+1)(n(n+4) - 1)(n-1)z_T + n(n+1)^3)}{n(n+1)^3 \sigma^2 (n+1 - z_T(n-1))^4}.\end{aligned}$$

Note that as $\gamma_T \rightarrow 0$ and hence $z_T \rightarrow \frac{n}{n-1}$, we have $\dot{\alpha}_T \rightarrow \frac{(n-1)\gamma_T}{(n+1)\sigma^2} > 0$ and $\dot{\beta}_T \rightarrow -\frac{n(n^2+n-2)\gamma_T}{(n+1)^3\sigma^2} < 0$. Finally, because $|\alpha_t|$ is bounded away from zero at all t , we have $\gamma_T \rightarrow 0$ as $T \rightarrow \infty$, and hence the derivatives have the desired signs for T large enough. (4.) That δ is eventually decreasing follows by evaluating (A.8) at $t = T$ using the boundary condition $\delta_T = \delta_T^m$ and signing the terms using part (1.).

(5.) If $r = 0$, (A.8) simplifies to

$$\dot{\delta}_t = \frac{(n-1)\alpha_t\beta_t\gamma_t(\delta_t - n\alpha_t(z_t - 1)((n+1)\delta_t - \bar{p}))}{n(n+1)\sigma^2} < 0,$$

since $\alpha_t < 0$ and $(n+1)\delta_t \geq \bar{p} = (n+1)\delta_t^m$ by part (1.).

Now consider the second time derivative $\ddot{\alpha}_t$, and evaluate it at a critical point of α_t . Solving $\dot{\alpha}_t = 0$ for g_0 and substituting into the second derivative, we obtain

$$\ddot{\alpha}_t = -\frac{\alpha_t^3\beta_t\gamma_t^2(n\alpha_t + 1)((n-1)n\alpha_t - 1)}{n^3\sigma^4} > 0,$$

since $n \geq 2$ and $\alpha_t \leq -1/2$.

Finally, we evaluate $\ddot{\beta}$ at a critical point of β . To this end, note that for $r = 0$,

$$\dot{\beta}_t = \frac{\alpha_t\beta_t\gamma_t}{n(n+1)\sigma^2} [n\alpha_t(1+n-z_t(n-1) - (n^2-1)\beta_t(z_t-1)) + (n-1)\beta_t].$$

At a critical point, the term in parentheses is nil. Since $\alpha_t < 0$, the second derivative $\ddot{\beta}_t$ is then proportional to

$$-\dot{\alpha}_t(1+n-z_t(n-1) - (n^2-1)\beta_t(z_t-1)) + \alpha_t\dot{z}_t(n-1 + (n^2-1)\beta_t).$$

We know z_t is strictly increasing, $\alpha_t < 0$, and the last term in parentheses is positive. Furthermore, $\dot{\beta}_t = 0$ implies $(1+n-z_t(n-1) - (n^2-1)\beta_t(z_t-1)) > 0$. Finally, $\delta_t = -\bar{p}(\alpha_t + \beta_t)$ from Theorem 1 implies that α_t is strictly increasing at a critical point of β_t . Therefore, both terms in $\ddot{\beta}_t$ are negative establishing quasiconcavity. \square

Proof of Proposition 2. (1.) The signaling components obviously vanish at T as then also the equilibrium play is myopic. Evaluate the slope of $\hat{\alpha}$ and $\hat{\alpha}^{br}$ at $t = T$.

We obtain

$$\dot{\hat{\alpha}}_T - \dot{\hat{\alpha}}_T^{br} = -\frac{\gamma_T(n-1)^2 z_T ((n-1)z_T - 2n)}{2n(n+1)^2 \sigma^2 (-(n-1)z_T + n+1)^3} > 0,$$

since $z_T \leq n/(n-1)$ implies both that the numerator is negative and that the denominator is positive. Because $\hat{\alpha}_T = \hat{\alpha}_T^{br}$, the signaling component $\hat{\alpha}_t - \hat{\alpha}_t^{br}$ is thus negative in a neighborhood of T . Now solve $\hat{\alpha}_t = \hat{\alpha}_t^{br}$ for z_t and substitute the resulting expression into $\dot{\hat{\alpha}}_t - \dot{\hat{\alpha}}_t^{br}$. We obtain,

$$\hat{\alpha}_t = \hat{\alpha}_t^{br} \Rightarrow \dot{\hat{\alpha}}_t - \dot{\hat{\alpha}}_t^{br} = \frac{(n-1)\alpha_t \beta_t \gamma_t ((n-1)\alpha_t - 1)}{2n(n+1)\sigma^2} > 0.$$

Thus, if $\hat{\alpha}_t - \hat{\alpha}_t^{br} = 0$ for some $t < T$, then the signaling component crosses zero from below at t , contradicting the fact that it is negative for all t close enough to T . We conclude that $\hat{\alpha}_t - \hat{\alpha}_t^{br} > 0$ for all $t < T$.

Now evaluate the slope of $\hat{\beta}$ and $\hat{\beta}^{br}$ at $t = T$. We obtain

$$\dot{\hat{\beta}}_T - \dot{\hat{\beta}}_T^{br} = -\frac{\gamma_T(n-1)^3 z_T}{2n(n+1)^2 \sigma^2 (n(-z_T) + n + z_T + 1)^3} < 0.$$

Because $\hat{\beta}_T = \hat{\beta}_T^{br}$, the signaling component $\hat{\beta}_t - \hat{\beta}_t^{br}$ is positive in a neighborhood of T . Solve $\hat{\beta}_t = \hat{\beta}_t^{br}$ for z_t and substitute the resulting expression into $\dot{\hat{\beta}}_t - \dot{\hat{\beta}}_t^{br}$. We obtain,

$$\hat{\beta}_t = \hat{\beta}_t^{br} \Rightarrow \dot{\hat{\beta}}_t - \dot{\hat{\beta}}_t^{br} = -\frac{(n-1)^2 \alpha_t^2 \beta_t \gamma_t}{2n(n+1)\sigma^2} < 0.$$

Thus, if the signaling component $\hat{\beta}_t - \hat{\beta}_t^{br}$ ever crosses zero it does so from above, contradicting the fact that it is positive at $t = T$.

Direct calculation yields $\delta_t - \delta_t^{br} = \frac{1}{2}((n+1)\delta_t - \bar{p}) \geq 0$, where the inequality follows since $\delta_t \geq \delta^m(\gamma_t) = \bar{p}/(n+1)$ by Proposition 1.1 and (10). Furthermore, by inspection of (A.8), $\dot{\delta}_t < 0$ if $\delta_t = \delta^m(\gamma_t)$, and thus $\delta_t > \bar{p}/(n+1)$ for all $t < T$.

(2.) Consider $\hat{\alpha}_t - \hat{\alpha}_t^{br}$, and suppose there exists a time t for which the signaling component has a slope of zero. Impose $r = 0$, solve $\dot{\hat{\alpha}}_t - \dot{\hat{\alpha}}_t^{br} = 0$ for β_t , and substitute into $\hat{\alpha}_t - \hat{\alpha}_t^{br}$. We obtain

$$\hat{\alpha}_t - \hat{\alpha}_t^{br} = \frac{(n-1)\alpha_t - 1}{2n(n+1)\alpha_t(z_t - 1) - 2} > 0,$$

contradicting our finding that $\hat{\alpha}_t \leq \hat{\alpha}_t^{br}$ for all t .

Likewise, we know the signaling component $\hat{\beta}_t - \hat{\beta}_t^{br}$ is decreasing at $t = T$. Now impose $r = 0$, and consider the slope $\dot{\hat{\beta}}_t - \dot{\hat{\beta}}_t^{br}$ at an arbitrary t . We obtain

$$\dot{\hat{\beta}}_t - \dot{\hat{\beta}}_t^{br} = -\frac{(n-1)\alpha_t \beta_t \gamma_t (n\alpha_t(z_t - 1) ((n+1)\beta_t + (n-1)\alpha_t z_t) - \beta_t)}{2n\sigma^2 z_t}.$$

If the slope of the signaling component satisfies $\dot{\hat{\beta}}_t \geq \dot{\hat{\beta}}_t^{br}$, then it must be that $(n+1)\beta_t + (n-1)\alpha_t z_t \leq 0$. However, the level of the signaling component is given by

$$\hat{\beta}_t - \hat{\beta}_t^{br} = \frac{(n+1)\beta_t + (n-1)\alpha_t z_t}{2z_t}.$$

Consider the largest t for which the signaling component has a slope of zero. Then the signaling component must be negative at that point. This contradicts our earlier finding that the signaling component is positive and decreasing in a neighborhood of T . Therefore, $\dot{\hat{\beta}}_t < \dot{\hat{\beta}}_t^{br}$ for all t .

Since $\delta_t - \delta_t^{br} = \frac{1}{2}((n+1)\delta_t - \bar{p})$, the claim follows by Proposition 1.4. \square

Proof of Proposition 3. (1.) The result follows from the properties of the expected total output established in the text before the proposition.

(2.) Firm i 's output on the equilibrium path is given by $Q_t^i = \alpha_t C^i + \beta_t \Pi_t + \delta_t$. Therefore, for any i and $j \neq i$, we have $Q_t^i - Q_t^j = \alpha_t (C^i - C^j)$. Proposition 1 shows that α is nonmonotone for T sufficiently large. \square

Proof of Proposition 4. We begin by constructing the distribution of Π_t under the true data-generating process. Substituting the equilibrium strategies into the law of motion for Π_t in Lemma A.1, we obtain $d\Pi_t = \lambda_t \alpha_t (n\Pi_t - \sum_i C^i) dt + \lambda_t \sigma dZ_t$, or

$$\begin{aligned} \Pi_t = \pi_0 \exp\left(\int_0^t n\lambda_s \alpha_s ds\right) - \sum_i C^i \int_0^t \lambda_s \alpha_s \exp\left(\int_s^t n\lambda_u \alpha_u du\right) ds \\ + \sigma \int_0^t \lambda_s \exp\left(\int_s^t n\lambda_u \alpha_u du\right) dZ_s. \end{aligned}$$

We conclude that conditional on C , Π_t is normally distributed with mean

$$\mathbb{E}[\Pi_t | C] = \pi_0 \exp\left(\int_0^t n\lambda_s \alpha_s ds\right) - \sum_i C^i \int_0^t \lambda_s \alpha_s \exp\left(\int_s^t n\lambda_u \alpha_u du\right) ds,$$

and variance

$$\text{Var}[\Pi_t | C] = \sigma^2 \int_0^t \lambda_s^2 \exp\left(2 \int_s^t n\lambda_u \alpha_u du\right) ds.$$

Recall also that $n\alpha_t \lambda_t = \dot{\gamma}_t / \gamma_t$, and hence $\exp(\int_s^t n\lambda_u \alpha_u du) = \gamma_t / \gamma_s$. We thus have

$$\mathbb{E}[\Pi_t | C] = \pi_0 \frac{\gamma_t}{\gamma_0} - \sum_i C^i \frac{1}{n} \int_0^t \frac{\dot{\gamma}_s}{\gamma_s} \frac{\gamma_t}{\gamma_s} ds = \pi_0 \frac{\gamma_t}{\gamma_0} - \frac{1}{n} \sum_i C^i \gamma_t \left(\frac{1}{\gamma_0} - \frac{1}{\gamma_t}\right),$$

and

$$\text{Var}[\Pi_t | C] = -\frac{1}{n^2} \int_0^t \dot{\gamma}_s \frac{\gamma_t^2}{\gamma_s^2} ds = \frac{1}{n^2} \gamma_t^2 \left(\frac{1}{\gamma_t} - \frac{1}{\gamma_0}\right).$$

Thus, conditional on the realized costs, firm i 's expected time- t flow profit is given by

$$\left(\bar{p} - C^i - \alpha_t \sum_{j=1}^n C^j - \beta_t n \mathbb{E}[II_t | C] - \delta_t n\right) (\alpha_t C^i + \beta_t \mathbb{E}[II_t | C] + \delta_t) - \beta_t^2 n \text{Var}[II_t | C].$$

Taking an expectation with respect to C , we obtain its *ex ante* expected time- t profit

$$W_t := \frac{\beta_t \gamma_t ((2\alpha_t + \beta_t)n + 1) - g_0 n (n(\alpha_t + (\alpha_t + \beta_t)^2) + \beta_t)}{n^2} - (\bar{p} - \pi_0)^2 (\alpha_t + \beta_t)(n(\alpha_t + \beta_t) + 1).$$

(1.) Recalling that $\gamma_0 = ng_0$, we have

$$W_0 - W^{\text{co}} = -g_0 \left[\frac{n^2 + n - 1}{(n+1)^2} + \alpha_t(\alpha_t + 1) \right] + (p - \pi_0)^2 \left[-(\alpha_t + \beta_t)(n(\alpha_t + \beta_t) + 1) - \frac{1}{(n+1)^2} \right].$$

Because $n \geq 2$ and $\alpha_t \leq 1/2$, the coefficient on g_0 is negative. The coefficient on $\bar{p} - \pi_0$ is negative as well because $\alpha_t + \beta_t \leq 1/(n+1)$. Similarly, using the terminal values of the equilibrium coefficients, we have

$$W_T - W^{\text{co}} = -\frac{g_0(n-1)n\gamma_T [g_0 n(2n^2 + n - 3) + (n+1)(n+3)\gamma_T]}{[g_0 n(n^2 - 1) + (n+1)^2 \gamma_T]^2},$$

which is negative because the coefficient on g_0 inside the brackets is positive for $n \geq 2$.

(2.) Since α is bounded away from zero, γ is a strictly decreasing function. We can thus use the standard change-of-variables formula to write the equilibrium coefficients and expected profits as a function of γ instead of t . Letting $\pi_0 = \bar{p}$ and $\Delta W(\gamma) := W(\gamma) - W^{\text{co}}$, this gives

$$\Delta W(\gamma) = -g_0 n^2 \left[\alpha(\gamma)(\alpha(\gamma) + 1) + \frac{n^2 + n - 1}{(n+1)^2} \right] - (g_0 n - \gamma)\beta(\gamma) (2n\alpha(\gamma) + 1 + n\beta(\gamma)).$$

To prove the claim, we show that $\Delta W(\gamma) > 0$ for some γ under the conditions stated in the proposition. Indexing equilibria by the terminal value of γ , we show first that for every $\gamma_T > 0$ sufficiently small (and thus every T large enough), there exists $\gamma^* > \gamma_T$ such that $\alpha^{\gamma_T}(\gamma^*) = -1$. Moreover, $\gamma^* \rightarrow 0$ as $\gamma_T \rightarrow 0$. To this end, note that each function α^{γ_T} satisfies the differential equation

$$\alpha'(\gamma) = \frac{\beta(\gamma)(\gamma + (n-1)n\alpha(\gamma)(g_0 n - \gamma) + g_0(n-1)n)}{\gamma n(\gamma + g_0(n-1)n)} + \frac{nr\sigma^2(\gamma + \alpha(\gamma)(\gamma + n(\gamma + g_0(n-1))) + g_0(n-1)n)}{\gamma^2 n\alpha(\gamma)(\gamma + g_0(n-1)n)}, \quad (\text{A.14})$$

with boundary condition $\alpha(\gamma_T) = \alpha^m(\gamma_T)$. The coefficient of $\beta(\gamma)$ is negative for $\gamma = \gamma_T$ sufficiently small, and it can cross zero only once. Hence we can obtain an upper bound on $\alpha(\gamma)$ for γ small by replacing $\beta(\gamma)$ with its myopic value $\beta^m(\gamma)$ defined in (10) and considering the resulting differential equation

$$\begin{aligned}\hat{\alpha}'(\gamma) &= \frac{\beta^m(\gamma)(\gamma + (n-1)n\hat{\alpha}(\gamma)(g_0n - \gamma) + g_0(n-1)n)}{\gamma n(\gamma + g_0(n-1)n)} \\ &\quad + \frac{nr\sigma^2(\gamma + \hat{\alpha}(\gamma)(\gamma + n(\gamma + g_0(n-1)))) + g_0(n-1)n}{\gamma^2 n\hat{\alpha}(\gamma)(\gamma + g_0(n-1)n)}\end{aligned}$$

with $\hat{\alpha}(\gamma_T) = \alpha^m(\gamma_T)$. Since $\beta^m(\gamma) \leq \beta(\gamma)$ for all $\gamma \geq 0$, we then have $\hat{\alpha}(\gamma) > \alpha(\gamma)$ for all γ such that

$$(\gamma + (n-1)n\alpha(\gamma)(g_0n - \gamma) + g_0(n-1)n) < 0. \quad (\text{A.15})$$

Now consider the derivative $\hat{\alpha}'(\gamma)$ and substitute in the value $\hat{\alpha}(\gamma) = -1$. The resulting expression is positive if and only if

$$\gamma \geq \bar{\gamma} := \frac{g_0(g_0(n-1)^3n - n(n^2-1)r\sigma^2)}{g_0(n-1)((n-1)n+1) + (n+1)^2r\sigma^2}.$$

The threshold $\bar{\gamma}$ is strictly positive if and only if

$$\frac{r\sigma^2}{g_0} < \frac{(n-1)^2}{n+1}. \quad (\text{A.16})$$

Moreover, then for every $\hat{\gamma} \in (0, \bar{\gamma})$, there exists $\hat{\gamma}_T > 0$ such that $\hat{\alpha}^{\hat{\gamma}_T}(\hat{\gamma}) = -1$. To see this, note that we can simply follow the ODE for $\hat{\alpha}$ to the left from the initial value $\hat{\alpha}(\hat{\gamma}) = -1$. As $\alpha^m(\hat{\gamma}) > -1$, the solution lies below the continuous function α^m at $\hat{\gamma}$. Moreover, it lies strictly above -1 at all $\gamma < \hat{\gamma}$. Since $\alpha^m(0) = -1$, there thus exists some $\gamma \in (0, \hat{\gamma})$ at which $\hat{\alpha}(\gamma) = \alpha^m(\gamma)$, which is our desired $\hat{\gamma}_T$. Finally, since the function $\hat{\alpha}^{\hat{\gamma}_T}$ is decreasing at $\hat{\gamma}$, the left-hand side of (A.15) is negative at $\hat{\gamma}$, which implies that $\hat{\alpha}^{\hat{\gamma}_T}$ bounds the corresponding $\alpha^{\hat{\gamma}_T}$ from above on $[\hat{\gamma}_T, \hat{\gamma}]$, and thus

$$-1 = \hat{\alpha}^{\hat{\gamma}_T}(\hat{\gamma}) > \alpha^{\hat{\gamma}_T}(\hat{\gamma}) \quad \text{for all } \hat{\gamma} < \bar{\gamma}.$$

But $\alpha^{\hat{\gamma}_T}(\hat{\gamma}_T) > -1$, so by the intermediate value theorem, there exists some $\gamma^* \in (\hat{\gamma}_T, \hat{\gamma})$ such that $\alpha^{\hat{\gamma}_T}(\gamma) = -1$. Moreover, since $\hat{\alpha}^{\hat{\gamma}_T}$ are solutions to the same ODE with different initial conditions, their paths cannot cross, which implies that $\hat{\gamma}$ is a monotone function of $\hat{\gamma}_T$. Letting $\hat{\gamma} \rightarrow 0$ then implies that $\gamma^* \rightarrow 0$ as $\hat{\gamma}_T \rightarrow 0$ as desired.

Now consider the difference $\Delta W(\gamma)$. Because $\Delta W(\gamma)$ is increasing in β , we use $\beta^m(\gamma)$ to bound $\beta(\gamma)$. We then substitute the value $\alpha(\gamma) = -1$ into $\Delta W(\gamma)$. By

the above argument, we can take the corresponding $\gamma > 0$ to be arbitrarily small by choosing T large enough if (A.16) holds. But evaluating the difference $\Delta W(\gamma)$ at $\gamma = 0$, we obtain

$$\Delta W(0) = \frac{g_0(n-1)^3 n (n^2 + n - 1)}{(n^2 - 1)^2} > 0,$$

so the result follows by continuity of $W(\gamma)$ in γ .

Finally, in the table below, we report the values of the upper bound \bar{r} defined by condition (A.16) and of the lower bound \underline{r} defined by the sufficient condition for existence, $g_0/\sigma^2 < r/\kappa(n)$, when $g_0/\sigma^2 = 1$. By inspection, $\underline{r} < \bar{r}$ for all $1 \leq n \leq 10$, confirming that the conditions in the proposition are compatible with our sufficient condition for existence.

n	2	3	4	5	6	7	8	9	10
\underline{r}	0.32	0.87	1.48	2.10	2.74	3.38	4.02	4.66	5.31
\bar{r}	0.33	1.00	1.80	2.66	3.57	4.50	5.44	6.40	7.36

(3.) A calculation analogous to the one for ΔW yields the difference between ex-ante expected consumer surplus and its complete-information level

$$\begin{aligned} \Delta CS := & \frac{1}{2} \left[g_0 n (\alpha_t + \beta_t)^2 - \frac{g_0 n}{(n+1)^2} - \beta_t \gamma_t (2\alpha_t + \beta_t) \right] \\ & + \frac{1}{2} (\bar{p} - \pi_0)^2 \left[n^2 (\alpha_t + \beta_t)^2 - \frac{n^2}{(n+1)^2} \right]. \end{aligned}$$

Because $\alpha + \beta \leq -1/(n+1)$, the coefficient on the term $\bar{p} - \pi_0$ is positive. Furthermore, we can use $\delta_t = -\bar{p}(\alpha_t + \beta_t)$ and $\delta_t \geq 1/(n+1)$ to bound β_t from above by $-\alpha_t - 1/(n+1)$ to obtain $\Delta CS \geq \frac{1}{2} \gamma_t [\alpha_t^2 - (n+1)^{-2}] > 0$. \square

A.5 Proofs for Section 6

We prove Proposition 5 in two main steps. First, we show that finite-horizon equilibria converge along a subsequence to a strategy profile of the infinite horizon game that is a solution to the corresponding HJB equation. Second, we show that the value under this limit strategy profile satisfies a transversality condition and hence constitutes a solution to each player's best response problem.

As a preliminary observation, we note that $g_0/\sigma^2 < 4r/(27n)$ strengthens the first case in (15) and hence α^T , β^T , δ^T , and γ^T are bounded uniformly in T (see beginning of Section A.4). Moreover, then $-n\alpha \geq \xi \geq 0$, and hence ξ^T is uniformly bounded as well. To see the last claim, note that $n\alpha^m + \xi^m = 0$. Therefore, for all T , we have $-n\alpha_T = \xi_T$. Now consider the sum $n\dot{\alpha}_t + \dot{\xi}_t$ and evaluate it at $n\alpha_t + \xi_t = 0$. We obtain

$$n\dot{\alpha}_t + \dot{\xi}_t = -\frac{n\alpha_t}{2\sigma^2 (g_0(n-1)n + \gamma_t)} (g_0(n-1)n (r\sigma^2 + \alpha_t\beta_t\gamma_t) + 2r\sigma^2\gamma_t).$$

Because the fraction is positive, we can bound γ_t in the term in parentheses with ng_0 and 0 respectively to bound the right-hand side from below. Thus, if $ng_0\alpha_t\beta_t+r\sigma^2 > 0$, then the function $-n\alpha$ crosses ξ from above only, and then $-n\alpha_T = \xi_T$ implies that $\xi < -n\alpha$ for all t . Because $\beta < -\alpha$, this clearly holds if $\alpha > a$ for some $a > -3/2$. The existence of such a constant a can be shown by first verifying that α is bounded from below by the solution to

$$\dot{y}_t = -ry_t(y_t + 1) + \frac{ng_0}{\sigma^2}y_t^4, \quad y_T = -1,$$

and then verifying that $y_t > -3/2$ when $g_0/\sigma^2 < 4r/(27n)$. We omit the details.

We adopt the notation introduced in the beginning of Section A.4, but redefine $f^T := (\alpha^T, \beta^T, \delta^T, \xi^T, \gamma^T)$ to include ξ^T . Note that Lemma A.6 continues to hold for f^T so redefined. Finally, note that each f^T satisfies $\dot{f}^T(t) = F(f^T(t))$ at every $t < T$, where $F : [-B, B]^5 \rightarrow \mathbb{R}^5$ is the continuous function on the right-hand side of our boundary value problem (written here including δ). By continuity, F is bounded on its compact domain implying that the functions $\{f^T\}$ are equi-Lipschitz.

Lemma A.7. *Any sequence $\{f^T\}$ of symmetric linear Markov equilibria contains a subsequence $\{f^{T_n}\}$ that converges uniformly to a continuously differentiable $f : [0, \infty) \rightarrow \mathbb{R}^5$ that satisfies $\dot{f} = F(f)$ and $\lim_{t \rightarrow \infty} f(t) = (\alpha^m(0), \beta^m(0), \delta^m(0), \xi^m(0), 0)$.*

Proof. The family $\{f^T\}$ is uniformly bounded and equi-Lipschitz and hence of locally bounded variation uniformly in T . Thus, Helly's selection theorem implies that there exists a subsequence of horizons $\{T_n\}_{n \in \mathbb{N}}$ with $T_n \rightarrow \infty$ such that f^T converges pointwise to some function f as $T \rightarrow \infty$ along the subsequence. We show that this convergence is in fact uniform.

Suppose to the contrary that there exists $\varepsilon > 0$ and a collection of times $\{T_k, t_k\}_{k \in \mathbb{N}}$ such that $\{T_k\}$ is a subsequence of $\{T_n\}$ and $\|f^{T_k}(t_k) - f(t_k)\| > \varepsilon$ for every k . By Lemma 2, there exists $t_\varepsilon < \infty$ such that for all $T_n \geq t \geq t_\varepsilon$, we have $\|f^{T_n}(t) - (x^*, 0)\| < \varepsilon/2$. Since $f^{T_n}(t) \rightarrow f(t)$ as $n \rightarrow \infty$, we then have $\|f^{T_n}(t) - f(t)\| < \varepsilon$ for all $T_n \geq t \geq t_\varepsilon$. This implies that t_k belongs to the compact interval $[0, t_\varepsilon]$ for all sufficiently large k , which in turn implies that no subsequence of $\{f^{T_k}\}$ converges uniformly on $[0, t_\varepsilon]$. But $\{f^{T_k}\}$ are uniformly bounded and equi-Lipschitz (and thus equicontinuous) and $[0, t_\varepsilon]$ is compact, so this contradicts the Arzela-Ascoli theorem. We therefore conclude that $\{f^{T_n}\}$ converges uniformly to f .

For differentiability of f , note first that uniform convergence of f^{T_n} to f implies that $\dot{f}^{T_n} = F(f^{T_n}) \rightarrow F(f)$ uniformly on every interval $[0, t]$, since F is continuous on a compact domain and hence uniformly continuous. Define $h : \mathbb{R}_+ \rightarrow \mathbb{R}^5$ by

$$h_i(t) := f_i(0) + \int_0^t F_i(f(s))ds, \quad i = 1, \dots, 5.$$

We conclude the proof by showing that $h = f$. As $f^{T_n} \rightarrow f$, it suffices to show that $f^{T_n} \rightarrow h$ pointwise. For $t = 0$ this follows by definition of h , so fix $t > 0$ and $\varepsilon > 0$. Choose N such that for all $n > N$, we have $\|f^{T_n}(0) - h(0)\| < \varepsilon/2$ and $\sup_{s \in [0, t]} \|\dot{f}^{T_n}(s) - F(f(s))\| < \varepsilon/(2t)$. Then for all $n > N$,

$$\|f^{T_n}(t) - h(t)\| \leq \|f^{T_n}(0) - h(0)\| + \left\| \int_0^t \dot{f}^{T_n}(s) ds - \int_0^t F(f(s)) ds \right\| < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Thus $f = h$ and $\dot{f} = \dot{h} = F(f)$. The limit of f as $t \rightarrow \infty$ follows by Lemma A.6. \square

Since the limit function $f = (\alpha, \beta, \delta, \xi, \gamma)$ satisfies the boundary value problem, we may construct a value function V of the form (9) as in the proof of Lemma A.3. Then the policy $(\alpha, \beta, \delta, \xi)$ satisfies the first-order condition (8) by construction and thus achieves the maximum on the right-hand side of (7). Hence, it remains to show that the transversality condition holds. In what follows, we use the fact that in the infinite-horizon game, a strategy Q is feasible if (i) $\mathbb{E}[\int_0^t Q_s^2 ds] < \infty$ for all $t \geq 0$, and (ii) the firm obtains a finite expected payoff. We need the following two lemmas.

Lemma A.8. *For any feasible strategy Q ,*

$$\lim_{t \rightarrow \infty} e^{-rt} v(t) \mathbb{E}[\Pi_t^Q] = \lim_{t \rightarrow \infty} e^{-rt} v(t) \mathbb{E}[\hat{\Pi}_t] = \lim_{t \rightarrow \infty} e^{-rt} v(t) \mathbb{E}[\hat{\Pi}_t^2] = 0$$

for any function v of polynomial growth. Also, $\limsup_{t \rightarrow \infty} e^{-rt} \mathbb{E}[(\Pi_t^Q)^2] < \infty$.

Proof. Regarding $\hat{\Pi}$, suppose that $(\Pi_0, \hat{\Pi}_0) = (\pi, \hat{\pi})$. Then, it is easy to see that

$$\hat{\Pi}_t = \hat{\pi} \hat{R}_{t,0} + c(1 - \hat{R}_{t,0}) + \int_0^t \hat{R}_{t,s} \sigma \lambda_s dZ_s.$$

where $\hat{R}_{t,s} := \exp(\int_s^t \lambda_u \alpha_u [1 + (n-1)(1-z_u)] du)$, $s < t$, is a discount factor (i.e., $\lambda_u \alpha_u [1 + (n-1)(1-z_u)] < 0$). In particular,

$$\mathbb{E}[\hat{\Pi}_t] = \hat{\pi} \hat{R}_{t,0} + c(1 - \hat{R}_{t,0}) < \max\{c, \hat{\pi}\}.$$

Also, by uniform boundedness,

$$\mathbb{E}\left[\left(\int_0^t \hat{R}_{t,s} \sigma \lambda_s dZ_s\right)^2\right] = \mathbb{E}\left[\int_0^t \hat{R}_{t,s}^2 \sigma^2 \lambda_s^2 ds\right] \leq K_1 t$$

for some $K_1 > 0$. Hence, $\mathbb{E}[\hat{\Pi}_t^2] \leq K_0 + K_1 t$. The limits for $\hat{\Pi}$ follow directly.

Regarding $(\Pi_t^Q)_{t \geq 0}$, letting $\tilde{R}_{t,s} := \exp(\int_s^t \lambda_u [n\alpha_u + \beta_u] du)$, we have that

$$\Pi_t^Q = \pi \tilde{R}_{t,0} + \int_0^t \tilde{R}_{t,s} \lambda_s [\delta_s - (n-1)\alpha_s (z\hat{\Pi}_s + (1-z_s)c)] ds + \int_0^t \tilde{R}_{t,s} \lambda_s Q_s ds + \int_0^t \tilde{R}_{t,s} \lambda_s \sigma dZ_s.$$

Defining $\mathbb{E}[I_t^1] := \int_0^t \tilde{R}_{t,s} \lambda_s \mathbb{E}[Q_s] ds$, Cauchy-Schwartz inequality implies

$$\mathbb{E}[I_t^1] \leq \left(\int_0^t \tilde{R}_{t,s}^2 \lambda_s^2 ds \right)^{1/2} \left(\int_0^t \mathbb{E}[Q_s]^2 ds \right)^{1/2} < Kt \left(\mathbb{E} \left[\int_0^t Q_s^2 ds \right] \right)^{1/2}.$$

Hence,

$$e^{-rt} \mathbb{E}[I_t^1] < e^{-rt/2} Kt \left(e^{-rt} \mathbb{E} \left[\int_0^t Q_s^2 ds \right] \right)^{1/2} < e^{-rt/2} Kt \left(\mathbb{E} \left[\int_0^\infty e^{-rs} Q_s^2 ds \right] \right)^{1/2},$$

where the last term is finite by feasibility of Q . Hence, $e^{-rt} \mathbb{E}[I_t^1] \rightarrow 0$. It is easy to verify that all other terms also converge to zero once discounted, and this also occurs when they are accompanied by v of polynomial growth. Thus, $e^{-rt} v(t) \mathbb{E}[\Pi_t^Q] \rightarrow 0$.

To conclude, in studying $e^{-rt} \mathbb{E}[(\Pi_t^Q)^2]$ the only non-trivial terms are

$$A_t := \left(\int_0^t \tilde{R}_{t,s} \lambda_s Q_s ds \right)^2 \quad \text{and} \quad B_t := \int_0^t \tilde{R}_{t,s} \lambda_s Q_s ds \int_0^t \tilde{R}_{t,s} \lambda_s \sigma dZ_s.$$

(For the others the limit exists and takes value zero.) Observe first that there is $\epsilon > 0$ such that $\tilde{R}_{t,s} < e^{-\epsilon \int_s^t \lambda_u du}$ for all $0 \leq t < \infty$; this follows from $n\alpha + \beta < 0$ and $\lim_{t \rightarrow \infty} n\alpha + \beta < 0$. Thus, from Cauchy-Schwartz and the fact that $\lambda < C$, some $C > 0$,

$$\begin{aligned} A_t &\leq \left(\int_0^t \tilde{R}_{t,s}^2 \lambda_s ds \right) \left(\int_0^t \lambda_s Q_s^2 ds \right) \leq C^2 \left(\int_0^t e^{-2\epsilon \int_s^t \lambda_u du} \lambda_s ds \right) \left(\int_0^t Q_s^2 ds \right) \\ &= C \frac{1 - e^{-2\epsilon \int_0^t \lambda_u du}}{2\epsilon} \left(\int_0^t Q_s^2 ds \right) < \tilde{C} \left(\int_0^t Q_s^2 ds \right). \end{aligned}$$

Consequently, $e^{-rt} \mathbb{E}[A_t] \leq \tilde{C} \mathbb{E} \left[e^{-rt} \int_0^t Q_s^2 ds \right] \leq \tilde{C} \mathbb{E} \left[\int_0^\infty e^{-rs} Q_s^2 ds \right] < \infty$, by feasibility. We conclude that $\limsup e^{-rt} \mathbb{E}[A_t] < \infty$.

Regarding B_t , by applying Cauchy-Schwartz again, we have

$$\mathbb{E}[B_t] \leq \mathbb{E} \left[\left(\int_0^t \tilde{R}_{t,s} \lambda_s Q_s ds \right)^2 \right]^{1/2} \mathbb{E} \left[\left(\int_0^t \tilde{R}_{t,s} \lambda_s \sigma dZ_s \right)^2 \right]^{1/2},$$

where the second term is bounded by some $(L_0 + L_1 t)^{1/2}$. Using the previous argument for A_t gives

$$e^{-rt} \mathbb{E}[A_t]^{1/2} \leq e^{-rt/2} v(t) \tilde{C}^{1/2} \mathbb{E} \left[e^{-rt} \int_0^t Q_s^2 ds \right]^{1/2} \leq e^{-rt/2} \tilde{C}^{1/2} \mathbb{E} \left[\int_0^\infty e^{-rs} Q_s^2 ds \right]^{1/2},$$

where the last term is finite by feasibility. Thus, $e^{-rt} \mathbb{E}[B_t] \leq e^{-rt} \mathbb{E}[A_t]^{1/2} (L_0 + L_1 t)^{1/2} \rightarrow 0$. It is easy to show that the rest of the terms in $\mathbb{E}[(\Pi_t^Q)^2]$ converge to zero using similar (and simpler) arguments. Hence, $\limsup e^{-rt} \mathbb{E}[(\Pi_t^Q)^2] < \infty$. \square

Lemma A.9. *Under the limit strategy $(\alpha, \beta, \delta, \xi)$, the system (A.9) admits on $[0, +\infty)$ a bounded solution for which $\lim_{t \rightarrow \infty} v_i(t)$ exists for each i , and the system (A.6) defines v_k ($k = 1, 4, 5, 8$) that have at most linear growth.*

Proof. Let $\theta := r + [2\alpha^2\gamma(n(1-z) + z)]/n\sigma^2$. Notice that because $z \leq n/(n-1)$, $\theta_t > r > 0$. It is easy to see that for $s > t$,

$$v_9(s)e^{-\int_0^s \theta_u du} - v_9(t)e^{-\int_0^t \theta_u du} = - \int_t^s e^{-\int_0^u \theta_v dv} [(n-1)\alpha_u z_u + \xi_u]^2 du.$$

We look for a solution such that $v_9(s) \exp(-\int_0^s \theta_u du) \rightarrow 0$ as $s \rightarrow \infty$. If it exists, then

$$v_9(t) = \int_t^\infty e^{-\int_t^s \theta_v dv} [(n-1)\alpha_s z_s + \xi_s]^2 ds.$$

Because $(n-1)\alpha_s z_s + \xi_s$ is uniformly bounded and $\theta > r > 0$, the right-hand side exists, and it is uniformly bounded. Hence, it corresponds to our desired solution. Moreover, the limit value of v_9 is, by L'Hopital's rule

$$\lim v_9(t) = \lim \frac{[(n-1)\alpha_s z_s + \xi_s]^2}{\theta_t} = \frac{[-n + n/2]^2}{r} = \frac{n^2}{4r}.$$

The other equations in (A.9) have similar solutions (i.e., taking the form of a net present value, with a finite limit value), and they can be found in an iterative fashion.

Solving $v_k(t)\alpha_t\lambda_t$ ($k = 1, 4, 5, 8$) as a function of the limit coefficients from (A.6) and using $\lim_{t \rightarrow \infty} f(t)$ from Lemma A.7, we see that $v_k(t)\alpha_t\lambda_t \rightarrow 0$. Because $\alpha_t \rightarrow -1$, $\gamma_t \in O(1/(a+bt))$, and $\lambda_t \propto \alpha_t\gamma_t$, this implies that $v_k(t)$ grows at most linearly. \square

We are now ready to show that the transversality condition holds (see, *e.g.*, Pham, 2009, Theorem 3.5.3).

Lemma A.10. *Under any feasible strategy Q , $\limsup_{t \rightarrow \infty} e^{-rt}\mathbb{E}[V(C, \Pi_t^Q, \hat{\Pi}_t, t)] \geq 0$. Moreover, under the limit strategy $(\alpha, \beta, \delta, \xi)$, the limit exists and it takes value zero.*

Proof. It obviously suffices to show the result conditional on any realized c . We first check the lim sup. Terms involving v_i , $i = 0, 1, 2, 3, 5, 6, 7, 9$ in V converge to zero by the last two lemmas. For the v_4 term, Cauchy-Schwartz implies

$$e^{-rt}v_4(t)\mathbb{E}[\Pi_t^Q \hat{\Pi}_t] \leq e^{-rt/2}v_4(t)\mathbb{E}[\hat{\Pi}_t^2]^{1/2} e^{-rt/2}\mathbb{E}[(\Pi_t^Q)^2]^{1/2},$$

where $e^{-rt/2}v_4(t)\mathbb{E}[\hat{\Pi}_t^2]^{1/2} \rightarrow 0$ as v_4 is at most $O(t)$ and $\mathbb{E}[\hat{\Pi}_t^2]$ is linear. By Lemma A.8, $\limsup e^{-rt}\mathbb{E}[(\Pi_t^Q)^2] < \infty$. Thus $e^{-rt}v_4(t)\mathbb{E}[\Pi_t^Q \hat{\Pi}_t] \rightarrow 0$ as $t \geq 0$. We deduce that the lim sup is non-negative by noticing that $e^{-rt}v_8(t)\mathbb{E}[(\Pi_t^Q)^2] \geq 0$ as $v_8 \geq 0$.

Since all terms except for $e^{-rt}v_8(t)\mathbb{E}[(\Pi_t^Q)^2]$ converge to zero under any feasible strategy, it remains to show that, under the limit strategy Q^* , $e^{-rt}v_8(t)\mathbb{E}[(\Pi_t^{Q^*})^2] \rightarrow 0$.

However, this is straightforward once we observe that

$$\begin{aligned} \Pi_t^{Q^*} &= \pi R_t - c R_t \int_0^t R_{t,s} \lambda_s \alpha_s [1 + (n-1)(1-z_s)] ds \\ &\quad + \int_0^t R_{t,s} \lambda_s \sigma dZ_s + \int_0^t R_{t,s} \lambda_s [\xi_s + (n-1)\alpha_s z_s] \hat{\Pi}_s ds. \end{aligned}$$

Indeed, because (i) $\mathbb{E}[(\int_0^t R_{t,s} \lambda_s \sigma dZ_s)^2]$ and $\mathbb{E}[\hat{\Pi}_t^2]$ grow at most linearly, (ii) the functions $(\alpha, \beta, \xi, z, \lambda)$ are all uniformly bounded, and (iii) $R_{t,s}$ is a discount rate, it is easy to verify that all terms in $\mathbb{E}[(\Pi_t^{Q^*})^2]$ decay to zero once discounted by e^{-rt} . \square

References

- ATHEY, S., AND K. BAGWELL (2008): “Collusion with Persistent Cost Shocks,” *Econometrica*, 76(3), 493–540.
- AUMANN, R., AND M. MASCHLER (1995): *Repeated Games with Incomplete Information*. The MIT press.
- BACK, K., C. H. CAO, AND G. A. WILLARD (2000): “Imperfect Competition among Informed Traders,” *Journal of Finance*, 55(5), 2117–2155.
- BERNHARDT, D., AND B. TAUB (2015): “Learning about Common and Private Values in Oligopoly,” *Rand Journal of Economics*, 46(1), 66–85.
- CISTERNAS, G. (2015): “Two-Sided Learning and Moral Hazard,” Discussion paper, MIT Sloan.
- DAVIS, M. H. A. (1979): “Martingale Methods in Stochastic Control,” in *Lecture Notes in Control and Information Sciences*, ed. by F. Allgöwer, and M. Morari, vol. 16, pp. 85–117. Springer.
- DAVIS, M. H. A., AND P. VARAIYA (1973): “Dynamic Programming Conditions for Partially Observable Stochastic Systems,” *SIAM Journal of Control*, 11(2), 226–261.
- DORASZELSKI, U., G. LEWIS, AND A. PAKES (2014): “Just Starting Out: Learning and Price Competition in a New Market,” Discussion paper, Harvard University.
- DORASZELSKI, U., AND M. SATTERTHWAIT (2010): “Computable Markov Perfect Industry Dynamics,” *The RAND Journal of Economics*, 41(2), 215–243.
- ERICSON, R., AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62(1), 53–82.

- ESCOBAR, J. F., AND J. TOIKKA (2013): “Efficiency in Games with Markovian Private Information,” *Econometrica*, 81(5), 1887–1934.
- FERSHTMAN, C., AND A. PAKES (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” *Quarterly Journal of Economics*, 127(4), 1611–1661.
- FOSTER, F. D., AND S. VISWANATHAN (1996): “Strategic trading when agents forecast the forecasts of others,” *The Journal of Finance*, 51(4), 1437–1478.
- FRIEDMAN, A. (1971): *Differential Games*. Courier Corporation.
- FUDENBERG, D., AND J. TIROLE (1986): “A “Signal-Jamming” Theory of Predation,” *RAND Journal of Economics*, 17(3), 366–376.
- GAL-OR, E. (1985): “Information Sharing in Oligopoly,” *Econometrica*, 53(2), 329–343.
- HÖRNER, J., AND S. LOVO (2009): “Belief-Free Equilibria in Games With Incomplete Information,” *Econometrica*, 77(2), 453–487.
- HÖRNER, J., S. TAKAHASHI, AND N. VIEILLE (forthcoming): “Truthful Equilibria in Dynamic Bayesian Games,” *Econometrica*.
- KALLIANPUR, G. (1980): *Stochastic Filtering Theory*. Springer-Verlag, New York.
- KYLE, A. S. (1985): “Continuous Auctions and Insider Trading,” *Econometrica*, 53(6), 1315–1335.
- (1989): “Informed Speculation with Imperfect Competition,” *Review of Economic Studies*, 56(3), 317–55.
- LAMBERT, N. S., M. OSTROVSKY, AND M. PANOV (2014): “Strategic Trading in Informationally Complex Environments,” Discussion paper, Stanford University.
- LIPTSER, R. S., AND A. N. SHIRYAEV (1977): *Statistics of random processes, I*. Springer, New York.
- MAILATH, G. J. (1989): “Simultaneous Signaling in an Oligopoly Model,” *Quarterly Journal of Economics*, 104(2), 417–427.
- MESTER, L. (1992): “Perpetual Signalling with Imperfectly Correlated Costs,” *RAND Journal of Economics*, 23(4), 548–563.

- MILGROM, P., AND J. ROBERTS (1982a): “Limit Pricing and Entry under Incomplete Information: an Equilibrium Analysis,” *Econometrica*, 50(2), 443–459.
- (1982b): “Predation, Reputation, and Entry Deterrence,” *Journal of Economic Theory*, 27(2), 280–312.
- MIRMAN, L. J., L. SAMUELSON, AND A. URBANO (1993): “Duopoly Signal Jamming,” *Economic Theory*, 3(1), 129–149.
- NYARKO, Y. (1997): “Convergence in Economic Models with Bayesian Hierarchies of Beliefs,” *Journal of Economic Theory*, 74, 266–296.
- PESKI, M. (2014): “Repeated Games with Incomplete Information and Discounting,” *Theoretical Economics*, 9(3), 651–694.
- PHAM, H. (2009): *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer-Verlag, Berlin.
- RAITH, M. (1996): “A General Model of Information Sharing in Oligopoly,” *Journal of Economic Theory*, 71(1), 260–288.
- RIORDAN, M. H. (1985): “Imperfect Information and Dynamic Conjectural Variations,” *RAND Journal of Economics*, 16(1), 41–50.
- SADZIK, T., AND C. WOOLNOUGH (2014): “Multidimensional Uncertainty and Hypersensitive Asset Prices,” Discussion paper, UCLA and NYU.
- SANNIKOV, Y. (2007): “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica*, 75(5), 1285–1329.
- TOWNSEND, R. M. (1983): “Forecasting the Forecasts of Others,” *Journal of Political Economy*, 91(4), 546–588.
- VIVES, X. (1984): “Duopoly Information Equilibrium: Cournot and Bertrand,” *Journal of Economic Theory*, 34(1), 71–94.
- (2011): “Strategic Supply Function Competition with Private Information,” *Econometrica*, 79(6), 1919–1966.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 76(6), 1375–1411.
- YONG, J., AND X. Y. ZHOU (1999): *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer-Verlag, New York.