

Dynamic Oligopoly with Incomplete Information

ALESSANDRO BONATTI

Sloan School of Management, MIT

GONZALO CISTERNAS

Sloan School of Management, MIT

and

JUUSO TOIKKA

MIT

First version received February 2016; final version accepted August 2016 (Eds.)

We consider learning and signalling in a dynamic 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 play depends on the history 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 costs and beliefs under the equilibrium strategies are non-monotone over time. We explain this by decomposing incentives into signalling and learning, and discuss implications for prices, quantities, and welfare.

Key words: Dynamic oligopoly, Asymmetric information, Learning, Signalling, Continuous time

JEL Codes: C73, D43, D82, L13

1. INTRODUCTION

In the theory of oligopoly, asymmetric information plays a central role in explaining anti-competitive practices such as limit pricing and predation (Milgrom and Roberts, 1982a, 1982b), or price rigidity in cartels (Athey and Bagwell, 2008). However, the fundamental question of how competition unfolds in a new market where firms start out with incomplete information about each other has received comparatively little attention. In such a market, a firm may be trying to simultaneously learn its competitors' types from the observation of market variables and influence their beliefs about its own type. This possibility of multi-sided learning and signalling makes the setting fascinating but creates a technical challenge responsible for the scarcity of results: the analyst must track the evolution of the firms' beliefs over time.

In this article, we make progress on the problem by using continuous-time methods to provide the first tractable analysis of Markov (perfect) equilibria in a dynamic oligopoly with incomplete information. More specifically, we study a stylized Cournot game where each firm privately knows its own cost and only observes the market price, which is subject to unobservable demand

shocks. The resulting equilibrium dynamics capture the jockeying for position among oligopolists before the market reaches its long-run equilibrium. We study how the firms' strategic behaviour is shaped by learning and signalling, and derive implications for the time paths of prices, profits, and consumer surplus.

To address the tractability of beliefs, we consider a linear-quadratic Gaussian environment: the market demand function and the firms' cost functions are linear in quantities, the constant marginal costs are drawn once and for all from a symmetric normal prior distribution, 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 along the equilibrium path, 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 behaviour is to allow current outputs to depend on the history only through the firms' beliefs about the costs. But when individual 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 a weighted sum of its own cost and the public posterior expectation about the average industry cost conditional on past prices. In other words, its own cost and the public belief are sufficient statistics for a firm's private belief. 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 resulting bias in the public belief.

The representation of beliefs yields a characterization of 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 in turn 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 non-linear 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 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: linear Markov equilibria fail to exist if the signalling incentive is sufficiently strong. On the other hand, we can say surprisingly much about the properties of such equilibria.

As quantities are strategic substitutes, each firm has an incentive to deviate up from the myopic output to manipulate its competitors' beliefs. This additional output has a deterministic component common to all cost types, which in equilibrium has no learning consequences.² However, the incentive to expand output is stronger the lower the firm's own cost: not only is it cheaper to do so, a low-cost firm also benefits more from the other firms scaling back their production. As a result, the equilibrium has the firms actively signalling their costs through the market price. We show that, in any symmetric linear Markov equilibrium, the equilibrium price carries enough statistical information about industry costs for the firms to asymptotically learn the average cost

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

2. This component is analogous to signal-jamming in environments with symmetric uncertainty. See, e.g., Holmström (1999), Riordan (1985), Fudenberg and Tirole (1986), or Mirman *et al.* (1993).

of their rivals. The identification problem caused by the *ex ante* symmetry of firms and the one-dimensional price prevents learning the costs of individual firms. But knowing the average is enough for the firms to play their complete information best response. Thus, equilibrium play converges asymptotically to the static complete information outcome for the realized costs.

We then show that the interplay of learning and signalling leads to a rich and interesting set of predictions. The key observation is that the equilibrium strategy assigns non-monotone weights to private and public information over time. In particular, the weight each firm assigns to its own cost is the largest (in absolute value) in the intermediate term. By decomposing the equilibrium strategy into a myopic and a forward-looking component, this can be understood as arising from two monotone effects. The myopic component, which only reflects learning, grows over time. Roughly, high-cost firms scale back their production further over time because they expect their rivals to be more aggressive as they become better informed. The forward-looking component, which captures signalling, decreases over time. This is because the firms' estimates of their competitors' costs become more precise, beliefs become less sensitive to price changes, and the incentive to signal diminishes.

Because the firms assign non-monotone weights to their own costs, the difference between any two firms' outputs is non-monotone over time conditional on their realized costs. In particular, efficient firms benefit the most from their cost advantage at intermediate times. From an industry standpoint, the higher sensitivity of individual outputs to costs improves the allocation of production, leading to higher profits. In fact, this allows us to even derive conditions under which the *ex ante* expectation of industry profits (as well as total surplus) is higher in the intermediate term than under complete information.

Finally, we observe that signalling drives the expected total quantity above the corresponding (static) complete information level. In turn, this implies that the expected market price is depressed below its complete information level. Moreover, we show that at any point in time, consumers only observing historical prices expect prices to increase in the future as the signalling incentive diminishes.

The above prediction about prices concerns an aggregate variable found in market-level data. In this vein, we also show that the volatility of total market output conditional on costs eventually decreases, but not necessarily monotonically so. Because of linear demand, this implies that the volatility of the average price eventually vanishes as well. As the drift also disappears, the price only exhibits variation due to demand shocks in the long run.

Some of our other predictions concern firm-level data as they involve individual costs and outputs. The result about the non-monotonicity of the equilibrium strategy is one such example. Namely, the coefficient on the firm's own cost in the equilibrium strategy determines the sensitivity of the firm's output to its cost. Thus, the shape of this coefficient is simply a prediction about the time pattern of the cost-sensitivity of individual firm output in a new market, which could readily be measured in firm-level data.

Our results are in line with some recent empirical evidence. In particular, a study of a newly deregulated British electricity market by Doraszelski *et al.* (2016) finds that (i) prices were low and increasing during an initial phase, but in the long run (ii) the price process settled down, and (iii) the market converged to a stable state akin to a static equilibrium. They interpret the findings to be the result of firms learning to play an equilibrium of a complete information game using some adaptive learning rule. Our analysis provides a complementary perspective, showing that the patterns can be qualitatively matched also by the equilibrium of a dynamic game of incomplete information where forward-looking firms learn about their competitors' types.³ However, our

3. The microstructure of the British market differs from that in our model. However, Cournot competition is often adopted in studies of electricity markets even when it is not descriptively accurate, see, for example,

model is unlikely to generate the later period of decreasing prices in the British market, but as the data pertain to a single realization of a time series, this could be random variation before convergence to static equilibrium. Alternatively, our model could be extended with a deterministic trend (in the demand or the costs), trumped by the initial signalling activity to still yield increasing prices early on.

Related Literature. Our work contributes to the literature on continuous-time games with Brownian information. In particular, our model adds incomplete information to the imperfect-monitoring games studied by Sannikov (2007). We share his focus on a fixed discount factor. However, Sannikov studies the entire set of public perfect equilibria, whereas we characterize a class of Markov equilibria.⁴ Prior work considering incomplete information includes Faingold and Sannikov (2011), who study reputation dynamics in the context of a population of small players facing a long-run player who may be a behavioural type, as well as Daley and Green (2012) and Dilmé (2014), who study one-sided signalling with a binary type and Brownian noise. Cisternas (2015) develops methods for games where uncertainty about the state of the world is symmetric in equilibrium but private beliefs arise after deviations. In contrast, in our game beliefs are private even on the equilibrium path.

The early literature on incomplete information in dynamic oligopoly considers issues such as limit pricing, predation, and reputation using models with one-sided information. See Milgrom and Roberts (1982a,b) and Fudenberg and Tirole (1986) among others. Mailath (1989) and Mester (1992) construct separating equilibria in two and three-period oligopoly games where all firms have private costs and actions are observable. More recently, Athey and Bagwell (2008) study collusion in a Bertrand oligopoly with persistent private costs. They identify conditions under which the best equilibrium for patient firms has all types pooling at the same price and no learning takes place. In contrast, we fix the discount rate, and learning and signalling are central to our analysis. There is also recent work on the role of information in static oligopoly; see Myatt and Wallace (2015) for Cournot competition, or Vives (2011) and Bernhardt and Taub (2015) for supply-function equilibria.

Finally, 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 multi-agent model by Foster and Viswanathan (1996) mentioned above, and its continuous-time version by Back *et al.* (2000). We share their focus on linear equilibria in a Gaussian environment. (Our results can be used to show that the *ad hoc* restriction on strategies in these works is equivalent to requiring them to be symmetric, linear, and Markov in our sense.) However, trading in a financial market with common values differs starkly from product market competition under strategic substitutes and private values. In the former, players limit trades in order to retain their informational advantage, whereas in the latter, they engage in excess production to signal low costs to discourage their rivals, leading to qualitatively different

Borenstein and Bushnell (1999) or Bushnell *et al.* (2008). Furthermore, Hortacsu and Puller (2008) show in an additively-separable supply-function model with private information that, in equilibrium, firms price optimally against residual demand as in Cournot competition.

4. The use of continuous-time methods and the focus on Markov equilibria also distinguishes our work from the literature on repeated Bayesian games with fully or partially persistent types. This literature has almost exclusively restricted attention to patient players, typically focusing on cooperative equilibria. See, for example, Aumann and Maschler (1995), Hörner and Lovo (2009), Escobar and Toikka (2013), Peşki (2014), or Hörner *et al.* (2015). There is also a literature on learning in repeated games of incomplete information under myopic play, see Nyarko (1997).

equilibrium behaviour. The differences between the games also result in the analysis being substantially different.⁵

Outline. We introduce the model in the next section and consider beliefs under linear strategies in Section 3. We then turn to Markov strategies and equilibria in Section 4, and discuss their properties in Section 5. We consider the infinite horizon case and discuss other possible extensions in Section 6. Section 7 concludes. 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 independent draws from a normal distribution with mean π_0 and variance g_0 .⁶

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

$$dY_t = (\bar{p} - \sum_i Q_t^i) dt + \sigma dZ_t, \quad Y_0 = 0, \tag{1}$$

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 (cf. Keller and Rady, 1999; Sannikov and Skrzypacz, 2007). Heuristically, the current market price dY_t/dt is given by a linear demand curve perturbed by additive i.i.d. noise. Thus, with slight abuse of terminology, we refer to the firms' observation of Y as the firms observing the market price.

A pure strategy for a firm determines current output as a function of the firm's cost, past revenues, and own past outputs. However, because of the noise in the revenue process, no firm can ever observe that another firm has deviated from a given strategy.⁷ 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. We thus define a strategy to be only a function of the firm's cost and revenues, leaving off-path behaviour unspecified. This notion of strategy extends 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) . 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 (weak) solution $Y \in L_2[0, T]$.⁸ We define the expected payoff of firm i under an admissible

5. In the finance models, the price is set by a market maker and the players' linear flow payoffs are determined by differences in beliefs. In contrast, we have an exogenous demand curve, quadratic payoffs, and cost levels matter. As a result, the equilibrium in the former is essentially characterized by a single ordinary differential equation for the market depth (see Back *et al.*, 2000) whereas it is not possible to further reduce our boundary value problem.

6. See Section 6 for the infinite horizon case and for a discussion of the symmetry, independence, and private value assumptions.

7. As 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 .

8. More precisely, the game takes place on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. The state space $\Omega = \mathbb{R}^n \times C[0, T]$ is the space of all possible cost realizations and (continuous) paths of Y . The filtration $\{\mathcal{F}_t\}$ is defined as follows. Let $\mathcal{B} = \mathcal{B}_1 \otimes \dots \otimes \mathcal{B}_n$ be the product sigma-algebra on \mathbb{R}^n generated by (C^1, \dots, C^n) , and let $\mathcal{B}^i = \{\emptyset, \mathbb{R}\} \otimes \dots \otimes \{\emptyset, \mathbb{R}\} \otimes \mathcal{B}_i \otimes \{\emptyset, \mathbb{R}\} \otimes \dots \otimes \{\emptyset, \mathbb{R}\}$. Let $\{\bar{\mathcal{F}}_t\}$ be the canonical filtration on $C[0, T]$, where each $\bar{\mathcal{F}}_t$ is generated by sets

strategy profile to be

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

where $r \geq 0$ is the common discount rate. The equality is by (1). It shows that the payoff can be thought of as the expected present value of the observable flow payoff $Q_t^i(dY_t - C^i dt)$. 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. 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 that

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

A profile of linear strategies is *symmetric* if the functions (α, f, δ) are the same for all firms.⁹ Our interest is in Nash equilibria in symmetric linear strategies that condition on the history only through its effect on beliefs about costs and calendar time. Such equilibria, defined formally below, are a natural extension of Markov perfect equilibrium to our setting.

As will be clear in what follows, the Gaussian information structure obtained under linear strategies is essential for the analysis. The key aspects of the model enabling this are the quadratic payoff function defined by the left-hand side of (2) and the monitoring technology defined by (1), under which the players observe a noisy public signal of the sum of everyone's actions. A similar analysis can be carried out for other quadratic stage games under this monitoring structure. Note that linear strategies require abstracting from corner solutions by allowing negative outputs, and the normal prior requires negative costs, explaining these two simplifying assumptions. However, with an appropriate choice of parameters, the probability of negative outputs and costs can be taken to be arbitrarily small because of the convergence of equilibrium play (see Corollary 2 below).¹⁰

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 $B_t(Y^t) := \int_0^t f_s^t dY_s + \delta_t$ depends only on public information.

$\{f \in C[0, T]: f_s \in \Gamma\}$ with $s \leq t$ and Γ a Borel set in \mathbb{R} . (Heuristically, $\{\bar{\mathcal{F}}_t\}$ amounts to observing the past of the process Y .) Now, define $\{\mathcal{F}_t\}$ by $\mathcal{F}_t = \mathcal{B} \otimes \bar{\mathcal{F}}_t$. A strategy for player i is a process Q^i progressively measurable with respect to $\{\mathcal{F}_t^i\}$, where $\mathcal{F}_t^i = \mathcal{B}^i \otimes \bar{\mathcal{F}}_t$. (Note that Q^i is also measurable with respect to $\{\mathcal{F}_t\}$.) Under an admissible strategy profile, the probability \mathbb{P} is the unique product measure on $\Omega = \mathbb{R}^n \times C[0, T]$ consistent with (1) and the normal prior on costs.

9. A 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 $q_t^i = \alpha_t c_i + \sum_{s < t} f_s^t (y_s - y_{s-1}) + \delta_t$. Equation (3) can be viewed as a limit of such strategies.

10. Requiring the updating to be Gaussian and the firms' best responses to be linear prevents us from using a non-linear transformation to keep quantities or costs non-negative.

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 revenue, and hence the relevant signal for firm i about the other firms' costs is

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

Therefore, firm i 's belief can be derived by applying the Kalman filter with Y^i as the signal and $C^{-i} := (C^1, \dots, C^{i-1}, C^{i+1}, \dots, C^n)$ 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 , the posterior belief of firm i at time $t \in [0, T]$ 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 non-increasing 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' behaviour 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.

More specifically, consider the posterior expectation about the average firm cost conditional on the revenue process Y under a symmetric linear strategy profile (α, f, δ) . This *public belief* is

defined as $\Pi_t := \frac{1}{n} \mathbb{E}[\sum_j C^j | \mathcal{F}_t^Y]$, with posterior variance $\gamma_t := \mathbb{E}[(\sum_j C^j - n\Pi_t)^2 | \mathcal{F}_t^Y]$.¹¹ 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, given α , the posterior variance is a decreasing function of time given by

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

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 non-decreasing function of t .

That is, on the path of play, 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 distributions, since under linear strategies the system is Gaussian. Moreover, since the variance 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 .¹²

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 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 $\hat{\Pi}_t^i$ is defined by

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

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

12. In 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.

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. (We give an interpretation for its law of motion in the context of the best-response analysis in Section 4.1.) 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 $(C^i, \hat{\Pi}^i, t)$.¹³

Special cases of Lemmas 2 and 3 were first derived in discrete time by Foster and Viswanathan (1996), who considered a restricted class of strategies; our results extend the argument to all symmetric linear strategy profiles.

4. MARKOV EQUILIBRIA

In games of complete information, a Markov (perfect) equilibrium allows behaviour to depend only on the payoff-relevant part of history. In our model, only 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 behaviour 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 novel characterization of symmetric linear Markov strategies.¹⁴

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.¹⁵

13. If 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.

14. The finance literature on insider trading simply assumes that strategies condition only on the initial private signal and the market makers' belief (which equals the market price). Lemma 4 can be applied to that setting to show that this assumption is equivalent to strategies being linear and Markov (in our sense).

15. Our strategies only prescribe behaviour on the path of play, so the observation in footnote 12 implies that if we replace "firm's belief" with "firm's hierarchy of beliefs" in the definition of a Markov strategy profile, then Lemma 4 continues to hold verbatim, as do all our other results. We have chosen to impose the Markov restriction in terms of the firms' (private) first-order beliefs to avoid having to introduce higher-order beliefs.

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

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.¹⁶ 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 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 follows by 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 follows. Let $V(c, \pi, \hat{\pi}^i, t)$ denote the optimal time- t continuation value of firm i with cost $C^i = c$, public belief $\Pi_t = \pi$, and counterfactual public belief $\hat{\Pi}_t^i = \hat{\pi}$.¹⁷ The Hamilton–Jacobi–Bellman (HJB) equation for the firm's problem is then

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

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 $z_t \hat{\pi} + (1 - z_t)c$. Note that of all the terms on the second line in (7), only the first one depends on q .

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 affect the drift of the public belief Π .

16. Noting 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.

17. We use upper case to denote random variables and lower case to denote their realizations (*i.e.* scalars).

It is worth noting that here continuous time greatly simplifies the analysis. Similar 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 term convex in q . Moreover, this term involves the value function—an endogenous object—making it hard to establish the existence and uniqueness of an optimum. In contrast, in continuous time q affects the drift of π , which in turn affects the value 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

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

for some differentiable $v_k: \mathbb{R} \rightarrow \mathbb{R}$, $k=0, \dots, 9$. Moreover, a linear optimal policy exists on and off the path of play.¹⁸ 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 in the form

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

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. Nevertheless, this argument shows that optimal off-path behaviour can be taken to be 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.

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

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

18. The 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).

In the proof of the following result, we show that these are the equilibrium coefficients for myopic players as a function of current posterior variance x . In particular, firm i 's time- T equilibrium best-response 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

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

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

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

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

with boundary conditions $\alpha_T = \alpha^m(\gamma_T)$, $\beta_T = \beta^m(\gamma_T)$, $\xi_T = \xi^m(\gamma_T)$, and $\gamma_0 = n g_0$.

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.¹⁹ We establish the result by direct substitution into the ODE for δ . Since this is an equilibrium relationship, it does not 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 non-linear boundary value problem defined by equations (11)–(14)

19. For 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 $d = -\bar{p}(a + b)$.

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 behaviour. On the upside, we obtain as a by-product a relatively complete description of equilibrium behaviour, 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 $\gamma_0 = n g_0$. 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 $\gamma_F \in (0, n g_0]$. 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 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 contradicts the existence of an equilibrium for all $g_0 < \bar{g}_0$.

The sufficient condition (15) for existence in Theorem 1 is satisfied if players are sufficiently impatient or if the horizon is sufficiently short. The condition is not tight; numerical analysis suggests that equilibria exist for parameters in a somewhat larger range. However, it 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 signalling incentive becoming too strong. Consistent with this interpretation, (15) becomes harder to satisfy if r decreases or T increases, either of which makes signalling more valuable, or if g_0/σ^2 increases, which increases the scope for signalling. To see why increasing the number of firms n is also problematic, note that under linear strategies, it is the sum of the firms' costs that enters into (1) and into the firms' payoffs through Q . Thus, the relevant initial variance is that of the sum of costs, or $n g_0$, which is increasing in n .

5. EQUILIBRIUM PROPERTIES

We then turn to the properties of linear Markov equilibria and derive implications for the firms' strategic behaviour, prices, quantities, and welfare.

We first summarize properties of the equilibrium coefficients.

Proposition 1. *Any symmetric linear Markov equilibrium satisfies the following properties:*

1. $(-\alpha_t, \beta_t, \delta_t) \geq (-\alpha^m(\gamma_t), \beta^m(\gamma_t), \delta^m(\gamma_t)) > 0$ for all t .




FIGURE 1
Equilibrium coefficients, $(r, \sigma, n, \bar{p}, T, g_0) = (0.1, 1, 2, 5, 5, 2)$

2. α is initially decreasing and if T is sufficiently large, it is eventually increasing.²⁰
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 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 its cost and increasing in the public belief.

The second and third part of Proposition 1 imply that the equilibrium coefficients on the cost, α , and on the public belief, β , are necessarily non-monotone for T sufficiently large. As we discuss below, this seemingly surprising pattern is a natural consequence of learning and signalling. In contrast, the myopic 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).²¹

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. 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 behaviour. 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

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

21. These coefficients correspond to an equilibrium of a dynamic game where players are myopic, but where γ evolves according to the actual equilibrium coefficient α . In contrast, the true equilibrium for myopic players (*i.e.* for $r = \infty$) is given by the system defined by (10) and (14). The coefficients are pointwise smaller in absolute value in the latter than in the former, because in the true myopic equilibrium γ decreases more slowly due to the smaller α and the myopic coefficients are monotone functions of variance by (10).

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 $\mathbb{P}[\|Q_t - q^N(C)\| < \varepsilon] > 1 - \varepsilon$.²²*

The key to the proof is the fact that under the sufficient condition for existence invoked in the statement, 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. The independence of t_ε from the horizon 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. (In Section 6 we show that as $T \rightarrow \infty$, our equilibria converge to a symmetric linear Markov equilibrium of the infinite horizon version of the model.)

5.1. Signalling and learning

In order to explain the qualitative properties of equilibrium strategies, we consider here how signalling 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: each firm has an incentive to increase its output above the myopic level to lower the price in an attempt to lead its competitors to underestimate its cost.²³ Indeed, compared to the myopic coefficient δ^m , which is constant, the equilibrium δ is greater 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

22. Here, \mathbb{P} denotes the joint law of (C, Q_t) under the equilibrium strategies in the game with horizon T . Without the condition on the parameters, play still converges to the static complete information Nash equilibrium, but our proof for that case allows the critical time t_ε to depend on the horizon T .

23. See, *e.g.*, Riordan (1985), Fudenberg and Tirole (1986), or Mirman *et al.* (1993).

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 signalling. The myopic time- t best response to the equilibrium strategy is found by setting $\partial V / \partial \pi \equiv 0$ in the second term in (8). Expressed 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

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

which equals the second term in (8), is then by construction only due to signalling. Accordingly, we refer to the coefficients on the right-hand side of (16) as *signalling components*.

Proposition 2. *In any symmetric linear Markov equilibrium, the signalling components satisfy the following properties:*

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 .²⁴

Armed with Proposition 2, we are now in a position to explain equilibrium signalling and the non-monotonicity of the equilibrium coefficients. Note first that the *ex ante* expected signalling 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 signalling quantity is positive as the firms are engaging in excess production in an effort to convince their rivals to scale back production. Moreover, if $r = 0$, the expected excess output is monotonically decreasing over time, reflecting the shorter time to benefit from an induced reduction in the rivals' output and the fact that beliefs are less sensitive to output when the firms' estimates of their rivals' costs are more precise.

The costs and benefits of signalling depend on the firm's own cost and its belief about its rivals' costs. A lower cost firm makes it cheaper to produce additional output, and then leads to higher

24. As 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 $r = 0$.




FIGURE 2

Learning and signalling incentives, $(r, \sigma, n, \bar{p}, T, g_0) = (0, 1, 2, 1, 4, 1, 2)$

profits from the expansion of market share when the other firms scale back their outputs. This is captured by the signalling component $\hat{\alpha}_t - \hat{\alpha}_t^{br}$ in (16) being negative. If $r = 0$, it is decreasing in absolute value over time for the same reasons why the expected signalling quantity discussed above is decreasing and vanishing at the end.

The existence of the strictly positive signalling 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 their costs are high and that it will hence be producing aggressively in the future. Producing more is cheaper when the belief is higher as the other firms are then expected to produce less. Moreover, as a firm with a higher belief also expects to produce more in the future, it expects a larger benefit from its rivals scaling back their output. Again, this signalling component is monotone decreasing over time when $r = 0$.

Turning to the non-monotonicity of the equilibrium coefficients, consider Figure 2, which illustrates the equilibrium coefficients $\hat{\alpha}$ and $\hat{\beta}$, the coefficients $\hat{\alpha}^{br}$ and $\hat{\beta}^{br}$ of the myopic best response to the equilibrium strategy, and the signalling components $\hat{\alpha} - \hat{\alpha}^{br}$ and $\hat{\beta} - \hat{\beta}^{br}$. The dashed curves depict the implied coefficients on C^i and M_t^i under the myopic coefficients from (10), assuming that the evolution of the posterior variance γ is driven by the corresponding myopic α^m . That is, the dashed curves correspond to the equilibrium of the dynamic game when players are myopic (*i.e.* with $r = \infty$).

In the myopic equilibrium (dashed), the weights on C^i and M_t^i are increasing in absolute value as the firms' information becomes more precise (*i.e.* as γ decreases over time). As the myopic equilibrium reflects only the effect of learning, its properties are 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. 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, the firms increasingly rely on them to form 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.

The myopic best reply to the equilibrium strategy reflects these forces, but it is also affected by the shape of the equilibrium coefficients. The fact that the equilibrium $\hat{\beta}$ is larger than the corresponding weight in the myopic equilibrium means in the context of our auxiliary static game that firm i 's opponent is responding more aggressively to its private signal. In response, firm i 's myopic best reply places a higher weight $\hat{\alpha}^{br}$ on its own cost. This explains why $\hat{\alpha}^{br}$ lies below the lower of the dashed curves, except at $t=0$ when there is no private history. Proposition 1 shows that β (and thus β^{br}) is eventually decreasing, which explains why $\hat{\alpha}^{br}$ is eventually slightly increasing in Figure 2. Similarly, as the equilibrium $\hat{\alpha}$ 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 signalling components $\hat{\alpha} - \hat{\alpha}^{br}$ and $\hat{\beta} - \hat{\beta}^{br}$, which are decreasing in absolute value by Proposition 2.

To summarize, both $\hat{\alpha}$ and $\hat{\beta}$ are the sum of a monotone signalling component, and of an almost monotone myopic component reflecting learning. Since $\hat{\alpha}$ and $\hat{\beta}$ are simply a regrouping of α and β , these two effects explain also the non-monotonicity of the latter.

5.2. Prices and quantities

The properties of the equilibrium coefficients have implications for the levels and time paths of prices and outputs. The relationship $\delta = -\bar{p}(\alpha + \beta)$ between the coefficients in Theorem 1 yields a simple expression for the expected total output 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 behaviour 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 output then carry over to the expected price with obvious sign reversals. We record these in the following proposition, which summarizes some properties of equilibrium prices and outputs.




FIGURE 3

Price and output paths, $(r, \sigma, n, \bar{p}, T, g_0, \pi_0, c^1, c^2) = (0.4, 0.75, 2, 100, 3.46, 2, 30, 20, 40)$

Proposition 3. *In any symmetric linear Markov equilibrium, prices and quantities satisfy the following properties:*

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 static 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 volatility of total output conditional on costs, $-(\alpha_t \beta_t \gamma_t)/\sigma$, is eventually decreasing in t . If $r = 0$, then the volatility decreases monotonically in t .*
3. *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, non-monotone.*

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 realized price and output paths for two firms with costs $(c^1, c^2) = (20, 40)$.

The second part of the proposition concerns the volatility of total output. Noting that $\sum_i Q_t^i = \alpha_t \sum_i C^i + n\beta_t \Pi_t + n\delta_t$, we see that this volatility is given by $n\beta_t \lambda_t \sigma$, where $\lambda_t \sigma$ is the volatility of the public belief Π_t . Recalling that $\lambda_t = -(\alpha_t \gamma_t)/(n\sigma^2)$ then gives the expression in Proposition 3. Since the volatility of output is driven by the volatility of the public belief Π_t , it eventually decreases as Π_t converges to the true average cost. Indeed, as $t \rightarrow \infty$, the total output converges to a constant by Corollary 2. If $r = 0$, this convergence is monotone, despite the non-monotonicity of α and β . Interestingly, however, the volatility of total output may be non-monotone, peaking at some intermediate time, when $r > 0$.

By inspection of (1), the drift of the revenue process is $\bar{p} - \sum_i Q_t^i$, and hence it simply mirrors the movements in total output. But the drift is just the expected price (where the expectation is over the time- t demand shock σdZ_t), and so the above discussion of output volatility also describes the volatility of the expected price. We thus see that the expected price may be the most volatile after the market has already operated for some time, but it will eventually settle down. Of course, realized prices continue to vary due to demand shocks even after outputs have converged (close) to complete information levels.

The non-monotonicity of the output difference $Q_t^i - Q_t^j$ in t in the third part of Proposition 3 can be clearly seen in Figure 3. The result follows simply by definition of Markov strategies and the non-monotonicity of α for T large. As we discuss below, it has implications for productive efficiency and hence for market profitability and market concentration.

5.3. Profits, consumer surplus, and market concentration

We now consider equilibrium profits and consumer surplus. In particular, we are interested in comparing them to the corresponding complete information benchmarks over time.

Of course, conditional on any history, the firms' profits as well as consumer surplus depend on the realized costs and prices as well as on the prior beliefs. Thus, in order to obtain a meaningful comparison, we consider the *ex ante* expected flow profits and flow consumer surplus. That is, we take expectations of these flows with respect to the vector of costs in our dynamic game. Similarly, we construct the complete information benchmark by taking an expectation of the static complete information Nash equilibrium profits and consumer surplus with costs drawn from the same prior distribution as in the dynamic game.²⁵

By the symmetry, the *ex ante* expectation of any firm i 's time- t flow profit is given by

$$W_t := \mathbb{E}[(\bar{p} - \sum_j Q_t^j - c^1)Q_t^1].$$

Routine calculations show that the expected static complete information Nash profit is

$$W^{co} := \frac{(\bar{p} - \pi_0)^2 + g_0(n^2 + n - 1)}{(n + 1)^2}.$$

The *ex ante* expected time- t flow consumer surplus, CS_t , and the expected static complete information Nash consumer surplus, CS^{co} , are defined analogously.

Proposition 4. *The expected profits and consumer surpluses satisfy the following properties:*

1. $CS_t > CS^{co}$ for all $t \in [0, T]$ in every symmetric linear Markov equilibrium.
2. $W_t < W^{co}$ for $t = 0$ and $t = T$ in every symmetric linear Markov equilibrium.
3. Suppose

$$\kappa(n) < \frac{r\sigma^2}{g_0} < \frac{n(2 - 3n + n^3)}{(n + 1)^3}. \tag{17}$$

25. Note that repetition of the static Nash equilibrium is the natural complete information analogue of our linear Markov equilibria. Indeed, it obtains as a limiting case as $g_0 \rightarrow 0$.

Then for all π_0 sufficiently close to \bar{p} and for all $\bar{T} > 0$, there exists $t > 0$ and $T > \bar{T}$ such that the T -horizon game has a symmetric linear Markov equilibrium where

$$\int_t^T e^{-rs}(W_s - W^{co})ds > 0 \quad \text{and} \quad \int_t^T e^{-rs}(TS_s - TS^{co})ds > 0,$$

where $TS_s := W_s + CS_s$ and $TS^{co} := W^{co} + CS^{co}$ are, respectively, the *ex ante* expected time- s total surplus and the expected static complete information Nash total surplus.

The first part of Proposition 4 shows that the *ex ante* expected flow consumer surplus always lies above the expected complete information consumer surplus. Heuristically, this follows since expected price is at all times below the complete information level if $\pi_0 \leq \bar{p}$ by Proposition 3 and any variation in the price conditional on the costs is only beneficial to consumers who can adjust their demands. (The actual proof allows for $\pi_0 > \bar{p}$.)

Turning to the profits, Figure 4 compares the expected flow profits under complete and incomplete information. The left and right panels contrast markets with a low and high mean of the cost distribution. There are two main forces at play. On the one hand, signal-jamming adds to the expected total output, relative to the complete information Nash equilibrium level. This wasteful spending drives down profits but (eventually) declines over time by Proposition 2. On the other hand, the firms' active signalling (*i.e.* α being more negative than its myopic value) increases the sensitivity of each firm's output to its own cost. This improves the allocation of production, holding fixed the total output level. The resulting higher productive efficiency leads to higher expected profits since, from the *ex ante* perspective, each firm receives an equal share of the industry profit.

Based on the first force alone, one would conjecture that expected flow profits always increase over time, as in the left panel of Figure 4. But recall from Proposition 1 that the sensitivity of output to cost, α , is non-monotone. In other words, signalling can result in the productive efficiency being highest in the medium run. As in the right panel of Figure 4, this can lead the expected flow profit W_t to surpass the expected complete information profit W^{co} at some intermediate time. The third part of Proposition 4 shows that the effect may be strong enough to cause even the *ex ante* expectation of the time- t continuation profits (*i.e.* $\int_t^T e^{-rs} W_s ds$) to be higher than under complete information. Moreover, since the consumer surplus is always above the complete information level, also the expected time- t continuation total surplus is then above the complete information level.

The conditions in the third part of Proposition 4 ensure (i) that the "average profitability of the market" (as captured by $\bar{p} - \pi_0$) is not too high relative to the variance of output, so that the gain in allocative efficiency of production is important enough to outweigh the effect of higher total output, and (ii) that each firm is sufficiently patient, so that the value of signalling is sufficiently large. The latter is guaranteed by the second inequality in (17). The first inequality in (17) is a technical condition. It guarantees the existence of our equilibrium for any T by Theorem 1, and implies that $(\alpha, \beta, \delta, \gamma)$ converge uniformly to well-defined limit functions as $T \rightarrow \infty$ (along some subsequence), which we use in the proof.²⁶

To summarize, after an initial unprofitable phase of high output levels, the combined effect of learning and signalling can improve the expected industry flow profits. This result suggests that some intermediate information structure would yield higher payoffs than both complete

26. For concreteness, we show at the end of the proof of Proposition 4 that the two inequalities in (17) define a non-empty interval of possible values for $r\sigma^2/g_0$ if $3 \leq n \leq 10$.




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

information and incomplete information with firms only sharing the common prior. In particular, the initial phase of wasteful spending could be avoided by releasing an exogenous public signal about the industry average cost. For a given precision level, such a signal induces the information structure (obtained by repeatedly observing the market price) that arises in equilibrium at some intermediate time.²⁷

The above discussion of equilibrium profits focused on *ex ante* expected profits, which by symmetry are the same across firms. For any cost realizations, however, the firms with the lowest costs earn the highest profits and have the largest market shares. A possible way to capture this is to use some standard measure of concentration such as the Herfindahl-Hirshman index, or HHI, which is defined as the sum of the squared market shares.

In a symmetric linear Markov equilibrium, the time-*t* HHI is given by

$$HHI_t := \sum_i \left(\frac{Q^i}{\sum_k Q^k} \right)^2 = \frac{\sum_i (\alpha_t C^i + \beta_t \Pi_t + \delta_t)^2}{(\alpha_t \sum_k C^k + n\beta_t \Pi_t + n\delta_t)^2}.$$

The presence of Π_t in the denominator makes it hard to study the behaviour of the HHI analytically, so we resort here to simulations in order to explore the evolution of market concentration over time. Figure 5 shows the simulated HHI, averaged over 500 runs of the market, for a two-firm industry with cost vector $(c^1, c^2) = (1, 4)$. As the figure illustrates, industry concentration can be non-monotone over time. This is driven by the non-monotonicity of the difference between the firms’ outputs over time (Proposition 3). The result is “overshooting”: the firm with the lowest cost captures relatively quickly a large market share that is larger than its share in the complete information outcome, to which the market eventually converges. We also see that for long horizons or low discount rates, the HHI can be decreasing for a long time.

27. In the static literature on *ex ante* information sharing in oligopoly (Vives, 1984; Gal-Or, 1985; Raith, 1996), output is most sensitive to costs under complete information, and the expected total quantity is constant in the precision of the information revealed. As a result, sharing full information about costs is optimal in Cournot models. Instead, in our dynamic model, total expected quantity is decreasing, but output is most sensitive to cost for some intermediate time, leading to the richer picture outlined above.




FIGURE 5

Simulated HHI, $(n, r, \sigma, \bar{p}, T, g_0, c^1, c^2, \pi_0) = (2, .65, .75, 10, 7.77, 2, 1, 4, 3)$

6. EXTENSIONS

6.1. Infinite horizon

We have assumed a finite horizon throughout. The main reason for this is a technical one: it allows us to derive and impose the relevant boundary conditions for the equilibrium system (11)–(14). However, we now show that our analysis carries over to an infinite horizon. More precisely, we show that our finite-horizon equilibria converge to an equilibrium of the infinite-horizon version of the model as $T \rightarrow \infty$ under a slight strengthening of the sufficient condition for existence (15). This result allows us to extend our findings about the equilibrium properties to the infinite-horizon case. It also provides a method for approximating an equilibrium of the infinite-horizon game using our boundary value problem.

For the formal statement, we use Theorem 1 to identify a Markov equilibrium of the T -horizon game with the tuple $(\alpha^T, \beta^T, \delta^T, \xi^T, \gamma^T)$. It is convenient to extend these 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.²⁸ Moreover, $\delta^* = -\bar{p}(\alpha^* + \beta^*)$ and $(\alpha^*, \beta^*, \xi^*, \gamma^*)$ is a solution to the system (11)–(14) 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 $g_0/\sigma^2 < 4r/(27n)$ strengthens the first case in (15) to ensure that all the functions are bounded uniformly in T , facilitating the convergence argument. In particular, if $g_0/\sigma^2 < r/\kappa(n)$,

28. There 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^* is bounded away from zero, and hence the limit posterior variance γ_t^* is strictly decreasing in t , implying that conditioning on t is equivalent to conditioning on γ_t^* .

then $(\alpha^T, \beta^T, \delta^T, \gamma^T)$ are uniformly bounded and converge uniformly (along a subsequence) as $T \rightarrow \infty$. The stronger condition allows us to also bound and show the convergence of ξ^T , and ultimately establish that the limit is an equilibrium of the infinite-horizon game by verifying a transversality condition.

Since beliefs and play converge in the limit of finite-horizon equilibria, this is immediately true of the infinite-horizon equilibrium we identify. Moreover, 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 as well. Finally, depending on the parameters, the coefficient 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 becomes vector-valued with an associated posterior covariance matrix, and the analysis of the resulting boundary value problem seems a daunting task. (See Lambert *et al.* (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 symmetric settings with positively or negatively correlated costs. To see this, note that the public belief simply tracks the average cost, so the laws of motion of Π and γ are unaffected by (symmetric) correlation in the prior; correlation only affects the initial condition for γ . Similarly, firm i 's private belief is still captured by M^i and γ^M ; only the initial condition for the matrix Γ in Lemma 1 is affected. Lemmas 2 and 3 then carry over as stated, except that the initial value of z will now be greater than 1 if costs are negatively correlated, and less than 1 if correlation is positive.

The above observation can be used to establish the following relationships among symmetric linear Markov equilibria across the different cases. For simplicity, suppose there are only two firms. Fix an equilibrium with independent cost draws. Then at any time t , the continuation equilibrium over the remaining horizon $[t, T]$ is an equilibrium of a game where the horizon is $[0, T - t]$ and the prior has negative correlation. (The latter equilibrium is characterized by the same differential equations as in the independent case, only the initial conditions are different.) The intuition for this result is as follows. Along the path of play, the outsider is learning the firms' average cost. Hence, from his perspective, individual costs are increasingly negatively correlated over time even though they were independent under the prior. Indeed, in the limit the average is known and costs are perfectly negatively correlated. On the other hand, if costs are *ex ante* symmetrically negatively correlated, then the outsider believes them to be negatively correlated already at $t = 0$. Moreover, each firm's own cost now provides a private signal about the others' cost. Thus, the resulting information structure is as if we had directly jumped to some time $t > 0$ in the independent case.

Similarly, when the prior distribution has symmetric positive correlation, the outsider is still just learning the average. As a result, from the outsider's perspective, the correlation of the costs is decreasing and eventually becomes zero at some time t , at which point the continuation equilibrium over $[t, T]$ is simply an equilibrium of a game where the horizon is $[0, T - t]$ and where the costs are independent according to the prior.

We can also introduce interdependent values, modelled as firm i 's cost being the sum $C^i + k \sum_{j \neq i} C^j$ for some $0 < k \leq 1$. The addition of an extra term in the payoff function changes the equations in our boundary value problem somewhat, but the derivation is analogous. Cost

interdependence reduces the incentive to signal, since any given firm having a lower cost implies that the other firms' costs 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 towards the end resembling the aggressive behaviour at the end of the horizon in models of insider trading in financial markets.

We have focused on the firms' uncertainty about their competitors' costs. A more general model of a new market would have both cost and demand uncertainty. A one-agent model with such two-dimensional uncertainty is studied by 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.

7. CONCLUDING REMARKS

We have analysed a stylized game of dynamic oligopolistic competition under incomplete information. In our game, firms must balance the intertemporal trade-off between flow-profit maximization and investment in manipulating their rivals' beliefs. The key to the tractability of our framework is the representation of symmetric linear Markov strategies in terms of the private costs and a public belief—a sufficient statistic based on public information only.

We have derived conditions for existence of a symmetric linear Markov equilibrium and characterized the time-varying equilibrium weights the firms assign to private and public information. In any such equilibrium, learning is partial—firms only learn the average of their rivals' costs—yet behaviour converges uniformly to the complete information Nash outcome. Finally, we have traced the rich implications of equilibrium behaviour for the patterns of relevant variables, such as prices, quantities, and industry profits.

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. 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 $B_t(Y^t) := \int_0^t f'_s dY_s + \delta_t$. 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*

$$d\Pi_t = -\frac{\alpha_t \gamma_t}{n\sigma^2} [dY_t - (\bar{p} - \alpha_t n\Pi_t - nB_t(Y^t))dt], \quad \Pi_0 = \pi_0,$$

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

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

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

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

where for Γ_t uniqueness is in the class of symmetric non-negative 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 $\Gamma^{ii} = \gamma_t^1, \Gamma^{ij} = \gamma_t^2, i \neq j$, 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 non-negative 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

$$x' \Gamma_t x = g_0 \|x\|_2^2 + \gamma_t^2 \left(\sum_i x_i\right)^2.$$

If $\gamma_t^2 \geq 0$, we are done. If $\gamma_t^2 < 0$, then, using the fact that $(\sum_i x_i)^2 \leq \|x\|_1^2$, we have

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

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 non-negative 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

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

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

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

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. ||

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.) \parallel

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

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

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 $z_t \lambda_t = \lambda_t^M$ gives²⁹

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

Therefore, we have

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

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 $\hat{\Pi}_t^i = \Pi_t$ 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

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

29. Observe 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 $\dot{\gamma}_t^M = -(\alpha_t \gamma_t^M / \sigma)^2$.

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.) \parallel

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 noise in the revenue 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.

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

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

for certain linear functionals L_t^i , $i = 1, 2, 3$.³⁰ 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). Furthermore, because the coefficient on $(Q_t^i)^2$ in the firm's flow payoff is invertible (as it simply equals -1), Yong and Zhou's Corollary 5.6 (p. 312) implies that the existence of a linear optimal policy is equivalent to the existence of a solution to the stochastic Hamiltonian system associated with the best-response problem. This Hamiltonian system is a linear forward-backward stochastic differential equation for which existence in our case follows by the result of Yong (2006). The form of the value function given in (9) then follows by the existence of a linear optimal policy.

We note then that the value function V is continuously differentiable in t and twice continuously differentiable in $(c, \pi, \hat{\pi})$.³¹ Thus it satisfies the HJB equation (7). This implies that the linear optimal policy $q = \alpha_t c + \beta_t \pi + \delta_t + \xi_t(\hat{\pi} - \pi)$,

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

31. Differentiability of each v_i in (9) can be verified using the fact that V is the value under the optimal policy $(\alpha, \beta, \delta, \xi)$, where (α, β, δ) are differentiable by assumption and ξ , which only enters V through an integral, can be taken to be continuous by Corollary 5.6 of Yong and Zhou (1999, p. 312).

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 , π , $\hat{\pi}$, 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 equation (7) we have

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

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 the derivatives $(\dot{v}_1(t), \dot{v}_4(t), \dot{v}_5(t), \dot{v}_8(t))$ 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

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

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). \parallel

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. Let $(\alpha, \beta, \xi, \gamma)$ be 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 9 on page 508). 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.³² 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). 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 $(v_0, v_2, v_3, v_6, v_7, v_9)$.

This system is stated here for future reference:

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

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 behaviour is given by (α, β, δ) as desired. \parallel

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 $G := [0, \tilde{\gamma}_F)$ be the maximal such interval with respect to set inclusion. (I.e. $\tilde{\gamma}_F = \sup\{\tilde{\gamma}_F \in \mathbb{R}_+ : \text{backward system has a solution for all } \gamma_F \in [0, \tilde{\gamma}_F)\}$.) Finally, define the function $\kappa : \mathbb{N} \rightarrow \mathbb{R}_{++}$ by

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

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 $\gamma_0 = n g_0$.*

32. For 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 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 $t = T$.

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 $-\bar{\rho}(\alpha + \beta) = \delta$ 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

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

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. 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

$$0 \geq \alpha_t \geq x_t = \frac{1}{\sqrt[3]{3 \frac{n g_0}{\sigma^2} (T-t) - 1}} \geq \frac{1}{\sqrt[3]{3 \frac{n g_0 T}{\sigma^2} - 1}} \geq -2.$$

Combining this with $0 \leq \gamma_t \leq \gamma_0 = n g_0$, 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 $g_0 \in (0, \bar{g}_0)$. 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 $g_0 < \bar{g}_0$. 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 $V(0, 0, 0, s) = v_0(s)$. The payoff $V(0, 0, 0, s)$ is the expected profit over $[s, T]$ under the equilibrium strategies conditional on $\Pi_s = \hat{\Pi}_s = C^i = 0$. Because $\Pi_s = \hat{\Pi}_s$, it is independent of ξ . Moreover, $(\alpha, \beta, \delta, \gamma)$ are bounded on $[s, T]$ uniformly over $g_0 \in [0, \bar{g}_0]$ by assumption. Hence, the equilibrium payoff $V(0, 0, 0, s)$ is bounded by some $B < \infty$ uniformly in g_0 .

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

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

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

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

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. But because $V(0, \pi, \hat{\pi}, \tau) \geq 0$ for all $(\pi, \hat{\pi})$, this implies that the time- s expectation of the deviation payoff tends to infinity as $g_0 \rightarrow \bar{g}_0$. Hence it dominates B (and thus $V(0, 0, 0, s)$) for g_0 close enough to \bar{g}_0 . But this contradicts the fact that a symmetric linear Markov equilibrium exist for all $g_0 < \bar{g}_0$. \parallel

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 $\gamma_i^T K$ for some $K < \infty$ independent of i and T .

Lemma A.6. *Let $g_0/\sigma^2 < r/\kappa(n)$. Then 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 $T \geq t \geq t_\eta$, giving 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 $T \geq t \geq t_\varepsilon$. \parallel

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

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

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 $\{|X_t| < \varepsilon\} \subset \{|\underline{Q}_t^i - q_i^N(C)| < \varepsilon\}$, and hence it follows that $\mathbb{P}[|\underline{Q}_t^i - q_i^N(C)| < \varepsilon] > 1 - \varepsilon$. \parallel

Proof of Proposition 1. (1.) Consider a symmetric linear Markov equilibrium (α, β, δ) with posterior variance γ . Denote the induced values of the myopic coefficients under γ 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 $\delta_T < 0$ 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

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

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

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

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. \parallel

Proof of Proposition 2. (1.) The signalling 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

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

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 signalling component $\hat{\alpha}_t - \hat{\alpha}_t^{br}$ is thus negative in a neighbourhood of T . Now solve $\hat{\alpha}_t = \hat{\alpha}_t^{br}$ for z_t and substitute the resulting expression into $\hat{\alpha}_t - \hat{\alpha}_t^{br}$. We obtain,

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

Thus, if $\hat{\alpha}_t - \hat{\alpha}_t^{br} = 0$ for some $t < T$, then the signalling 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 signalling component $\hat{\beta}_t - \hat{\beta}_t^{br}$ is positive in a neighbourhood 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 signalling 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 signalling component has a slope of zero. Impose $r=0$, solve $\hat{\alpha}_t - \hat{\alpha}_t^{br} = 0$ for β_t , and substitute into $\dot{\hat{\alpha}}_t - \dot{\hat{\alpha}}_t^{br}$. We obtain

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

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

Likewise, we know the signalling 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 signalling 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 signalling 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 signalling component has a slope of zero. Then the signalling component must be negative at that point. This contradicts our earlier finding that the signalling component is positive and decreasing in a neighbourhood 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. \parallel

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

(2.) By Lemma A.1, the volatility of the public belief Π_t is $-(\alpha_t \gamma_t)/(n\sigma^2)\sigma = \lambda_t \sigma$. Thus the total output $\sum_i Q_t^i = \alpha_t \sum_i C^i + n\beta_t \Pi_t + n\delta_t$ has volatility $n\beta_t \lambda_t \sigma = -(\alpha_t \beta_t \gamma_t)/\sigma$. Differentiating $\alpha_t \beta_t \gamma_t$ with respect to t , setting $t=T$, and using (11), (12), and (14) gives

$$\left. \frac{d}{dt}(\alpha_t \beta_t \gamma_t) \right|_{t=T} = \frac{g_0^2(n-1)^3 n^3 \gamma_T^2 (g_0(n-1)n + \gamma_T)^2 (g_0 n(3n(n+1) - 2) + (-n^2 + n + 2)\gamma_T)}{(n+1)^3 \sigma^2 (n(g_0(n-1) + \gamma_T) + \gamma_T)^5}.$$

The last term in the numerator is positive for all n , because $\gamma_T \leq ng_0$ implies

$$g_0 n(3n(n+1) - 2) + (-n^2 + n + 2)\gamma_T > 2g_0 n^2(2+n) > 0.$$

Thus $-\alpha_t \beta_t \gamma_t$ is eventually decreasing.

In the undiscounted case, we obtain

$$\frac{d}{dt}(\alpha_t \beta_t \gamma_t) = \frac{\alpha_t^2 \beta_t \gamma_t^2 ((n-1)n\alpha_t(-2(n+1)\beta_t(z_t - 1) - z_t) - 2\beta_t)}{n(n+1)\sigma^2}.$$

Because $z_t \geq 1$ and $-\alpha_t > \beta_t > 0$ we rewrite the terms in parentheses in the numerator as

$$(1-n)n\alpha_t z_t - 2\beta_t (n(n^2 - 1)\alpha_t(z_t - 1) + 1) > (1-n)n\alpha_t z_t - 2\beta_t > (n-1)n\beta_t - 2\beta_t.$$

Therefore, we obtain the following bound:

$$\frac{d}{dt}(\alpha_t \beta_t \gamma_t) > \frac{(n-2)\alpha_t^2 \beta_t^2 \gamma_t^2}{n\sigma^2} > 0.$$

This shows that $-\alpha_t \beta_t \gamma_t$ is strictly decreasing in t .

(3.) Firm i 's output on the equilibrium path is given by $Q_t^i = \alpha_i C^i + \beta_i \Pi_t + \delta_t$. Therefore, for any i and $j \neq i$, we have $Q_t^i - Q_t^j = \alpha_i(C^i - C^j)$. Proposition 1 shows that α is non-monotone for T sufficiently large. \parallel

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

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

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

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

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

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

A similar derivation yields the *ex ante* expected time- t consumer surplus

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

(1.) Subtracting the expected complete information static Nash consumer surplus CS^{co} from the *ex ante* expected time- t consumer surplus CS_t gives

$$\frac{1}{2} \left[g_0 n (\alpha_t + \beta_t)^2 - \frac{g_0 n}{(n+1)^2} - \beta_t \gamma_t (2\alpha_t + \beta_t) \right] + \frac{1}{2} (p - \pi_0)^2 \left[n^2 (\alpha_t + \beta_t)^2 - \frac{n^2}{(n+1)^2} \right].$$

We claim that this expression is positive. Since $\delta_t \geq \delta^m$, we know that $\alpha_t + \beta_t < -1/(n+1)$, and hence the second term is positive. Consider the first term. The sum of the first two terms inside the brackets is again positive, and the last term is positive as $-\alpha_t > \beta_t$ for all t .

(2.) 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

$$W_T - W^{\text{co}} = - \frac{g_0(n-1)n\gamma_T [g_0n(2n^2+n-3) + (n+1)(n+3)\gamma_T]}{[g_0n(n^2-1) + (n+1)^2\gamma_T]^2},$$

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

(3.) By part (1.), it suffices to establish the result for flow profits. We assume $\pi_0 - \bar{p} = 0$; the result for $\pi_0 - \bar{p}$ small enough follows by continuity of the profits in π_0 .

For any T , fix a symmetric linear Markov equilibrium with coefficients (α^T, β^T) and posterior variance γ^T . Throughout the proof we restrict attention to a strictly increasing sequence of horizons T such that (i) $(\alpha^T, \beta^T, \gamma^T)$, viewed as functions on $[0, \infty)$ by setting $(\alpha_t^T, \beta_t^T, \gamma_t^T) = (\alpha_t^T, \beta_t^T, \gamma_t^T)$ for all $t > T$, converge uniformly to well-defined limits $(\alpha^*, \beta^*, \gamma^*)$ as $T \rightarrow \infty$, and (ii) $\gamma_t^T \rightarrow 0$ monotonically as $T \rightarrow \infty$. The existence of a sequence satisfying (i) can be established as in Lemma A.9 below because $(\alpha^T, \beta^T, \gamma^T)$ are bounded uniformly in T when $\kappa(n) < r\sigma^2/g_0$ (see the proof of Lemma A.4). Moreover, the argument shows that the limits $(\alpha^*, \beta^*, \gamma^*)$ satisfy (11), (12), and (14) on $[0, \infty)$ with $(\alpha_t^*, \beta_t^*, \gamma_t^*) \rightarrow (\alpha^m(0), \beta^m(0), 0)$ as $t \rightarrow \infty$. Since $\gamma_t^T \rightarrow 0$ by Corollary 1, we can then take a further subsequence of horizons in order to satisfy (ii).

For $T \geq t \geq 0$, let $\Delta W_t^T := W_t^T - W^{\text{co}}$, where W_t^T is the *ex ante* expectation of the time- t equilibrium flow profit in the game with horizon T . Given the convergence of $(\alpha^T, \beta^T, \gamma^T)$, ΔW^T (viewed as a function on $[0, \infty)$ by setting $\Delta W_t^T = \Delta W_t^T$ for all $t > T$) converges uniformly to some ΔW^* as $T \rightarrow \infty$ (along our sequence). It suffices to show that $\Delta W_t^* > 0$ for all t large enough. Indeed, then $\int_t^\infty e^{-rs} \Delta W_s^* ds > 0$ for any t large enough, and hence by uniform convergence of ΔW^T to ΔW^* , we have $\int_t^T e^{-rs} \Delta W_s^T ds > \int_t^\infty e^{-rs} \Delta W_s^* ds > 0$ for any sufficiently large T (along our sequence; hence the need for \bar{T} in the statement of the result), where the first inequality follows because $W_s^T = W_s^T < 0$ for $s > T$ by the first part of Proposition 4.

It turns out to be convenient to change variables. Note that each α^T is bounded away from zero, and so each γ^T is a strictly decreasing function. We can thus invert γ^T and write the equilibrium coefficients and expected profits as a function of the posterior variance instead of time. We denote these functions $\alpha^{\gamma T}$ and $\beta^{\gamma T}$, indexed by the terminal posterior variance $\gamma_T := \gamma_T^T$ instead of the horizon (e.g., $\alpha^{\gamma T}(\gamma) = \alpha^T((\gamma^T)^{-1}(\gamma))$ for $\gamma \in [\gamma_T, n g_0]$). By the uniform convergence of $(\alpha^T, \beta^T, \gamma^T)$, the functions $\alpha^{\gamma T}$ and $\beta^{\gamma T}$ have well-defined pointwise limits as $\gamma_T \rightarrow 0$ (along our subsequence). Abusing notation, we denote these limits α and β . The corresponding limit flow-profit difference written in terms of γ is then

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

Noting that $\Delta W(\gamma) = \Delta W^*((g^*)^{-1}(\gamma))$ and $\lim_{t \rightarrow \infty} \gamma_t^* \rightarrow 0$, it then suffices to show that $\Delta W(\gamma) > 0$ for all $\gamma > 0$ small enough.

The rest of the argument proceeds as follows: We show first that α is strictly decreasing (and hence strictly less than $\alpha(0) = -1$) in a neighbourhood of 0 by constructing a sequence of linear upper bounds for the family of functions $\alpha^{\gamma T}$ for γ_T small. We use this to bound β from below by its complete information level in a neighbourhood of 0. This in turn allows deriving a better upper bound for α and a lower bound for ΔW . The latter satisfies an ODE, which we use to establish the result (after another change of variables).

Lemma A.7. *The limit function α is strictly decreasing in a neighbourhood of $\gamma = 0$ if*

$$\frac{r\sigma^2}{g_0} < \frac{(n-1)^2}{n+1}. \quad (\text{A.14})$$

Proof. Given the change of variable, each $\alpha^{\gamma T}$ satisfies the differential equation

$$\begin{aligned} \alpha^{\gamma T}(\gamma) &= \frac{\beta(\gamma)(\gamma + (n-1)n\alpha^{\gamma T}(\gamma)(g_0n - \gamma) + g_0(n-1)n)}{\gamma n(\gamma + g_0(n-1)n)} \\ &+ \frac{nr\sigma^2(\gamma + \alpha^{\gamma T}(\gamma)(\gamma + n(\gamma + g_0(n-1)n))) + g_0(n-1)n}{\gamma^2 n \alpha^{\gamma T}(\gamma)(\gamma + g_0(n-1)n)} \end{aligned} \quad (\text{A.15})$$

with boundary condition $\alpha^{\gamma T}(\gamma_T) = \alpha^m(\gamma_T)$. Replacing $\alpha^{\gamma T}(\gamma)$ with $\alpha^m(\gamma)$ in the first term of (A.15), we obtain an upper bound on its numerator. In particular, the coefficient on $\beta^{\gamma T}(\gamma)$ is negative for all $\gamma \leq \bar{\gamma} := n\gamma_0(n-1)^2/(n^2+1)$. Hence, we obtain an upper bound on $\alpha^{\gamma T}$ by replacing $\beta^{\gamma T}(\gamma)$ with its myopic value $\beta^m(\gamma)$, which was defined in (10). This

bound applies over the interval $[\gamma_T, \bar{\gamma}]$ for all γ_T sufficiently small. It is given by the ODE

$$\hat{\alpha}^{\gamma_T}(\gamma) = \frac{\beta^m(\gamma)(\gamma + (n-1)n\hat{\alpha}^{\gamma_T}(\gamma)(g_0n - \gamma) + g_0(n-1)n)}{\gamma n(\gamma + g_0(n-1)n)} + \frac{nr\sigma^2(\gamma + \hat{\alpha}^{\gamma_T}(\gamma)(\gamma + n(\gamma + g_0(n-1))) + g_0(n-1)n)}{\gamma^2 n\hat{\alpha}^{\gamma_T}(\gamma)(\gamma + g_0(n-1)n)} \tag{A.16}$$

with $\hat{\alpha}^{\gamma_T}(\gamma_T) = \alpha^m(\gamma_T)$. Since $\beta^m(\gamma) \leq \beta(\gamma)$ for all $\gamma \geq 0$, we then have $\hat{\alpha}^{\gamma_T}(\gamma) > \alpha^{\gamma_T}(\gamma)$ for all γ small enough. Moreover, as these bounds are ODEs, their paths for different initial values γ_T cannot cross.

We now study the right-hand side of (A.16) to construct a linear upper bound on $\hat{\alpha}^{\gamma_T}$. The right-hand side of (A.16) is strictly concave in $\hat{\alpha}^{\gamma_T}$, and for γ small enough, it is strictly decreasing in $\hat{\alpha}^{\gamma_T}$ when $\hat{\alpha}^{\gamma_T} = -1$ and strictly increasing in γ . Furthermore, under condition (A.14), there exists $\hat{\gamma} > 0$ such that the right-hand side of (A.16) is strictly negative when $\hat{\alpha}^{\gamma_T} = -1$ and $\gamma < \hat{\gamma}$.

These properties imply that there exist $\tilde{\gamma} \in (0, \hat{\gamma})$ and $\tilde{\gamma}_T < \tilde{\gamma}$ such that $\hat{\alpha}^{\tilde{\gamma}_T}(\tilde{\gamma}) = -1$. Furthermore, for all $\gamma_T \leq \tilde{\gamma}_T$ and $(\gamma, \hat{\alpha}^{\gamma_T}) \in [0, \tilde{\gamma}] \times [-1, \alpha^m(\tilde{\gamma}_T)]$, the right-hand side of (A.16) is bounded from above by $\hat{\alpha}^{\tilde{\gamma}_T}(\tilde{\gamma})$. (This slope is obtained by substituting the values $\hat{\alpha} = -1$ and $\gamma = \tilde{\gamma}$.) Therefore, for all $\gamma_T \leq \tilde{\gamma}_T$, $\hat{\alpha}^{\gamma_T}$ is bounded from above on $(0, \tilde{\gamma}]$ by a linear function with slope $\hat{\alpha}^{\tilde{\gamma}_T}(\tilde{\gamma})$ that takes value $\alpha^m(\gamma_T)$ at γ_T . Because this holds for all $\gamma_T \in (0, \tilde{\gamma}_T)$, the function $\hat{\alpha}$ obtained by letting $\gamma_T \rightarrow 0$ is also strictly decreasing for γ small enough. The lemma now follows as, by construction, $\hat{\alpha}(0) = \alpha(0) = \alpha^m(0)$ and $\hat{\alpha} \geq \alpha$. \square

Lemma A.8. Under condition (17), the limit function β lies above $\beta^m(0)$ in a neighbourhood of $\gamma = 0$.

Proof. Consider the ODE for β , which is given by

$$\beta'(\gamma) = \frac{n(n+1)r\sigma^2\beta(\gamma)(\gamma(n+1) + \gamma_0(n-1)n) - \gamma(n-1)\beta(\gamma)^2(\gamma + \gamma_0(n-1)n)}{\gamma^2 n(n+1)\alpha(\gamma)(\gamma + \gamma_0(n-1)n)} - \frac{\beta(\gamma)(\gamma((n^2-1)\beta(\gamma) + n+1) + \gamma_0 n(\beta(\gamma) + n^2(-\beta(\gamma)) + n-1))}{\gamma(n+1)(\gamma + \gamma_0(n-1)n)} - \frac{\gamma_0(n-1)n^3 r\sigma^2}{\gamma^2 n(n+1)\alpha(\gamma)(\gamma + \gamma_0(n-1)n)} \tag{A.17}$$

with $\beta(0) = \beta^m(0)$. Now assume $r\sigma^2(n+1)^3 > ng_0(n-1)^2$, which is implied by the first inequality in (17). The right-hand side of (A.17) is then strictly decreasing in α for γ sufficiently small. Because, as $\gamma \rightarrow 0$, the coefficient α approaches -1 from below, we use $\alpha = -1$ to bound $\beta'(\gamma)$ in a neighbourhood of $\gamma = 0$. This gives a lower bound $\hat{\beta}$ on β . Let $\beta^{\text{co}} := \beta^m(0)$. One can verify that if the second inequality in (17) holds, there exists $\tilde{\gamma} > 0$ small enough such that $\hat{\beta}$ is strictly increasing in γ whenever $\hat{\beta}(\gamma) = \beta^{\text{co}}$ and $\gamma \leq \tilde{\gamma}$. Therefore, $\beta(\gamma) \geq \hat{\beta}(\gamma) \geq \beta^{\text{co}}$ for all γ close enough to 0. \square

We now turn to ΔW . Notice first that

$$\frac{\partial \Delta W(\gamma)}{\partial \beta} = (-ng_0 - \gamma)(1 + 2n(\alpha(\gamma) + \beta(\gamma))).$$

Proposition 1 implies that $\alpha + \beta \leq -1/(n+1)$. Thus, $\Delta W(\gamma)$ is locally increasing in β around the equilibrium values of our coefficients. Furthermore, $\Delta W(\gamma)$ is strictly concave in β . We can therefore use $\beta^{\text{co}} \leq \beta$ to construct a lower bound on $\Delta W(\gamma)$ for γ sufficiently small. Furthermore, there exists $\bar{\alpha} < -1$ such that the resulting bound on $\Delta W(\gamma)$ is strictly decreasing in α for all $\alpha > \bar{\alpha}$. Because $\alpha(0) = -1$ and all coefficients are continuous functions, there exists a neighbourhood of $\gamma = 0$ in which $\Delta W(\gamma)$ is decreasing in α .

We now define the family of functions $\tilde{\alpha}^{\gamma_T}$ as solutions to the differential equation in (A.16), where $\beta^m(\gamma)$ is replaced by the constant β^{co} . (Identical steps to those in Lemma A.7 establish that the limiting bound $\tilde{\alpha}$ (as $\gamma_T \rightarrow 0$) is strictly decreasing for γ small enough.) We then define a lower bound on the profit difference by setting $\beta(\gamma) = \beta^{\text{co}}$ and $\alpha(\gamma) = \tilde{\alpha}(\gamma)$ in $\Delta W(\gamma)$. This gives

$$w(\gamma) := \frac{n(\gamma(2(n+1)n\tilde{\alpha}(\gamma) + n^2 + n + 1) - \gamma_0 n(n+1)(\tilde{\alpha}(\gamma) + 1)((n+1)\tilde{\alpha}(\gamma) + 2n))}{(n+1)^2}.$$

To finish the proof, we solve $w(\gamma) = w$ for γ and substitute the resulting expression into the derivative $w'(\gamma)$. This gives $w'(\gamma)$ in terms of w and $\tilde{\alpha}$ only. Since $\tilde{\alpha}$ is strictly decreasing for γ small enough, we can make the change of variables $\gamma \mapsto \tilde{\alpha}$ (i.e. we divide by $\tilde{\alpha}'(\gamma)$) to obtain a differential equation for $w(\tilde{\alpha})$. Finally, evaluating the expression for

$w'(\tilde{\alpha})$ at $w=0$ and $\tilde{\alpha}=-1$, we have $w'(0)=-n^2(n-1)\gamma_0/(n+1)<0$. Therefore, w is strictly decreasing in $\tilde{\alpha}$ at $\gamma=0$. Since $\tilde{\alpha}$ is strictly decreasing in γ , this implies that $\Delta W(\gamma)\geq w(\gamma)>0$ for $\gamma>0$ small enough as desired.

Finally, in the table below, we report the values of the left-hand side and the right-hand side of condition (17) for different values of n :

n	2	3	4	5	6	7	8	9	10
LHS of (17)	0.32	0.87	1.48	2.10	2.74	3.38	4.02	4.66	5.31
RHS of (17)	0.30	0.94	1.73	2.59	3.50	4.43	5.38	6.37	7.30

By inspection, LHS of (17) is strictly smaller than RHS of (17) for $3\leq n\leq 10$, confirming that for such n , there exists a non-empty open interval of possible values of $r\sigma^2/g_0$ for which the assumptions of the second part of Proposition 4 are satisfied. \parallel

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 $-\alpha \geq \xi \geq 0$, 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 $-\alpha_T = \xi_T$. Now consider the sum $n\dot{\alpha}_t + \dot{\xi}_t$ and evaluate it at $n\alpha_t + \xi_t = 0$. We obtain

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

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 $ng_0\alpha_t\beta_t + r\sigma^2 > 0$ for all t , then the function $-\alpha$ crosses ξ from above only, and then $-\alpha_T = \xi_T$ implies $\xi_t < -\alpha_t$ for $t < T$. 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.9. *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 $\sup_{s \in [0,t]} \|f^{T^n}(s) - F(f(s))\| < \varepsilon/(2t)$. Then for all $n > N$,

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

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

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 admissible if (i) $\mathbb{E}[\int_0^t Q_s^2 ds] < \infty$ for all $t \geq 0$, (ii) revenue process (1) has a unique solution, and (iii) firms' expected payoffs are finite. We need the following two lemmas.

Lemma A.10. *For any admissible 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 $\hat{\Pi}_t$, 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 $\hat{\Pi}_t$ 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-Schwarz inequality implies

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

Hence,

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

where the last term is finite by admissibility 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

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

(For the others the limit exists and takes value zero.) Observe first that there is $\varepsilon > 0$ such that $\tilde{R}_{t,s} < e^{-\varepsilon \int_s^t \lambda_u du}$ 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-Schwarz and the fact that $\lambda < C$, some $C > 0$,

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

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 admissibility. We conclude that $\limsup e^{-rt}\mathbb{E}[A_t] < \infty$.

Regarding B_t , by applying Cauchy–Schwarz 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 admissibility. 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, $\limsup e^{-rt}\mathbb{E}[(\Pi_t^Q)^2] < \infty$. \parallel

Lemma A.11. *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

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

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.9, 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. \parallel

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

Lemma A.12. *Under any admissible 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 limsup. 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–Schwarz 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.10, $\limsup e^{-rt}\mathbb{E}[(\Pi_t^Q)^2] < \infty$. Thus $e^{-rt}v_4(t)\mathbb{E}[\Pi_t^Q \hat{\Pi}_t] \rightarrow 0$ as $t \geq 0$. We deduce that the limsup 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 admissible 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

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

Indeed, because (i) $\mathbb{E}\left[\left(\int_0^t R_{t,s}\lambda_s \sigma dZ_s\right)^2\right]$ 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} . \parallel

Acknowledgments. We thank Johannes Hörner, Chris Knittel, Greg Pavlov, Ariel Pakes, Larry Samuelson, Andy Skrzypacz, Philipp Strack, Mike Whinston, and Alexander Wolitzky for useful discussions, and the editor Botond Kőszegi, three anonymous referees, and the audiences at various seminars and conferences for comments.

REFERENCES

- ATHEY, S. and BAGWELL, K. (2008), “Collusion with Persistent Cost Shocks”, *Econometrica*, **76**, 493–540.
- AUMANN, R. and MASCHLER, M. (1995), *Repeated Games with Incomplete Information* (Cambridge, MA: The MIT Press).
- BACK, K., CAO, C. H. and WILLARD, G. A. (2000), “Imperfect Competition among Informed Traders”, *Journal of Finance*, **55**, 2117–2155.
- BERNHARDT, D. and TAUB, B. (2015), “Learning about Common and Private Values in Oligopoly”, *Rand Journal of Economics*, **46**, 66–85.
- BORENSTEIN, S. and BUSHNELL, J. (1999), “An Empirical Analysis of the Potential for Market Power in California’s Electricity Industry”, *The Journal of Industrial Economics*, **47**, 285–323.
- BUSHNELL, J. B., MANSUR, E. T. and SARAVIA, C. (2008), “Vertical Arrangements, Market Structure, and Competition: An Analysis of Restructured US Electricity Markets”, *American Economic Review*, **98**, 237–66.
- CISTERNAS, G. (2015), “Two-Sided Learning and Moral Hazard” (Discussion paper, MIT Sloan).
- DALEY, B. and GREEN, B. (2012), “Waiting for News in the Market for Lemons”, *Econometrica*, **80**, 1433–1504.
- DILMÉ, F. (2014), “Dynamic Quality Signaling with Hidden Actions” (Discussion paper, University of Bonn).
- DORASZELSKI, U., LEWIS, G. and PAKES, A. (2016), “Just Starting Out: Learning and Price Competition in a New Market” (Discussion paper, University of Pennsylvania).
- ESCOBAR, J. F. and TOIKKA, J. (2013), “Efficiency in Games with Markovian Private Information”, *Econometrica*, **81**, 1887–1934.
- FAINGOLD, E. and SANNIKOV, Y. (2011), “Reputation in Continuous-Time Games”, *Econometrica*, **79**, 773–876.
- FOSTER, F. D. and VISWANATHAN, S. (1996), “Strategic Trading When Agents Forecast the Forecasts of Others”, *The Journal of Finance*, **51**, 1437–1478.
- FRIEDMAN, A. (1971), *Differential Games* (New York: John Wiley & Sons).
- FUDENBERG, D. and TIROLE, J. (1986), “A “Signal-Jamming” Theory of Predation”, *RAND Journal of Economics*, **17**, 366–376.
- GAL-OR, E. (1985), “Information Sharing in Oligopoly”, *Econometrica*, **53**, 329–343.
- HOLMSTRÖM, B. (1999), “Managerial Incentive Problems: A Dynamic Perspective”, *The Review of Economic Studies*, **66**, 169–182.
- HÖRNER, J. and LOVO, S. (2009), “Belief-Free Equilibria in Games With Incomplete Information”, *Econometrica*, **77**, 453–487.
- HÖRNER, J., TAKAHASHI, S. and VIEILLE, N. (2015), “Truthful Equilibria in Dynamic Bayesian Games”, *Econometrica*, **83**, 1795–1848.
- HORTACSU, A. and PULLER, S. L. (2008), “Understanding Strategic Bidding in Multi-unit Auctions: A Case Study of the Texas Electricity Spot Market”, *RAND Journal of Economics*, **39**, 86–114.
- KALLIANPUR, G. (1980), *Stochastic Filtering Theory* (New York: Springer-Verlag).
- KELLER, G. and RADY, S. (1999), “Optimal Experimentation in a Changing Environment”, *Review of Economic Studies*, **66**, 475–507.
- KYLE, A. S. (1985), “Continuous Auctions and Insider Trading”, *Econometrica*, **53**, 1315–1335.
- (1989), “Informed Speculation with Imperfect Competition”, *Review of Economic Studies*, **56**, 317–55.
- LAMBERT, N. S., OSTROVSKY, M. and PANOV, M. (2014), “Strategic Trading in Informationally Complex Environments” (Discussion paper, Stanford University).
- LIPTSER, R. S. and SHIRYAEV, A. N. (1977), *Statistics of random processes, I* (New York: Springer).
- MAILATH, G. J. (1989), “Simultaneous Signaling in an Oligopoly Model”, *Quarterly Journal of Economics*, **104**, 417–427.
- MESTER, L. (1992), “Perpetual Signaling with Imperfectly Correlated Costs”, *RAND Journal of Economics*, **23**, 548–563.
- MILGROM, P. and ROBERTS, J. (1982a), “Limit Pricing and Entry under Incomplete Information: an Equilibrium Analysis”, *Econometrica*, **50**, 443–459.
- MILGROM, P. and ROBERTS, J. (1982b), “Predation, Reputation, and Entry Deterrence”, *Journal of Economic Theory*, **27**, 280–312.
- MIRMAN, L. J., SAMUELSON, L. and URBANO, A. (1993), “Duopoly Signal Jamming”, *Economic Theory*, **3**, 129–149.
- MYATT, D. P. and WALLACE, C. (2015), “Cournot Competition and the Social Value of Information”, *Journal of Economic Theory*, **158B**, 466–506.
- 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**, 651–694.
- PHAM, H. (2009), *Continuous-Time Stochastic Control and Optimization with Financial Applications* (Berlin: Springer-Verlag).
- RAITH, M. (1996), “A General Model of Information Sharing in Oligopoly”, *Journal of Economic Theory*, **71**, 260–288.

- RIORDAN, M. H. (1985), "Imperfect Information and Dynamic Conjectural Variations", *RAND Journal of Economics*, **16**, 41–50.
- SADZIK, T. and WOOLNOUGH, C. (2014), "Multidimensional Uncertainty and Hypersensitive Asset Prices" (Discussion paper, UCLA).
- SANNIKOV, Y. and SKRZYPACZ, A. (2007), "Impossibility of Collusion Under Imperfect Monitoring with Flexible Production", *American Economic Review*, **97**, 1794–1823.
- SANNIKOV, Y. (2007), "Games with Imperfectly Observable Actions in Continuous Time", *Econometrica*, **75**, 1285–1329.
- TOWNSEND, R. M. (1983), "Forecasting the Forecasts of Others", *Journal of Political Economy*, **91**, 546–588.
- VIVES, X. (1984), "Duopoly Information Equilibrium: Cournot and Bertrand", *Journal of Economic Theory*, **34**, 71–94.
- VIVES, X. (2011), "Strategic Supply Function Competition with Private Information", *Econometrica*, **79**, 1919–1966.
- YONG, J. (2006), "Linear Forward-Backward Stochastic Differential Equations with Random Coefficients", *Probability Theory and Related Fields*, **135**, 53–83.
- YONG, J. and ZHOU, X. Y. (1999), *Stochastic Controls: Hamiltonian Systems and HJB Equations* (New York: Springer-Verlag).