

Collusion with Persistent Cost Shocks

Susan Athey and Kyle Bagwell

First Draft: March, 2003; This Draft: July, 2006

Abstract

We consider a dynamic Bertrand game, in which prices are publicly observed and each firm receives a privately observed cost shock in each period. Although cost shocks are independent across firms, within a firm costs follow a first-order Markov process. We analyze the set of collusive equilibria available to firms, emphasizing the best collusive scheme for the firms at the start of the game. In general, there is a tradeoff between productive efficiency, whereby the low-cost firm serves the market in a given period, and high prices. We show that when costs are perfectly correlated over time within a firm, if the distribution of costs is log concave and firms are sufficiently patient, then the optimal collusive scheme entails price rigidity: firms set the same price and share the market equally, regardless of their respective costs. Productive efficiency can be achieved in equilibrium under some circumstances, but such equilibria are not optimal. When serial correlation of costs is imperfect, partial productive efficiency is optimal. For the case of two cost types, first-best collusion is possible if the firms are patient relative to the persistence of cost shocks, but not otherwise. We present numerical examples of first-best collusive schemes.

*Athey: Harvard University (Department of Economics) and NBER. Bagwell: Columbia University (Kelvin J. Lancaster Professor of Economic Theory in the Department of Economics, and School of Business) and NBER. We are grateful to Nageeb Ali, Joe Harrington, David Levine, Alberto Martin, Ichiro Obara, two anonymous referees and seminar participants at the Berkeley 2004 Summer Institute in Competitive Strategy, Columbia, FTC, Georgetown, Gerzensee, Harvard, Northwestern, MIT, Norwegian School of Economics and Business Administration, NYU, Oslo, Princeton, Stanford, UBC, UCLA and UCSD for helpful discussions. We thank the National Science Foundation (Athey: SES-9983820, SES-0351500; Bagwell: SES-0214021, SES-0518802) for generous financial support.

1. Introduction

A rich literature analyzes the problem of collusion in repeated games. In one important class of models, the price or output decisions of individual firms are imperfectly observed by other firms. Green and Porter (1982) and Abreu, Pearce and Stacchetti (1986) analyze the case where individual firm output decisions are not observed; instead, a noisy signal (the market price) of the firms' output decisions is publicly observed. Imperfect observability of actions is an important concern in some settings, as when the product is an intermediate input and the customers are firms who negotiate prices individually. In other settings, however, the relevant firm behavior is publicly observed. For example, in a government procurement auction, the bids are usually publicly available; and many consumer goods are sold at publicly posted prices.

Even when firm behavior is publicly observed, colluding firms may confront significant informational problems. Firms' production costs often have important components that are private information, due to variations in supply contracts and process innovations. In an ideal collusive scheme, firms would communicate truthfully about their respective costs, so that, at each point in time, they could both maintain high prices *and* assign all production to the firm(s) with the lowest production costs. Such a scheme is possible, however, only if firms have incentives to communicate truthfully and accept the corresponding market-share assignments.

This discussion motivates consideration of a different class of collusion models, in which firms are privately informed as to their cost types and take actions that are publicly observed. Aoyagi (2003), Athey and Bagwell (2001), Athey, Bagwell and Sanchirico (2004) and Skryzpacz and Hopenhayn (2004) develop models of this kind, where firms play a repeated Bertrand pricing game, or equivalently, act as bidders in a series of procurement auctions. Importantly, these papers focus on the case where cost types are independent over time. This assumption is not always plausible. In some procurement auction settings, for example, individual firms may enjoy persistent sources of cost advantage. Similarly, for a firm selling products in a market, there are often many components of cost. Over time, different parts of the production process may see improvement, or the firm may sign new contracts for inputs, where the contracts may last several periods. Although firms may have a general understanding of the overall cost structure of their competitors, they may lack specific knowledge of the factors that change over time and affect competitor costs.¹

In this paper, we study collusion among firms that are privately informed as to their respective cost types, where cost shocks are persistent over time. The case of perfect persistence is a special case of our analysis.² Formally, we model firms as interacting over an infinite horizon

¹ The relevance of private information about persistent aspects of production cost has been emphasized by observers of important recent collusive agreements. See, e.g., Connor's (2001, p. 17) discussion of the lysine cartel. Likewise, McMillan (1991) reports that private information as to costs was an important consideration among colluding construction firms in Japan.

² LaCasse (2001) and Chakrabarti (2001) analyze collusion among firms when costs are perfectly persistent. These papers construct examples of separating and pooling equilibria, but do not analyze the optimal collusive

in a dynamic game of Bertrand competition, where in any period each firm privately observes its cost shock before publicly selecting its price. So as to allow the greatest possible scope for collusion, we assume that colluding firms can communicate before setting their prices; this allows us to analyze the extent to which communication is useful for colluding firms, an important question for antitrust policy. We assume as well that firms can allocate market share unequally among themselves, so long as they charge the same price. Although cost shocks are independent across firms, within a firm costs follow a first-order Markov process.³ Under the assumption that demand is inelastic, we analyze the set of collusive equilibria available to firms, emphasizing the best collusive scheme for the firms at the start of the game. Given that cost types are private information and persistent, our game is not a repeated game but rather a dynamic game with hidden state variables. We are interested in what we call Perfect Public Bayesian Equilibria (PPBE) of this game. This is a refinement of Perfect Bayesian Equilibrium, and it is the natural extension of the Perfect Public Equilibrium (PPE) (Fudenberg, Levine, and Maskin, 1994) solution concept for the dynamic Bayesian game that we study.⁴ Relatively little prior work has analyzed dynamic games where players have (perfectly or imperfectly) persistent private information, and thus we believe that the techniques developed here may be useful in other contexts as well.⁵

Whether costs shocks are persistent or not, firms face important incentive constraints when attempting to achieve productive efficiency, whereby the lowest-cost firm serves the market in each period. Productive efficiency requires that lower-cost firms expect higher market share, and this creates a temptation for higher-cost firms to mimic lower-cost firms. Such mimicry can be dissuaded if a firm expects that, following a period where it reports lower costs, it will

equilibria. LaCasse (2001) studies the case of perfect substitutes and downward-sloping demand, with two cost types; Chakrabarti (2001) analyzes Cournot oligopoly with discrete types.

³ When types are correlated across firms, collusion may be easier. First, the loss associated with failing to allocate production to the lowest-cost firm is smaller, since firm costs are typically similar. Second, it is less costly to provide incentives, since a firm may be punished for reported cost types that differ from those of its competitors. See Aoyagi (2003) for discussion of collusion when types are correlated across players.

⁴ Cole and Kocherlakota (2001) develop an analogous concept for dynamic games with hidden actions and hidden state variables; they refer to their concept as Markov-Private Perfect Equilibrium.

⁵ For example, the folk theorem of Fudenberg, Levine, and Maskin (1994) does not apply to dynamic games with persistent private information. Persistent private information has received attention in other contexts. The early studies of reputation formation (Kreps and Wilson, 1982; Milgrom and Roberts, 1982) analyze finite-horizon dynamic games in which there is a small probability that a player is a perfectly persistent and irrational type. Recently, Abreu and Gul (2000) and Abreu and Pearce (2003) extend and apply this approach to dynamic bargaining games. Our work differs in that all types are rational, occur with substantial probability, and may change over time; further, we focus on cooperative equilibria, and revelation of types in each period is necessary for efficiency. Kennan (2001) analyzes a dynamic bargaining game between a buyer and a seller, in which the buyer's private information may change over time and follows a two-state Markov process. Our work differs in that it allows for all agents to have private information; furthermore, we focus on cooperative equilibria. Watson (2002) examines the role of perfectly persistent private information in a model of dynamic relationship formation wherein the stakes of the relationship evolve over time. In the dynamic contracting literature, Laffont and Tirole (1988) highlight a ratchet effect, when the agent's type is perfectly persistent and only short-term contracts are feasible. Fernandez and Phelan (2000) numerically analyze a dynamic principal-agent problem where the agent's type follows a Markov process; and Battaglini (2005) studies long-term contracts that induce separation in similar settings. We discuss the dynamic contracting literature further in Section 3.

receive lower market share or endure lower prices. For the cartel, the benefit of productive efficiency in the current period thus must be balanced against the costs of possibly inefficient production and/or low prices in future periods. Future productive inefficiency would arise, if future market share were withheld from a firm that continues to enjoy low costs.

At an intuitive level, it is not immediately clear whether collusion becomes easier or harder when cost types are correlated over time. On the one hand, when cost types are persistent, it becomes possible for a firm to make inferences about rival's costs, using past price observations. Accordingly, one might expect that persistence would give firms more instruments with which to reveal information and assign production efficiently. On the other hand, persistence creates a new incentive to mimic another cost type, since mimicry today influences rival's beliefs in the future. Indeed, when cost types exhibit persistence, repeated play of the equilibrium of a one-shot game is no longer an equilibrium in the dynamic game, and so simply constructing any equilibrium (let alone the best or worst) is non-trivial.

We begin our formal analysis by constructing the best pooling equilibrium. In pooling equilibria, strategies do not depend on firms' past or present cost types. These equilibria are appealing given the considerations just described, since there is no incentive to manipulate opponent beliefs. In the best pooling equilibrium, firms share the market equally at the customer's reservation price in every period. Firms thus achieve high prices but no productive efficiency. In the case of a deviation, they switch to a "punishment" pooling equilibrium that takes a "carrot-stick" form: firms start the punishment by sharing the market at a low price and after each period switch with some probability to an equilibrium where all firms share the market at the reservation price. The best pooling equilibrium exists if firms are sufficiently patient.

Under the assumption that the (initial) distribution of costs is log-concave, we next establish that the best pooling equilibrium is actually the best (unconstrained) equilibrium in the limiting case where costs are perfectly persistent, and is close to optimal when costs are close to perfectly persistent. The case with perfectly persistent costs raises some interesting issues. At first, it might seem that the dynamic game would immediately collapse to a static one when costs do not change over time. But this intuition is incomplete: even though costs do not change, firms can still make different choices at different points in time. For example, there exist equilibria where firms use an initial signaling phase to reveal cost types, followed by a phase where prices are higher but market shares are allocated according to the early signals. Under the log-concavity assumption, however, the equilibrium that maximizes ex ante cartel payoffs is extremely simple, when firms are sufficiently patient: all firms share the market at the customer's reserve price.⁶

⁶ Interestingly, the optimal equilibrium is the same as the optimal equilibrium in the model studied by Athey, Bagwell, and Sanchirico (2004), in which costs are independent over time but firms use Strongly Symmetric Perfect Public Equilibria (SPPE), whereby at the beginning of each period, strategies and expected payoffs are symmetric. The SPPE solution concept is appropriate when the winning price, or even the vector of firm prices, is publicly observed, but a firm is not able to identify the individual behavior of any firm other than itself. In contrast, in this paper we do not use the strong symmetry restriction, and in general firms will not view one another symmetrically after the beginning of the game.

In contrast, when the log-concavity assumption is relaxed, we show that equilibria with an initial signaling phase may yield greater profits than the best pooling equilibria. Intuitively, equilibria with an initial signaling phase entail a tradeoff between the cost of low prices in the signaling phase and the benefit of greater productive efficiency in the future; we provide conditions under which the benefit dominates when log-concavity fails. In such cases, the equilibrium construction must respect the inability of firms to commit to their response to the revelation of information.

We next return to the case of imperfectly persistent costs and establish that some productive efficiency is optimal when firms are sufficiently patient. Indeed, if firms are sufficiently patient to enforce the best pooling equilibrium, then they can enforce as well a simple “odd-even” equilibrium that achieves partial productive efficiency and yields higher expected profits. An odd-even equilibrium employs a simple two-period rotation scheme. In odd periods, firms announce their types, and the low-cost firm gets a high market share. In even periods, if a firm received the high market share in the previous period, it now receives a reduced market share, where the amount of the reduction is chosen to deter mimicry in the previous (odd) period. If costs are independent over time, this scheme would induce some productive efficiency in odd periods, while in even periods the market would be served by an average-cost firm. This scheme is less effective, however, when costs are close to perfectly persistent, since high-cost firms are then likely to serve the market in even periods.

Given that persistence compromises the effectiveness of the odd-even scheme, it is natural to ask whether it is possible to attain first-best collusion. We know that when the serial correlation is equal to zero, the model is equivalent to that of Athey and Bagwell (2001), who show that first-best collusion is possible. But it is not immediately clear that first-best can be attained or approximated in a PPBE once cost types are persistent. With even slight persistence, the game changes in fundamental ways: firms have new incentives to signal and manipulate their opponents’ beliefs in ways that may advantage them in the future.

We show that, despite these complications, when the persistence of types is not too high relative to the patience of firms, first-best collusion is possible. The collusive equilibrium calls for firms to announce their cost types, and firms that have announced low cost in the recent past give up market share in states of the world where the firms have the same cost type. If firms are sufficiently patient, they can be induced to wait for such states. We present numerical examples of first-best collusive schemes. Although these schemes involve subtle incentives, they can be described fairly easily and computed analytically for specific parameter values.

Taken together, our results indicate that when patience is high relative to the persistence of cost types, the best equilibrium entails productive efficiency, high prices, and market shares that are less positively correlated over time than are the cost types. This type of equilibrium is non-stationary and fairly complex. In contrast, when persistence of cost types is large relative to patience of firms, but firms are still moderately patient, the best equilibrium is very simple: it entails productive inefficiency, high prices, and stable market shares. A variety of empirical

evidence and descriptive accounts of collusion establish an association between collusion and rigid prices and stable market shares (Athey, Bagwell and Sanchirico, 2004); yet, there are also case studies of collusive schemes that were quite sophisticated and where market share was exchanged intertemporally among co-conspirators (Athey and Bagwell, 2001).

Finally, we consider whether equilibria with productive efficiency in every period exist at all (even with low prices) when cost types are very persistent. We show that in the continuum-type case, it is sometimes possible to construct an equilibrium that entails productive efficiency in every period, but any such equilibrium yields per-period profits equal to those in the one-shot Nash equilibrium. The equilibrium we construct has above-competitive prices in the first period but lower prices subsequently. Intuitively, a firm gains when it is perceived to have higher costs, since rivals then price less aggressively in future periods. If firms are too patient, the incentive for a lower-cost type to mimic the price of a higher-cost type in the initial signaling phase is overwhelming, and no equilibrium with productive efficiency exists.

The paper proceeds as follows. Section 2 introduces the model. The next two sections analyze pooling equilibria and separating equilibria, respectively. In Section 5, we consider punishment equilibria that entail productive efficiency. Section 6 concludes.

2. The Model

In this section, we introduce the model. There are I firms that meet in periods $t = 1, \dots, \infty$. Throughout, we use the following notational conventions. If \mathcal{X}_i represents a set, then $\mathcal{X} = (\mathcal{X}_1, \dots, \mathcal{X}_I)$. Given a sequence $\{x_{i,t}\}_{t=1}^{\infty}$, where $x_{i,t} \in \mathcal{X}_i$, \mathbf{x}_t denotes $(x_{1,t}, \dots, x_{I,t})$, x_i^t denotes $(x_{i,1}, \dots, x_{i,t})$, and $\mathbf{x}^t = (\mathbf{x}_1, \dots, \mathbf{x}_t)$.

We posit I *ex ante* identical firms that meet in periods $t = 1, \dots, \infty$ to engage in Bertrand competition for sales in a homogenous-good market. Firms discount the future at rate $\delta \in (0, 1)$. We assume that in each period, demand is inelastic, and there is a unit mass of identical consumers with a fixed reservation price r , where $r > \bar{\theta}$, with $\bar{\theta}$ being the highest possible cost for firms. Thus, demand is stationary over time, and consumers are not strategic players. Let $p_{i,t} \in \mathbb{R}_+$ denote the price chosen by firm i .

Firm i 's "cost type" in time t , $\theta_{i,t}$, follows a first-order Markov processes with support $\Theta_i \subseteq [\underline{\theta}, \bar{\theta}]$. The commonly known distribution function is $F(\cdot|\theta_{i,t-1})$, where $\theta_{i,t-1}$ is the firm's cost type in period $t - 1$. Let $\theta_{i,1}$ be drawn from the prior $F_0(\cdot)$ (where this prior is the stationary distribution of the Markov process). To avoid the need for special notation, we use the convention that $F(\cdot|\theta_{i,0}) = F_0(\cdot)$ throughout (even though there may be no such $\theta_{i,0}$ that returns the prior). To represent vectors of cost types and probability distributions, let $\mathbf{F}(\cdot|\boldsymbol{\theta}_{t-1}) = (F(\cdot|\theta_{1,t-1}), \dots, F(\cdot|\theta_{I,t-1}))$. We emphasize that the cost shocks are not affected by any actions the firms may take.

We refer to two special cases of this model throughout the paper.

Model Definitions:

Model 1: $\Theta_i = \{\underline{\theta}, \bar{\theta}\} \in \mathbb{R}^2$, with $\underline{\theta} < \bar{\theta}$. Both $F_0(\cdot)$ and $F(\cdot|\theta_{i,t-1})$ have full support on Θ_i for all $\theta_{i,t-1} \in \Theta_i$. We let $\underline{\theta} = L$ and $\bar{\theta} = H$ and use the notation for types interchangeably. We let $\lambda_{\theta_{i,t-1}} = F(L|\theta_{i,t-1})$. We focus on the case of positive serial correlation, whereby $1 > \lambda_L > \lambda_H > 0$.

Model 2: $\Theta_i = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$. Cost types are perfectly persistent, so that $F(\cdot|\theta_{i,t-1})$ places all of the probability weight on $\theta_{i,t-1}$. F_0 has a strictly positive density over its support.

Thus, Model 1 is a two-type model in which types are imperfectly persistent, whereas Model 2 allows for a continuum of types and focuses on the case of perfect persistence.

Note that if the game ends after period 1, and firms play the Bertrand pricing game described above, the model is equivalent to a one-shot, first-price procurement auction with posted reservation price r , where types are drawn from distribution $F_0(\cdot)$. For both models, the static equilibria when all firms draw types from $F_0(\cdot)$ entail productive efficiency and prices that are always weakly less than the highest cost type.⁷

Each period t of the dynamic game follows a timeline, in which each firm i : (i) privately observes a new cost shock, $\theta_{i,t} \in \Theta_i$; (ii) may engage in “cheap talk,” whereby it publicly announces its cost type $a_{i,t} \in \mathcal{A}_i$ (announcements are simultaneous), where $|\mathcal{A}_i| \geq |\Theta_i|$; (iii) simultaneously selects a price $p_{i,t} \in \mathcal{P}_i = \mathbb{R}_+$ and a maximum quantity it is willing to sell, $q_{i,t} \in \mathcal{Q}_i = [0, 1]$, both of which are publicly observed; (iv) receives market share $\varphi_i(\mathbf{p}_t; \mathbf{q}_t)$, where $\varphi_i : \mathcal{P} \times \mathcal{Q} \rightarrow [0, 1]$ is an exogenous, stationary rationing rule such that if $N(\mathbf{p}_t)$ is the set of firms charging the lowest price, $q_{i,t} \geq \varphi_i(\mathbf{p}_t; \mathbf{q}_t) \geq \min(1/|N(\mathbf{p}_t)|, q_{i,t})$. Thus, if the quantity restrictions for the low-price firms sum to one, each will get exactly its quantity restriction. Also, if each low-price firm selects $q_{i,t} \geq 1/N(\mathbf{p}_t)$, then each low-price firm receives a market share allocation of $1/N(\mathbf{p}_t)$. If $\sum_{i \in N(\mathbf{p}_t)} q_{i,t} < 1$, an analogous rationing rule is used to allocate the residual market share to the set of firms charging the next-lowest price, and so on until the market share is exhausted or there are no firms remaining with non-zero quantity restrictions. The details of the rationing rule are not important for any of our results.

We pause to provide some additional explanation for the announcements and quantity restrictions. They are included to allow firms to share the market in unequal proportions as a function of firm cost types. An equilibrium might specify that following some history, if two firms both announce low costs and tie for the lowest price, then one firm gets only 1/3 of the market. If that firm deviates from the agreement and takes more than 1/3, the deviation will be observable by opponents. If announcements were eliminated, then quantity restrictions might lead firms to underproduce in some states of the world. However, if the model was modified

⁷ Athey and Bagwell (2001) and Riley and Samuelson (1981) characterize the static equilibria for Models 1 and 2, respectively. When firms can be asymmetric (i.e., have different cost distributions), there is typically productive inefficiency in Model 2. Athey and Bagwell (2004) characterize the static equilibrium for asymmetric firms for Model 1, while the static equilibrium for Model 2 is characterized by Maskin and Riley (2000).

so that quantity restrictions came after pricing decisions, prices could take over some of the coordination role of the quantity restrictions. In the absence of both announcements and quantity restrictions, the set of possible firm market shares would be quite limited. Despite this, below we show that in many cases, announcements and quantity restrictions are not necessary to achieve optimal equilibria. In other cases, eliminating these features would lead to reductions in efficiency for some parameter values. Including them allows us to analyze whether communication is useful, an important question for antitrust policy.

Summarizing, within a period, after announcements, prices, and quantity restrictions are determined as $(\mathbf{a}_t, \mathbf{p}_t, \mathbf{q}_t)$, a firm will receive profits $(p_{i,t} - \theta_{i,t})\varphi_i(\mathbf{p}_t; \mathbf{q}_t)$. Notice that announcements do not directly affect profits; they simply influence the firms' choices of prices and quantity restrictions.

Let $\mathcal{Z}_i = \mathcal{A}_i \times \mathcal{P}_i \times \mathcal{Q}_i$ and $\mathcal{Z} = (\mathcal{Z}_1, \dots, \mathcal{Z}_I)$. The set of possible ‘‘period strategies’’ in a given period t for firm i is given by $\mathcal{S}_i \subset \{s_{i,t} | s_{i,t} : \mathcal{A} \times \Theta_i \rightarrow \mathcal{Z}_i\}$, where the elements $s_{i,t}$ of \mathcal{S}_i must be decomposable into three component functions, $s_{i,t} = (\alpha_{i,t}, \rho_{i,t}, \psi_{i,t})$, satisfying the following restrictions. The announcement function $\alpha_{i,t}$ depends only on the firm's own cost type. Since announcements precede pricing and quantity decisions, the latter choices depend on both the firm's own type as well as the announcements of others. In other words, when types are given by $\boldsymbol{\theta}_t$, firm i first announces $a_{i,t} = \alpha_{i,t}(\theta_{i,t})$. Then each firm i observes $\mathbf{a}_{-i,t} = \boldsymbol{\alpha}_{-i,t}(\boldsymbol{\theta}_{-i,t})$. Next, each firm i sets price $p_{i,t} = \rho_{i,t}(\mathbf{a}_t, \theta_{i,t})$ and chooses quantity restriction $q_{i,t} = \psi_{i,t}(\mathbf{a}_t, \theta_{i,t})$. Finally, market shares are determined by $\varphi_{i,t}(\mathbf{p}_t; \mathbf{q}_t)$.

In this paper, we will occasionally analyze a situation where one firm wishes to undercut another. Selectively, in those cases, we will use the convention that $\varepsilon > 0$ is the smallest price increment, so that a firm undercuts by charging ε or 2ε less than its opponent. However, we will not necessarily include ε in any profit calculations. Thus, when we say that ‘‘firm j charges $p_i - \varepsilon$,’’ this should be interpreted as saying that firm j charges a negligible amount less than p_i , thereby winning the market at (essentially) price p_i .

2.1. Interim Profits

We now introduce notation for the ‘‘interim’’ payoffs for firm i , after firm i knows its cost type in period t but before firm i acts. Let ν_i be a probability distribution over Θ_i . Given a function g , we let $\mathbb{E}_{\theta_{i,t}}[g(\theta_{i,t})|\nu_i]$ denote the expectation over values of $\theta_{i,t}$ taking ν_i as the probability distribution over $\theta_{i,t}$. When $\nu_i = F(\cdot; \theta_{i,t-1})$, we simply write $\mathbb{E}_{\theta_{i,t}}[g(\theta_{i,t})|\theta_{i,t-1}]$.

At the interim stage, for any given period strategy $s_{i,t} = (\alpha_{i,t}, \rho_{i,t}, \psi_{i,t})$, a firm can deviate from this strategy in several ways. The firm might choose a deviant announcement ($a_{i,t} \neq \alpha_{i,t}(\theta_{i,t})$); it might choose prices and quantity restrictions that are inconsistent with the set of realized announcements or its own type ($p_{i,t} \neq \rho_{i,t}(\mathbf{a}_t, \theta_{i,t})$ or $q_{i,t} \neq \psi_{i,t}(\mathbf{a}_t, \theta_{i,t})$); or it might do some combination of these things. All of these possible deviations can be represented by an alternative strategy $\tilde{s}_{i,t} \neq s_{i,t}$.

However, there is one particular type of deviation, termed an “on-schedule deviation,” that will play a special role in the analysis. In this type of deviation, $\tilde{s}_{i,t}$ specifies that type $\theta_{i,t}$ “mimics” $\hat{\theta}_{i,t} \neq \theta_{i,t}$: that is, $\tilde{\alpha}_{i,t}(\theta_{i,t}) = \alpha_{i,t}(\hat{\theta}_{i,t})$, and for all \mathbf{a}_t , $\tilde{\rho}_{i,t}(\mathbf{a}_t, \theta_{i,t}) = \rho_{i,t}(\mathbf{a}_t, \hat{\theta}_{i,t})$ and $\tilde{\psi}_{i,t}(\mathbf{a}_t, \theta_{i,t}) = \psi_{i,t}(\mathbf{a}_t, \hat{\theta}_{i,t})$. Although this deviation can be represented directly through the strategy $\tilde{s}_{i,t}$, it will be more convenient to introduce direct notation for mimicry. Formally, if firm i ’s beliefs about opponent types at the start of period t are given by $\boldsymbol{\nu}_{-i}$, then the following expressions represent interim expected market share and expected profits for firm i when firm i has type $\theta_{i,t}$ but mimics the behavior that type $\hat{\theta}_{i,t}$ would use given period strategy $s_{i,t}$:

$$\bar{m}_i(\hat{\theta}_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i}) = \mathbb{E}_{\boldsymbol{\theta}_{-i,t}} \left[\varphi_i \left(\boldsymbol{\rho}_t(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), (\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t})); \boldsymbol{\psi}_t(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), (\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t})) \right) \middle| \boldsymbol{\nu}_{-i} \right],$$

$$\bar{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i}) = \mathbb{E}_{\boldsymbol{\theta}_{-i,t}} \left[\begin{array}{c} (\rho_{i,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), \hat{\theta}_{i,t}) - \theta_{i,t}) \cdot \\ \varphi_i \left(\boldsymbol{\rho}_t(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), (\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t})); \boldsymbol{\psi}_t(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), (\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t})) \right) \end{array} \middle| \boldsymbol{\nu}_{-i} \right].$$

We emphasize that the notation is redundant (that is, given $\tilde{s}_{i,t}$ as defined above, $\bar{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i}) = \bar{\pi}_i(\theta_{i,t}, \theta_{i,t}; (\tilde{s}_{i,t}, \mathbf{s}_{-i,t}), \boldsymbol{\nu}_{-i})$), and we do not in any way limit the set of possible deviations to these “on-schedule” deviations.

We distinguish on-schedule deviations from “off-schedule” deviations. In an off-schedule deviation, a firm chooses an action or a series of actions that no cost type should have chosen in equilibrium. Since off-schedule deviations are observable and should never happen in equilibrium, they can be severely punished so as to deter the deviation. In contrast, on-schedule deviations are not observable as deviations. Thus, any future punishment associated with mimicking $\hat{\theta}_{i,t}$ must be borne in equilibrium, when type $\hat{\theta}_{i,t}$ actually occurs.

We say that firm i ’s announcement is *uninformative* if there exists some constant c_i such that $\alpha_{i,t}(\theta_{i,t}) = c_i$ for all $\theta_{i,t}$. In the case where all firms’ announcements are uninformative,

$$\bar{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i}) = (\rho_{i,t}(\mathbf{c}, \hat{\theta}_{i,t}) - \theta_{i,t}) \cdot \bar{m}_i(\hat{\theta}_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i}); \quad (2.1)$$

that is, expected profits are equal to price minus cost times expected market share. In many of the collusive equilibria we construct, the price will not depend directly on the announcements, as it will be constant at r ; then, (2.1) can be used.

2.2. Evolution of Beliefs in the Dynamic Game

We now develop notation for the evolution of firms’ beliefs using Bayes’ rule. Let $\Delta\Theta_i$ be the set of probability distributions over Θ_i . The “product structure” of the game implies that firm j ’s private information about its history of cost types and past deviations does *not* provide firm j with any relevant information about firm i . We thus impose that the beliefs of each firm $j \neq i$ about firm i ’s cost shocks evolve in the same way along the equilibrium path. We can then say that “firms’ beliefs about opponents in period t are given by $\boldsymbol{\mu}_t$,” by which we mean that

for each firm $j \neq i$, firm j has belief $\mu_{i,t}$ about $\theta_{i,t}$. Note that $\mu_{i,t}$ may differ from $F(\cdot|\theta_{i,t-1})$, which is the belief that firm i has about $\theta_{i,t}$ at the start of period t .

Let $\mu_{i,t}^p$ be the posterior belief that firms $j \neq i$ hold about $\theta_{i,t}$ at the end of period t , after observing \mathbf{z}_t , given a conjectured period strategy $s_{i,t}$ and the period's prior belief $\mu_{i,t}$. We say that \mathbf{z}_t is *compatible* with $s_{i,t}$ and $\mu_{i,t}$ if there exists some cost type $\theta_{i,t}$ such that $\theta_{i,t}$ is in the support of $\mu_{i,t}$ and $z_{i,t} = s_{i,t}(\mathbf{a}_t, \theta_{i,t})$. In this case, the posterior $\mu_{i,t}^p$ is determined using Bayes' rule. On the other hand, if \mathbf{z}_t is not compatible with $s_{i,t}$ and $\mu_{i,t}$, then Bayes' rule does not pin down the posterior beliefs $\mu_{i,t}^p$. The posterior belief is then specified by the analyst.

We let $\tilde{T} : \Delta\Theta_i \times \mathcal{S}_i \times \mathcal{Z} \rightarrow 2^{\Delta\Theta_i}$ denote the correspondence that gives the set of possible period $t+1$ beliefs about $\theta_{i,t}$ given $(\mu_{i,t}, s_{i,t}, \mathbf{z}_t)$. For all \mathbf{z}_t that are compatible with $s_{i,t}$ and $\mu_{i,t}$, \tilde{T} is single-valued. It is the belief that, as a Bayesian (and given its knowledge of the stochastic process for the evolution of costs) firm j should have about firm i at the beginning of period $t+1$, given that firm j started period t with belief $\mu_{i,t}$, conjectured that firm i used period strategy $s_{i,t}$, and observed a compatible vector of public actions \mathbf{z}_t . Formally, $\tilde{T}(\mu_{i,t}, s_{i,t}, \mathbf{z}_t) = \mathbb{E}_{\theta_{i,t}} \left[F(\cdot|\theta_{i,t}) | \mu_{i,t}^p \right]$.

Consider now $(\mu_{i,t}, s_{i,t}, \mathbf{z}_t)$ such that \mathbf{z}_t is not compatible with $s_{i,t}$ and $\mu_{i,t}$. Then,

$$\tilde{T}(\mu_{i,t}, s_{i,t}, \mathbf{z}_t) = \left\{ \mu_{i,t+1} \in \Delta\Theta_i : \exists \mu_{i,t}^p \in \Delta\Theta_i \text{ s.t. } \mu_{i,t+1} = \mathbb{E}_{\theta_{i,t}} \left[F(\cdot|\theta_{i,t}) | \mu_{i,t}^p \right] \right\}.$$

The period $t+1$ belief about $\theta_{i,t+1}$ is then selected from this set, with the particular selection corresponding to the posterior that is specified for period t . Note that even though Bayes' rule no longer determines the posterior belief $\mu_{i,t}^p$, the updating rule still places restrictions on beliefs about $\theta_{i,t+1}$. For example, if the posterior $\mu_{i,t}^p$ places all of the probability weight on $\theta_{i,t} = \theta'$, then $\mu_{i,t+1} = F(\cdot|\theta_{i,t} = \theta')$; thus, if F has full support, $\mu_{i,t+1}$ must have full support as well.

Henceforth, we let $T(\mu_{i,t}, s_{i,t}, \mathbf{z}_t) \in \tilde{T}(\mu_{i,t}, s_{i,t}, \mathbf{z}_t)$ be a selection from \tilde{T} (recalling that the selection is pre-determined for \mathbf{z}_t that is compatible with $s_{i,t}$ and $\mu_{i,t}$). Then, we may describe the evolution of beliefs in a simple, recursive form: $\mu_{i,t+1} = T(\mu_{i,t}, s_{i,t}, \mathbf{z}_t)$. We may summarize the updating function for all firms at once as $\mathbf{T}(\boldsymbol{\mu}_t, \mathbf{s}_t, \mathbf{z}_t) = (T(\mu_{1,t}, s_{1,t}, \mathbf{z}_t), \dots, T(\mu_{I,t}, s_{I,t}, \mathbf{z}_t))$.

Consider some examples of the evolution of beliefs about opponents. If $s_{i,t}$ assigns a unique vector of public actions to each type, then firm j can infer perfectly from \mathbf{z}_t the value of $\theta_{i,t}$. In that case, $\mu_{i,t+1}(\cdot) = F_i(\cdot|\theta_{i,t})$. On the other hand, if $\mu_{i,t} = F_0$, and $s_{i,t}$ always assigns the same actions $z_{i,t}$ to all cost types, then after \mathbf{z}_t is observed, $\mu_{i,t+1} = F_0$, since nothing is learned in period t and F_0 is the stationary distribution of the Markov process.

2.3. Extensive-Form Strategies in the Dynamic Game

To analyze this game, we need to specify a solution concept. We will consider Perfect Bayesian Equilibria (PBE) that satisfy a further refinement, namely that player strategies depend on public histories and their own payoff-relevant private information, but not otherwise on their history of private information. We refer to the solution concept we impose here as Perfect Public

Bayesian Equilibrium (PPBE). As will be clear, PPBE is a natural extension of Perfect Public Equilibrium (PPE) to dynamic Bayesian (rather than repeated) games. Given that firms are ex ante symmetric, we impose the further restriction that firms use ex ante symmetric strategies; that is, strategies are exchangeable as a function of public and private histories. This implies that asymmetries in firm strategies starting in period t arise as a result of asymmetric cost realizations and behavior in the past. This restriction does not affect the qualitative nature of our results, but it greatly simplifies the exposition.

We begin with some notation. The observable public history in the dynamic game is an infinite sequence of realized reports, prices, and quantity restrictions, $h^t = \{\mathbf{a}^t, \mathbf{p}^t, \mathbf{q}^t\}$. Let \mathcal{H}^t be the set of possible public histories at date t . In addition to the public history, at the beginning of period t , firm i knows θ_i^{t-1} , its past history of cost types. We also assume that after every period, firms can observe the realization of some public randomization device and then select continuation equilibria on this basis. To ease the notational burden, we do not introduce explicit notation for the randomization process.

We consider “public strategies,” whereby firm i ’s strategy maps any history with the same h^{t-1} and the same $\theta_{i,t}$ to the same actions. Firm i ’s strategy in period t does *not* depend on the private history θ_i^{t-1} . It makes sense that firm i would not care about past cost types once it learns its current cost type: past cost types do not influence its own beliefs (given that types are Markov); and they were not observed by any opponents, so they do not affect opponent beliefs either. Thus, in our game, when all opponents use public strategies, each firm has a best response that is public. Given the stochastic process for cost types, our definition of public strategies is the analog of the standard definition for repeated games.⁸

Firm i ’s strategy in the extensive-form game can be described by a sequence $\sigma_{i,t}$, where $\sigma_{i,t} : \mathcal{H}^{t-1} \rightarrow \mathcal{S}_i$ and the dependence of actions on current cost types is incorporated in the period strategies $s_{i,t} \in \mathcal{S}_i$. The full (public and private) history of the game can be described by $\{h^t, \theta^t\}_{t=1}^\infty$. Given $\{\theta_t, \sigma_t\}_{t=1}^\infty$, a full path of play, $\{\mathbf{a}_t, \mathbf{p}_t, \mathbf{q}_t, \theta_t\}_{t=1}^\infty$, is induced. Payoffs for firm i at time 1 thus may be written as:

$$\sum_{t=1}^\infty \delta^{t-1} (p_{i,t} - \theta_{i,t}) \varphi_i(\mathbf{p}_t; \mathbf{q}_t),$$

where $s_{i,t} = \sigma_{i,t}(h^{t-1}) \forall i$, and $(a_{i,t}, p_{i,t}, q_{i,t}) = s_{i,t}(\boldsymbol{\alpha}_t(\boldsymbol{\theta}_t), \theta_{i,t})$.

It is more useful to develop notation for the expected payoffs for firm i at time τ , which can be written as a function of beliefs about opponents $\boldsymbol{\mu}_{-i,\tau}$, public history $h^{\tau-1}$, and firm i ’s

⁸ Our focus on PPBE thus excludes from consideration PPE that employ private strategies. See Kandori and Obara (2006) for further discussion.

current cost shock $\theta_{i,\tau}$, where T is a selection from \tilde{T} , as defined in Section 2.2:

$$\begin{aligned} \tilde{v}_i(\{\sigma_t\}_{t=\tau}^\infty, \boldsymbol{\mu}_{-i,\tau}, h^{\tau-1}, \theta_{i,\tau}) &= \mathbb{E}_{\{\boldsymbol{\theta}^t\}_{t=\tau}^\infty} \left[\sum_{t=\tau}^\infty \delta^{t-\tau} \mathbb{E}_{\theta_{i,t}} [\bar{\pi}_i(\theta_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}) | \theta_{i,t-1}] | \boldsymbol{\mu}_{-i,\tau}, \theta_{i,\tau} \right], \\ &\text{where, for all } t \geq \tau, \\ s_{j,t} &= \sigma_{j,t}(h^{t-1}) \text{ for all } j \in \{1, \dots, I\}, \\ z_{j,t} &= (a_{j,t}, p_{j,t}, q_{j,t}) = s_{j,t}(\boldsymbol{\alpha}_t(\boldsymbol{\theta}_t), \theta_{j,t}) \text{ for all } j \in \{1, \dots, I\}, \text{ and} \\ \boldsymbol{\mu}_{-i,t+1} &= \mathbf{T}_{-i}(\boldsymbol{\mu}_{-i,t}, \mathbf{s}_{-i,t}, \mathbf{z}_t). \end{aligned}$$

When taking the expectation in this definition, the distribution over future actions and cost types is induced by the strategies and current beliefs about cost types. Using this notation, a PPBE is a collection of public strategies $\{\boldsymbol{\sigma}_t^*\}_{t=1}^\infty$ such that $\sigma_i^* = \{\sigma_{i,t}^*\}_{t=1}^\infty$ is weakly optimal at every information set, together with initial beliefs $\boldsymbol{\mu}_1$ and a belief updating function $T \in \tilde{T}$ that determines firms' beliefs about opponents at each date t so that, for all i , h^{t-1} , and \mathbf{z}_t ,

$$\mu_{i,t+1} = T(\mu_{i,t}, \sigma_{i,t}^*(h^{t-1}), \mathbf{z}_t). \quad (2.2)$$

Formally, for all $\theta_{i,\tau}$, $h^{\tau-1}$, and firm i 's beliefs about the firms' opponents $\boldsymbol{\mu}_{-i,\tau}$ generated through (2.2) by $h^{\tau-1}$ and $\{\boldsymbol{\sigma}_{-i,t}^*\}_{t=1}^{\tau-1}$,

$$(\sigma_{i,t}^*)_{t=\tau}^\infty = \arg \max_{(\sigma_{i,t})_{t=\tau}^\infty} \tilde{v}_i(\{\sigma_{i,t}, \boldsymbol{\sigma}_{-i,t}^*\}_{t=\tau}^\infty, \boldsymbol{\mu}_{-i,\tau}, h^{\tau-1}, \theta_{i,\tau}). \quad (2.3)$$

It is straightforward to verify that our definitions imply that a PPBE is a PBE.⁹

3. Pooling Equilibria

This section focuses on pooling equilibria, whereby strategies do not depend on firms' cost shocks. In order to evaluate the efficiency of equilibria, we introduce the concept of ‘‘partial productive efficiency.’’ In a given period, if firms sell to the entire unit-mass of consumers, then the firms achieve partial productive efficiency if the expected industry production cost (weighted across firms by market shares) is less than that which would be achieved were instead each firm assigned an equal share of the market. The firms achieve productive efficiency, if the lowest-cost firm (or firms) always receives all market share. We also introduce the notion of a ‘‘scheme.’’ We use this word to refer to a set of strategies with particular properties. This terminology is useful, since particular strategies may or may not represent equilibrium strategies, or may represent equilibrium strategies only for particular parameter values. Given a [name] scheme, we use the term ‘‘[name] equilibrium’’ to refer to a PPBE in which a [name] scheme is used.

⁹ See Fudenberg and Tirole (1991, pp. 331-3) for the requirements of PBE. One restriction that we have not explicitly imposed on our selection T from \tilde{T} is that firm i 's belief about firm j 's type is consistent with the understanding that the behavior of any other firm cannot signal this information. It would be cumbersome to formally specify this requirement given that there are two stages in each period of the game, and the requirement does not affect any of the analysis we conduct in this paper. However, in general this requirement should also be imposed on the selection T .

3.1. Pooling Equilibria Exist when Firms Are Patient

Pooling equilibria may take several forms. In one class of stationary pooling equilibria, along the equilibrium path, the strategies specify that firms always share the market equally at a particular price. We refer to such strategies as a *rigid-pricing scheme*. In the *best rigid-pricing scheme*, firms share the market at the price r as long as no firms deviate.

Rigid-Pricing Scheme: A set of strategies where, on the equilibrium path, firms share the market equally at a fixed price p' . Along the equilibrium path, for all i and t , $p_{i,t} = p'$ and $q_{i,t} = 1$ (non-binding quantity restrictions), and announcements are uninformative. In a best rigid-pricing scheme, $p' = r$.

Another class of pooling equilibria use non-stationary strategies that take a “carrot-stick” form. Formally, we define a *carrot-stick scheme* as follows:

Carrot-Stick Scheme: A set of strategies where, on and off of the equilibrium path, announcements are uninformative and $q_{i,t} = 1$ (non-binding quantity restrictions). There are two states. In the war state, all firms choose price $p_w < r$; and in the reward state, all firms charge price equal to r . The firms begin in the war state. In the war state, if all firms choose price p_w in a given period, the firms switch to the reward state with probability $\chi \in [0, 1]$, and return to the war state with probability $1 - \chi$. In the reward state, if all firms choose price r in a given period, the firms remain in the reward state with probability 1. In each period, if any firm charges a price other than the assigned price, the firms switch to the war state with probability 1. In a *worst carrot-stick scheme*, in the initial period, a firm with the highest type, $\bar{\theta}$, is indifferent between choosing the price p_w and charging a higher price, where by charging a higher price the firm would sell nothing and restart the scheme with probability 1.

In a carrot-stick scheme, firms may be induced to price below cost in the present period, when they anticipate the reward of getting a high price in the future. But no such scheme can be used in a PPBE if the highest type, $\bar{\theta}$, is unwilling to charge p_w in the war state. In particular, the war state prevails in the initial period, and the highest type then cannot prefer a deviation in which it selects a higher price, sells nothing and restarts the war state in the second period. This motivates our definition of the worst carrot-stick scheme. As confirmed in the Appendix, in a worst carrot-stick scheme, if a firm begins the game with the highest type, then it earns a payoff, $v^{cs}(\bar{\theta})$, that is positive in Model 1 and zero in Model 2. Intuitively, in Model 1, if the payoff to such a firm is too low, then this firm would prefer to “sit out” for a period and re-start the war state in the next period, when its type may be $\underline{\theta}$ instead. We show further that, in a worst carrot-stick scheme, the payoff to a firm with cost type $\theta_{i,t}$ is

$$v^{cs}(\theta_{i,t}) = v^{cs}(\bar{\theta}) + \frac{(\gamma_{\bar{\theta}} - \gamma_{\theta_{i,t}})}{I(1 - \delta)},$$

where $\frac{\gamma_{\theta_{i,t}}}{1-\delta} = \mathbb{E} \left[\sum_{s=t}^{\infty} \delta^{s-t} \theta_{i,s} \mid \theta_{i,t} \right]$ is the expected discounted unit cost for firm i , given that firm i currently has cost type $\theta_{i,t}$.

Within the class of strategies characterized by pooling, payoffs are clearly maximized when the best rigid-pricing scheme is used. Thus, if firms are sufficiently patient to enforce the best rigid-pricing scheme in a PPBE, then the resulting best rigid-pricing equilibrium is a *best pooling equilibrium*. As is well known, it is easier to enforce an equilibrium when a more severe punishment is employed following any deviation. We impose here that the punishment takes the form of a worst carrot-stick equilibrium.¹⁰ The following result establishes the existence of a critical discount factor above which a best (rigid-pricing) pooling equilibrium can be supported by a worst carrot-stick pooling equilibrium.

Proposition 1. *Consider either Model 1 or Model 2. Let $\delta_c \equiv \frac{(r-\theta)(I-1)}{(r-\theta)I-(\theta-\theta)} \in (\frac{I-1}{I}, 1)$. For all $\delta \in [\delta_c, 1)$, there exists a worst carrot-stick equilibrium. Thus, a best rigid-pricing equilibrium exists where following a deviation, firms switch to a worst carrot-stick equilibrium.*

Proof: See Appendix.

We observe that the critical discount factor is independent of the degree of persistence of cost types. The proof shows that the critical discount factor is determined by the constraint that a firm with the lowest cost type, $\underline{\theta}$, must not gain from deviating by undercutting the price r in the reward state. This constraint reflects a tradeoff for such a firm between the immediate benefit from undercutting and thus serving the entire market and the future cost of triggering the war state and thereby diminishing the continuation value. The cost outweighs the benefit if the firm is sufficiently patient. Intuitively, the persistence of the firm's cost type enters by affecting the continuation values that it expects after cooperation and deviation, respectively. But whether the firm cooperates or deviates will not affect its expected discounted unit cost level, $\gamma_{\underline{\theta}}$. Furthermore, in a carrot-stick equilibrium, the market share that a firm enjoys is independent of whether the firm deviated in the past. Thus, while a deviation reduces a firm's expected future revenue, it does not alter the firm's expected future costs. Consequently, the extent to which a deviation causes a diminishment in the continuation value is independent of the degree of persistence in cost types.

Pooling equilibria are appealingly simple, in that both on and off the equilibrium path, no messages are sent and no new information about types is revealed. If initial beliefs are equal to the prior, firms' beliefs about opponents remain fixed at the prior. Thus, a key feature of our constructed pooling equilibria is that beliefs about opponents do not play an important role. Carrot-stick equilibria can thus serve as punishments supporting other equilibria without modification, even when information about types has been revealed.

¹⁰ In general, there will be more severe punishments than the worst carrot-stick scheme; in particular, punishments with some separation of types will typically be worse. See Section 5 for an example.

3.2. The Optimality of Pooling Equilibria with Extreme Persistence

In this subsection, we establish that pooling is (approximately) optimal when the persistence of cost shocks is (near-) perfect, if the prior distribution F_0 is log-concave or if r is large enough. We show that partial pooling may be optimal when these conditions are not satisfied.

3.2.1. Pooling Is Optimal with Perfect Persistence

We focus here on Model 2 and thus analyze the special case where firm cost types are perfectly persistent. In particular, we characterize the best collusive equilibrium from the perspective of firms at the beginning of the game, before they learn their cost types. As a punishment equilibrium, we use the worst carrot-stick equilibrium of Proposition 1. Our main result is robust to allowing for small probabilities that types change over time, and thus the result can be thought of as a limiting case.

Recall now the distinction between on-schedule (unobservable) deviations and off-schedule (observable) deviations. Under perfect persistence, mimicking one type throughout the entire game is an on-schedule deviation, and any other deviation is an off-schedule deviation. This is true because mimicking one type, and then another type, reveals that the firm must have misrepresented at one point, given that cost types do not change throughout the game.

The main result of this subsection establishes that the best rigid-pricing equilibrium described in Proposition 1 is optimal ex ante under fairly mild parameter restrictions. To prove this result, we proceed in two steps. First, we maximize ex ante firm profits in a relaxed setting, where we choose market share and revenue functions directly and require only that on-schedule deviations are unattractive. When F_0 is log-concave or r is large enough, we show that the optimal market share and revenue functions can be achieved when firms use the best rigid-pricing scheme. Second, we consider whether the solution in the relaxed setting is also immune to off-schedule deviations. Here, we recall from Proposition 1 that the best rigid-pricing scheme is actually immune to all deviations, including off-schedule deviations, when $\delta \geq \delta_c$.

Proposition 2. *Consider Model 2. Suppose that $\delta \geq \delta_c$, for δ_c defined in Proposition 1. Then, if F_0 is log-concave or if r is large enough, an (ex ante) optimal PPBE is the best rigid-pricing equilibrium described in Proposition 1.*

Proof: First, we maximize ex ante firm profits in a relaxed setting, in which we choose market share and revenue functions subject to the on-schedule incentive constraint. Let $\check{R}_i(\hat{\theta}_{i,1})$ and $\check{M}_i(\hat{\theta}_{i,1})$ denote the expected future discounted revenues and market shares that firm i anticipates if it mimics type $\hat{\theta}_{i,1}$ throughout the game. If firm i 's type is $\theta_{i,1}$, then the present discounted value of profits for firm i can be represented as $U_i(\hat{\theta}_{i,1}, \theta_{i,1}) \equiv \check{R}_i(\hat{\theta}_{i,1}) - \theta_{i,1}\check{M}_i(\hat{\theta}_{i,1})$. The on-schedule incentive-compatibility constraint requires that, for all i ,

$$U_i(\theta_{i,1}, \theta_{i,1}) \geq U_i(\hat{\theta}_{i,1}, \theta_{i,1}) \text{ for all } \hat{\theta}_{i,1}, \theta_{i,1}. \quad (3.1)$$

By standard arguments, (3.1) holds only if

$$\check{M}_i(\theta_{i,1}) \text{ is nonincreasing in } \theta_{i,1} \quad (3.2)$$

and (by the envelope theorem)

$$U_i(\theta_{i,1}, \theta_{i,1}) = U_i(\bar{\theta}, \bar{\theta}) + \int_{\check{\theta}_i=\theta_{i,1}}^{\bar{\theta}} \check{M}_i(\check{\theta}_i) d\check{\theta}_i. \quad (3.3)$$

Using (3.3) and integration by parts, (3.1) implies that

$$\mathbb{E}_{\theta_{i,1}} [U_i(\theta_{i,1}, \theta_{i,1}) | F_0] = \mathbb{E}_{\theta_{i,1}} \left[\check{R}_i(\bar{\theta}) - \bar{\theta} \cdot \check{M}_i(\bar{\theta}) + \frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})} \check{M}_i(\theta_{i,1}) \middle| F_0 \right]. \quad (3.4)$$

We now define the relaxed program: Choose a value $\check{R}_i(\bar{\theta})$ and a function $\check{M}_i(\theta_{i,1})$ to maximize (3.4) subject to (3.2),

$$\mathbb{E}_{\theta_{i,1}} [\check{M}_i(\theta_{i,1}) | F_0] = \frac{1}{1-\delta} \frac{1}{I}, \text{ and} \quad (3.5)$$

$$\check{R}_i(\bar{\theta}) \leq r \cdot \check{M}_i(\bar{\theta}), \quad (3.6)$$

where (3.5) is imposed since firms (and strategies) are assumed ex ante symmetric¹¹ and (3.6) is imposed since type $\bar{\theta}$ cannot sell its market share at a price higher than r . We emphasize that the constraints of the relaxed program are substantially less restrictive than those that would be imposed were we to maximize over strategies that form a PPBE.

If F_0 is log-concave or if r is large enough, we may now use Proposition 5 of Athey, Bagwell and Sanchirico (2004) to show that the relaxed program is solved when $\check{M}_i(\theta_{i,1}) \equiv \frac{1}{1-\delta} \frac{1}{I}$ and $\check{R}_i(\bar{\theta}) = r \cdot \check{M}_i(\bar{\theta})$. In particular, their result shows that among all functions $\check{M}_i(\theta_{i,1})$ that satisfy (3.2) and (3.5), the last term of (3.4) is maximized when $\check{M}_i(\theta_{i,1}) \equiv \frac{1}{1-\delta} \frac{1}{I}$ if F_0 is log-concave. The central idea is easily understood: log-concavity is equivalent to $\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}$ nondecreasing, and so the last term of (3.4) is maximized when \check{M}_i puts as much weight as possible on high realizations of $\theta_{i,1}$, subject to the constraints that \check{M}_i is nonincreasing and must achieve a given average value (i.e., subject to (3.2) and (3.5)). This is accomplished when \check{M}_i is constant in $\theta_{i,1}$.¹² Consider now the first term in (3.4). Observe that $\check{R}_i(\bar{\theta}) - \bar{\theta} \cdot \check{M}_i(\bar{\theta}) \leq [r - \bar{\theta}] \check{M}_i(\bar{\theta}) \leq \frac{r - \bar{\theta}}{(1-\delta)I}$, where the first inequality follows from (3.6) and the second inequality follows from (3.2) and

¹¹ If we allowed for ex ante asymmetric strategies (and thus market shares), we would need to modify the arguments to require that total ex ante expected market share (across all firms) is $1/(1-\delta)$, but we would allow the ex ante expectations to differ across firms. Given any vector of firm-specific ex ante expected market shares, pooling will be optimal under the conditions stated in the proposition, except that the critical discount factor must be increased due to the increased temptation of the disadvantaged firm to cheat.

¹² Formally, since $\int_{\bar{\theta}}^{\bar{\theta}} I(1-\delta) \cdot \check{M}_i(\theta_{i,1}) \cdot f_0(\theta_{i,1}) d\theta_{i,1} = 1$, $g(\theta_{i,1}; \check{M}_i) = I(1-\delta) \cdot \check{M}_i(\theta_{i,1}) \cdot f_0(\theta_{i,1})$ is a probability density. Since F_0/f_0 is nondecreasing, the expected value of F_0/f_0 with respect to g is increased if the associated distribution $G(\theta_{i,1}; \check{M}_i)$ is shifted by First-Order Stochastic Dominance (FOSD). But if $M_i^r(\theta_{i,1}) \equiv 1/(I(1-\delta))$, then $G(\theta_{i,1}; M_i^r)$ dominates (FOSD) $G(\theta_{i,1}; \check{M}_i)$ for all valid market-share functions $\check{M}_i(\theta_{i,1})$.

(3.5). Thus, the first term in (3.4) is maximized when $\check{M}_i(\bar{\theta}) \equiv \frac{1}{1-\delta} \frac{1}{\bar{I}}$ and $\check{R}_i(\bar{\theta}) = r \cdot \check{M}_i(\bar{\theta})$. Likewise, if r is large enough, then the first term of (3.4) dominates the second, so that again the maximum is achieved when $\check{M}_i(\bar{\theta}) \equiv \frac{1}{1-\delta} \frac{1}{\bar{I}}$ and $\check{R}_i(\bar{\theta}) = r \cdot \check{M}_i(\bar{\theta})$.

We now observe that the revenue value and market share function that solves the relaxed program, in fact, can be achieved by strategies in the extensive form game. In particular, the best rigid-pricing scheme delivers ex ante market share and revenue to each firm of $\frac{1}{1-\delta} \frac{1}{\bar{I}}$ and $r \cdot \frac{1}{1-\delta} \frac{1}{\bar{I}}$, respectively, and so the solution to the relaxed program is achieved. We have already established in Proposition 1 that when $\delta \geq \delta_c$, the best rigid-pricing scheme is used in a PPBE of the game. Thus, the best rigid-pricing equilibrium described in Proposition 1 must be an optimal equilibrium, when $\delta \geq \delta_c$ and F_0 is log-concave or r is large enough. ■

Before proceeding, we comment on the approach of the proof of Proposition 2. We considered the solution to a relaxed problem, where strategies only needed to be immune to a restricted class of deviations, whereby one cost type mimics another for all time. It was fortuitous that the solution took such a simple form, and that the solution was immune to all types of deviations.¹³ A different solution to the relaxed problem might not have been an equilibrium to the original game. Indeed, this situation occurs when we relax the parameter restrictions of Proposition 2, as we show in Section 3.2.3.

It might seem that it would be advantageous for firms to signal their cost types at the beginning of the game and then enjoy the benefits of collusion with complete information about costs for the remainder of the game. However, Proposition 2 shows that such an equilibrium is not optimal if firms are sufficiently patient. Part of the intuition follows from the fact that the signaling costs must be proportional to the gains from signaling. If a low-cost firm expects a market-share advantage for the remainder of the game, it will be willing to expend great costs signaling at the beginning. A more subtle intuition highlights the role of log-concavity of F_0 . This is an important condition: as we confirm below in Section 3.2.3, if r is close to $\bar{\theta}$ and F_0 is not log-concave, the best rigid-pricing scheme may no longer be optimal. Log-concavity is equivalent to requiring that $\frac{F_0}{f_0}(\theta_{i,1})$ is nondecreasing. Giving more market share to higher-cost types, as in a rigid-pricing scheme, allows lower-cost types to get higher utility in equilibrium, without inducing mimicry from higher-cost types. The term $\frac{F_0}{f_0}(\theta_{i,1})$ measures the magnitude of this effect. When it is higher for higher types, it is optimal to allocate more market share to higher types, in spite of the resulting efficiency losses.

This result is closely related to a finding of McAfee and McMillan (1992). They show that in a static cartel, when no monetary transfers are allowed and there is no future to provide incentives, the optimal mechanism for the cartel is a rigid-pricing strategy, if it can somehow be enforced. Athey, Bagwell, and Sanchirico (2004) provide a similar result in a repeated game model of a self-enforcing cartel, where costs are i.i.d. over time. In that context, non-stationary

¹³ Indeed, we have proved something stronger than stated in the proposition: rigid pricing is optimal among all Bayesian Nash equilibria that are immune only to deviations from on-path histories.

equilibria are possible, and the prospect of a reduced continuation value may thus deter a high-cost type from mimicking a low-cost type. Athey, Bagwell, and Sanchirico (2004) focus on equilibria that are strongly symmetric, however, and so any use of future continuation values to provide incentives is necessarily wasteful from the perspective of the cartel, as it must take the form of an industry-wide price war. The availability of an industry-wide price war would correspond to a limited form of transfers in a static model. In this context, they find that the optimal equilibrium does not use industry-wide price wars to provide incentives. Instead, the best rigid-pricing scheme is optimal.

In the context of collusion with perfectly persistent cost types, Proposition 2 shows that the best rigid-pricing scheme is also optimal. Importantly, we obtain this result while allowing for asymmetric (and non-stationary) equilibria; in particular, utility can be transferred from one firm to another through future play (as well as “wasted” through industry-wide price wars). However, because cost types do not change over time, if a low-type firm transfers utility to a high-type firm by giving up future market share, this transfer must be inefficient. This contrasts with the first-best result presented in Athey and Bagwell (2001), who study a model where cost types are discrete and i.i.d. over time, and asymmetric equilibria are allowed. It is then possible for a firm that has low costs today to make an efficient future transfer to a firm that has high costs today. The firms achieve an efficient future transfer by waiting for a date in which they have the same cost type.

Our result is also related to a well-known finding (see, e.g., Fudenberg and Tirole, 1991, pp. 299-301) in the dynamic contracting literature. This finding states that, in a dynamic principal-agent model with persistent types, if the principal has commitment power, then it is impossible for the principal to improve upon a mechanism that replicates the optimal static mechanism in each period. Under the assumptions of log-concavity and sufficient patience, we also find that the optimal dynamic solution entails the repeated play of the optimal static solution. Despite this apparent similarity, our analysis differs in several respects. There is no “principal” in our model who commits to allocations and transfers as a function of reports. Indeed we analyze schemes that are optimal from the ex ante perspective of “agents” (i.e., firms). While the principal-agent model allows monetary transfers to “settle up” transfers of utility immediately, we assume that direct sidepayments between firms are infeasible, and so state-contingent transfers of utility may need to be implemented through subsequent play, play that must respect incentive constraints given current information. In other words, in the PPBE of our dynamic game, agents cannot make commitments with respect to their future behavior. In Section 3.2.3, we analyze an example in which log concavity fails and find that, even if agents are very patient, the optimal mechanism with commitment cannot be achieved as a PPBE of the dynamic game.

Our result that the best rigid-pricing scheme is optimal builds from two key assumptions: inelastic demand and linear costs. We now discuss how the results change when these assump-

tions are relaxed. First, we note that rigid pricing has two important features. First, the allocation of market share is inefficient. Second, there is pooling: all types receive the same outcomes. We argue that some productive inefficiency is a robust feature of our model but that full pooling is not. Yet, pooling is not a knife-edge result: it strictly dominates the alternatives in our baseline model; and for cost functions and demand specifications sufficiently close to our baseline model, pooling remains optimal.

First, consider variations on the specification of demand. In the supplementary material, we briefly analyze models in which (i) products are perfect substitutes but demand is downward-sloping, (ii) products are imperfect substitutes, and (iii) firms compete in quantities. In case (i), we show that when demand is sufficiently inelastic, the best rigid-pricing scheme is still optimal. In general, forces in favor of at least partial pooling remain present in all of the alternative models.

Second, consider costs that are nonlinear in output. Variation in market share for a given cost type can then either increase or decrease expected costs, depending on whether the cost function is concave or convex. In addition, it is possible that a firm's expected market share can be non-monotonic in its cost type without violating incentive compatibility, since not only the mean but also the variance of the market share for a particular cost type affect payoffs for that type. Despite these complications, we show in the supplementary material that when costs are concave in output, a best rigid-pricing equilibrium offers strictly higher payoffs than any equilibrium in which the highest-cost type receives less than $1/I$ market share in each period (the market share under rigid pricing). However, we have not ruled out the (seemingly pathological) possibility that profits can be increased by resorting to an even more extreme form of inefficiency with some separation of types, where the highest-cost type gets even greater than average market share.

3.2.2. Pooling Is Approximately Optimal with Near-Perfect Persistence

Our next result shows that the best pooling equilibrium is approximately optimal in Model 1, when persistence gets high enough. This confirms that the result of Proposition 2 can be taken as a limiting case. As mentioned in the Introduction, incentives as well as the possible evolution of beliefs may be distinctly different in models where types are fixed with probability 1 and models where types can change with positive probability. However, in pooling equilibria, these distinctions are less important.

To establish our limiting result, we consider a sequence $\{\lambda_L^n, \lambda_H^n\}$ such that $\lim_{n \rightarrow \infty} \{\lambda_L^n, \lambda_H^n\} = \{1, 0\}$, and such that $F_0^n = F_0$, where F_0^n is the prior (stationary) distribution given $\{\lambda_L^n, \lambda_H^n\}$. Although our proof does not make use of the particular form of the sequence, for concreteness we define

$$\lambda_L^n = 1 - (1 - F_0(\underline{\theta}))/n, \quad \lambda_H^n = F_0(\underline{\theta})/n. \quad (3.7)$$

Corollary 1. Consider Model 1. Fix $\delta > \delta_c$. Let the sequence $\{\lambda_L^n, \lambda_H^n\}$ be defined in (3.7). If $(r - \bar{\theta})/(\bar{\theta} - \underline{\theta}) > 1 - F_0(\underline{\theta})$, then the limit as $n \rightarrow \infty$ of the payoff achieved by a given firm in the optimal ex ante PPBE approaches $\frac{1}{1-\delta}(r - \mathbb{E}_{\theta_{i,1}}[\theta_{i,1}|F_0])$, which is also the payoff attained by the best rigid-pricing scheme given F_0 and perfectly persistent types.

Proof: See Appendix.

This result can be understood in two parts. First, the parameter restriction implies that pooling is optimal when types are perfectly persistent, so that the analog of Proposition 2 holds for a two-type model. Second, we show that payoffs are upper hemi-continuous in the persistence parameter.

3.2.3. Partial Pooling May Be Optimal when Log-Concavity Fails

So far, we have focused on the case where the prior distribution F_0 is log-concave or r is large enough. We now consider how the analysis would change if these conditions fail. In that event, it is direct to show that in a static mechanism design problem, or in a repeated game with cost types that are i.i.d. over time, the optimal scheme typically entails partial pooling. However, a number of subtleties arise when analyzing the repeated game with perfectly persistent types.

To focus our discussion, consider a prior distribution with the following properties: there exists $z \in (\underline{\theta}, \bar{\theta})$ such that $F_0(\theta_{i,1})/f_0(\theta_{i,1})$ is strictly increasing on $[\underline{\theta}, z)$ and on $(z, \bar{\theta}]$, but $F_0(\theta_{i,1})/f_0(\theta_{i,1})$ decreases discretely at z . A specific example follows (where $z = 2/3$):

Example 1: The distribution is piecewise uniform and described by the following density function: $f_0(\theta_{i,1}) = \lambda_1$ for $\theta_{i,1} \in [0, 1/3)$, $f_0(\theta_{i,1}) = \lambda_2$ for $\theta_{i,1} \in [1/3, 2/3)$, and $f_0(\theta_{i,1}) = 3 - \lambda_1 - \lambda_2$ for $\theta_{i,1} \in [2/3, 1]$, and $f_0(\theta_{i,1}) = 0$ elsewhere, where $0 < \lambda_1 + \lambda_2 < 3$. If λ_2 is small enough, $\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}$ will jump up at $1/3$ and then jump down at $2/3$.

As a benchmark, let us begin by considering the optimal collusive scheme subject only to the on-schedule incentive constraints in the first period. (Recall that we took this approach in the proof of Proposition 2.) We can use similar arguments to Proposition 2 to show that the scheme which maximizes $\mathbb{E}_{\theta_{i,1}}[U_i(\theta_{i,1}, \theta_{i,1})|F_0]$ subject to the first-period on-schedule incentive constraints is characterized by pooling *within* each of the intervals $[\underline{\theta}, z)$ and $[z, \bar{\theta}]$, since F_0 is log-concave within those intervals. Consider now the comparison between two possible schemes: one with pooling throughout $[\underline{\theta}, \bar{\theta}]$ (specifically, the best rigid-pricing scheme); and one with pooling within each interval $[\underline{\theta}, z)$ and $[z, \bar{\theta}]$, but separation between the two steps.

The following class of schemes has separation between two steps. How much the market shares differ for the two regions of types is a parameter of the scheme.

Two-Step Scheme: A set of strategies where, on and off of the equilibrium path, announcements are uninformative. We may then suppress notation and represent the pricing strategy

within a period as $\rho_{i,t}(\theta_{i,1})$. For $\theta_{i,1} \in [\underline{\theta}, z)$, $\rho_{i,1}(\theta_{i,1}) = p_L$ and $\rho_{i,t}(\theta_{i,1}) = r$ for all $t > 1$, while for $\theta_{i,1} \in [z, \bar{\theta}]$, $\rho_{i,t}(\theta_{i,1}) = r$ for all t . In period 1, each player makes the (nonbinding) quantity restriction $q_{i,t} = 1$, so that the market share allocation is $1/I$ if all choose the same price, and otherwise the low-price firms serve the market. If N firms tied for the lowest price in period 1, then in all subsequent periods each of these N firms sets $q_{i,t} = (1 - (I - N)k)/N$, while all other firms set $q_{i,t} = k$, where $k \in [0, 1/I)$. Any off-schedule deviation results in a switch to the worst carrot-stick scheme analyzed in Proposition 1. The price p_L is determined by the on-schedule incentive constraint in the first period, and so it will depend on k ; for simplicity, we allow that p_L may be negative.

Simple Two-Step Scheme: A Two-Step Scheme with $k = 0$.

We show below that the simple two-step scheme does *not* satisfy all relevant off-schedule incentive constraints, but we begin by analyzing it as a benchmark. To compare the best rigid-pricing scheme with the simple two-step scheme, let y be the expected discounted market share for types on $[\underline{\theta}, z)$. Using (3.4), ex ante expected profits as a function of y are

$$\begin{aligned} \mathbb{E}_{\theta_{i,1}} [U_i(\theta_{i,1}, \theta_{i,1}) | F_0] &= (r - \bar{\theta}) \frac{1}{1 - F_0(z)} \left(\frac{1}{I(1 - \delta)} - F_0(z)y \right) \\ &+ y \int_{\underline{\theta}}^z F_0(\theta_{i,1}) d\theta_{i,1} + \frac{1}{1 - F_0(z)} \left(\frac{1}{I(1 - \delta)} - F_0(z)y \right) \int_z^{\bar{\theta}} F_0(\theta_{i,1}) d\theta_{i,1}. \end{aligned}$$

This expression is increasing in y if

$$\frac{\int_{\underline{\theta}}^z F_0(\theta_{i,1}) d\theta_{i,1}}{F_0(z)} > r - \bar{\theta} + \int_{\underline{\theta}}^{\bar{\theta}} F_0(\theta_{i,1}) d\theta_{i,1}. \quad (3.8)$$

Then, when (3.8) holds, profits will be strictly higher with the simple two-step scheme, which gives greater market share to types on $[\underline{\theta}, z)$.

Example 1, cont.: Condition (3.8) holds if $\lambda_1, \lambda_2 \in (0, 3)$, and

$$\lambda_1 > (\lambda_1 + \lambda_2) \frac{(9(r - 1) + \lambda_1 + \lambda_2)}{3 - (\lambda_1 + \lambda_2)}.$$

This requires that $4/3 > r$ and $\lambda_1 + \lambda_2 < (3 - 9(r - 1))/2$. If $r = 1.2$ and $\lambda_1 + \lambda_2 = 1/5$, (3.8) holds if $1/5 > \lambda_1 > 1/7$.

So, when (3.8) holds, the simple two-step scheme would dominate the best rigid-pricing scheme, if the simple two-step scheme were a PPBE. But it is not. To see why, suppose $\theta_{1,1} < z$ and $\bar{\theta} > \theta_{2,1} > z$. Starting in the second period, firm 2 makes 0 profit under the simple two-step scheme. As firm 2 makes positive profit under the worst carrot-stick equilibrium, it would make an off-schedule deviation (e.g., undercutting r in period 2) and trigger the punishment.

We can salvage an equilibrium with partial pooling, however, by specifying that firms 1 and 2 share the market, unequally, after period 1. We define $M^{cs}(\theta_{i,t})$ as the minimum market

share that a firm must receive to deter an off-schedule deviation in periods $t > 1$ when firms price at r . This must satisfy

$$\frac{M^{cs}(\theta_{i,t})}{(1-\delta)}(r - \theta_{i,t}) = r - \theta_{i,t} + \delta v^{cs}(\theta_{i,t}) = r - \theta_{i,t} + \frac{\delta}{1-\delta} \frac{1}{I}(\bar{\theta} - \theta_{i,t}),$$

where the second equality follows since, in Model 2, $v^{cs}(\bar{\theta}) = 0$ and $\gamma_{\bar{\theta}} - \gamma_{\theta_{i,t}} = \bar{\theta} - \theta_{i,t}$. For all $\theta_{i,t} \in [\underline{\theta}, \bar{\theta}]$, it is straightforward to verify that $M^{cs}(\theta_{i,t})$ is decreasing, and that $M^{cs}(\theta_{i,t}) \leq 1/I$ if also $\delta \in [\delta_c, 1)$. Recall that the worst carrot-stick equilibrium can be supported over this discount-factor interval. Now, consider the following modification of the simple two-step scheme.

Market-Sharing Two-Step Scheme: A two-step scheme with $k = M^{cs}(z)$.

The price p_L is determined by the on-schedule incentive constraint in the first period, and so it will take on a different value than in the simple two-step scheme.

By construction, the market-sharing two-step scheme satisfies on-schedule constraints in period 1. After period 1, strategies specify behavior as a function of first-period observed behavior, and there are no additional on-schedule constraints. Since $M^{cs}(\theta_{i,1}) \leq 1/I$, all firms receive market share greater than or equal to $M^{cs}(\theta_{i,1})$, and so off-schedule constraints are satisfied for all firms in periods $t > 1$. It remains only to consider the off-schedule constraints in period 1.

Proposition 3. *Consider Model 2.*

- (i) *For all $\delta \in [\delta_c, 1)$, there exists a market-sharing two-step equilibrium.*
- (ii) *If (3.8) holds, the market-sharing two-step equilibrium yields greater ex ante expected profits than the best rigid-pricing equilibrium.*

Proof: See Appendix. ■

The first part of this result shows that the market-sharing two-step equilibrium exists whenever the best rigid-pricing equilibrium exists ($\delta \geq \delta_c$). To establish that no additional patience is required to support the market-sharing two-step scheme, we show that it is the lowest-cost type who is most tempted to deviate from the best rigid-pricing scheme. With the two-step scheme, low-cost types expect to receive greater market share in the first period, and so gain less from undercutting either their assigned price (p_L) or the price assigned to higher-cost types (r).¹⁴ The second part of Proposition 3 shows that it is possible to improve upon the best-rigid pricing scheme, when r is not too large and F_0 is not log-concave. We have not, however, proved that the market-sharing two-step equilibrium is optimal. A scheme with multiple steps would achieve further separation of types in the first period and thereby allocate lower market share

¹⁴ Thus, if a punishment were available that required less patience to implement than the worst carrot-stick equilibrium, it would be possible to support the market-sharing two-step scheme with a lower range of discount factors than that required to support the best rigid-pricing scheme. As defined, however, the market-sharing two-step scheme relies on the worst carrot-stick equilibrium, and we require $\delta \geq \delta_c$ for that equilibrium to exist.

to the higher cost types in future periods. The effect of additional steps on ex ante expected profit, however, is sensitive to the precise way in which F_0 fails to be log-concave.

Before proceeding, we pause to interpret the market-sharing two-step equilibrium. In this equilibrium, there is an initial signalling phase, where some firms set low prices (and thus receive a higher market share than the remaining firms). In all subsequent periods, firms that initially cut prices receive more market share than those who did not, but all firms have some market share. Notice that empirically, this behavior might appear to entail an initial “price war” followed by a “collusive phase,” where the market shares during the price war phase determine the market shares of the “collusive phase.” Roos (2003) argues that a price war in the lysine industry started by Archer Daniels Midland in the early 1990s had these characteristics.

4. Equilibria with Some Productive Efficiency

In this section, we study equilibria with at least some productive efficiency. First, we establish that when firms are sufficiently patient, there exist simple equilibria with partial productive efficiency and pricing at r . However, as persistence of types grows, the degree of productive efficiency approaches zero, even if firms are very patient. Second, we construct more sophisticated equilibria with productive efficiency and pricing at r on the equilibrium path. These first-best equilibria exist when firms are patient relative to the persistence of cost types.

4.1. Partial Productive Efficiency with Arbitrary Persistence

Consider Model 1. We show that so long as δ is greater than δ_c , the critical discount factor for supporting the best rigid-pricing equilibrium, it is possible to improve upon the best rigid-pricing equilibrium. To accomplish this, we introduce a simple *odd-even scheme*, whereby firms report their types in odd periods and implement partial productive efficiency. In even periods, the firms that received higher market share in the prior period must reduce their market share. The even-period market shares are determined so as to provide incentives for truthful revelation in the odd periods.

Odd-Even Scheme: A set of strategies described as follows. Consider two functions, the odd-period market share allocation function $\kappa_i^o : \mathcal{A} \rightarrow [0, 1]$, and the even-period market share allocation function $\kappa_i^e : \mathcal{A} \rightarrow [0, 1]$, where $\sum_i \kappa_i^o(\mathbf{a}_t) = 1$ and $\sum_i \kappa_i^e(\mathbf{a}_t) = 1$ for all $\mathbf{a}_t \in \mathcal{A}$. If t is an odd period, firms announce their cost types, so that $\mathbf{a}_t = \boldsymbol{\theta}_t$. Then, the firms choose $p_{i,t} = r$, but they share the market unequally, as determined by $q_{i,t} = \kappa_i^o(\mathbf{a}_t)$. In period $t + 1$, an even period, announcements are uninformative. All firms choose $p_{i,t+1} = r$, and $q_{i,t+1} = \kappa_i^e(\mathbf{a}_t)$. In all periods, any off-schedule deviation is punished by switching to a worst carrot-stick scheme.

In the Appendix proof of Proposition 4, we formally define a value function for an odd-even equilibrium, as well as the on- and off-schedule constraints. Here, we provide some intuition

about the role of the constraints and when they will bind, focusing for simplicity on the case where $I = 2$. In an on-schedule deviation, a firm that just observed its cost type $\theta_{i,t}$ in odd period t mimics type $\hat{\theta}_{i,t}$ throughout periods t and $t+1$ (since the strategies call for uninformative announcements in period $t + 1$). If low-cost types receive more market share in odd periods ($\kappa_i^o(L, H) > 1/2$), the binding on-schedule incentive constraint will be the constraint that deters a high-cost type from mimicking a low-cost type in an odd period. This deviation will be deterred if $\kappa_i^e(L, H)$ is sufficiently low and $\kappa_i^e(H, L)$ is sufficiently high. The most tempting off-schedule deviation has a firm undercut the collusive price of r . Such a deviation may be tempting in either an odd period or an even period; in an even period $t + 1$, such a deviation yields a short-term gain of $(1 - \kappa_i^e(\theta_{i,t}, \theta_{j,t}))(r - \theta_{i,t+1})$. The even-period off-schedule constraint is especially likely to bind if $(\theta_{i,t}, \theta_{j,t}) = (L, H)$ and $\theta_{i,t+1} = L$, since (in order to respect the on-schedule constraints when $\kappa_i^o(L, H) > 1/2$) we will have $\kappa_i^e(L, H) < 1/2$.

For $\delta > \delta_c$, we now establish that an odd-even equilibrium can be constructed that achieves a strictly higher ex ante payoff for firms than does the best rigid-pricing equilibrium. An attractive feature of the constructed scheme is that the on-schedule incentive constraints are satisfied pointwise (so that announcing truthfully is *ex post* incentive compatible); thus, the specific manner in which beliefs are formed about the evolution of rivals' costs is unimportant. The proof builds from two insights. First, starting from the best rigid-pricing equilibrium, it is possible to maintain on-schedule incentive compatibility and raise the ex ante payoff for firms, by making an intertemporal exchange in market shares under which a firm that reveals itself to have low (high) costs in the odd period experiences an increase (decrease) in its odd-period market share and a decrease (increase) in its market share in the subsequent even period. Intuitively, market share is then redistributed from known high-cost firms to a known low-cost firm in the odd period, while in the subsequent period market share is redistributed from a firm that, under imperfect persistence, probably has low costs to firms that probably have high costs. Second, when the intertemporal exchange in market shares is small, the off-schedule constraints are sure to be satisfied provided that $\delta > \delta_c$, so that these constraints are slack under the best rigid-pricing equilibrium. This follows since the odd-even value function is continuous in the market shares and the odd-even equilibrium is identical to the best rigid-pricing equilibrium in the limit, when $\kappa_i^e(\mathbf{a}_t) = \kappa_i^o(\mathbf{a}_t) = 1/I$ for all \mathbf{a}_t . For some parameter values (e.g. when persistence is low), the improvement in the future value of cooperation attained by increasing productive efficiency relaxes the off-schedule constraints enough so that large intertemporal exchanges of market share become possible even at $\delta = \delta_c$.

Proposition 4. *Consider Model 1. If $\delta > \delta_c$, then there exists an odd-even equilibrium, where this equilibrium achieves strictly higher ex ante payoffs than the best rigid-pricing equilibrium.*

Proof: See Appendix.

We next consider optimal equilibria within the odd-even class. For given parameter val-

ues, the market share choices that maximize ex ante expected profits in the class of odd-even equilibria can be computed using standard linear programming techniques.

Example 2: *Suppose that $I = 2$, $\lambda_L = .95$, $\lambda_H = .025$, $L = 1$, $H = 2$, $r = 2.1$, and $\delta = \delta_c + .001 = .918$. Then, the odd-even equilibrium that maximizes ex ante expected profits uses the following values: $\kappa_i^o(L, L) = \kappa_i^o(H, H) = .5$, $\kappa_1^o(L, H) = 1 - \kappa_1^o(H, L) = .59$, $\kappa_i^e(L, L) = \kappa_i^e(H, H) = .5$, and $\kappa_1^e(L, H) = 1 - \kappa_1^e(H, L) = .42$. The downward on-schedule constraint (whereby type H is tempted to mimic L) is binding in all odd periods t . If t is odd and $\theta_t = (L, H)$, firm 1's market share is .59 in period t and .42 in period $t+1$. The off-schedule constraint limits productive efficiency: to increase $\kappa_1^o(L, H)$, a decrease in $\kappa_1^e(L, H)$ is required to respect on-schedule constraints, but the off-schedule constraint for firm 1's low type in even periods prevents a large decrease.*

Even though market shares fluctuate in the odd-even equilibrium of Example 2, very little efficiency is gained. Although production is fairly efficient in each odd period, production is fairly inefficient in even periods. With persistent types, a realization of (L, H) is often followed by another realization of (L, H) . In the even period, the low-cost firm then serves *less* than half of the market. Indeed, using the prior distribution of costs, the average cost over an average two-period cycle is 1.663, just slightly below the average cost in a pooling equilibrium of 1.667, and far above the prior expected value of the minimum of the two firms' costs, 1.444.

Numerical calculations confirm that for more moderate persistence than Example 2, off-schedule constraints are typically slack at $\delta = \delta_c$, and there is productive efficiency in odd periods. However, even periods have substantial inefficiency.

The odd-even scheme has very limited history-dependence on the equilibrium path. That limitation may be quite important. To gain some intuition, let us modify the odd-even scheme to allow κ_1^e to depend on \mathbf{a}_{t-2} as well as \mathbf{a}_t .

Example 3: *Consider the parameter values of Example 2. Then, the (modified) odd-even equilibrium that maximizes ex ante expected profits has $\kappa_i^o(L, L) = \kappa_i^o(H, H) = .5$, $\kappa_1^o(L, H) = 1 - \kappa_1^o(H, L) = 1$. That is, there is productive efficiency in odd periods. The values of $\kappa_i^e(\theta_1, \theta_2, \mathbf{a}_{t-2})$ in even periods are given in the following table:*

Table 1: Market Shares in Even Periods

	$\kappa_i^e(L, L, \mathbf{a}_{t-2})$	$\kappa_i^e(L, H, \mathbf{a}_{t-2})$	$\kappa_i^e(H, H, \mathbf{a}_{t-2})$	$\kappa_i^e(H, L, \mathbf{a}_{t-2})$
\mathbf{a}_{t-2}				
(L, L)	.5	.501	.5	.499
(L, H)	.205	.117	0	0
(H, L)	.795	1	1	.882
(H, H)	.5	.622	.5	.378

The market shares in even periods are set to maximize productive efficiency while still respecting

the on-schedule incentive constraints, which bind in odd periods. Suppose that $t - 2$ is odd, and that $\mathbf{a}_{t-2} = (L, H)$, in which case firm 1 serves the entire market in period $t - 2$. Firm 1 must receive reduced market share in a period following $t - 2$ in order to respect the on-schedule incentive constraint. However, it is likely that firm 1 will also be low cost in period $t - 1$. In order to keep market share for firm 1 as high as possible in period $t - 1$, it is useful to put off some of the “punishment” of firm 1 to a future period. We see from Table 1 that $\kappa_i^e(L, L, L, H)$ is equal to .205. Given that $\boldsymbol{\theta}_t = (L, L)$, there is no expected efficiency loss in the even period $t + 1$ from reducing market share for firm 1. But, this low market share helps relax the on-schedule incentive constraint for firm 1 in period $t - 2$.

In the equilibrium of Example 3, expected cost is 1.62, 2.6% lower than in Example 2. However, the low productive efficiency in even periods still limits the productive efficiency of the mechanism. Even when δ approaches 1, the average cost never gets lower than 1.588. The example illustrates how greater history-dependence can help: incentives for firm 1 to admit high cost in period $t - 2$ are provided by granting additional market share to firm 1 in period $t + 1$, in the event that both firms have the same cost in period t .

More generally, numerical calculations confirm the following regularities. For any level of persistence, there is an upper bound on the efficiency gain from the odd-even scheme. As the persistence of cost types grows, this upper bound approaches zero.

4.2. First-Best with Moderate Persistence

In this subsection, we focus on Model 1 with two firms ($I = 2$) and show that it is possible to attain first-best collusion for some parameter values (values where patience is high relative to persistence). The collusive scheme that delivers these payoffs is referred to as a *first-best scheme* and is a generalization of that proposed by Athey and Bagwell (2001). In each period, firms announce their cost types. If one firm has high cost and the other has low cost, the low-cost firm serves the market. If both have the same cost, the firms split the market, typically unevenly, where the splits are a function of past reports and are constructed to favor firms who have announced high costs in the past.

First-Best Scheme: A set of strategies in which, along the equilibrium path, firms announce their cost types and price at r in each period, so that $\mathbf{a}_t = \boldsymbol{\theta}_t$ and $p_{i,t} = r$. Further, quantity restrictions satisfy productive efficiency: if $a_{i,t} = L$ and $a_{j,t} = H$, then $q_{i,t} = 1$ and $q_{j,t} = 0$. In all periods, any off-schedule deviation is punished by switching to the worst carrot-stick scheme. It remains to specify \mathbf{q}_t when $a_{i,t} = a_{j,t}$ as a function of history. On the equilibrium path, play depends on \mathbf{a}_{t-1} as well as which of the two firms is “favored.” We represent this using “states,” $\omega_1(\boldsymbol{\theta}_{t-1})$ and $\omega_2(\boldsymbol{\theta}_{t-1})$, where firm 1 prefers $\omega_1(\boldsymbol{\theta}_{t-1})$ to $\omega_2(\boldsymbol{\theta}_{t-1})$. Let the set of states that may be reached on the equilibrium path be denoted

$$\Omega^e = \{\omega_1(L, L), \omega_2(L, L), \omega_1(L, H), \omega_2(L, H), \omega_1(H, L), \omega_2(H, L), \omega_1(H, H), \omega_2(H, H)\}.$$

For each state $\omega_j(\boldsymbol{\theta}_{t-1})$, we define period strategies and transitions to subsequent states as a function of observables. Formally, let $\Theta \equiv \Theta_i^2$ and

$$\tilde{q}_1(\cdot; j, \boldsymbol{\theta}_{t-1}) : \Theta \rightarrow [0, 1] \text{ and } \tilde{g}(\cdot; j, \boldsymbol{\theta}_{t-1}) : \Theta \rightarrow [0, 1], \quad (4.1)$$

so in state $\omega_j(\boldsymbol{\theta}_{t-1})$, following announcements $\boldsymbol{\theta}_t$, firm 1 gets market share $\tilde{q}_1(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$, firm 2 gets market share $1 - \tilde{q}_1(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$, and they subsequently move to state $\omega_1(\boldsymbol{\theta}_t)$ with probability $\tilde{g}(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$ and to state $\omega_2(\boldsymbol{\theta}_t)$ with probability $1 - \tilde{g}(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$.

For concreteness, consider the following example.

Example 4: Let $r = 2.1, H = 2, L = 1, \lambda_L = .7, \lambda_H = .5$ and $\delta = .92 > \delta_c = .917$. Using the representation just described, the following table can be used to construct strategies for a first-best scheme.

Table 2: First-Best Scheme Description

	$\tilde{q}_1(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$				$\tilde{g}(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1})$			
Today's cost: $\boldsymbol{\theta}_t$	(L, L)	(L, H)	(H, L)	(H, H)	(L, L)	(L, H)	(H, L)	(H, H)
State: $(j, \boldsymbol{\theta}_{t-1})$								
(1, L, L)	.5	1	0	1	0	0	1	.5
(2, L, L)	.5	1	0	0	1	0	1	.5
(1, H, L)	1	1	0	.5	.5	0	1	.5
(2, L, H)	0	1	0	.5	.5	0	1	.5
(1, H, H)	.5	1	0	1	.5	0	1	.5
(2, H, H)	.5	1	0	0	.5	0	1	.5

The entries in the table incorporate productive efficiency: in every state, $\tilde{q}_1(L, H; j, \boldsymbol{\theta}_{t-1}) = 1 - \tilde{q}_1(H, L; j, \boldsymbol{\theta}_{t-1}) = 1$. The entry $\tilde{g}(\boldsymbol{\theta}_t; j, \boldsymbol{\theta}_{t-1}) = 0$ specifies that following $\boldsymbol{\theta}_t = (L, H)$, the firms transition to state $\omega_2(L, H)$ (the state that favors firm 2) with probability 1. Similarly, following $\boldsymbol{\theta}_t = (H, L)$, the firms transition to the state that favors firm 1, $\omega_1(H, L)$, with probability 1. The transitions following “ties” ($\boldsymbol{\theta}_t = (L, L)$ or $\boldsymbol{\theta}_t = (H, H)$) favor each firm equally in each state. In some states, market shares in the case of “ties” are unequal: the favored firm receives more market share. Note the states $\omega_1(L, H)$ and $\omega_2(H, L)$ are not used.

An interesting feature of this example is that communication and binding quantity constraints are not necessary to implement the first-best scheme. The market share divisions can be accomplished by judicious choices of prices. For example, in state (1, L, L), firm 1’s high-cost type prices at $r - \varepsilon$ and its low-cost type prices at $r - 2\varepsilon$, while firm 2’s high-cost type prices at r and its low-cost type prices at $r - 2\varepsilon$. However, in general, the set of parameters for which first-best equilibria exist is smaller when communication and quantity constraints are not available, since eliminating these instruments places additional constraints on the scheme.

How does one verify that a first-best scheme is in fact a PPBE? In a repeated game, the dynamic programming tools of Abreu, Pearce, and Stacchetti (1986, 1990) can be applied to construct “self-generating” sets of equilibrium values. In the supplementary material for this article, we describe this in detail. Here, we show that it is possible to verify whether a first-best scheme is a PPBE by solving a system of linear equations that define a firm’s value in each state as a function of its cost type, and then checking appropriately defined on-schedule and off-schedule incentive constraints.

The first step is to construct extensive form strategies that specify play following any public history. Given an initial state, we can determine the state of the game following any history (where the history includes the outcome of public randomization) simply by tracking through the transitions specified by the strategies. If following a particular history the state is determined to be $\omega_j(\boldsymbol{\theta}_{t-1}) \in \Omega^e$, period strategies are given by

$$(\alpha_{i,t}(\theta_{i,t}), \rho_{i,t}(\mathbf{a}_t, \theta_{i,t}), \psi_{i,t}(\mathbf{a}_t, \theta_{i,t})) = (\theta_{i,t}, r, \tilde{q}_i(\mathbf{a}_t; j, \boldsymbol{\theta}_{t-1})) \equiv \tilde{s}_i(\mathbf{a}_t, \theta_{i,t}; j, \boldsymbol{\theta}_{t-1}). \quad (4.2)$$

Following every history up to time t , if announcements in period t are given by \mathbf{a}_t , beliefs in period $t + 1$ are given by $\mathbf{F}(\cdot; \mathbf{a}_t)$.¹⁵

Let \mathcal{V} be the set of functions $\mathbf{v} = (v_1, \dots, v_I)$ such that $v_i : \Theta_i \rightarrow \mathbb{R}$. This is the set of possible “type-contingent payoff functions,” specifying payoffs to each firm as a function of its true type in a particular period. If \mathbf{V} maps from \mathcal{Z} to \mathcal{V} , we can interpret V_i as an analog of a continuation value function, but one that maps from today’s public outcomes to a vector of future values, one for each type that firm i could be in the next period. Given such a V_i , and given beliefs about opponents $\boldsymbol{\mu}_t$ and posited strategies \mathbf{s}_t , the following represents firm i ’s expected payoffs in period t , before announcements are made, when firm i has type $\theta_{i,t}$ and mimics type $\hat{\theta}_{i,t}$:

$$\begin{aligned} \bar{u}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}, V_i) &= \bar{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}) \\ &+ \delta \mathbb{E}_{\theta_{i,t+1}, \theta_{-i,t}} \left[V_i \left(\mathbf{s}_t(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), (\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t})) \right) (\theta_{i,t+1}) \mid \boldsymbol{\mu}_{-i,t}, \theta_{i,t} \right]. \end{aligned}$$

The motivation for considering type-contingent payoff functions becomes clear when we see that future values depend on today’s “reported type” $\hat{\theta}_{i,t}$ through the induced public actions that serve to select a continuation equilibrium, while firm i constructs its beliefs about the future distribution of $\theta_{i,t+1}$ based on its true type.

We now define two functions, $\tilde{\mathbf{v}} \in \mathcal{V}$ and $\tilde{\mathbf{V}}$ mapping from \mathcal{Z} to \mathcal{V} , through the following

¹⁵ If a player makes an announcement that is not in Θ_i , the players proceed to the carrot-stick equilibrium where beliefs are forever after irrelevant, and so without loss of generality we can let the beliefs be as if the deviant player announced $\theta_{i,t} = L$. Given this, we simplify the discussion by focusing on the case where only announcements in Θ_i are possible for player i .

system of equations for each firm i .

$$\begin{aligned}
& \text{For all } (j, \boldsymbol{\theta}_{t-1}) \in \{1, 2\} \times \Theta : & (4.3) \\
\tilde{V}_i(\mathbf{a}_t, \mathbf{p}_t, \mathbf{q}_t; j, \boldsymbol{\theta}_{t-1}) &= \tilde{g}(\mathbf{a}_t; j, \boldsymbol{\theta}_{t-1}) \tilde{v}_i(\cdot; 1, \mathbf{a}_t) + (1 - \tilde{g}(\mathbf{a}_t; j, \boldsymbol{\theta}_{t-1})) \tilde{v}_i(\cdot; 2, \mathbf{a}_t) \\
& \text{whenever } \mathbf{a}_t \in \Theta, \mathbf{p}_t = (r, r), \text{ and } \mathbf{q}_t = \tilde{\mathbf{q}}_t(\hat{\boldsymbol{\theta}}_t; j, \boldsymbol{\theta}_{t-1}), \\
& \text{and } \tilde{V}_i(\mathbf{a}_t, \mathbf{p}_t, \mathbf{q}_t; j, \boldsymbol{\theta}_{t-1}) = v^{cs} \text{ otherwise, where} \\
& \tilde{v}_i(\theta_{i,t}; j, \boldsymbol{\theta}_{t-1}) = \bar{u}_i(\theta_{i,t}, \theta_{i,t}, \mathbf{s}, \boldsymbol{\mu}_{-i}, V_i) \\
& \text{with } \mathbf{s} = \tilde{\mathbf{s}}(\cdot; j, \boldsymbol{\theta}_{t-1}), \boldsymbol{\mu}_{-i} = F(\cdot; \boldsymbol{\theta}_{-i,t-1}), V_i = \tilde{V}_i(\cdot; j, \boldsymbol{\theta}_{t-1}).
\end{aligned}$$

Standard arguments (e.g. Abreu, Pearce, and Stacchetti (1986)) can be used to verify that, if we consider a particular $\omega_j(\boldsymbol{\theta}) \in \Omega^e$, if initial beliefs are $\mathbf{F}(\cdot; \boldsymbol{\theta})$ and extensive-form strategies and beliefs are specified using the construction outlined above, $\tilde{v}_i(\theta'_i; j, \boldsymbol{\theta})$ is indeed the expected discounted future payoff that player i expects after observing type θ'_i .

It remains to specify incentive constraints that ensure that a set of strategies are best responses given beliefs following every history. We have already verified that starting from any initial beliefs, the carrot-stick strategy profiles form a PPBE for discount factors above a critical level, so we focus on $\omega_j(\boldsymbol{\theta}) \in \Omega^e$. It is straightforward to verify that some types of deviations are dominated. For example, announcing a type outside of Θ_i is dominated by a deviation with an announcement in Θ_i followed by a price of $r - \varepsilon$. This deviation also dominates deviations involving quantity restrictions or lower prices. The following constraint guarantees that all off-schedule deviations are deterred for firm i in state $\omega_j(\boldsymbol{\theta}_{t-1}) \in \Omega^e$:

$$\tilde{v}_i(\theta_{i,t}; j, \boldsymbol{\theta}_{t-1}) \geq r - \theta_{i,t} + \delta v^{cs}. \quad (4.4)$$

On-schedule deviations are deterred for firm i in state $\omega_j(\boldsymbol{\theta}_{t-1}) \in \Omega^e$ if

$$\forall \theta_{i,t}, \hat{\theta}_{i,t}, \tilde{v}_i(\theta_{i,t}; j, \boldsymbol{\theta}_{t-1}) \geq \bar{u}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \tilde{\mathbf{s}}(\cdot; j, \boldsymbol{\theta}_{t-1}), F(\cdot; \boldsymbol{\theta}_{-i,t-1}), \tilde{V}_i(\cdot; j, \boldsymbol{\theta}_{t-1})). \quad (4.5)$$

Note that initial beliefs at the start of the game may not be among those in the first-best scheme ($\{\boldsymbol{\mu} : \boldsymbol{\mu} = \mathbf{F}(\cdot; \boldsymbol{\theta}_t) \text{ for some } \boldsymbol{\theta}_t\}$). The scheme can be augmented to allow for initial beliefs, but we do not introduce notation for that here.

Summarizing this analysis, we have

Proposition 5. *Fix $I = 2$ and consider the two-type model with imperfect persistence, with primitives δ, r, L, H , and \mathbf{F} , with $\delta \geq \delta_c$. Fix the specification of a first-best scheme \tilde{g} and \tilde{q}_1 , and define the corresponding $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{V}}$ as in (4.2) and (4.3). If for each $(i, j, \boldsymbol{\theta}_{t-1}) \in \{1, 2\}^2 \times \Theta$, the on-schedule and off-schedule constraints, (4.5) and (4.4), hold, then this first-best scheme is a PPBE set that yields first-best profits in every period.*

Proposition 5 can be applied to Example 4. It identifies the incentive constraints that can be checked to guarantee that a first-best scheme is a PPBE. However, the result does not give any indication of whether, for a given set of primitives, there will exist a specification for (4.1) that can satisfy the incentive constraints. It turns out that it is cumbersome to give necessary and sufficient conditions for specifications such that the first-best equilibrium exists, since that exercise involves searching over a large-dimensional space. However, we have constructed examples with a wide range of specifications; a general theme, not surprisingly, is that patience needs to be large relative to persistence.

5. Productive Efficiency with Perfect Persistence

As discussed in the last section, when patience is high relative to persistence, it may be possible to construct a first-best equilibrium. We showed in Section 3 that when persistence is extreme, pooling is optimal or approximately optimal so long as the distribution of cost types is log-concave. In this section, we ask whether equilibria with productive efficiency exist at all when cost types are perfectly persistent. We show that even if we allow for low prices, productive efficiency requires that patience be *low* enough. In addition, the equilibrium we construct requires a severe form of punishment in the case of off-schedule deviations.

5.1. Separating Equilibria with Productive Efficiency

We analyze here Model 2. Any proposed equilibrium with productive efficiency in each period must be immune to deviations whereby one type mimics another type in every period. This incentive compatibility requirement in turn implies

$$U_i(\theta_{i,1}, \theta_{i,1}) = U_i(\bar{\theta}, \bar{\theta}) + \int_{\tilde{\theta}_i=\theta_{i,1}}^{\bar{\theta}} \check{M}_i(\tilde{\theta}_i) d\tilde{\theta}_i = \frac{1}{1-\delta} \int_{\tilde{\theta}_i=\theta_{i,1}}^{\bar{\theta}} (1 - F_0(\tilde{\theta}_i))^{I-1} d\tilde{\theta}_i, \quad (5.1)$$

following the logic and using the notation from the proof of Proposition 2. That is, each player must expect per-period profits equal to those of the static Nash equilibrium. Thus, an equilibrium with productive efficiency would not be very profitable. Indeed, it can be shown that if $(\bar{\theta} - \underline{\theta})/I \geq \int_{\underline{\theta}}^{\bar{\theta}} (1 - F(\tilde{\theta}))^{I-1} d\tilde{\theta}$, then all types other than $\bar{\theta}$ would earn strictly less in an equilibrium with productive efficiency than in the worst carrot-stick equilibrium.¹⁶

However, it remains to analyze whether an equilibrium exists that delivers productive efficiency. The static Nash equilibrium is no longer an equilibrium in the dynamic game, since beliefs change after first-period play. Focusing on the case where $I = 2$, we now consider a *productive efficiency scheme*, in which firms do not communicate but rather set prices in a way that ensures productive efficiency in each period. Clearly, in the first period of such a scheme,

¹⁶ Simply compare the profits of the two equilibria for each type, and note that the difference is convex. The stated inequality holds (with equality) for the uniform distribution, for example.

the firms can achieve productive efficiency only if they use a pricing strategy that is strictly increasing in costs and symmetric across firms.

Productive Efficiency Scheme: A set of strategies such that, in each period, announcements are uninformative and market share proposals are not binding ($q_{i,t} \geq 1$). The first-period pricing strategy of firm i is denoted $\rho_{i,1}(\theta_{i,1})$, and is strictly increasing and symmetric across firms. Each firm infers the other firm’s cost once first-period prices are observed. Let θ_w and θ_l denote, respectively, the inferred cost of the “winner” (the lower-cost firm) and “loser” (the higher-cost firm) in the period-one pricing contest. Each firm adopts a stationary price along the equilibrium path in periods $t > 1$. Let $\beta(\theta_w, \theta_l)$ denote the price selected by the winner in periods $t > 1$, and we suppose that the loser charges ε more. We restrict attention to $\beta(\theta_w, \theta_l) \in [\theta_w, \theta_l]$.¹⁷ In all periods, any off-schedule deviation induces the belief threat punishment, as described below in Section 5.2.

When firms use a productive efficiency scheme, an off-schedule deviation may become apparent due to an inconsistency between a firm’s first-period and (say) second-period prices. For example, suppose firm i has type $\theta_{i,1}$ and undertakes an on-schedule deviation in period 1 by mimicking the price of a higher type, $\hat{\theta}_{i,1} > \theta_{i,1}$. Suppose firm j ’s type is lower than $\hat{\theta}_{i,1}$, so that firm j wins the first-period pricing contest and enters period 2 with the belief that $\hat{\theta}_{i,1} = \theta_l > \theta_w = \theta_{j,1}$. If the scheme specifies a period-2 price for firm j such that $\theta_{i,1} < \beta(\theta_w, \theta_l)$, then firm i will charge the price $\beta(\theta_w, \theta_l) - \varepsilon$ in period 2. Firm i ’s period 2 behavior then reveals its first-period deviation, and in period 3 the firms proceed to the belief threat punishment.

Productive efficiency equilibria are difficult to construct. Separation in the first period must be achieved, even though the first-period price may affect beliefs and thereby future profits. A subtlety arises because of a potential non-differentiability of payoffs in the first period for a firm of type $\theta_{i,1}$ at $\rho_{i,1}(\theta_{i,1})$. If firm i charges $\rho_{i,1}(\hat{\theta}_{i,1})$ for $\hat{\theta}_{i,1} > \theta_{i,1}$ in the first period, it is possible that $\hat{\theta}_{i,1} > \theta_{j,1} > \theta_{i,1}$, in which case firm i will undercut firm j ’s period-2 price, $\beta(\theta_{j,1}, \hat{\theta}_{i,1})$. On the other hand, if firm i charges $\rho_{i,1}(\hat{\theta}_{i,1})$ for $\hat{\theta}_{i,1} < \theta_{i,1}$, it is possible that $\hat{\theta}_{i,1} < \theta_{j,1} < \theta_{i,1}$, in which case firm i would not select the period-2 price $\beta(\hat{\theta}_{i,1}, \theta_{j,1})$ but would instead set a higher price (e.g., above r) and earn zero profit. Thus, payoffs change at different rates for “upward” deviations than for “downward” deviations. However, if $\beta(\theta_w, \theta_l)$ is strictly increasing in both arguments at appropriate rates, it is possible to exactly equalize the incentive to deviate upward with the incentive to deviate downward.

Strict monotonicity of $\beta(\theta_w, \theta_l)$ in turn requires that the first-period pricing schedule places each firm type above its static reaction curve (i.e., at a price such that first-period expected profit would be higher if a slightly lower price were selected). Intuitively, when a firm contemplates an increase in its first-period price, it then foresees a loss in its first-period expected

¹⁷ This restriction is required if a productive efficiency scheme is to be used in an PPBE: if $\beta(\theta_w, \theta_l) < \theta_w$, the winner would deviate (e.g., price above r) in period 2; and if $\beta(\theta_w, \theta_l) > \theta_l$, the loser would deviate and undercut the winner in period 2.

profit, and this loss is balanced against the benefit of the higher future price, $\beta(\theta_w, \theta_l)$, that the firm would enjoy were it to win the first-period pricing contest.

In the supplementary material for this article, we establish that a productive efficiency equilibrium exists if two conditions hold. The first condition is that

$$\inf_{\theta'_{i,1} > \theta''_{i,1}} \frac{f_0(\theta'_{i,1})}{f_0(\theta''_{i,1})} > \frac{\delta}{1 - \delta(1 - \delta)} \quad \text{and} \quad \inf_{\theta'_{i,1} < \theta''_{i,1}} \frac{f_0(\theta'_{i,1})}{f_0(\theta''_{i,1})} > \delta, \quad (5.2)$$

and the second condition is that δ is sufficiently small that, for all $\theta_{i,1}$,

$$\frac{2f_0(\theta_{i,1})}{(1 - F_0(\theta_{i,1}))^2} \int_{\tilde{\theta}_i = \theta_{i,1}}^{\bar{\theta}} (1 - F_0(\tilde{\theta}_i)) d\tilde{\theta}_i > \delta. \quad (5.3)$$

For any $\delta < 1$, the conditions hold when F_0 is sufficiently close to uniform. As well, for any F_0 , the conditions hold if δ is sufficiently small.

Proposition 6. *Consider Model 2 and suppose $I = 2$. If (5.2) and (5.3) are satisfied, then there exists a productive efficiency equilibrium. Specifically, in the first period, each firm i uses the following strategy:*

$$\rho_{i,1}(\theta_{i,1}) = \theta_{i,1} + \frac{2}{2 - \delta} \frac{1}{1 - F_0(\theta_{i,1})} \int_{\tilde{\theta}_i = \theta_{i,1}}^{\bar{\theta}} (1 - F_0(\tilde{\theta}_i)) d\tilde{\theta}_i.$$

Let $\theta_w = \min(\theta_{1,1}, \theta_{2,1})$, while $\theta_l = \max(\theta_{1,1}, \theta_{2,1})$. If firm i is the low-cost firm in period 1, then for all $t > 1$, firm i sets price

$$p_{i,t} = \beta(\theta_w, \theta_l) = \frac{1 - \delta}{2 - \delta} \theta_w + \frac{1}{2 - \delta} \theta_l,$$

while firm $j \neq i$ sets price $p_{j,t} = \beta(\theta_w, \theta_l) + \varepsilon$ for $\varepsilon > 0$.

Conditions (5.2) and (5.3) are satisfied in a rich parameter space; however, when they are not satisfied, a productive efficiency equilibrium may fail to exist. Intuitively, the highest-cost type ($\bar{\theta}$) gets no future profit and thus prices at cost in the first period. All other types, however, distort their first-period prices upward, in an attempt to signal higher costs and thereby secure a higher future price. If firms are very patient, the benefit of a higher future price is significant, and greater distortions in the first-period price are incurred. It is then possible that higher-cost types may price above $\bar{\theta}$ and thus the first-period price of the highest-cost type. This implies a non-monotonicity in the first-period pricing function, in contradiction to the hypothesis of a separating equilibrium.

We conclude that separating equilibria with productive efficiency exist under certain conditions.¹⁸ Such equilibria are characterized by strategic signaling in the first period. They thus repre-

¹⁸ As we show in the supplementary material section, when $I = 2$, if the parameters of the model satisfy conditions (5.2) and (5.3), then the cross-partial of the profit function for type $\theta_{i,1}$ that mimics type $\hat{\theta}_{i,1}$ is globally positive. Standard arguments then ensure that global incentive compatibility is implied by local incentive compatibility. When $I > 2$, we can show that an analogous productive efficiency equilibrium exists if δ is sufficiently small; however, a different proof is required, since the cross-partial of the profit function for type $\theta_{i,1}$ that mimics $\hat{\theta}_{i,1}$ becomes negative for $\hat{\theta}_{i,1}$ sufficiently close to $\bar{\theta}$. For $I > 2$, we use numerical methods to verify directly that global incentive compatibility is satisfied if δ is sufficiently low for the uniform distribution.

sent the Bertrand counterpart to the separating equilibria constructed by Mailath (1989) for a two-period model with differentiated products and perfectly persistent cost types.

5.2. Belief Threat Punishment

We now consider punishments that are not themselves equilibria at the start of the game, because they rely on beliefs that may only arise following a deviation from equilibrium. We seek to identify the most severe punishment of this sort. To this end, we employ the *belief threat punishment*: a deviant firm is forever after believed to have the lowest cost and is thus expected to charge a low price, regardless of the subsequent path of play, which in turn makes it rational for non-deviating firms to punish with their own low prices.

Belief Threat Punishment: Suppose that firm i engages in an off-schedule deviation in period τ . All firms $j \neq i$ thereafter believe that firm i has the lowest costs, $\underline{\theta}$, and they set the price $p_{j,t} = \underline{\theta} + 2\varepsilon$ in all future periods $t > \tau$, regardless of the evolution of play.

Now, if firm i indeed did have cost $\underline{\theta}$, then its best response against the belief threat punishment following its own deviation would in fact be to set $p_{i,t} = \underline{\theta} + \varepsilon$. If firm i does not have low cost, it chooses any price greater than $p_{j,t}$. This behavior is sequentially rational - each firm is doing its best from any point forward, given its beliefs and the equilibrium strategies of other firms. Furthermore, this is the most severe possible punishment outcome, since a deviant firm earns zero profit in the continuation game, independent of the discount factor.

While the belief threat punishment serves as a useful benchmark, it is not entirely plausible. An immediate objection to the construction just presented is that all firm j 's adopt dominated strategies (pricing below cost, for all histories) in the continuation. This objection can be handled easily, however, if we modify the above strategies to include a carrot-stick component.

Carrot-Stick Belief Threat Punishment: Suppose that firm i engages in an off-schedule deviation in period τ . The firms then impose a belief threat punishment with the modification that, in period $t > \tau$, if the deviant firm i plays $p_{i,t} = \underline{\theta} + \varepsilon$ and each firm $j \neq i$ plays $p_{j,t} = \underline{\theta} + 2\varepsilon$, then with some probability $\chi \in (0, 1)$ the firms switch to the best rigid-pricing equilibrium. Otherwise, they continue with the described punishment strategies.

For χ sufficiently low, the deviant firm still earns approximately zero profit. But it is now a strict best response for a non-deviant firm j to select $p_{j,t} = \underline{\theta} + 2\varepsilon$ throughout the punishment phase: this strategy induces a distribution over zero and positive profits, whereas any other strategy induces zero or negative profit in the current period and serves only to delay the eventual escape to the collusive continuation. Thus, the described strategies are no longer dominated. In this case, the continuation play itself requires a discount factor that is sufficiently high, since firms must be dissuaded from undercutting r in the punishment phase when $\chi > 0$.

The (carrot-stick) belief threat punishment implies a new critical discount factor for the best rigid-pricing equilibrium. Formally:

Proposition 7. *Consider Model 2 and suppose $\delta > (I - 1)/I$. Then, there exists a best rigid-pricing equilibrium. If firm i deviates, the continuation entails a carrot-stick belief threat punishment, and so firms $j \neq i$ price at $\underline{\theta} + 2\varepsilon$ in subsequent periods, and firm i prices above $\underline{\theta} + 2\varepsilon$ unless its cost type is less than $\underline{\theta} + 2\varepsilon$.*

Proof: See Appendix. ■

The critical discount factor $(I - 1)/I$ is strictly less than δ_c , and so we now have a lower critical discount factor for supporting the best rigid-pricing equilibrium. We note that $(I - 1)/I$ is also the standard critical discount factor for Bertrand supergames with complete information. Thus, if we are willing to impose the (carrot-stick) belief threat punishment, then incomplete information does not necessitate a higher discount factor in order to support the optimal collusive arrangement (under log-concavity).

While the equilibrium of Proposition 7 entails undominated strategies, one may object that the non-deviating firms might relinquish their worst-case beliefs after a deviation, if the deviant firm consistently did not price at $\underline{\theta} + \varepsilon$. Our specification requires a dogged pessimism: even if the deviant firm i hasn't priced at $\underline{\theta} + \varepsilon$ yet, each firm $j \neq i$ remains sure that firm i will do so tomorrow. Standard refinements also do not eliminate this equilibrium. The belief threat punishment as stated, however, is not robust to the possibility of imperfect persistence.

6. Conclusion

We analyze a dynamic Bertrand game (or equivalently, a series of procurement auctions), in which prices are publicly observed and each firm is privately informed as to its costs. Costs are independent across firms, but each firm's cost exhibits persistence over time. We characterize the set of collusive equilibria, giving particular emphasis to the collusive scheme that is optimal for firms at the start of the game. When costs are perfectly persistent, if the distribution of costs is log concave and firms are sufficiently patient, then the optimal collusive scheme entails price rigidity. While it is possible in some circumstances for firms to implement separating equilibria with productive efficiency, these equilibria are not optimal. When costs can take two types and are imperfectly persistent, some productive efficiency is typically optimal. First-best collusion is possible if the firms are sufficiently patient relative to the degree of persistence.

With the basic modeling framework now established, a number of exciting extensions may be considered. For example, it would be interesting to examine the model when firms face fixed costs and participation in the market is endogenous. A tension might then arise between collusion and predation, with the latter option perhaps having particular appeal to a firm that believes its unit cost is relatively low. Likewise, it would be interesting to include an investment

process, whereby firms could endogenously influence their respective cost distributions.¹⁹

At a methodological level, our analysis is novel in that we characterize optimal cooperation in a dynamic game with persistent, private information. In many applications, agents seek a self-enforcing cooperative relationship, and private information is important and persistent. The techniques developed here should be useful for such applications.

7. Appendix

Proof of Proposition 1: Consider the strategies described in the text. We refer to the equilibrium starting with price p_w as being in the “war” state. Let $\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi)$ be the discounted profit to firm i with cost type $\theta_{i,t}$ in period t if the firms are in the war state.

In order to use the same expressions for Models 1 and 2, we use the $\{\underline{\theta}, \bar{\theta}\}$ notation throughout the proof. For Model 1, we use the notation that

$$\begin{aligned}\gamma_{\underline{\theta}} &= (1 - \delta)\mathbb{E}\left[\sum_{s=t}^{\infty} \delta^{s-t}\theta_{i,t} \middle| \theta_{i,t} = \underline{\theta}\right] = \frac{(1 - \delta(1 - \lambda_{\bar{\theta}})\underline{\theta}) + \delta(1 - \lambda_{\underline{\theta}})\bar{\theta}}{1 - \delta(\lambda_{\underline{\theta}} - \lambda_{\bar{\theta}})}, \\ \gamma_{\bar{\theta}} &= (1 - \delta)\mathbb{E}\left[\sum_{s=t}^{\infty} \delta^{s-t}\theta_{i,t} \middle| \theta_{i,t} = \bar{\theta}\right] = \frac{(1 - \delta\lambda_{\underline{\theta}})\bar{\theta} + \delta\lambda_{\bar{\theta}}\underline{\theta}}{1 - \delta(\lambda_{\underline{\theta}} - \lambda_{\bar{\theta}})}.\end{aligned}$$

For Model 2, in the work below, we use the convention that $\gamma_{\theta_{i,t}} = \theta_{i,t}$ and $\lambda_{\underline{\theta}} = 1 > 0 = \lambda_{\bar{\theta}}$.

We begin by representing the continuation values in collusive and war states. In both models, for the collusive state, we have

$$v_i^r(\theta_{i,t}) = \frac{1}{I(1 - \delta)}(r - \gamma_{\theta_{i,t}}). \quad (7.1)$$

In Model 1, starting in the war state, the profit earned over the game by type $\theta_{i,t}$ is

$$\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi) = \frac{p_w - \theta_{i,t}}{I} + \delta \left[\begin{array}{l} \lambda_{\theta_{i,t}}(\chi v_i^r(\underline{\theta}) + (1 - \chi)\tilde{v}_i^{cs}(\underline{\theta}; p_w, \chi)) \\ + (1 - \lambda_{\theta_{i,t}})(\chi v_i^r(\bar{\theta}) + (1 - \chi)\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi)) \end{array} \right].$$

Likewise, in Model 2, starting in the war state, the game profit for type $\theta_{i,t}$ is

$$\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi) = \frac{p_w - \theta_{i,t}}{I} + \delta[\chi v_i^r(\theta_{i,t}) + (1 - \chi)\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi)].$$

For both models, the solution to the system of equations can be represented as

$$\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi) = \frac{1}{I(1 - \delta)} \left(\frac{\delta\chi r + (1 - \delta)p_w}{1 - \delta(1 - \chi)} - \gamma_{\theta_{i,t}} \right). \quad (7.2)$$

We next characterize further the war-state continuation value function. Using (7.2), we observe that

$$\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi) - \tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi) = \frac{(\gamma_{\bar{\theta}} - \gamma_{\theta_{i,t}})}{I(1 - \delta)}. \quad (7.3)$$

¹⁹ See Fershtman and Pakes (2000, 2004) for numerical analyses of endogenous investment among colluding firms in related settings.

Thus, $\tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi)$ is simply $\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi)$ plus a type-dependent function of model primitives, and so the most severe war state is achieved (pointwise) if $\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi)$ is reduced to the minimum value possible in a carrot-stick pooling equilibrium. To determine this minimum value, we consider a particular off-schedule incentive constraint in the war state; namely, type $\bar{\theta}$ must be deterred from setting a price above p_w . With such a deviation, type $\bar{\theta}$ would earn zero profit in the current period and induce the war state in the next period. The relevant constraint is

$$\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi) \geq \delta[\lambda_{\bar{\theta}}\tilde{v}_i^{cs}(\underline{\theta}; p_w, \chi) + (1 - \lambda_{\bar{\theta}})\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi)], \quad (7.4)$$

where we recall that $\lambda_{\bar{\theta}} = 0$ in Model 2. Using (7.3), we may rewrite (7.4) and determine the desired minimum value:

$$\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi) \geq \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}. \quad (7.5)$$

Notice that when $\tilde{v}_i^{cs}(\bar{\theta}; p_w, \chi) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}$, we have from (7.3) that

$$v_i^{cs}(\theta_{i,t}) \equiv \tilde{v}_i^{cs}(\theta_{i,t}; p_w, \chi) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2} + \frac{(\gamma_{\bar{\theta}} - \gamma_{\theta_{i,t}})}{I(1 - \delta)}. \quad (7.6)$$

We may use (7.2) to confirm that $v_i^{cs}(\bar{\theta}) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}$ is achieved if and only if (p_w, χ) satisfy

$$\chi = \frac{1 - \delta}{\delta} \frac{\bar{\theta} - p_w}{r - \bar{\theta}}. \quad (7.7)$$

Given (7.7), we note that $\chi \in [0, 1]$ and $p_w \in [0, r)$ if and only if

$$\max\{0, \frac{\bar{\theta} - \delta r}{1 - \delta}\} \leq p_w \leq \bar{\theta}. \quad (7.8)$$

Fixing $v_i^{cs}(\bar{\theta}) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}$, we next consider the remaining off-schedule incentive constraints in the war state. First, if $p_w \leq \theta_{i,t}$, then type $\theta_{i,t}$ must be deterred from deviating to a price higher than p_w . The relevant constraint is $v_i^{cs}(\theta_{i,t}) \geq \delta\mathbb{E}_{\theta_{i,t+1}}[v_i^{cs}(\theta_{i,t+1})|\theta_{i,t}]$. When $\theta_{i,t} = \bar{\theta}$, this constraint is equivalent to (7.4) and thus holds here with equality. For $\theta_{i,t} < \bar{\theta}$, the constraint is strictly satisfied, since $\delta < 1$ and $v_i^{cs}(\theta_{i,t}) > v_i^{cs}(\bar{\theta}) \geq 0$ follows from (7.6). Second, if $p_w > \theta_{i,t}$, then type $\theta_{i,t}$ must be deterred from undercutting p_w . The relevant constraint is then $v_i^{cs}(\theta_{i,t}) \geq p_w - \theta_{i,t} + \delta\mathbb{E}_{\theta_{i,t+1}}[v_i^{cs}(\theta_{i,t+1})|\theta_{i,t}]$. It is straightforward to verify that if this constraint holds for $\underline{\theta}$, then it holds for all higher types. Using (7.6), we may then derive that this constraint is satisfied if and only if $p_w \leq \underline{\theta} + \frac{(\bar{\theta} - \underline{\theta})}{I} \equiv \bar{p}_w \in (\underline{\theta}, \bar{\theta})$. Summarizing, when $v_i^{cs}(\bar{\theta}) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}$, the remaining off-schedule constraints hold in the war state if and only if the following condition holds:

$$\text{if } p_w > \underline{\theta}, \text{ then } p_w \leq \underline{\theta} + \frac{(\bar{\theta} - \underline{\theta})}{I} \equiv \bar{p}_w \in (\underline{\theta}, \bar{\theta}). \quad (7.9)$$

What is the critical discount factor above which we can satisfy $v_i^{cs}(\bar{\theta}) = \delta\lambda_{\bar{\theta}}\frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1 - \delta)^2}$ while also satisfying (7.9)? As (7.8) reveals, when δ is low, the former may require the selection of

a high p_w ; however, a high p_w may lead to a failure of the latter. We determine the critical discount factor by setting $p_w = \bar{p}_w$ and $\chi = 1$, and then solving (7.7) for δ . The solution is

$$\delta_w = \frac{\bar{\theta} - \bar{p}_w}{r - \bar{p}_w} = \frac{(\bar{\theta} - \underline{\theta})(I - 1)}{(r - \underline{\theta})I - (\bar{\theta} - \underline{\theta})} \in (0, 1). \quad (7.10)$$

For $\delta < \delta_w$, we cannot satisfy all of the off-schedule constraints for a carrot-stick pooling equilibrium. For $\delta \geq \delta_w$, we may specify p_w and χ so as to satisfy all off-schedule constraints while generating $v_i^{cs}(\bar{\theta}) = \delta \lambda_{\bar{\theta}} \frac{(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}})}{I(1-\delta)^2}$. To this end, we select $p_w = \bar{p}_w$ and $\chi = \bar{\chi}$, where $\bar{\chi} \equiv \frac{1-\delta}{\delta} \frac{\bar{\theta} - \bar{p}_w}{r - \bar{\theta}}$. For $\delta \geq \delta_w$, \bar{p}_w satisfies (7.8), and thus $\bar{\chi} \in [0, 1]$.

We come next to the off-schedule constraints in the collusion state, taking as given the strategies and payoffs for carrot-stick as just described. For Model 1, the off-schedule incentive constraint is now

$$v_i^r(\theta_{i,t}) \geq r - \theta_{i,t} + \delta[\lambda_{\theta_{i,t}} v_i^{cs}(\underline{\theta}) + (1 - \lambda_{\theta_{i,t}}) v_i^{cs}(\bar{\theta})]. \quad (7.11)$$

In Model 2, the off-schedule incentive constraint takes the form

$$v_i^r(\theta_{i,t}) \geq r - \theta_{i,t} + \delta v_i^{cs}(\theta_{i,t}). \quad (7.12)$$

For each model, straightforward calculations confirm that the off-schedule constraint is most difficult to satisfy when $\theta_{i,t} = \underline{\theta}$. Thus, allowing now for either model, we may use (7.1) and (7.6) and conclude that the off-schedule constraint holds in the collusion state if and only if

$$r - \gamma_{\underline{\theta}} \geq (1 - \delta) I (r - \underline{\theta}) + \delta(\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}}) \left(\frac{\delta \lambda_{\bar{\theta}}}{1 - \delta} + \lambda_{\underline{\theta}} \right).$$

This constraint holds if and only if $\delta \geq \delta_c$, where δ_c is the unique $\delta \in (0, 1)$ that satisfies:

$$\delta = \frac{(r - \underline{\theta})I - (r - \gamma_{\underline{\theta}})}{(r - \underline{\theta})I - (\gamma_{\bar{\theta}} - \gamma_{\underline{\theta}}) \left(\frac{\delta \lambda_{\bar{\theta}}}{1 - \delta} + \lambda_{\underline{\theta}} \right)}. \quad (7.13)$$

Solving (7.13), we obtain that

$$\delta_c = \frac{(r - \underline{\theta})(I - 1)}{(r - \underline{\theta})I - (\bar{\theta} - \underline{\theta})}. \quad (7.14)$$

Finally, given that $r > \bar{\theta}$, a simple comparison of (7.10) and (7.14) reveals that $1 > \delta_c > \delta_w > 0$. Thus, in both models, the off-schedule constraint for the collusion state determines the critical discount factor, and that discount factor is δ_c as defined in (7.14). Hence, if $\delta \geq \delta_c$, we may enforce collusion at r using a carrot-stick punishment scheme, with $p_w = \bar{p}_w$ and $\chi = \bar{\chi}$. ■

Proof of Corollary 1: First, we fix $\delta > \delta_c$ for δ_c given in (7.14), and observe that the rigid-pricing equilibrium exists, when cost types are sufficiently persistent (i.e., as $\lambda_L \rightarrow 1$ and $\lambda_H \rightarrow 0$). This follows since δ_c is independent of the degree of persistence.

Second, if cost types are perfectly persistent but there are two possible cost types, then to show that the ex ante optimal PPBE is the best rigid-pricing scheme, we may mimic the proof approach used for Proposition 2, except that we now need to find the condition analogous to log-concavity for the two-type model.²⁰ Analogous to the proof of Proposition 2, we consider a relaxed program, whereby we choose $\check{R}_i(\theta_{i,1})$ and $\check{M}_i(\theta_{i,1})$ to maximize $\mathbb{E}_{\theta_{i,1}} \left[\check{R}_i(\theta_{i,1}) - \theta_{i,1} \check{M}_i(\theta_{i,1}) | F_0 \right]$ subject to the following set of constraints: $\check{R}_i(\theta_{i,1}) \leq r \cdot \check{M}_i(\theta_{i,1})$ for each $\theta_{i,1}$ (pricing constraint), $\mathbb{E}_{\theta_{i,1}} \left[\check{M}_i(\theta_{i,1}) \right] = 1/(I(1-\delta))$ (ex ante symmetry constraint), $\check{R}_i(\bar{\theta}) - \bar{\theta} \check{M}_i(\bar{\theta}) \geq \check{R}_i(\underline{\theta}) - \bar{\theta} \check{M}_i(\underline{\theta})$ (IC-down), and $\check{R}_i(\underline{\theta}) - \underline{\theta} \check{M}_i(\underline{\theta}) \geq \check{R}_i(\bar{\theta}) - \underline{\theta} \check{M}_i(\bar{\theta})$ (IC-up). It is straightforward to show that (IC-down) and (IC-up) imply that $\check{M}_i(\theta_{i,1})$ is nonincreasing (market share monotonicity), and given that, (IC-up) must be slack unless $\check{M}_i(\underline{\theta}) = \check{M}_i(\bar{\theta})$, and further $\check{R}_i(\bar{\theta}) = r \cdot \check{M}_i(\bar{\theta})$ (if not, $\check{R}_i(\bar{\theta})$ could be increased). In turn, it can be shown that (IC-down) must bind. Then, substituting in the (IC-down) constraint, the objective can be written $\check{M}_i(\bar{\theta})(r - \bar{\theta}) + F_0(\underline{\theta}) \check{M}_i(\underline{\theta})(\bar{\theta} - \underline{\theta})$. Substituting in the ex ante symmetry constraint, the problem becomes to choose $\check{M}_i(\bar{\theta}) \leq 1/(I(1-\delta))$ (where the bound comes from ex ante symmetry and market share monotonicity) to maximize

$$\check{M}_i(\bar{\theta})(r - \bar{\theta}) + F_0(\underline{\theta}) \frac{1}{F_0(\underline{\theta})} \left(\frac{1}{I(1-\delta)} - \check{M}_i(\bar{\theta})(1 - F_0(\underline{\theta})) \right) (\bar{\theta} - \underline{\theta}).$$

This expression is linear in $\check{M}_i(\bar{\theta})$, and it is increasing in $\check{M}_i(\bar{\theta})$ if and only if $(r - \bar{\theta})/(\bar{\theta} - \underline{\theta}) > 1 - F_0(\underline{\theta})$, as desired. Thus, under the parameter restriction, the solution is $\check{M}_i(\bar{\theta}) = \check{M}_i(\underline{\theta}) = 1/(I(1-\delta))$ and $\check{R}_i(\theta_{i,1}) = r \cdot \check{M}_i(\theta_{i,1})$, which can be implemented as a PPBE using the best rigid-pricing policy.

Third, define the following objects. For the purposes of this proof, we modify our notational convention so that $\mathbb{E}_{\{\boldsymbol{\theta}^t\}_{t=1}^{\infty}} [g(\{\boldsymbol{\theta}^t\}_{t=1}^{\infty}) | \mathbf{F}_{-i,0}, \{\lambda'_L, \lambda'_H\}]$ indicates that the expectation is taken using transition probabilities $\{\lambda'_L, \lambda'_H\}$. Further, we modify the definition of \tilde{v}_i to include the arguments $\{\lambda'_L, \lambda'_H\}$, so that $\tilde{v}_i(\boldsymbol{\sigma}, \mathbf{F}_{-i,0}, h^0, \theta_{i,1}, \{\lambda'_L, \lambda'_H\})$ is calculated using transition probabilities $\{\lambda'_L, \lambda'_H\}$, where h^0 denotes the null history. Let $\boldsymbol{\sigma} = \{\boldsymbol{\sigma}_t\}_{t=\tau}^{\infty}$. Define $\tilde{v}_i^n(\boldsymbol{\sigma}) \equiv \mathbb{E}_{\theta_{i,1}} [\tilde{v}_i(\boldsymbol{\sigma}, \mathbf{F}_{-i,0}, h^0, \theta_{i,1}, \{\lambda'_L, \lambda'_H\}) | F_0]$. Define the present discounted values of market share and revenue given that firm i of type $\theta_{i,1}$ mimics type $\hat{\theta}_{i,1}$ throughout the game:

$$\begin{aligned} M_i(\hat{\theta}_{i,1}; \{\boldsymbol{\sigma}_t\}_{t=1}^{\infty}, \{\lambda'_L, \lambda'_H\}) &\equiv \mathbb{E}_{\{\boldsymbol{\theta}^t\}_{t=1}^{\infty}} \left[\sum_{t=1}^{\infty} \delta^{t-1} \bar{m}_i(\hat{\theta}_{i,1}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}) \Big| \mathbf{F}_{-i,0}, \{\lambda'_L, \lambda'_H\} \right] \\ R_i(\hat{\theta}_{i,1}; \{\boldsymbol{\sigma}_t\}_{t=1}^{\infty}, \{\lambda'_L, \lambda'_H\}) &\equiv \mathbb{E}_{\{\boldsymbol{\theta}^t\}_{t=1}^{\infty}} \left[\sum_{t=1}^{\infty} \delta^{t-1} \rho_{i,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,1}, \boldsymbol{\theta}_{-i,1}), \hat{\theta}_{i,1}) \cdot \bar{m}_i(\hat{\theta}_{i,1}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}) \Big| \mathbf{F}_{-i,0}, \{\lambda'_L, \lambda'_H\} \right] \end{aligned}$$

²⁰ See Athey and Bagwell (2001) for a related argument when cost types that are i.i.d. over time.

where, for all $t \geq 1$,

$$\begin{aligned} s_{j,t} &= (\alpha_{j,t}, \rho_{j,t}, \psi_{j,t}) = \boldsymbol{\sigma}_{j,t}(h^{t-1}) \text{ for all } j \in \{1, \dots, I\}, \\ z_{j,t} &= (a_{j,t}, p_{j,t}, q_{j,t}) = s_{j,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,1}, \boldsymbol{\theta}_{-i,1}), \theta_{j,1}) \text{ for all } j \neq i, \\ z_{i,t} &= (a_{i,t}, p_{i,t}, q_{i,t}) = s_{i,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,1}, \boldsymbol{\theta}_{-i,1}), \hat{\theta}_{i,1}), \text{ and} \\ \boldsymbol{\mu}_{-i,t+1} &\in \mathbf{T}_{-i}(\boldsymbol{\mu}_{-i,t}, \mathbf{s}_{-i,t}, \mathbf{z}_t). \end{aligned}$$

Let

$$\check{v}_i^*(\boldsymbol{\sigma}) \equiv \mathbb{E}_{\theta_{i,1}}[\check{v}_i(\boldsymbol{\sigma}, \mathbf{F}_{-i,0}, \{\}, \theta_{i,1}, \{1, 0\}) | F_0] = \mathbb{E}_{\theta_{i,1}}[R_i(\theta_{i,1}; \boldsymbol{\sigma}, \{1, 0\}) - \theta_{i,1}M_i(\theta_{i,1}; \boldsymbol{\sigma}, \{1, 0\})].$$

Let $\boldsymbol{\sigma}^R \equiv \{\boldsymbol{\sigma}_t^R\}_{t=1}^\infty$ be the extensive-form strategies for the best rigid-pricing scheme. For a given $\boldsymbol{\sigma}$, let $c(\boldsymbol{\sigma}, \hat{\theta}_i)$ denote the strategy derived from $\boldsymbol{\sigma}$ by following the behavior that $\boldsymbol{\sigma}$ would assign if the firm had realizations of type $\hat{\theta}_i$ in every period. Finally, define a strategy $\boldsymbol{\sigma}^n$ to be individually rational given $\{\lambda_L^n, \lambda_H^n\}$, if, starting from each period t (and given each possible h^t and induced beliefs $\boldsymbol{\mu}_{-i,t}$), each firm i expects average per-period payoffs of at least zero from that point on. Since each firm can always guarantee profits of 0 by charging a price sufficiently high, only individually rational strategies can be used in a PPBE.

Fourth, we need to prove the following claim: for all $\varepsilon > 0$, there exists n^* such that for all $n > n^*$, $\check{v}_i^n(\boldsymbol{\sigma}) - \check{v}_i^n(\boldsymbol{\sigma}^R) < \varepsilon$ for all $\boldsymbol{\sigma}$ such that $\boldsymbol{\sigma}$ is individually rational given $\{\lambda_L^n, \lambda_H^n\}$, and such that $\check{v}_i^n(\boldsymbol{\sigma}) \geq \check{v}_i^n(c(\boldsymbol{\sigma}, \hat{\theta}_i))$ for $\hat{\theta}_i \in \{L, H\}$. If this claim is true, we would conclude that for high enough n , the best rigid-pricing policy approximates the payoffs that could be attained using the best strategy that is immune to a deviation where a firm pretends to be type L in every period, as well as to all deviations where a firm pretends to be type H in every period. Clearly, the best strategy that is immune to this limited class of deviations yields profits at least as high as the best PPBE strategy.

To establish the claim, we proceed by contradiction. In particular, consider the following hypothesis, which we refer to as the no-convergence hypothesis, or NC: that there exists $\varepsilon > 0$, such that for all n^* , there exists $n' > n^*$ and a corresponding $\boldsymbol{\sigma}^{n'}$ such that $\check{v}_i^{n'}(\boldsymbol{\sigma}^{n'}) - \check{v}_i^{n'}(\boldsymbol{\sigma}^R) \geq \varepsilon$, $\boldsymbol{\sigma}^{n'}$ is individually rational given $\{\lambda_L^{n'}, \lambda_H^{n'}\}$, and $\check{v}_i^{n'}(\boldsymbol{\sigma}^{n'}) \geq \check{v}_i^{n'}(c(\boldsymbol{\sigma}^{n'}, \hat{\theta}_i))$ for $\hat{\theta}_i \in \{L, H\}$.

For each $\boldsymbol{\sigma}^{n'}$ and each $\hat{\theta}_{i,1}$, we can calculate $M_i(\hat{\theta}_{i,1}; \boldsymbol{\sigma}^{n'}, \{1, 0\})$ and $R_i(\hat{\theta}_{i,1}; \boldsymbol{\sigma}^{n'}, \{1, 0\})$. Since these are real numbers drawn from compact subsets of the real line, there exists a convergent subsequence. Call the limits $M_i^*(\hat{\theta}_{i,1})$ and $R_i^*(\hat{\theta}_{i,1})$, and restrict attention to that subsequence.

We then argue that $\check{v}_i^{n'}(\boldsymbol{\sigma}^{n'})$ converges to $\mathbb{E}_{\theta_{i,1}}[R_i^*(\theta_{i,1}) - \theta_{i,1}M_i^*(\theta_{i,1}) | F_0]$. To do so, we argue that $\check{v}_i^*(\boldsymbol{\sigma}^{n'}) - \check{v}_i^{n'}(\boldsymbol{\sigma}^{n'}) + \mathbb{E}_{\theta_{i,1}}[R_i^*(\theta_{i,1}) - \theta_{i,1}M_i^*(\theta_{i,1}) | F_0] - \check{v}_i^*(\boldsymbol{\sigma}^{n'})$ converges to zero. Start by considering the last two terms, which can be rewritten

$$\mathbb{E}_{\theta_{i,1}} \left[R_i^*(\theta_{i,1}) - \theta_{i,1}M_i^*(\theta_{i,1}) - \left[R_i(\theta_{i,1}; \boldsymbol{\sigma}^{n'}, \{1, 0\}) - \theta_{i,1}M_i(\theta_{i,1}; \boldsymbol{\sigma}^{n'}, \{1, 0\}) \right] \middle| F_0 \right].$$

This converges to zero by definition of $R_i^*(\theta_{i,1})$ and $M_i^*(\theta_{i,1})$.

To show that $\check{v}_i^*(\boldsymbol{\sigma}^{n'}) - \check{v}_i^{n'}(\boldsymbol{\sigma}^{n'})$ converges to zero, recall that \check{v}_i^* and $\check{v}_i^{n'}$ differ at a given strategy profile $\boldsymbol{\sigma}$ only because they place different weights on the probability that different histories are realized; in particular, \check{v}_i^* is calculated assuming that costs do not change over time, while $\check{v}_i^{n'}$ is calculated assuming transition probabilities $\{\lambda_L^{n'}, \lambda_H^{n'}\}$. Since $\boldsymbol{\sigma}^{n'}$ is assumed individually rational, following all histories, firm i 's expected discounted future payoffs must be between 0 and $\frac{r-\theta}{1-\delta}$. For a given $\boldsymbol{\sigma}$, $\check{v}_i^*(\boldsymbol{\sigma})$ and $\check{v}_i^{n'}(\boldsymbol{\sigma})$ differ only following a period where some firm experiences a cost change, in which case the difference between expected payoffs from that point on (computed using $\{\lambda_L^n, \lambda_H^n\}$) is at most $\frac{r-\theta}{1-\delta}$. Consider the case where $\lambda_L^n \leq 1 - \lambda_H^n$ for each n (other cases are analogous). Since given transition probabilities $\{\lambda_L^{n'}, \lambda_H^{n'}\}$, $(\lambda_L^{n'})^{I(t-1)}$ is an lower bound on the probability that up until period t , no firm experienced a cost change, and $(1 - (\lambda_L^{n'})^I)$ is an upper bound on the probability that a cost change occurs in period t ,

$$\left| \check{v}_i^*(\boldsymbol{\sigma}) - \check{v}_i^{n'}(\boldsymbol{\sigma}) \right| \leq \sum_{t=1}^{\infty} \delta^t (\lambda_L^{n'})^{I(t-1)} (1 - (\lambda_L^{n'})^I) \frac{r-\theta}{1-\delta} = \frac{\delta}{1-\delta} (r-\theta) \frac{1 - (\lambda_L^{n'})^I}{1 - \delta(\lambda_L^{n'})^I}.$$

This bound does not depend on $\boldsymbol{\sigma}$, and so $\check{v}_i^*(\boldsymbol{\sigma}) - \check{v}_i^{n'}(\boldsymbol{\sigma})$ converges to zero uniformly as $n' \rightarrow \infty$ (and $\lambda_L^{n'} \rightarrow 1$). This implies that $\check{v}_i^*(\boldsymbol{\sigma}^{n'}) - \check{v}_i^{n'}(\boldsymbol{\sigma}^{n'})$ converges to 0.

Finally, note that as $n' \rightarrow \infty$, $\check{v}_i^{n'}(\boldsymbol{\sigma}^R)$ converges to $\mathbb{E}_{\theta_{i,1}} \left[\frac{1}{1-\delta} \frac{1}{I} (r - \theta_{i,1}) \right]$. So, by the maintained hypothesis NC, $\mathbb{E}_{\theta_{i,1}} [R_i^*(\theta_{i,1}) - \theta_{i,1} M_i^*(\theta_{i,1}) - \frac{1}{1-\delta} \frac{1}{I} (r - \theta_{i,1}) | F_0] \geq \varepsilon$, and

$$R_i^*(\underline{\theta}) - \underline{\theta} M_i^*(\underline{\theta}) \geq R_i^*(\bar{\theta}) - \bar{\theta} M_i^*(\bar{\theta}), \quad R_i^*(\bar{\theta}) - \bar{\theta} M_i^*(\bar{\theta}) \geq R_i^*(\underline{\theta}) - \underline{\theta} M_i^*(\underline{\theta}). \quad (7.15)$$

Since $M_i^{n'}(\theta_{i,1}) \equiv M_i(\theta_{i,1}; \boldsymbol{\sigma}^{n'}, \{\lambda_L', \lambda_H'\})$ is nonincreasing (a consequence of the assumption $\check{v}_i^{n'}(\boldsymbol{\sigma}^{n'}) \geq \check{v}_i^{n'}(c(\boldsymbol{\sigma}^{n'}, \hat{\theta}_i))$ for $\hat{\theta}_i \in \{L, H\}$), $\mathbb{E}_{\theta_{i,1}} [M_i^{n'}(\theta_{i,1}) | F_0] = 1/(I(1-\delta))$ (by ex ante symmetry), and since $M_i^{n'}(\theta_{i,1}) \in [0, 1/(1-\delta)]$ for all $\theta_{i,1}$, $M_i^*(\theta_{i,1})$ has these properties as well. Similarly, since $R_i(\theta_{i,1}; \boldsymbol{\sigma}^{n'}, \{1, 0\}) \in [0, r/(1-\delta)]$, $R_i^*(\theta_{i,1}) \in [0, r/(1-\delta)]$.

But, we already argued that $R_i^*(\theta_{i,1}) = \frac{r}{I(1-\delta)}$, $M_i^*(\theta_{i,1}) = 1/(I(1-\delta))$ maximizes $\mathbb{E}_{\theta_{i,1}} [R_i^*(\theta_{i,1}) - \theta_{i,1} M_i^*(\theta_{i,1})]$ subject to (7.15) and the latter set of requirements on R_i^* and M_i^* , a contradiction. ■

Proof of Proposition 3: Under the specified two-step scheme,

$$M^L(z) \equiv \sum_{N=1}^I \binom{I-1}{N-1} F_0(z)^{N-1} (1 - F_0(z))^{I-N} \left(\frac{1}{N} (1 - (I-N) M^{cs}(z)) \right)$$

is the market share that a type on $[\underline{\theta}, z)$ expects to receive in all periods after period 2. To satisfy the first-period on-schedule constraint, p_L must satisfy

$$(p_L - z) \left(\frac{1 - (1 - F_0(z))^I}{F_0(z)I} \right) + (r - z) \frac{\delta}{1 - \delta} M^L(z) = \frac{r - z}{1 - \delta} \left((1 - F_0(z))^{I-1} \frac{1}{I} + \delta (1 - (1 - F_0(z))^{I-1}) M^{cs}(z) \right),$$

or

$$p_L - z = \frac{r - z}{\frac{1 - (1 - F_0(z))^I}{F_0(z)I}} \left((1 - F_0(z))^{I-1} \frac{1}{I} - \frac{\delta}{1 - \delta} \left(M^L(z) - (1 - (1 - F_0(z))^{I-1}) M^{cs}(z) - (1 - F_0(z))^{I-1} \frac{1}{I} \right) \right).$$

Note that $p_L < r$, and that p_L decreases with δ , because the expected market share in periods $t > 1$ is greater when p_L is chosen.

By construction, if $\delta \geq \delta_c$, then the off-schedule constraints are satisfied for periods $t > 1$. Now consider the off-schedule constraints in period 1. There is no gain to making a deviant announcement, since announcements are uninformative. The most profitable price deviations entail undercutting either p_L or r . All possible deviations in quantity restrictions in the first period are dominated by price deviations: undercutting a price of p_L or r guarantees that the firm wins the market, and undercutting r guarantees a positive profit for all types.

Let us begin by considering the temptation to undercut p_L in the first period. Note that if $p_L < \theta_{i,1}$, which will hold as δ approaches 1, then the incentive to undercut p_L disappears. So, this constraint only arises for moderate δ . On-schedule incentive compatibility tells us that types on $[z, \bar{\theta}]$ prefer to use their assigned strategies rather than mimic types below z . Thus, it is sufficient to check that all types prefer to follow the behavior assigned to types on $[\underline{\theta}, z]$ rather than undercut p_L . The latter constraint can be represented as follows:

$$\left(\frac{1 - (1 - F_0(z))^I}{F_0(z)I} \right) (p_L - \theta_{i,1}) + \frac{\delta}{1 - \delta} (r - \theta_{i,1}) M^L(z) \geq p_L - \theta_{i,1} + \frac{\delta}{1 - \delta} \frac{1}{I} (\bar{\theta} - \theta_{i,1}). \quad (7.16)$$

We begin by comparing this constraint to the off-schedule constraint imposed by the best rigid-pricing scheme:

$$\frac{1}{I} (r - \theta_{i,1}) + \frac{\delta}{1 - \delta} (r - \theta_{i,1}) \frac{1}{I} \geq r - \theta_{i,1} + \frac{\delta}{1 - \delta} \frac{1}{I} (\bar{\theta} - \theta_{i,1}).$$

The first-period benefits from a deviation are smaller with the market-sharing two-step scheme, because $p_L < r$ and $\frac{1 - (1 - F_0(z))^I}{F_0(z)I} > \frac{1}{I}$; the future benefits to cooperating are higher because $M^L(z) > \frac{1}{I}$. Thus, if $\delta \geq \delta_c$, so that the best rigid-pricing scheme is feasible, then (7.16) will be satisfied as well, and the critical discount factor such that (7.16) is satisfied is less than δ_c .

Now consider the temptation to undercut the price of r . First consider this temptation for types $\theta_{i,1} \in [\underline{\theta}, z]$. The off-schedule constraint is more difficult to satisfy as $\theta_{i,1}$ increases on $[\underline{\theta}, z]$, because in all periods the market share is higher from following the equilibrium strategies than from engaging in the deviation and then switching to the carrot-stick equilibrium. Thus, lower cost types find the equilibrium strategy relatively more appealing. The first-period off-schedule constraint that deters this deviation for all $\theta_{i,1} \in [\underline{\theta}, z]$ is then

$$\begin{aligned} \left(\frac{1 - (1 - F_0(z))^I}{F_0(z)I} \right) (p_L - z) + \frac{\delta}{1 - \delta} (r - z) M^L(z) &\geq (r - z) (1 - F_0(z))^{I-1} + \frac{\delta}{1 - \delta} \frac{1}{I} (\bar{\theta} - z) \\ &= (r - z) \left(\frac{1}{1 - \delta} M^{cs}(z) - \left(1 - (1 - F_0(z))^{I-1} \right) \right) \end{aligned}$$

where the equality follows by the definition of $M^{cs}(z)$. Substituting in for p_L yields

$$\frac{\delta}{1 - \delta} (1 - F_0(z))^{I-1} (1 - M^{cs}(z)) \geq M^{cs}(z) - \frac{(1 - F_0(z))^{I-1}}{1 - \delta} \left[\frac{1}{I} - \delta \right] - \left(1 - (1 - F_0(z))^{I-1} \right). \quad (7.17)$$

For $\delta \geq \delta_c$, we know $M^{cs}(z) < \frac{1}{I}$. The left-hand side is thus strictly greater than the value, \underline{L} , it achieves when $M^{cs}(z)$ is replaced with $\frac{1}{I}$; likewise, the right-hand side is strictly lower than the value, \bar{R} , it achieves when $M^{cs}(z)$ is replaced with $\frac{1}{I}$. Comparing, we find that $\underline{L} > \bar{R}$, and so (7.17) is satisfied as a strict inequality. Intuitively, the prospect of future market share $M^{cs}(z)$ is enough to deter a deviation in period 2, when the gain in market share from undercutting r is greater than it is in period 1 because some opponent types will choose p_L in period 1; thus, the promise of $M^{cs}(z)$ in the future is more than enough to deter a deviation in period 1.

Now consider the temptation to undercut r for $\theta_{i,1} \in [z, \bar{\theta}]$. The relevant constraint is

$$\frac{r - \theta_{i,1}}{1 - \delta} \left((1 - F_0(z))^{I-1} \cdot \frac{1}{I} + \delta \left(1 - (1 - F_0(z))^{I-1} \right) M^{cs}(z) \right) \geq (r - \theta_{i,1}) (1 - F_0(z))^{I-1} + \frac{\delta}{1 - \delta} \frac{1}{I} (\bar{\theta} - \theta_{i,1}).$$

We compare the critical discount factor for this constraint to δ_c . Rewriting, we have

$$\begin{aligned} & \frac{\delta}{1 - \delta} (r - \theta_{i,1}) \frac{1}{I} - (r - \theta_{i,1}) \left(1 - \frac{1}{I} \right) - \frac{\delta}{1 - \delta} \frac{1}{I} (\bar{\theta} - \theta_{i,1}) \\ & \geq \frac{\delta}{1 - \delta} \left(\frac{1 - (1 - F_0(z))^{I-1}}{(1 - F_0(z))^{I-1}} \right) \left(\frac{1}{I} (\bar{\theta} - \theta_{i,1}) - M^{cs}(z) (r - \theta_{i,1}) \right). \end{aligned} \quad (7.18)$$

The left-hand side of this expression is positive for all $\theta_{i,1}$ if $\delta \geq \delta_c$. The right-hand side is decreasing in $\theta_{i,1}$, since $\frac{1}{I} > M^{cs}(z)$. So we look for conditions under which the right-hand side is negative when $\theta_{i,1} = z$. Substituting in for $M^{cs}(z)$, we use the fact that

$$\frac{1}{I} (\bar{\theta} - z) - \left(1 - \delta + \delta \frac{1}{I} \frac{\bar{\theta} - z}{r - z} \right) (r - z) = \left(\frac{1}{I} (\bar{\theta} - z) - (r - z) \right) (1 - \delta),$$

which is negative, to establish that the right-hand side of (7.18) is negative. Thus, the critical discount factor such that (7.18) holds is less than δ_c . ■

Proof of Proposition 4: We begin by representing beliefs and payoffs under the odd-even scheme. Given \mathbf{a}_{t-2} , firm j 's belief about $\theta_{i,t}$ for $i \neq j$ at the start of period t is given by $\mu_{i,t}(a_{i,t-2}) = \lambda_H (1 - \lambda_{a_{i,t-2}}) + \lambda_L \lambda_{a_{i,t-2}}$. Given \mathbf{a}_t , firm j 's belief about $\theta_{i,t+1}$ for $i \neq j$ at the start of period $t+1$ is given by $\mu_{i,t+1}(a_{i,t}) = \lambda_{a_{i,t}}$. We let $v_i^{oe}(\hat{\theta}_{i,t}, \theta_{i,t}, \mathbf{a}_{t-2})$ be the expected discounted value of payoffs for firm i in an odd period t , given \mathbf{a}_{t-2} and at the point where firm i has just observed its cost type $\theta_{i,t}$, if it mimics type $\hat{\theta}_{i,t}$ throughout periods t and $t+1$, but expects to report truthfully from period $t+2$ onwards. This function can be defined recursively for each firm i , as follows:

$$v_i^{oe}(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{a}_{t-2}) = \mathbb{E}_{\theta_{-i,t}} \left[\begin{array}{l} \kappa_i^o(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}) \cdot (r - \theta_{i,t}) \\ + \delta \kappa_i^e(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}) (r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t}]) \\ + \delta^2 \mathbb{E}_{\theta_{i,t+2}} [v_i^{oe}(\theta_{i,t+2}, \theta_{i,t+2}; \hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}) | \theta_{i,t}] \end{array} \middle| \boldsymbol{\mu}_{-i,t}(\mathbf{a}_{-i,t-2}) \right].$$

With this definition in place, we observe that on-schedule deviations can be deterred if

$$v_i^{oe}(\theta_{i,t}, \theta_{i,t}; \mathbf{a}_{t-2}) \geq v_i^{oe}(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{a}_{t-2}) \text{ for all } \mathbf{a}_{t-2}, \hat{\theta}_{i,t}, \theta_{i,t}.$$

Note that for firm i , whether firm i 's report in period $t - 2$ was truthful or not does not affect expected payoffs or beliefs.

Off-schedule deviations can happen in either odd periods or even periods. Since announcements do not affect the prices of opponents, off-schedule deviations in announcements are not tempting. The most tempting off-schedule deviation for firm i in an odd period is to price at $r - \varepsilon$, which leads firm i to capture the whole market. Such a deviation would be followed by reversion to the worst carrot-stick equilibrium. This deviation is deterred if

$$v_i^{oe}(\theta_{i,t}, \theta_{i,t}; \mathbf{a}_{t-2}) \geq r - \theta_{i,t} + \delta \mathbb{E}_{\theta_{i,t+1}} [v_i^{cs}(\theta_{i,t+1}) | \theta_{i,t}] \text{ for all } \theta_{i,t}, \mathbf{a}_{t-2}.$$

A similar deviation must be deterred in even periods, after firm i learns $\theta_{i,t+1}$:

$$\begin{aligned} & \kappa_i^e(\theta_{i,t}, \boldsymbol{\theta}_{-i,t}) (r - \theta_{i,t+1}) + \delta \mathbb{E}_{\theta_{i,t+2}} [v_i^{oe}(\theta_{i,t+2}, \theta_{i,t+2}; \theta_{i,t}, \boldsymbol{\theta}_{-i,t}) | \theta_{i,t+1}] \\ & \geq r - \theta_{i,t+1} + \delta \mathbb{E}_{\theta_{i,t+2}} [v_i^{cs}(\theta_{i,t+2}) | \theta_{i,t+1}] \text{ for all } \theta_{i,t}, \boldsymbol{\theta}_{-i,t}, \mathbf{a}_{t-2}, \theta_{i,t+1}. \end{aligned}$$

If this latter constraint holds, then no matter what happened in the odd period (in particular, whether firm i was honest then), firm i does not wish to deviate in the subsequent even period.

We now proceed to construct the scheme described in the proposition. In our constructed odd-even scheme, prices are always equal to r on the equilibrium path. In addition, $\kappa_i^o(\boldsymbol{\theta}_t) = 1/I$ unless $\theta_{j,t} = L$ and $\boldsymbol{\theta}_{-j,t} = (H, \dots, H) = \mathbf{H}$ for some j . Further, when firm i alone reports low costs, for $j \neq i$, $\kappa_j^o(\theta_{i,t} = L, \boldsymbol{\theta}_{-i,t} = \mathbf{H}) = (I - 1)\phi + 1/I$, $\kappa_j^o(\theta_{i,t} = L, \boldsymbol{\theta}_{-i,t} = \mathbf{H}) = 1/I - \phi$,

$$\kappa_i^e(\theta_{i,t} = L, \boldsymbol{\theta}_{-i,t} = \mathbf{H}) = 1/I - (I - 1)\phi \frac{1}{\delta} \cdot \frac{r - H}{r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t} = H]}, \quad (7.19)$$

and

$$\kappa_j^e(\theta_{i,t} = L, \boldsymbol{\theta}_{-i,t} = \mathbf{H}) = 1/I + \phi \frac{1}{\delta} \cdot \frac{r - H}{r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t} = H]}. \quad (7.20)$$

We first show that the constructed odd-even scheme satisfies all on-schedule incentive constraints. Reports in odd period t only affect play in periods t and $t + 1$ on the equilibrium path. Further, a sufficient condition for the on-schedule constraints to hold is that they hold pointwise in $\boldsymbol{\theta}_{-i,t}$. Thus, the on-schedule constraints are satisfied if, for every $\hat{\theta}_{i,t}$, $\theta_{i,t}$ and $\boldsymbol{\theta}_{-i,t}$, we have

$$\begin{aligned} & \left(\kappa_i^o(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}) - \kappa_i^o(\theta_{i,t}, \boldsymbol{\theta}_{-i,t}) \right) \cdot (r - \theta_{i,t}) \\ & \leq \delta \left(\kappa_i^e(\theta_{i,t}, \boldsymbol{\theta}_{-i,t}) - \kappa_i^e(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}) \right) (r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t}]). \end{aligned} \quad (7.21)$$

Suppose that $\boldsymbol{\theta}_{-i,t}$ has no component equal to L . It is straightforward to check that (7.19) is defined so that (7.21) is exactly binding for firm i when $\hat{\theta}_{i,t} = L$, $\theta_{i,t} = H$ and $\boldsymbol{\theta}_{-i,t} = \mathbf{H}$.

Simple calculations reveal that

$$\begin{aligned}
& (r - L)(r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t} = H]) - (r - H)(r - \mathbb{E}_{\theta_{i,t+1}} [\theta_{i,t+1} | \theta_{i,t} = L]) \\
& \hspace{20em} (7.22) \\
& = (H - L)r[1 + \lambda_H - \lambda_L] - H[H - \lambda_L(H - L)] + L[H - \lambda_H(H - L)] \\
& > (H - L)H[1 + \lambda_H - \lambda_L] - H[H - \lambda_L(H - L)] + L[H - \lambda_H(H - L)] \\
& \hspace{10em} = (H - L)^2 \lambda_H > 0,
\end{aligned}$$

where the first inequality uses $r > H$ and both inequalities use imperfect persistence ($\lambda_H > 0$). Given (7.22), it is direct to verify that (7.21) is slack when $\hat{\theta}_{i,t} = H$, $\theta_{i,t} = L$, and $\boldsymbol{\theta}_{-i,t} = \mathbf{H}$. Suppose next that $\boldsymbol{\theta}_{-i,t}$ has at least two components equal to L . By construction, $\kappa_i^o(L, \boldsymbol{\theta}_{-i,t}) = \kappa_i^o(H, \boldsymbol{\theta}_{-i,t}) = \kappa_i^e(L, \boldsymbol{\theta}_{-i,t}) = \kappa_i^e(H, \boldsymbol{\theta}_{-i,t}) = 1/I$, so that (7.21) holds. Finally, suppose that $\boldsymbol{\theta}_{-i,t}$ has exactly one component equal to L . Then $\kappa_i^o(L, \boldsymbol{\theta}_{-i,t}) = 1/I > \kappa_i^o(H, \boldsymbol{\theta}_{-i,t}) = 1/I - \phi$, and $\kappa_i^e(L, \boldsymbol{\theta}_{-i,t}) = 1/I$, while $\kappa_i^e(H, \boldsymbol{\theta}_{-i,t})$ is equal to the right-hand side of (7.20). Then, it is straightforward to verify that (7.21) holds exactly for this $\boldsymbol{\theta}_{-i,t}$ when $\theta_{i,t} = H$ and $\hat{\theta}_{i,t} = L$. Further, using (7.22), it is direct to verify that the constraint is slack for this $\boldsymbol{\theta}_{-i,t}$ when $\hat{\theta}_{i,t} = H$ and $\theta_{i,t} = L$. Thus, all on-schedule constraints hold for this scheme.

Next, we represent payoffs for the constructed odd-even scheme, and we verify that this scheme improves expected profits relative to the best rigid-pricing scheme, wherein all market shares are equal to $1/I$ and prices equal to r on the equilibrium path. To simplify notation, consider $i = 1$ and take t odd. First, observe that unless $\boldsymbol{\theta}_t = (L, \mathbf{H})$ or $\boldsymbol{\theta}_t = (H, \boldsymbol{\theta}_{-1,t})$ where $\boldsymbol{\theta}_{-1,t}$ has exactly one component equal to L , all market shares in t and $t+1$ are equal to $1/I$, and there is then no difference between the constructed odd-even scheme and the best rigid-pricing scheme. If $\boldsymbol{\theta}_t = (L, \mathbf{H})$, firm i expects profits over the next two periods equal to

$$\begin{aligned}
& \kappa_1^o(L, \mathbf{H}) \cdot (r - L) + \delta \kappa_1^e(L, \mathbf{H}) (r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | \theta_{1,t} = L]) \\
& = (r - L)/I + \delta (r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | L]) / I \\
& \quad + (I - 1)\phi \cdot \left((r - L) - \frac{(r - H)(r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | \theta_{1,t} = L])}{r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | \theta_{1,t} = H]} \right) \\
& > (r - L)/I + \delta (r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | L]) / I,
\end{aligned}$$

where the final term gives the two-period expected profit under the best rigid-pricing scheme, and the inequality follows from (7.22). Finally, consider the case where exactly one component of $\boldsymbol{\theta}_{-1,t}$ is equal to L , and $\theta_{1,t} = H$. Then, firm 1 expects profits over the next two periods equal to

$$(1/I - \phi) \cdot (r - H) + \delta \left(\frac{1}{I} + \phi \cdot \frac{1}{\delta} \frac{r - H}{r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | \theta_{1,t} = H]} \right) (r - \mathbb{E}_{\theta_{1,t+1}} [\theta_{1,t+1} | \theta_{1,t} = H]).$$

Relative to a scheme with all market shares equal to $1/I$, the difference is $-\phi \cdot (r - H) + \phi \cdot (r - H) = 0$. Thus, for all $\boldsymbol{\theta}_t$, firm i 's expected profits over t and $t + 1$ are sometimes higher and never lower under the constructed odd-even scheme.

Finally, we observe that the off-schedule constraints hold with slack when $\delta > \delta_c$ and all market shares are equal to $1/I$, by the equivalence of the best rigid-price scheme and the odd-even scheme in that case. Since equilibrium payoffs in the odd-even scheme are continuous in market shares, if $\delta > \delta_c$, then for $\phi > 0$ small enough, the off-schedule constraints will be satisfied in the odd-even scheme with partial pooling as well. ■

Proof of Proposition 7: We established above that the carrot-stick belief threat punishment does not entail the use of weakly dominated strategies. Let $\xi(\theta_{i,1})$ be the present discounted value a deviant firm expects in the carrot-stick belief threat punishment. For $\theta_{i,1} \leq \underline{\theta} + 2\varepsilon$, this value is approximately $\xi(\theta_{i,1}) = \chi\delta\frac{1}{I}(r - \theta_{i,1}) + (1 - \chi)\delta\xi(\theta_{i,1})$, or $\xi(\theta_{i,1}) = \frac{\chi\delta}{1 - (1 - \chi)\delta}\frac{1}{I}(r - \theta_{i,1})$. Higher types price above $\underline{\theta} + 2\varepsilon$ and thus receive $\xi(\theta_{i,1}) = 0$. For any $\theta_{i,1}$, firm i does not gain by deviating from pricing at r in each period, if the following off-schedule constraint holds:

$$\frac{r - \theta_{i,1}}{I} \frac{1}{1 - \delta} \geq r - \theta_{i,1} + \delta\xi(\theta_{i,1}). \quad (7.23)$$

Rewriting, we obtain $\frac{(I-1)(r-\theta_{i,1})+I\delta\xi(\theta_{i,1})}{I(r-\theta_{i,1})+I\delta\xi(\theta_{i,1})} \leq \delta$. The left-hand side is increasing in $\xi(\theta_{i,1})$. Thus, for χ sufficiently small, $\xi(\theta_{i,1})$ is arbitrarily close to zero for all $\theta_{i,1}$, and we are sure to satisfy (7.23) if $\delta > (I - 1)/I$. ■

8. References

- Abreu, D. and F. Gul, “Bargaining and Reputation,” *Econometrica*, 68(1), January 2000, 85-117.
- Abreu, D. and D. Pearce, “A Behavioral Model of Bargaining with Endogenous Types,” working paper, Princeton University, November 2003.
- Abreu, D., D. Pearce and E. Stacchetti, “Optimal Cartel Equilibria with Imperfect Monitoring,” *Journal of Economic Theory*, 39(1), June 1986, pp. 251-69.
- Abreu, D., D. Pearce and E. Stacchetti, “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58.5, September 1990, pp. 1041-63.
- Aoyagi, M., “Bid Rotation and Collusion in Repeated Auctions,” *Journal of Economic Theory*, 112(1), September 2003, 79-105.
- Athey, S., K. Bagwell, and C. Sanchirico, “Collusion and Price Rigidity,” *The Review of Economic Studies*, 71(2), April 2004, 317-49.
- Athey, S. and K. Bagwell, “Optimal Collusion with Private Information,” *RAND Journal of Economics*, 32(3), Autumn 2001, 428-65.
- Athey, S. and K. Bagwell, “Dynamic Auctions with Persistent Private Information,” Mimeo, Stanford University, 2004.
- Battaglini, M., “Long-Term Contracting with Markovian Consumers,” *American Economic Review* 95(3), June 2005, 637-58.
- Chakrabarti, S., “Information Revelation and Collusion in Oligopolies with Unknown Costs,” Working Paper, Cornell University, 2001.
- Cole, H. and N. Kocherlakota, “Dynamic Games with Hidden Actions and Hidden States,” *Journal of Economic Theory* 98, 2001, 114-26.

- Connor, J. M., *Global Price Fixing: Our Customers Are the Enemy*, Boston, Mass: Kluwer Academic Press, 2001.
- Fernandez, A. and C. Phelan, "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, 2000, 91, 223-47.
- Fershtman, C. and A. Pakes, "A Dynamic Oligopoly with Collusion and Price Wars," *RAND Journal of Economics*, 31(2), 2000, 207-36.
- Fershtman, C. and A. Pakes, "Finite State Dynamic Games with Asymmetric Information: A Computational Framework," Mimeo, Harvard, July 2004.
- Fudenberg, D., D. J. Levine and E. Maskin, "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5) September 1994, 997-1039.
- Fudenberg, D. and J. Tirole, *Game Theory*, Cambridge, Mass. and London: MIT Press, 1991.
- Green, E. J. and R. H. Porter, "Noncooperative Collusion under Imperfect Price Information," *Econometrica*, 52(1) January 1984, 87-100.
- Kandori, M. and I. Obara, "Efficiency in Repeated Games Revisited: the Role of Private Strategies," *Econometrica*, 74, March 2006, 499-519.
- Kennan, J., "Repeated Bargaining with Persistent Private Information," *Review of Economic Studies*, 68, October 2001, 719-55.
- Kreps, D. and R. Wilson, "Reputation and Imperfect Information," *Journal of Economic Theory*, 1982, 27, 253-79.
- LaCasse, C., "Price-Fixing with Incomplete Information," Mimeo, University of Alberta, 1999.
- Laffont, J.-J. and J. Tirole, "The Dynamics of Incentive Contracts," *Econometrica*, 1988, 56, 1153-75.
- Mailath, G., "Simultaneous Signaling in an Oligopoly Model," *The Quarterly Journal of Economics*, May 1989, 104, 417-27.
- Maskin, E. and J. G. Riley, "Asymmetric Auctions," *Review of Economic Studies*, July 2000, 67(3), 413-438.
- McAfee, P. and J. McMillan, "Bidding Rings," *American Economic Review*, 82(3), June 1992, 579-599.
- McMillan, J., "Dango: Japan's Price-Fixing Conspiracies," *Economics and Politics*, 3, November, 201-218.
- Milgrom, P. and J. Roberts, "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27, 280-312.
- Milgrom, P. and I. Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70(2), March 2002, 583-601.
- Riley, J. G. and W. F. Samuelson, "Optimal Auctions," *American Economic Review*, 71(3), 381-92.
- Roos, N., "Examining Models of Collusion: the Market for Lysine," Mimeo, University of Sydney, 2003.
- Skryzpacz, A. and H. Hopenhayn, "Tacit Collusion in Repeated Auctions," *Journal of Economic Theory*, 114(1), January 2004, 153-69.
- Watson, J., "Starting Small and Commitment," *Games and Economic Behavior*, 38, 2002, 176-99.