

Collusion and Price Rigidity

SUSAN ATHEY

Stanford University and NBER

KYLE BAGWELL

Columbia University and NBER

and

CHRIS SANCHIRICO

University of Pennsylvania

First version received December 2000; final version accepted November 2002 (Eds.)

We consider an infinitely repeated Bertrand game, in which prices are publicly observed and each firm receives a privately observed, i.i.d. cost shock in each period. We focus on symmetric perfect public equilibria, wherein any “punishments” are borne equally by all firms. We identify a tradeoff that is associated with collusive pricing schemes in which the price to be charged by each firm is strictly increasing in its cost level: such “fully sorting” schemes offer efficiency benefits, as they ensure that the lowest-cost firm makes the current sale, but they also imply an informational cost (distorted pricing and/or equilibrium-path price wars), since a higher-cost firm must be deterred from mimicking a lower-cost firm by charging a lower price. A rigid-pricing scheme, where a firm’s collusive price is independent of its current cost position, sacrifices efficiency benefits but also diminishes the informational cost. For a wide range of settings, the optimal symmetric collusive scheme requires (i) the absence of equilibrium-path price wars and (ii) a rigid price. If firms are sufficiently impatient, however, the rigid-pricing scheme cannot be enforced, and the collusive price of lower-cost firms may be distorted downward in order to diminish the incentive to cheat. When the model is modified to include i.i.d. public demand shocks, the downward pricing distortion that accompanies a firm’s lower-cost realization may occur only when current demand is high.

1. INTRODUCTION

In the standard model of collusion, symmetric firms interact in an infinitely repeated Bertrand game in which past prices are publicly observed. The standard model offers a number of insights, but it presumes an unchanging market environment. This is an important limitation, since the scope for testing a theory of collusion is greater when the theory offers predictions concerning the manner in which collusive prices vary with underlying market conditions.

This limitation is partially addressed in two celebrated extensions. Rotemberg and Saloner (1986) introduce publicly observed demand shocks that are i.i.d. over time.¹ When the demand shock is large, the incentive to cheat (undercut the collusive price) is acute, and collusion becomes more difficult to enforce. Rather than forego collusive activity altogether, firms then reduce the collusive price and thereby diminish the incentive to cheat. Thus, markups are countercyclical. Like the standard model, their model does not predict that actual “price wars” occur on the equilibrium path; rather, the success of collusion varies along the equilibrium path with the demand shocks that are encountered.

Following the seminal work of Stigler (1964), a second literature stresses that a firm may be unable to monitor perfectly the behaviour of its rivals. Green and Porter (1984) explore this possibility in an infinitely repeated Cournot model. They assume that a firm cannot observe

1. Haltiwanger and Harrington (1991) and Bagwell and Staiger (1997) consider further extensions.

the output choices of rivals but that all firms observe a public signal (the market price) that is influenced both by output choices and an unobserved demand shock.² A colluding firm that witnesses a low market price then faces an inference problem, as it is unclear whether the low-price outcome arose as a consequence of a bad demand shock or a secret output expansion by a rival. The Green and Porter (1984) model thus represents collusion in the context of a repeated moral-hazard (hidden-action) model, and a central feature of their analysis is that wars occur along the equilibrium path following bad demand shocks.

In this paper, we explore a third extension of the standard collusion model. We consider an infinitely repeated Bertrand game, in which each firm is privately informed of its unit cost level in each period, there is a continuum of possible costs, and the cost realization is i.i.d. across firms and time. Current price selections (but not cost realizations) are publicly observed before the beginning of the next period. We thus represent collusion in the context of a repeated adverse-selection (hidden-information) model with publicly observed actions (prices).³

Our model is well designed to contribute to a long-standing issue in Industrial Organization concerning the relationship between collusion and price rigidity in the presence of cost shocks. Empirical studies by Mills (1927), Means (1935) and Carlton (1986, 1989) conclude that prices are more rigid in concentrated industries, suggesting that collusion is associated with a greater tendency toward price rigidity. In addition, over the past several decades, antitrust enforcement has uncovered numerous price-fixing agreements in which firms co-ordinate on a particular price and enforce stable market shares over time. In these examples, colluding firms adjust price occasionally in response to changes in overall market conditions, but they sacrifice the efficiency advantages that could be gained by allowing a firm with a temporary cost advantage to serve a larger market share.⁴

At the same time, the Industrial Organization literature has not provided a satisfactory theory that links price rigidity with collusion. The best known theory is the “kinked demand curve” theory offered by Sweezy (1939) and Hall and Hitch (1939). As Scherer (1980) and Tirole (1988) discuss, however, this theory has important shortcomings. Industrial Organization economists have thus gravitated toward the more informal view that price rigidity is appealing to collusive firms, because a rigid-price collusive scheme prevents mistrust and reduces the risk of a price war. Carlton (1989) explains:

The property of the kinked demand curve theory that price is unresponsive to some cost fluctuations is preserved in most discussions of oligopoly theory whether or not based on the kinked demand curve. The reasoning is that in oligopolies prices fluctuate less in response to cost changes (especially small ones) than they would otherwise in order not to disturb existing oligopolistic discipline. Anytime a price change occurs in an oligopoly, there is a risk that a price war could break out. Hence, firms are reluctant to change price (Carlton, 1989, pp. 914–915).

We develop here a rigorous evaluation of this informal reasoning. Focusing on the *private* cost fluctuations that firms experience, we explore the extent to which “mistrust” limits colluding

2. While the Green and Porter (1984) model is developed in the context of Cournot competition, the main insights can be captured in a repeated Bertrand setting, as Tirole (1988) shows. The Green–Porter (1984) model is further extended by Porter (1983), Abreu, Pearce and Stacchetti (1986, 1990) and Fudenberg, Levine and Maskin (1994).

3. Athey and Bagwell (2001), a sequel to the present effort, explores a related model. We discuss this paper in what follows. Our model is also related to recent work that extends the Green and Porter (1984) model to allow for privately observed demand signals. Compte (1998) and Kandori and Matshushima (1998) suppose that firms publicly choose “messages” after privately observing their respective demand signals. In our setting, firms are privately informed as to their respective costs and the public action is a (payoff-relevant) price choice.

4. For example, a European cartel of cartonboard producers set stable market shares, but adjusted the fixed price every 6 months in response to changing demand conditions (European Commission, 1994). Additional examples are discussed in Business Week (1975, 1998) and Scherer (1980).

firms' ability to respond to their respective cost positions. The costs of mistrust are formalized in terms of the price wars and pricing distortions that are required to dissuade firms from misrepresenting their private information.

We begin with a formal analysis of the static Bertrand game with inelastic demand and private cost information.⁵ This game constitutes the stage game of our repeated-game model. In the unique Nash equilibrium of the stage game, the symmetric pricing strategy is strictly increasing in the firm's cost level. An advantage of Nash pricing is that sales in the current period are allocated to the firm with the lowest cost. This is the *efficiency benefit* of a "fully sorting" (*i.e.* strictly increasing) pricing scheme. Of course, from the firms' perspective, Nash pricing also has an important limitation: sales are allocated at low prices.

We turn next to the repeated-game model and explore whether firms can then support better-than-Nash profits. We focus on the class of *symmetric perfect public equilibria (SPPE)*. An SPPE collusive scheme at a given point in time can be described by (i) a price for each cost type and (ii) an associated equilibrium continuation value for each vector of current prices, where the continuation value is symmetric across firms. In an SPPE, therefore, colluding firms move symmetrically through any co-operative or price-war phases.

We observe that a collusive scheme must satisfy two kinds of incentive constraints. First, for every firm and cost level, the short-term gain from cheating with an *off-schedule deviation* (*i.e.* with a price that is not assigned to any cost type and that thus represents a clear deviation) must be unattractive, in view of the (off-the-equilibrium-path) price war that such a deviation would imply. As is usual in repeated-game treatments of collusion, this constraint is sure to be met if firms are sufficiently patient. Second, the proposed conduct must also be such that no firm is ever attracted to an *on-schedule deviation*, whereby a firm of a given cost type misrepresents its private information and selects a price intended for a different cost type.

To characterize the optimal SPPE, we build on the dynamic-programming techniques put forth by Abreu *et al.* (1986, 1990) (APS) and Fudenberg *et al.* (1994). We draw an analogy between our repeated hidden-information game and the static mechanism design literature, in which the on-schedule incentive constraint is analogous to the standard incentive-compatibility constraint, the off-schedule incentive constraint serves as a counterpart to the traditional participation constraint, and the continuation values play the role of "transfers". However, unlike a standard mechanism design problem in which transfers are unrestricted, the set of feasible continuation values is limited and endogenously determined. In particular, we may associate a price war with a transfer that is borne symmetrically by all firms.

We break our analysis of optimal SPPE into two parts. We suppose first that firms are patient, so that the off-schedule constraint is met. The on-schedule constraint then captures the *informational costs* of collusion that confront privately informed firms. The central problem is that the scheme must be constructed so that a higher-cost firm does not have an incentive to misrepresent its costs as lower, by selecting a lower price and thereby securing for itself a higher expected market share. In an SPPE, the informational costs of collusion may be manifested in two ways. First, the prices of lower-cost firms may be distorted to sub-monopoly levels. This is a potentially effective means of eliciting truthful cost information, since higher-cost firms find lower prices less appealing. Second, following the selection of lower prices, the collusive scheme may sometimes call for a future equilibrium-path price war. The current-period benefit of a lower price then may be of sufficient magnitude to compensate for the future cost of a price war only if the firm truly has lower costs in the current period.

5. The Bertrand model assumes homogeneous goods, which is a common feature of collusive markets (Hay and Kelley (1974), Scherer (1980, p. 203)). Collusion is also often associated with inelastic demand (Eckbo, 1976).

A rich array of collusive schemes fit within the SPPE category. One possibility is that firms incur the informational costs of collusion purely in terms of distorted pricing. An example is the *Nash-pricing scheme*, in which firms repeatedly play the Nash equilibrium of the static game. Another possibility is that firms initially achieve full sorting and adopt higher-than-Nash prices. In this case, some of the informational costs of collusion must be reflected in the future cost of a price war: such a scheme satisfies the on-schedule constraint only if equilibrium-path wars sometimes follow the selection of prices assigned to lower-cost firms. A further possibility is that firms may neutralize the informational costs of collusion altogether, by adopting a *rigid-pricing scheme*, in which each firm selects the same price in each period, whatever its current cost position. The downside of the rigid-pricing scheme is that it sacrifices efficiency benefits: one firm may have lower costs than its rivals, and yet the firms share the market. These schemes highlight the central tradeoff between efficiency benefits and informational costs that colluding firms must reconcile. More generally, collusive pricing schemes may be strictly increasing over some intervals of costs and rigid over other regions, with wars following some pricing realizations.

Our first main finding is that firms fare poorly under any SPPE collusive scheme that insists upon full sorting. In fact, considering the entire set of fully sorting SPPE, we find that firms can do no better than the Nash-pricing scheme. We next consider the full class of SPPE collusion schemes and report a second main finding: if firms are sufficiently patient, then an optimal SPPE collusive scheme can be achieved without recourse to equilibrium-path price wars (*i.e.* with stationary strategies). This finding contrasts interestingly with the predictions of the Green and Porter (1984) model.

Armed with these findings, we add some additional structure and provide a characterization of the optimal SPPE collusive scheme. When firms are patient and the distribution of cost types is log-concave, we establish a third main finding: optimal SPPE collusion is characterized by a rigid-pricing scheme, in which firms select the same price (namely, the reservation price of consumers) in each period, whatever their cost levels. We thus offer an equilibrium interpretation of the association between price rigidity and collusion described above.

We then turn to the second part of our analysis and consider impatient firms. We show that impatience creates an additional disadvantage to price wars: a scheme with high prices today sustained by wars in the future makes a deviation especially profitable today, while simultaneously reducing the value of co-operation in the future. Our second (no-wars) finding therefore continues to hold when firms are impatient. Next, we observe that the off-schedule constraint is particularly demanding for lower-cost types. Intuitively, when a firm draws a lower-cost type, the temptation to cheat and undercut the assigned price is severe, since the resulting market-share gain is then especially appealing. For impatient firms, a collusive scheme thus must ensure that lower-cost types receive sufficient market share and select sufficiently low prices in equilibrium, so that the gains from cheating are not too great.

This logic is reminiscent of the argument made by Rotemberg and Saloner (1986), although here it is private cost shocks (as opposed to public demand shocks) that necessitate modification of the collusive scheme. We confirm this logic with our fourth main finding: if firms are not sufficiently patient to enforce the rigid-pricing scheme, they may still support a partially rigid collusive scheme, in which the price of lower-cost types is reduced in order to mitigate the incentive to cheat. This finding suggests that symmetric collusion between impatient firms may be marked by occasional (and perhaps substantial) price reductions by individual firms. These departures occur when a firm receives a favourable cost shock, and they represent a permitted "escape clause" (*i.e.* an opportunity to cut prices and increase market share without triggering retaliation) within the collusive scheme. More generally, we establish conditions for impatient firms under which, if better-than-Nash profits can be achieved, then optimal SPPE collusion is

characterized by a stationary pricing scheme in which prices are rigid over intervals of costs (*i.e.* the optimal pricing scheme is a weakly increasing step function).

To further develop the relationship between our theory and that of Rotemberg and Saloner (1986), we next extend our model to include public i.i.d. demand shocks. The off-schedule constraint is then most difficult to satisfy when market demand is high and a firm's cost shock is low. We thus offer a fifth main finding: in an extended model with public i.i.d. demand shocks, if firms are not sufficiently patient to enforce the rigid-pricing scheme, optimal SPPE collusion may be characterized by a stationary pricing scheme, in which an individual firm charges a lower price in high-public-demand and low-private-cost states. The Rotemberg and Saloner (1986) prediction of countercyclical pricing is thus robust to private cost information. Our model has the further prediction that prices are more variable when today's demand is high.

Throughout, we restrict attention to symmetric schemes. This restriction is important. Asymmetric schemes allow one firm to enjoy a more profitable continuation value than another. Such schemes thus facilitate transfers *from* one firm *to* another. Athey and Bagwell (2001) analyse optimal asymmetric schemes. In the repeated game that they feature, the stage game allows that firms communicate cost information and make market-share proposals before setting prices. Under the assumption that unit costs are high or low, they construct an *asymmetric perfect public equilibrium (APPE)* that delivers first-best profits when the patience of firms exceeds a critical (finite) level.⁶ In this construction, when a firm announces that its costs are high, it is favoured with greater market share in future periods. The resulting equilibrium play is highly non-stationary.

The present paper differs from that of Athey and Bagwell (2001), both in terms of the models employed and the predictions derived. The present paper considers a model with a continuum of types and studies optimal symmetric schemes, while allowing that firms may be impatient and that demand may be volatile. In contrast, Athey and Bagwell (2001) use a two-type model and study optimal asymmetric schemes, with the central (first-best) finding applying when firms are sufficiently patient. At a predictive level, the present paper provides a formal interpretation for the traditional view that standard collusion entails fixed prices and stable market shares over time. Athey and Bagwell (2001) show that sophisticated collusion, in which firms track and reward individual firm behaviour over time, communicate and allocate market shares, is characterized by market shares that are unstable over time. While we do not propose a theory of how firms co-ordinate upon an equilibrium, the two papers together suggest an intriguing possibility: stationary SPPE are appealing simple and may be descriptive of less formal (and perhaps tacit) collusive ventures, while optimal APPE are quite sophisticated and may be most plausible in the presence of a small number of well-organized conspirators that interact frequently and communicate explicitly. Finally, SPPE may be the only available option if firms cannot observe individual firm behaviour. This occurs, *e.g.* in procurement auctions with more than two bidders, if the winning bid—but not the name of the winner—is announced.

As just implied, a special case of our model is a repeated procurement auction. We may thus relate our findings to those developed by McAfee and McMillan (1992), in their analysis of bidding rings. They describe evidence that fixed-price schemes (*i.e.* “identical bidding”) are widely used. In a static model, they show that a fixed price is the optimal strategy for bidding cartels in first-price auctions for a single object, when the cartels are “weak” (*i.e.* firms are unable to make transfers). In our analysis of the optimal SPPE, we generalize the weak-cartel model, since the static mechanism we analyse is directly derived from a repeated game and allows for a restricted class of transfers (corresponding to symmetric price wars). Our rigid-pricing finding

6. For a family of repeated private-information games, Fudenberg *et al.* (1994) show that first-best payoffs can be reached in the limit as players become infinitely patient.

thus provides additional theoretical support for the practice of identical bidding. We also extend the analysis to incorporate impatient firms.

We describe the static and repeated games in Sections 2 and 3, respectively. The latter is related to the mechanism-design approach in Section 4. We present our findings for SPPE among patient firms in Section 5, and Section 6 considers impatient firms. These sections close with a discussion of an extended model with downward-sloping demand. Section 7 concludes.

2. THE STATIC GAME

We begin with a static game of Bertrand competition in which firms possess private information. This game illustrates the immediate tradeoffs that confront firms in determining their pricing policies and serves as a foundation on which our subsequent dynamic analysis builds.

We posit n *ex ante* identical firms that engage in Bertrand competition for sales in a homogeneous-good market. Following Spulber (1995), we modify the standard Bertrand model with the assumption that each firm is privately informed as to its unit cost level. Cost levels are i.i.d. across firms. In particular, firm i 's "type" θ_i is determined by an independent draw from the support $[\underline{\theta}, \bar{\theta}]$ according to the commonly known distribution function $F(\theta)$. We assume that the corresponding density $f(\theta) \equiv F'(\theta)$ is strictly positive on $[\underline{\theta}, \bar{\theta}]$. After the firms learn their respective cost types, they simultaneously choose prices. Let $\rho_i \in \mathbb{R}_+$ denote the price chosen by firm i , with $\boldsymbol{\rho} \equiv (\rho_1, \dots, \rho_n)$ then representing the associated price profile. We assume a unit mass of identical consumers, each of whom has an inelastic demand for one unit up to some reservation price r , where $r \geq \bar{\theta}$.

A price strategy for firm i is a function $p_i(\theta_i)$ mapping from the set of cost types, $[\underline{\theta}, \bar{\theta}]$, to the set of possible prices, \mathbb{R}_+ . For simplicity, the function p_i is assumed continuously differentiable, except perhaps at a finite number of points (so as to allow for schedules with jumps). A price strategy profile is thus a vector $\mathbf{p}(\boldsymbol{\theta}) \equiv (p_i(\theta_i), \mathbf{p}_{-i}(\boldsymbol{\theta}_{-i}))$, where $\boldsymbol{\theta} \equiv (\theta_i, \boldsymbol{\theta}_{-i})$ is the vector of cost types and $\mathbf{p}_{-i}(\boldsymbol{\theta}_{-i})$ is the profile of rival price strategies. Each firm chooses its price strategy with the goal of maximizing its expected profit, given its cost type. To represent a firm's expected profit, we require two further definitions. First, we define $\pi(\rho, \theta) \equiv \rho - \theta$ as the profit that a firm receives when it sets the price ρ and has cost type θ and "wins" the entire unit mass of consumers. Second, we specify a Bertrand market-share-allocation function, $m_i(\boldsymbol{\rho})$, that indicates firm i 's market share when the vector of realized prices is $\boldsymbol{\rho}$. This function allocates consumers evenly among firms that select the lowest price in the market.

We may now represent firm i 's *interim* profit, which is the expected profit for firm i when it has cost type θ_i , selects the price ρ_i and anticipates that rival prices will be determined by the rival pricing strategy profile, $\mathbf{p}_{-i}(\boldsymbol{\theta}_{-i})$. With $\bar{m}(\rho_i; \mathbf{p}_{-i}) \equiv E_{\theta_{-i}}[m_i(\rho_i, \mathbf{p}_{-i}(\boldsymbol{\theta}_{-i}))]$, firm i 's interim profit function may be written as

$$\bar{\pi}(\rho_i, \theta_i; \mathbf{p}_{-i}) \equiv \pi(\rho_i, \theta_i) \bar{m}(\rho_i; \mathbf{p}_{-i}). \quad (2.1)$$

When firms adopt a symmetric pricing strategy, $p(\cdot)$, we use the notation $\bar{m}(\rho_i; p)$ and $\bar{\pi}(\rho_i, \theta_i; p)$.

We now describe the essential tradeoff that confronts a firm when it sets its price. As illustrated in (2.1), when a firm of a given type considers whether to lower its price, it must weigh the effect of an increase in the chance of winning (through $\bar{m}_i(\rho_i; p)$) against the direct effect of the price reduction on profit-if-win (through $\pi(\rho_i, \theta_i)$). An important feature of the model is that different types feel differently about this tradeoff. In particular, the interim profit function satisfies a *single-crossing property*: lower types find the expected-market-share increase that accompanies a price reduction relatively more appealing than do higher types, since lower

types have lower total costs and thus higher profit-if-win. The single-crossing property implies that higher-cost firms always select higher prices (*i.e.* $p_i(\theta_i)$ is non-decreasing).

The stage game may be analysed using standard techniques from the auction literature.⁷

Proposition 1. *The static game has a unique Nash equilibrium. Moreover, the unique Nash equilibrium pricing strategy (1) is symmetric: $p_i \equiv p^e, \forall i$; (2) is continuously differentiable and strictly increasing over $\theta \in (\underline{\theta}, \bar{\theta})$; (3) is below the monopoly price: $p^e(\theta) < r, \forall \theta < \bar{\theta}$; and (4) yields positive interim profit for all types but the highest θ , who never wins and whose price $p^e(\bar{\theta}) = \bar{\theta}$ would yield zero profit even if it did win.*

Notice that the symmetric equilibrium pricing strategy p^e is continuous and strictly increasing. Further, the price always falls at or below $\bar{\theta}$, no matter how high is r . For future reference, we let π^{NE} denote a firm’s expected profit in the Nash equilibrium of the static game.

3. THE REPEATED GAME

In this section, we define the repeated game. We also present a “Factored Program” and establish a relationship between solutions to this program and optimal SPPE.

3.1. The model

Imagine that firms meet period after period to play the stage game described in the previous section, each with the objective of maximizing its expected discounted stream of profit. Assume further that, upon entering a period of play, a firm observes only the history of: (i) its own cost draws, (ii) its own pricing schedules, and (iii) the realized prices of its rivals. Thus, we assume that a firm does not observe rival types or rival price schedules.

Formally, we describe the repeated game in the following terms. A full path of play is an infinite sequence $\{\theta^t, \mathbf{p}^t\}$, with a given pair in the sequence representing a vector of types and price schedules at date t . The infinite sequence implies a public history of realized price vectors, $\{\rho^t\}$, and pathwise payoffs for firm i may be thus defined as

$$u_i(\{\theta^t, \mathbf{p}^t\}) = \sum_{t=1}^{\infty} \delta^{t-1} \pi(\rho_i^t, \theta_i^t) m_i(\rho^t).$$

At the close of period τ , firm i possesses an information set, which may be written as $h_i = \{\theta_i^t, p_i^t, \rho_{-i}^t\}_{t=1}^{\tau}$. (The null history is the firm’s information set at the beginning of the first period.) A (pure) strategy for firm i , $s_i(h_i)(\theta_i)$, associates a price schedule with each information set h_i . Each strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ induces a probability distribution over play paths $\{\theta^t, \mathbf{p}^t\}$ in the usual manner. The expected discounted payoff from \mathbf{s} is thus the expectation $\bar{u}_i(\mathbf{s}) = E[u_i(\{\theta^t, \mathbf{p}^t\})]$ taken with respect to this measure on play paths.

3.2. A dynamic programming approach

Under our assumptions, firm types are i.i.d. across time (and firms), and so the repeated game has a recursive structure. It is therefore natural to follow Fudenberg *et al.* (1994) (FLM) and employ a recursive solution concept: we focus upon sequential equilibria in which each firm’s strategy conditions only upon the publicly observed history of realized prices. Such strategies are called *public strategies* and such sequential equilibria are called *perfect public equilibria (PPE)*. A public strategy may thus be abbreviated as a map from finite public histories $\{\rho^t\}_{t=1}^{\tau}$ to

7. See also Spulber (1995), who establishes Proposition 1 for general demand functions.

price schedules. We further restrict attention to SPPE, whereby following every public history, firms adopt symmetric price schedules: $s_i(\rho^1, \dots, \rho^\tau) = s_j(\rho^1, \dots, \rho^\tau), \forall i, j, \tau, \rho^1, \dots, \rho^\tau$. Symmetry means that all firms suffer future punishments and rewards together on an industry-wide basis.

Drawing on the work of Abreu *et al.* (1986, 1990), we apply the tools of dynamic programming to this setting. Let $\mathcal{V}_s \subset \mathbb{R}$ denote the set of SPPE continuation values and write $\underline{\mathcal{V}}_s \equiv \inf \mathcal{V}_s$ and $\overline{\mathcal{V}}_s \equiv \sup \mathcal{V}_s$. Note, initially, that with a continuum of possible pricing strategies there is no *a priori* basis from which to argue that either $\overline{\mathcal{V}}_s \in \mathcal{V}_s$ or $\underline{\mathcal{V}}_s \in \mathcal{V}_s$;⁸ if $\overline{\mathcal{V}}_s \in \mathcal{V}_s$, then we say that $\overline{\mathcal{V}}_s$ is an *optimal* SPPE value. Following APS, any symmetric public strategy profile $\mathbf{s} = (s, \dots, s)$ can be *factored* into a first-period price schedule p and a *continuation payoff function* $v : \mathbb{R}_+^n \rightarrow \mathbb{R}$. The continuation payoff function describes the repeated-game payoff $v(\rho)$ enjoyed by all firms from the perspective of period 2 onward after each first-period price realization $\rho = (\rho_1, \dots, \rho_n) \in \mathbb{R}_+^n$. We define $\bar{v}(\rho_i; \mathbf{p}_{-i}) \equiv E_{\theta_{-i}}[v(\rho_i, \mathbf{p}_{-i}(\theta_{-i}))]$ as the expected continuation payoff when a firm selects ρ_i and expects other firms to price according to \mathbf{p}_{-i} . In view of our symmetry restriction, we may simply write $\bar{v}(\rho_i; p)$, and similarly each firm's expected payoff from \mathbf{s} can be written as $E_{\theta_i}[\bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p)]$.

We now consider the Factored Program, in which we choose factorizations directly in order to maximize a firm's expected payoff, subject to (i) the feasibility constraint that the continuation payoff rests always in the SPPE value set and (ii) the incentive constraint that a firm cannot gain by deviating to an alternative pricing schedule (given the continuation payoff function and under the assumption that other firms follow the pricing schedule):

The Factored Program. Choose price schedule p and continuation payoff function v to maximize

$$\begin{aligned} & E_{\theta_i}[\bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p)] \\ & \text{subject to: } \forall \rho \in \mathbb{R}_+^n, v(\rho) \in \mathcal{V}_s, \text{ and} \\ & \forall \tilde{p}, E_{\theta_i}[\bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p)] \geq E_{\theta_i}[\bar{\pi}(\tilde{p}(\theta_i), \theta_i; p) + \delta \bar{v}(\tilde{p}(\theta_i); p)]. \end{aligned}$$

Lemma 1. Consider any symmetric public strategy profile $\mathbf{s}^* = (s^*, \dots, s^*)$ with the corresponding factorization (p^*, v^*) . Then, \mathbf{s}^* is an optimal SPPE if and only if (p^*, v^*) solves the Factored Program.

This lemma, standard in the literature, establishes that we may characterize the set of optimal SPPE by solving the Factored Program.

It is also possible to analyse the set of equilibria. If firms can randomize over continuation equilibria (*e.g.* using a public randomization device), then the set of SPPE values is convex and is thus fully characterized when the best and worst SPPE are found, where the worst equilibrium value is attained by minimizing rather than maximizing the objective in the Factored Program. In Section 6, we analyse the worst SPPE; until then, we focus on optimal SPPE.

3.3. The Interim Program for games with private information

We next reformulate the Factored Program so that it can be analysed using existing tools from the (static) mechanism design literature. We begin by observing that a SPPE in a repeated game with private information must be immune to two kinds of current-period deviations.

8. We could attempt to prove compactness of the set in general terms. Instead, we establish compactness in the process of characterizing best and worst SPPE values.

A firm deviates “off-schedule” when it chooses a price not specified for *any* cost realization (*i.e.* a price not in the range of p). When a firm prices in this manner, it has unambiguously deviated: the deviant price is “off the equilibrium path”. As is standard, such deviations are most effectively deterred when firms use the worst available punishment as a threat. In contrast, a firm deviates “on-schedule” when it chooses a price that is assigned under p to *some* cost level, but not its own. For example, a firm may be tempted to choose the lower price assigned to a lower-cost realization in order to increase its chances of winning the market. Importantly, a rival can not be sure that the deviating firm was not truly of the cost type that it is imitating: the deviant price is “on the equilibrium path”. An on-schedule deviation of this kind can be prevented if the collusive scheme imposes a punishment when low prices are chosen. But such a punishment would be costly, since it would occur along the equilibrium path of play, whenever firms actually realize low costs.

With this distinction at hand, we take the constraints of the Factored Program, put them in interim form and parse them into two groups, and rewrite this program as follows:

The Interim Program. Choose price schedule p and continuation payoff function v to maximize

$$E_{\theta_i} [\bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p)]$$

subject to:

Off-Schedule Constraints: $\forall \rho' \notin p([\underline{\theta}, \bar{\theta}]),$

(IC-off1) $\forall \theta_{-i}, \quad v(\rho', \mathbf{p}_{-i}(\theta_{-i})) \in \mathcal{V}_s$

(IC-off2) $\forall \theta_i, \quad \bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p) \geq \bar{\pi}(\rho', \theta_i; p) + \delta \bar{v}(\rho'; p)$

On-Schedule Constraints: $\forall \hat{\theta}_i,$

(IC-on1) $\forall \theta_{-i}, \quad v(p(\hat{\theta}_i), \mathbf{p}_{-i}(\theta_{-i})) \in \mathcal{V}_s$

(IC-on2) $\forall \theta_i, \quad \bar{\pi}(p(\theta_i), \theta_i; p) + \delta \bar{v}(p(\theta_i); p) \geq \bar{\pi}(p(\hat{\theta}_i), \theta_i; p) + \delta \bar{v}(p(\hat{\theta}_i); p).$

Notice how the on-schedule constraints are written in “direct” form: for given p and v , (IC-on2) requires that a firm with type θ_i does better by “announcing” that its type is θ_i than by announcing some other type, $\hat{\theta}_i$, when other firms are presumed to announce truthfully. This suggests that (IC-on2) may correspond to a “truth-telling” constraint in an appropriate mechanism design formulation.

4. COLLUSION AMONG PATIENT FIRMS AND MECHANISM DESIGN

In this section, we build on this suggestion, showing that when firms are patient, the Interim Program can be relaxed to yield a new program that we call the Mechanism Design Program. We also use existing tools from the mechanism design literature to begin our characterization of the optimal SPPE.

4.1. The Mechanism Design Program

Our general approach has two steps. First, we relax the Interim Program, by dropping off-schedule constraints and allowing continuation payoff functions beyond those that are actually feasible in the repeated game. With some notational adjustment, we then arrive at the Mechanism Design Program. Second, we provide conditions under which the solution in the relaxed setting corresponds to a factorization that continues to satisfy all of the constraints of the Interim Program. In this way, we identify conditions under which optimal SPPE may be characterized

by solving the Mechanism Design Program. The Mechanism Design Program is useful because we can apply existing tools to it directly; for example, in what follows, we make repeated use of the revenue equivalence theorem.

Our first step is to relax the constraints of the Interim Program as follows: (i) the off-schedule constraints are ignored and (ii) (IC-on1) is replaced with a relaxed constraint

$$\forall \hat{\theta}_i, \bar{v}(p(\hat{\theta}_i); p) \leq \bar{V}_s. \quad (\text{IC-on1}')$$

The first relaxation is without loss of generality if firms are sufficiently patient, since then off-schedule deviations are anyway not tempting. To appreciate the second relaxation, we recall that under (IC-on1), for every on-schedule vector of prices, the continuation value is drawn from the SPPE set, \mathcal{V}_s ; in contrast, (IC-on1') requires only that firm i 's *expected* continuation value does not exceed the *supremum* of the SPPE set, \bar{V}_s .⁹

We next introduce direct-form notation. Let $\Pi(\hat{\theta}, \theta; p) \equiv \bar{\pi}(p(\hat{\theta}), \theta; p)$ denote the current-period profit that a firm of type θ would expect were it to announce that its type is $\hat{\theta}$. We define as well a general “transfer” or “punishment” function, $T(\hat{\theta})$, which a firm expects to incur when it announces $\hat{\theta}$. We now define:

The Mechanism Design Program. Choose price schedule p and a punishment function T to maximize

$$E_{\theta}[\Pi(\theta, \theta; p) - T(\theta)]$$

subject to: For all θ , $T(\theta) \geq 0$;

$$(\text{IC-onM}) \forall \hat{\theta}, \theta, \Pi(\theta, \theta; p) - T(\theta) \geq \Pi(\hat{\theta}, \theta; p) - T(\hat{\theta}).$$

Suppose (p, v) satisfies the constraints of the Interim Program. We may then translate (p, v) into (p, T) , according to $T(\hat{\theta}) \equiv \delta[\bar{V}_s - \bar{v}(p(\hat{\theta}); p)]$. It is direct that (p, T) satisfies the constraints of the Mechanism Design Program; further, using this translation, the objectives of the Mechanism Design and Interim Programs rank (p, v) pairings in the same order. The Mechanism Design Program is thus a relaxed version of the Interim Program.

The meaning of T warrants emphasis. For a given SPPE, if a firm that announces $\hat{\theta}$ expects a continuation value below the supremum of the SPPE set (*i.e.* if $\bar{V}_s > \bar{v}(p(\hat{\theta}); p)$), then we may interpret the SPPE as specifying (in expectation) a “war”. There is then “no war” if the expected continuation value equals the supremum of the SPPE set (*i.e.* if $\bar{V}_s = \bar{v}(p(\hat{\theta}); p)$). Using the translation $T(\hat{\theta}) \equiv \delta[\bar{V}_s - \bar{v}(p(\hat{\theta}); p)]$, we thus may associate $T(\hat{\theta}) > 0$ ($T(\hat{\theta}) = 0$) with a future that follows a firm’s announcement of $\hat{\theta}$ and in which there is a war (no war).

We come now to the second step in our approach, where we provide the conditions under which a solution (p^*, T^*) to the Mechanism Design Program can be translated back into a factorization (p^*, v^*) that satisfies all of the constraints of the Interim Program. The following proposition identifies an important set of conditions of this kind:¹⁰

Proposition 2 (Stationarity). *Suppose (p^*, T^*) solves the Mechanism Design Program and $T^* \equiv 0$. Then $\exists \hat{\delta} \in (0, 1)$ such that, for all $\delta \geq \hat{\delta}$, there exists an optimal SPPE which is stationary, wherein firms adopt p^* after all equilibrium-path histories, and p^* solves the following program: maximize $E_{\theta} \Pi(\theta, \theta; p)$ subject to $\forall \hat{\theta}, \theta, \Pi(\theta, \theta; p) \geq \Pi(\hat{\theta}, \theta; p)$.*

9. Note particularly that (IC-on1') allows $v \equiv \bar{V}_s$, even though we as yet have no assurance that $\bar{V}_s \in \mathcal{V}_s$.

10. In our working paper (Athey, Bagwell and Sanchirico, 1998), we show that the general approach also extends to cases where T^* can be strictly positive (in which case the corresponding equilibrium is non-stationary). $T^* > 0$ may be the unique solution when demand is downward-sloping, but not, as we will show, for inelastic demand.

To establish this result, we develop two implications of the maintained assumption that $(p^*, T^* \equiv 0)$ is optimal in the relaxed environment. First, using our translation above, it then follows that $(p^*, v^* \equiv \bar{V}_s)$ is a solution to the Interim Program, if it satisfies the additional constraints of this program. In turn, this implies that $(p^*, v^* \equiv \bar{V}_s)$ is (weakly) superior to any SPPE factorization. Formally, $E_\theta[\Pi(\theta, \theta; p^*) + \delta \bar{V}_s] \geq \bar{V}_s$.

Second, we claim that if firms are sufficiently patient, it then follows that the repeated play of p^* (coupled with appropriate off-schedule punishments) is an SPPE. It is straightforward that this pattern of play satisfies the on-schedule incentive constraint: since p^* satisfies (IC-onM) when $T^* \equiv 0$, each firm will follow p^* when future play does not vary with the on-schedule price. Further, the repeated play of the static Nash equilibrium, p^e , is always an equilibrium of the repeated game; therefore, when firms are sufficiently patient, the threat of Nash reversion deters an off-schedule deviation (where a firm chooses a price not assigned to any cost type by p^*).¹¹ The claim is thus established. In turn, it implies that $E_\theta[\Pi(\theta, \theta; p^*)]/(1 - \delta) \leq \bar{V}_s$.

Combining our two inequalities, we obtain the desired result: $\bar{V}_s = E_\theta[\Pi(\theta, \theta; p^*)]/(1 - \delta)$. In words, if the Mechanism Design Program is solved with $(p^*, T^* \equiv 0)$, and if firms are sufficiently patient, then an optimal SPPE is easily characterized: firms adopt the pricing schedule p^* in each period, where p^* is the solution to the static program stated in Proposition 2.

It is now possible to preview the analysis that follows. In Section 4.2, we characterize the set of (p, T) that satisfy (IC-onM). This puts us in position to solve the Mechanism Design Program, using tools standard in the mechanism design literature. In Section 5, we establish that there is always a solution in which $T^* \equiv 0$. For patient firms, Proposition 2 then implies that an optimal SPPE is characterized by the stationary adoption of the accompanying pricing schedule, p^* . Under the assumption that the distribution function $F(\theta)$ is log-concave, we find that the optimal pricing schedule takes a simple form: $p^*(\theta) \equiv r$. We thus report conditions under which for sufficiently patient firms the optimal SPPE is stationary and requires all firms to charge the same price r , regardless of their costs. This rigid-pricing scheme is supported by the threat that if any other price is observed, the firms will revert to the worst SPPE, which delivers continuation value \underline{V}_s . When $F(\theta)$ is log-concave, we show in Section 6 that the worst SPPE is attained through Nash reversion (*i.e.* the Nash-pricing scheme is used in each period). We also establish there several additional predictions that arise when firms are less patient.

4.2. Consequences of on-schedule incentive compatibility

We begin our analysis of the Mechanism Design Program by characterizing the implications of the on-schedule constraint (IC-onM). We do this in the following lemma (where we use the notation $\Pi_\theta(\theta, \theta; p) = \frac{\partial}{\partial \theta} \Pi(\hat{\theta}, \theta; p)|_{\hat{\theta}=\theta}$):

Lemma 2 (Constraint reduction). (p, T) satisfies (IC-onM) if and only if (p, T) satisfies:

- (i) $p(\theta)$ is weakly increasing, and
- (ii) $\Pi(\theta, \theta; p) - T(\theta) = \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} \Pi_\theta(\tilde{\theta}, \tilde{\theta}; p) d\tilde{\theta}$.

This result is standard in the mechanism design literature, and it follows from the single-crossing property. We develop next an interpretation of Lemma 2, in order to provide the intuitive foundation for many of the findings that follow, and in particular our result that often the optimal SPPE uses the rigid-pricing scheme.

11. We establish below in Proposition 5(ii) that p^* yields higher expected profit than does the static Nash equilibrium, p^e .

The situation analysed here contrasts with the usual Principal–Agent formulation, since now the “agents” (*i.e.* firms) design their *own* schedule, with the goal of generating as much profit as possible. The expression in (ii) thus can be interpreted as reflecting the profit that can be distributed to θ , without inducing mimicry by other types. As (ii) reveals, the interim profit (inclusive of wars), $\Pi(\theta, \theta; p) - T(\theta)$, that is “left” for type θ , after incentive-compatibility constraints are considered, consists of two terms: the “profit-at-the-top” and the “information” or “efficiency” rents earned by the types between θ and $\bar{\theta}$.

To interpret these terms, let us consider a type θ_k that is just below $\bar{\theta}$. How much can this type earn, without inducing mimicry by type $\bar{\theta}$? Type θ_k can earn the same profit as does type $\bar{\theta}$ plus a bit extra, where the extra portion is attributable to the greater efficiency (*i.e.* lower costs) that θ_k actually enjoys. Similarly, a type θ_{k-1} that is slightly lower than θ_k can earn the same profit as does θ_k plus a bit extra. Pulling these points together, we now have a direct interpretation of (ii): the profit for any type θ equals the profit-at-the-top plus the accumulated efficiency rents of higher types (note that $\Pi_\theta < 0$). An important implication is that when the profit-at-the-top is increased, the highest type has less incentive to misrepresent itself as a lower type, and this relaxation in the incentive constraints in turn permits lower types to earn higher profits.

What determines the magnitude of the efficiency rents? To answer this, let us define the market share expected by a firm when it announces $\hat{\theta}$ as $M(\hat{\theta}; p) \equiv \bar{m}(p(\hat{\theta}); p)$, where the expectation is over the announcements of other firms (assumed truthful). We observe that

$$-\int_{\theta}^{\bar{\theta}} \Pi_{\theta}(\bar{\theta}, \bar{\theta}; p) d\bar{\theta} = \int_{\theta}^{\bar{\theta}} M(\bar{\theta}; p) d\bar{\theta}. \quad (4.1)$$

The magnitude of the efficiency rents is thus determined by the allocation of market shares across types.

We note that the firms have two instruments, prices and wars, with which to sort between types. It is useful to consider whether the availability of the war instrument expands the set of incentive-compatible market-share allocations. To explore this issue, we introduce a simple restriction. Consider a scheme (p, T) and let θ_K denote the lowest θ for which $p(\theta) = p(\bar{\theta})$. We restrict attention to schemes (p, T) for which $\Pi(\theta_K, \theta_K; p) - T(\theta_K) \geq 0$. It is straightforward to show that only schemes that satisfy the restriction will be optimal, and maintaining the restriction simplifies the exposition of our findings.¹² We can now establish that the use of wars does *not* expand the range of sorting alternatives available to the firm. Formally:

Lemma 3. *Given a scheme (p, T) that satisfies (IC-onM) and $T \geq 0$, and associated market-share allocation $M(\theta; p)$, there exists an alternative scheme $(\tilde{p}, \tilde{T} \equiv 0)$ which also satisfies (IC-onM), and such that $M(\theta; p) = M(\theta; \tilde{p})$ and $\Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = \Pi(\bar{\theta}, \bar{\theta}; \tilde{p})$.*

In short, given an original incentive-compatible scheme and market-share allocation, we may construct an alternative incentive-compatible scheme that delivers the same market-share allocation without using wars, while also providing the same profit-at-the-top. This construction requires that the prices are adjusted away from their original levels, and so the lemma does not determine which firm types (if any) prefer the alternative scheme. We explore this issue next.

12. For any (p, T) that fails the restriction, we may modify the argument associated with Lemma 3 in the following way. First, note that since (IC-onM) implies that overall profit is decreasing in θ , $\Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) < 0$. Next, rather than constructing the alternative scheme from (p, T) , we instead follow the steps in the Appendix and construct $(\tilde{p}, \tilde{T} \equiv 0)$ from (\check{p}, \check{T}) , where $\check{p}(\theta) = \bar{p}$ and $\check{T}(\theta) = 0$ for $\theta \in [\theta_K, \bar{\theta}]$ and $(\check{p}, \check{T}) = (p, T)$ elsewhere. The resulting $(\tilde{p}, \tilde{T} \equiv 0)$ satisfies (IC-onM), even though (\check{p}, \check{T}) does not, and it also satisfies $M(\cdot; p) = M(\cdot; \tilde{p}) = M(\cdot; \check{p})$, and $\Pi(\bar{\theta}, \bar{\theta}; \tilde{p}) - \tilde{T}(\bar{\theta}) = \Pi(\bar{\theta}, \bar{\theta}; \check{p}) > \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta})$. Since our goal is to establish that $(\tilde{p}, \tilde{T} \equiv 0)$ provides higher profit than does (p, T) , the latter inequality strengthens the results that follow.

Using Lemma 2 and equation (4.1), we observe that type θ 's interim profit is determined as

$$\Pi(\theta, \theta; p) - T(\theta) = \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} M(\tilde{\theta}; p) d\tilde{\theta}. \tag{4.2}$$

The next result, which is well known from auction theory, follows directly:

Lemma 4 (Revenue equivalence theorem). *Consider any (p, T) which satisfies (IC-onM). Then any other (\tilde{p}, \tilde{T}) which satisfies (IC-onM), $M(\theta; p) = M(\theta; \tilde{p})$ and $\Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = \Pi(\bar{\theta}, \bar{\theta}; \tilde{p}) - \tilde{T}(\bar{\theta})$ must also satisfy $\Pi(\theta, \theta; p) - T(\theta) = \Pi(\theta, \theta; \tilde{p}) - \tilde{T}(\theta)$ for all θ .*

Intuitively, suppose that firms start with the scheme (p, T) and that they then consider an alternative scheme (\tilde{p}, \tilde{T}) which is on-schedule incentive compatible and delivers the same profit-at-the-top. If in addition the alternative scheme maintains the original market-share allocation, then the efficiency rents are also preserved for every type. As the alternative scheme alters neither the profit-at-the-top nor the efficiency rents, it follows from (4.2) that this scheme maintains as well the original interim profit for all types.

At this point, we have extracted three lessons. First, after accounting for incentive compatibility, firms may be attracted to pricing schemes that raise the profit-at-the-top. Second, for a given amount of sorting, firms are free to choose whether to implement the corresponding market-share allocation with wars. Third, once the profit-at-the-top and the market-share allocation are determined, interim profit is fixed for all types. With these lessons in place, we are prepared now to characterize optimal SPPE for patient firms.

5. OPTIMAL COLLUSION AMONG PATIENT FIRMS

In this section, we analyse optimal symmetric collusion among firms that are patient. We present our central points in five steps. First, we consider symmetric collusive schemes that are *fully sorting* (i.e. $p(\theta)$ is strictly increasing). Second, we explore whether an optimal SPPE for patient firms requires equilibrium-path wars (i.e. $T(\theta) > 0$ for some θ). Third, we add further structure and characterize the optimal symmetric pricing scheme. Fourth, we identify the relationships between our findings and those in the broader literature. Finally, we discuss the manner in which our results extend when demand is downward-sloping.

5.1. Fully sorting pricing schemes

The set of fully sorting schemes includes a variety of candidates. One possibility is that the firms employ the static Nash equilibrium in each period. Alternatively, the firms may attempt to sort in the first period with higher prices, perhaps near the monopoly level. Such schemes satisfy on-schedule incentive constraints only if they include equilibrium-path wars.

Consider what profits are possible in an SPPE where the pricing function is strictly increasing in the first period. Since the highest type makes no sales, profit-at-the-top is simply $-T(\bar{\theta})$. Because a firm wins the market if and only if all other firms announce higher types, efficiency rents are uniquely determined for the class of fully sorting schemes, with $M(\theta; p) = [1 - F(\theta)]^{n-1}$. Thus, by Lemma 4, any two fully sorting schemes which satisfy (IC-onM) differ only if the profit-at-the-top, $-T(\bar{\theta})$, differs. The best profit-at-the-top is achieved when $T(\bar{\theta}) = 0$. But the Nash-pricing function, p^e , is fully sorting, and $(p^e, T \equiv 0)$ satisfies (IC-onM). Recall that $T \equiv 0$ corresponds to $\bar{v}(p(\theta); p) = \bar{V}_s$ for all θ . Therefore, the best collusive scheme that is fully sorting in the first period yields expected payoffs equal to $\pi^{NE} + \delta \bar{V}_s$,

the profit from the static Nash equilibrium plus the discounted profit from the best available continuation SPPE.

Now consider imposing a stronger requirement, namely that the pricing scheme is strictly increasing in every period of play. Let \mathcal{V}_s^{FS} be the set of such SPPE and let \bar{V}_s^{FS} be the supremum of the set. The logic of the previous paragraph implies that \bar{V}_s^{FS} is attained by a scheme that delivers the static Nash profit plus the best available continuation value in the restricted set \mathcal{V}_s^{FS} . Thus, $\bar{V}_s^{FS} = \pi^{NE} + \delta \bar{V}_s^{FS}$, or $\bar{V}_s^{FS} = \pi^{NE} / (1 - \delta)$. It follows that an optimal SPPE under full sorting is simply the repeated play of the static Nash equilibrium; further, this holds for any discount factor, since the Nash-pricing scheme satisfies all off-schedule constraints as well. Summarizing:

Proposition 3. *Among the class of SPPE where the pricing function is strictly increasing in every period, for any distribution function F and discount factor δ , an optimal SPPE is the repeated play of the static Nash equilibrium after all histories.*

5.2. No wars on the equilibrium path

For the class of fully sorting pricing functions, the analysis in the previous section shows there is no benefit in supporting higher prices with on-schedule wars. We now extend this argument for any initial market-share allocation.

Consider any original scheme (p, T) that entails wars somewhere. Lemma 3 guarantees the existence of an alternative no-war scheme $(\tilde{p}, \tilde{T} \equiv 0)$ that is incentive compatible, induces the same market-share allocation, and generates the same profit-at-the-top. Lemma 4 then implies that the alternative no-war scheme gives the same interim profit (inclusive of wars) as did the original scheme. As shown in the Appendix (proof of Lemma 3), the alternative schedule achieves this profit by exchanging any war in the original schedule for a lower price. Applying this argument, together with our stationarity result (Proposition 2), we conclude that:¹³

Proposition 4. *Allow for any distribution function F . If (p^*, T^*) is a solution to the Mechanism Design Program, then there exists as well a solution (\tilde{p}, \tilde{T}) with $\tilde{p}(\theta) \leq p(\theta)$ and $\tilde{T}(\theta) \equiv 0$. Thus, if firms are sufficiently patient, there then exists an optimal SPPE that is stationary: firms use the pricing scheme $\tilde{p}(\theta)$ following every history along the equilibrium path, and $E_\theta \Pi(\theta, \theta; \tilde{p}) / (1 - \delta) = \bar{V}_s$.*

We see from Proposition 4 that wars have no value: for any distribution F , if there exists an optimal SPPE that uses wars, then there exists as well an optimal SPPE that does not.

5.3. Optimal pricing

We are now prepared to determine the optimal SPPE pricing scheme when firms are patient. Given the “no-wars” finding from the previous subsection, we seek the price strategy $p^*(\theta)$ that solves the program presented in our stationarity proposition. With some additional structure, this pricing scheme is easily characterized:¹⁴

13. We include here a restriction on the discount factor, so that the stationarity proposition may be used. In what follows, in Proposition 8, we develop the further argument that the off-schedule constraints are relaxed in moving from a scheme with wars to a no-war scheme, and we are thus able to remove this restriction.

14. This result also holds in a model with discrete types, except that an additional parameter restriction is required, one that depends on the gap between r and the highest type. The restriction is satisfied when the distance between types becomes small enough. Details are available from the authors.

Proposition 5. For δ sufficiently large:

- (i) If either (a) F is log-concave, or (b) $r - \bar{\theta}$ is sufficiently large, then the equilibrium path of the optimal SPPE is characterized by price rigidity ($p^*(\theta) \equiv r$) and no wars ($T^*(\theta) \equiv 0$).
- (ii) In any optimal SPPE that is stationary, there exists an open interval of cost types over which pricing is rigid, and per-period profit greater than that in the static Nash equilibrium is attained: $\bar{V}_s > \pi^{NE}/(1 - \delta)$.

For patient firms, if the distribution function is log-concave or r is large enough, the optimal SPPE is described as follows: firms select the price r in each period, whatever their private cost realizations, so long as all firms have selected the price r in all previous periods.¹⁵ Further, firms can always exceed Nash payoffs, if they are sufficiently patient.

The rigid-pricing scheme $p(\theta) \equiv r$ has benefits and costs. An important benefit of this scheme is that it satisfies the on-schedule incentive constraint without recourse to equilibrium-path wars. Furthermore, the price is as high as possible. However, an evident cost of the rigid-pricing scheme is that it sacrifices efficiency benefits: it may be that one firm has a low cost while another firm has a high cost, but under the rigid-pricing scheme each of these firms sells to $1/n$ -th of the market. The content of the proposition (part (i)) is that the benefits of the rigid-pricing scheme exceed the costs, provided that the distribution function is log-concave, or the reservation price is high enough.

As this proposition is central to the paper, we include the proof in the text. By Proposition 4, we may focus on solutions to the Mechanism Design Program for which $T(\theta) \equiv 0$. Using (4.2), we may now write a firm's expected profit as

$$E_\theta[\Pi(\theta, \theta; p)] = E_\theta \left[\pi(p(\bar{\theta}), \bar{\theta}) \cdot M(\bar{\theta}; p) + \int_\theta^{\bar{\theta}} M(\tilde{\theta}; p) d\tilde{\theta} \right].$$

Next, employing a standard trick from the literature on optimal auctions (Myerson (1981), Bulow and Roberts (1989)), we may integrate by parts and rewrite our objective function as

$$E_\theta[\Pi(\theta, \theta; p)] = E_\theta \left[\pi(p(\bar{\theta}), \bar{\theta}) \cdot M(\bar{\theta}; p) + \frac{F}{f}(\theta) \cdot M(\theta; p) \right]. \tag{5.1}$$

Consider the first term of this expression. Since the single-crossing property implies that p is non-decreasing, profit-at-the-top is highest when all firms set the same price, so that the highest type is never underpriced. The most profitable rigid-pricing scheme is the one in which firms fix the price at r . Thus, $p(\theta) \equiv r$ maximizes profit-at-the-top, given that downward-sloping pricing schemes are not on-schedule incentive compatible.

Consider now the second term. Log-concavity plays an important role here. Intuitively, the "contribution" of an increase in a given type's profit to the firm's expected profit is governed by the fraction of types below it (which enjoy a relaxed incentive constraint and thus earn higher profit), conditional on the "probability" that the given type will actually arise. We interpret $F(\theta)/f(\theta)$ as a measure of the contribution of an increase in type θ 's profit to the firm's expected profit. Log-concavity ensures that this measure is non-decreasing, so that allocating market share to higher types, and away from lower (and more efficient) types, improves cartel profit. This suggests that the rigid-pricing scheme $p(\theta) \equiv r$ also maximizes the expected efficiency rent term, $E_\theta[\frac{F}{f}(\theta) \cdot M(\theta; p)]$.

15. We have assumed $r \geq \bar{\theta}$. If this assumption were relaxed, then the optimal scheme would entail that firms with cost types greater than r sit out rather than endure negative profits.

To formalize this suggestion, observe first that for all p , $E_\theta[M(\theta; p)] = 1/n$: before a firm learns its type, it has an expected market share of $1/n$. Define

$$\Phi(\theta; p) = n \int_{\underline{\theta}}^{\theta} M(\tilde{\theta}; p) f(\tilde{\theta}) d\tilde{\theta},$$

and note that $\Phi(\cdot; p)$ is a probability distribution for each p . Let p^R be a rigid-pricing scheme. Given that $M(\theta; p)$ is non-increasing when p is non-decreasing, $M(\theta; p^R) \equiv 1/n$ must cross $M(\theta; p)$ from below for any p . Thus, $\Phi(\cdot; p^R)$ dominates $\Phi(\cdot; p)$ by first-order stochastic dominance (FOSD) for all p . In words, a rigid-pricing scheme puts the most weight on high-cost types. By the definition of FOSD, it follows that if $F(\theta)/f(\theta)$ is non-decreasing,

$$\begin{aligned} E_\theta \left[\frac{F}{f}(\theta) \cdot M(\theta; p^R) \right] &= \frac{1}{n} \int_{\underline{\theta}}^{\bar{\theta}} \frac{F}{f}(\theta) d\Phi(\theta; p^R) \geq \frac{1}{n} \int_{\underline{\theta}}^{\bar{\theta}} \frac{F}{f}(\theta) d\Phi(\theta; p) \\ &= E_\theta \left[\frac{F}{f}(\theta) \cdot M(\theta; p) \right] \end{aligned}$$

for all p non-decreasing. Thus, the second term in (5.1) is maximized with a rigid-pricing scheme. Finally, *only* a rigid-pricing scheme can be optimal because, given our assumption that $f > 0$, $0 = \frac{F}{f}(\underline{\theta}) < \frac{F}{f}(\theta)$ for $\theta > \underline{\theta}$. Any non-rigid scheme would place more weight on $\frac{F}{f}(\underline{\theta})$.

Since $p(\theta) \equiv r$ maximizes profit-at-the-top and is the uniquely optimal rigid-pricing scheme that does so, we conclude that rigid pricing at r strictly dominates all other incentive-compatible pricing schemes when F is log-concave. Under log-concavity, if firms are sufficiently patient, the optimal SPPE thus entails rigid pricing at r in each period.

The assumption that the distribution is log-concave is common in the contracting literature, and many distributions satisfy this assumption. But we can also consider the optimal market-share allocation even if F/f is decreasing on some intervals. If r is close enough to $\bar{\theta}$, the optimal pricing rule maximizes the second term of (5.1); clearly, it is desirable to put more weight on the higher values of F/f , subject to p non-decreasing. This may entail intervals of both sorting and pooling. Under our assumption that $f(\underline{\theta}) > 0$, strict log-concavity always holds in a neighbourhood of $\underline{\theta}$. Thus, it is always optimal for sufficiently patient firms to use rigid pricing at the bottom of the pricing function, and from this part (ii) of the proposition follows: sufficiently patient firms can sustain SPPE per-period profit strictly above the static Nash equilibrium (which entails sorting at the bottom). Finally, observe that if $r - \bar{\theta}$ is sufficiently large, so that profit-at-the-top is great enough, the first term of (5.1) dominates the second. Then, the benefits from pooling market share with high-cost types outweigh those from allocating market share amongst types with the highest F/f , and the rigid-pricing scheme is again optimal.

We close this subsection with an instructive example. Suppose that there are two firms and consider the family of two-step pricing schedules. Such schedules are characterized by a low price ρ_1 , a high price ρ_2 , and a breakpoint θ_2 , such that all types below (above) θ_2 select the low (high) price. A rigid-pricing scheme is then an extreme case, in which $\theta_2 = \underline{\theta}$ or $\theta_2 = \bar{\theta}$. For $\theta_2 \in (\underline{\theta}, \bar{\theta})$, the single-crossing property ensures that the on-schedule incentive-compatibility constraints are satisfied if and only if a firm of type θ_2 is indifferent between the bottom and top steps: $\pi(\rho_1, \theta_2)M(\underline{\theta}; p) = \pi(\rho_2, \theta_2)M(\bar{\theta}; p)$, where $M(\underline{\theta}; p) = F(\theta_2)/2 + [1 - F(\theta_2)]$ and $M(\bar{\theta}; p) = [1 - F(\theta_2)]/2$. Given ρ_2 and θ_2 , incentive compatibility thus determines ρ_1 and thereby expected profit. Suppose further that the distribution of costs is uniform over the unit interval. Direct calculations then yield $E[\Pi(\theta, \theta; p)] = [2\rho_2(1 - \theta_2) + \theta_2 + (\theta_2)^2 - 1]/4$. Expected profit is thus maximized when $\rho_2 = r$ and θ_2 is set at a corner. With $r \geq \bar{\theta}$, the maximum is achieved when $\theta_2 = \theta$. Thus, a rigid-pricing scheme with $p(\theta) \equiv \rho_2 = r$ is optimal within the two-step family. We may also compare the rigid-pricing and Nash-pricing schemes.

Expected profit under the Nash-pricing scheme is $1/6$, which is less than the value $(2r - 1)/4$ that is generated by the rigid-pricing scheme.

5.4. Related literature

We now compare our findings to those established in related papers. We recall first the McAfee and McMillan (1992) study of bidding rings. Working with a static model of procurement auctions, McAfee and McMillan show that bidders in a “weak cartel” (where transfers are not allowed) collude best if they agree to bid the same price, r . Our Proposition 5 is closely related; indeed, their model can be mapped into the Mechanism Design Program by setting $T \equiv 0$. The analysis here extends their results, however, by formally connecting the static results to the repeated-game context and demonstrating that a rigid-pricing scheme is optimal even when schemes that sustain sorting using “wasteful” transfers ($T > 0$) are allowed. In the next section, we further generalize the analysis to consider the possibility of impatient firms.

Our findings are also related to work on cartel design under private information. As Cramton and Palfrey (1990) and Kihlstrom and Vives (1992) show, if firms can design a mechanism where they communicate their cost types and make side-payments to one another, then the firms can achieve full efficiency benefits without any pricing distortion (*i.e.* production is allocated to the lowest-cost firm, who sets its monopoly price), although the participation constraints may fail for some cost types.¹⁶ In the repeated-game context, as Athey and Bagwell (2001) show, firms may use asymmetric continuation values to mimic side-payments from one firm to another. In the model considered here, however, the SPPE restriction prevents such cross-firm transfers. But it is still possible for firms to achieve full efficiency benefits, if they use fully sorting pricing schemes. In this case, however, informational costs, manifested as pricing distortions and/or future price wars, must be experienced. The costs exceed the benefits when F is log-concave.

Interestingly, communication would not have any value in our model. The realization of full efficiency benefits does not require communication, while the informational costs remain, with or without communication, so long as firms cannot make side-payments (or use asymmetric continuation values). In short, the optimal direct revelation mechanism without side-payments is characterized by Proposition 5, and no communication is required to implement it.

Our findings are also related to those presented in the hidden-action collusion literature. A central feature of the Green and Porter (1984) and Abreu *et al.* (1986, 1990) papers is that collusive conduct involves periodic reversions to price wars. Our model can be placed within their hidden-action modelling framework, if we think of a firm’s strategy, $p(\theta)$, as its hidden action and the resulting price, $p = p(\theta)$, as the public signal, where the distribution of this public signal is then determined by the pricing function itself and the distributional properties of θ . The main difference is that we allow for an endogenous support of the public signal. Put differently, our model may be understood as a hidden-action model with *endogenous imperfect monitoring*. To see this, recall that in the Green and Porter (1984) and Abreu *et al.* (1986, 1990) modelling framework, the support of the publicly observed market price is independent of the private output selections made by firms. In our model, by contrast, the support of the signal is itself determined in equilibrium. In particular, if firms employ a rigid-pricing schedule in which they choose r under all cost realizations, then in equilibrium the support of the public signal is degenerate, as rival firms expect to observe the price r , no matter what cost realization the firm experiences. This in turn enables firms to limit wars to off-equilibrium-path events.

16. Their analysis builds on earlier work by Roberts (1985), who shows that a scheme with full efficiency benefits may not be incentive compatible, when firms can communicate but are unable to make side-payments. See also the McAfee and McMillan (1992) analysis of “strong cartels”, in which bidders can make transfers to one another.

We assume that firms' costs are i.i.d. over time. An alternative assumption is that each firm's privately observed cost is constant over time. LaCasse (1999) and Chakrabarti (2001) analyse this possibility with downward-sloping demand and discrete cost types. They provide conditions under which there exist equilibria that are either separating or pooling. Our results can be used to characterize optimal equilibria of a game with constant costs, when firms are patient. In particular, if F is log-concave or r is sufficiently high, then optimal equilibria entail rigid pricing at r .¹⁷ In the next section, we extend our model to the case of downward-sloping demand. Our findings again characterize optimal equilibria of a game with constant costs and patient firms. Optimal equilibria then generally entail at least partial rigidity.

5.5. Downward-sloping demand

We consider now the case of downward-sloping demand. While collusive behaviour is often associated with inelastic-demand markets, most demand curves have at least some elasticity, and so it is important to investigate the robustness of our results.

To this end, we modify the model as follows. Maintaining the assumption that goods are perfect substitutes, we now define the profit-if-win function as $\pi(\rho, \theta) \equiv (\rho - \theta)D(\rho)$, where D is a twice-continuously differentiable market demand function that satisfies $D > 0 > D'$ over the relevant range. We assume that $\pi(\rho, \theta)$ is strictly quasiconcave in ρ , with a unique maximizer, $p^m(\theta)$, where $p^m(\bar{\theta}) \geq \bar{\theta}$. The monopoly price, $p^m(\theta)$, is strictly increasing in θ .

With this modification, the interim profit function continues to be given by $\Pi(\hat{\theta}, \theta; p) = \pi(p(\hat{\theta}), \theta)M(\hat{\theta}; p)$. We next define $q(\rho; p)$ as the quantity a firm expects to sell when it sets price ρ and opponents use pricing function p . In the present context, $q(\rho; p) = D(\rho) \cdot E_{\theta_{-i}}[m_i(\rho, p_{-i}(\theta_{-i}))]$, and so $q(p(\hat{\theta}); p) = D(p(\hat{\theta}))M(\hat{\theta}; p)$. The function Π thus may be alternatively expressed as $\Pi(\hat{\theta}, \theta; p) = (p(\hat{\theta}) - \theta)q(p(\hat{\theta}); p)$. We note that Lemma 2 remains valid and implies that

$$\Pi(\theta, \theta; p) - T(\theta) = \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + \int_{\bar{\theta}}^{\theta} q(p(\tilde{\theta}); p) d\tilde{\theta}. \quad (5.2)$$

This expression highlights similarities and differences with the inelastic-demand case. As in the case of inelastic demand, if type θ expects to sell a higher quantity, then all cost types below θ can enjoy higher profit, since efficiency rents are greater at higher quantities. But now the magnitude of efficiency rents is determined by the shape of the pricing function (*i.e.* the allocation of market shares across types) *and* the level of prices (*i.e.* the size of market demand). As a consequence, Lemma 4 no longer holds: if firms construct an alternative collusive scheme through a change in p and T that preserves profit-at-the-top and incentive compatibility at the original market-share allocation, the interim profit of lower types is *not* preserved. Just like in an auction model with risk-averse firms, revenue equivalence breaks down when the level of price interacts directly with the firm's true cost type. Thus, we cannot immediately conclude that price wars are not needed. Indeed, our working paper (Athey *et al.*, 1998) provides a detailed discussion showing when price wars might be used and when they can be ruled out. Here, we focus our attention on price rigidity.

To understand the benefits and costs of rigid pricing, we take the expectation of profit from (5.2) and integrate by parts, to derive an expression analogous to (5.1):

$$E_{\theta}[\Pi(\theta, \theta; p) - T(\theta)] = \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + E_{\theta} \left[q(p(\theta); p) \frac{F}{f}(\theta) \right]. \quad (5.3)$$

17. Building on the present paper, Athey and Bagwell (2002) establish this finding.

Recall that in the case of inelastic demand, a rigid price at r maximizes both profit-at-the-top and the expected efficiency rent term. This is no longer true when demand is downward-sloping. Instead, the two terms are in conflict about the level of the price. To maximize the profit-at-the-top, it would be optimal to have a rigid price, as the market share for type $\bar{\theta}$ is thereby made as large as possible; furthermore, the best rigid price would be $p^m(\bar{\theta})$. To maximize the efficiency rent term, it would also be optimal to have a rigid price, following an argument analogous to the one we used for inelastic demand; however, the best rigid price would now be as low as possible, in order to make the quantity produced (and thus the efficiency rents) as large as possible. The conflict between the two terms about the optimal price level implies that price rigidity is not necessarily optimal.

At the same time, (5.3) also indicates that there are robust forces in favour of at least partial price rigidity. To see this, suppose that p is strictly increasing on $[x, y] \subset [\underline{\theta}, \bar{\theta}]$. Now define $\check{p}(\theta)$ implicitly so that it agrees with p outside of $[x, y]$ but is constant on $[x, y]$ at a price \bar{p} such that

$$E_{\theta}[q(\check{p}(\theta); \check{p}) \mid \theta \in [x, y]] = q(\bar{p}; \check{p}) = E_{\theta}[q(p(\theta); p) \mid \theta \in [x, y]]. \tag{5.4}$$

Then, just as in the inelastic-demand case, we can define a probability distribution

$$\Phi(\theta; p, x, y) = \frac{1}{q(\bar{p}; \check{p})} \int_x^{\theta} q(p(\theta); p) dF(\theta \mid \theta \in [x, y]).$$

Since $\Phi(\theta; \check{p}, x, y)$ dominates $\Phi(\theta; p, x, y)$ by FOSD, if $F(\theta)/f(\theta)$ is non-decreasing, then

$$E_{\theta} \left[q(\check{p}(\theta); \check{p}) \frac{F}{f}(\theta) \mid \theta \in [x, y] \right] \geq E_{\theta} \left[q(p(\theta); p) \frac{F}{f}(\theta) \mid \theta \in [x, y] \right],$$

and a force in favour of rigidity is thus illustrated. With downward-sloping demand, however, the new scheme does not in general satisfy (IC-onM) unless T is modified. While we can always find some \check{T} that satisfies (IC-onM), this \check{T} might violate the constraint $\check{T}(\theta) \geq 0$, and it is also possible that $\check{T}(\bar{\theta}) > T(\bar{\theta})$, implying a reduction in profit-at-the-top. Despite these complications, we establish conditions under which partial or full rigidity is optimal:

Proposition 6. *Suppose that demand is downward-sloping.*

- (i) *Among the class of SPPE where the pricing function is strictly increasing in every period, for any distribution function F and discount factor δ , the optimal SPPE is the repeated play of the static Nash equilibrium.*
- (ii) *For δ sufficiently large, if F is log-concave and demand is sufficiently inelastic, then the equilibrium path of the optimal SPPE is characterized by price rigidity and no wars.*

To interpret this result, observe that Proposition 6(i) represents a strengthening of Proposition 3: when demand is downward-sloping, the repeated play of the static Nash equilibrium is uniquely optimal among fully sorting SPPE. Downward-sloping demand introduces new forces in favour of schemes with low (*i.e.* static Nash) prices and no equilibrium-path price wars, as opposed to schemes with high (*e.g.* monopoly) prices and equilibrium-path wars. With lower prices, the level of market demand is greater, and the efficiency rents are thus higher.

An implication of Proposition 6(i) is that, if firms are to improve upon the static Nash equilibrium, then some rigidity is required. With downward-sloping demand, as Spulber (1995) shows, the static Nash equilibrium is fully sorting with $p^e(\bar{\theta}) = \bar{\theta}$. It is straightforward to show that it is possible to achieve profit greater than in the static Nash equilibrium if, for example, $p^m(\bar{\theta}) > \bar{\theta}$, by introducing some rigidity at the top of the pricing schedule.

Further, Proposition 6(ii) establishes that a rigid-pricing scheme is optimal when demand is sufficiently inelastic. Additional characterizations of the optimal scheme with downward-sloping demand are provided in Athey *et al.* (1998).

We now consider briefly the possibility that firms sell imperfect substitutes. In this case, $q(\rho; p) = E_{\theta_{-i}}[D(\rho, p_{-i}(\theta_{-i}))]$, and expected profit is still given by (5.3). This case is more complex, because the expected demand for each cost type depends on the entire pricing function. Recall our construction of \check{p} as in (5.4). As a result of changing the pricing function from p to \check{p} , the expected profit for types $\theta \notin [x, y]$ may increase or decrease, depending on the shape of the demand curve. If demand is linear, however, then firms care only about the expected price of opponents. In this event, p and \check{p} yield the same expected price and $\Pi(\theta, \theta; p) = \Pi(\theta, \theta; \check{p})$ for $\theta \notin [x, y]$. But it still remains to show that there exists a feasible \check{T} that maintains (IC-onM) and does not decrease profit-at-the-top. The analysis of this problem is beyond the scope of the present paper; here, we simply note that it can be shown that if F and $1 - F$ are log-concave, and demand is linear, then any scheme that improves upon the static Nash equilibrium is at least partially rigid. A similar result holds when firms compete in quantities rather than prices.¹⁸

In summary, the result that fully sorting pricing schemes are not optimal is quite robust, and the forces in favour of pooling remain in more general models, although new forces that oppose rigidity may lead to optimal schemes that are only partially rigid.

6. OPTIMAL COLLUSION AMONG IMPATIENT FIRMS

We now consider collusion among impatient firms. We proceed in two general steps. First, we characterize the worst SPPE (the most severe punishment) and then determine the critical patience level above which the rigid-pricing scheme can be enforced. Second, we consider less patient firms, who are unable to enforce the rigid-pricing scheme, and explore how they best collude. Impatience creates an additional motivation for the avoidance of price wars. In addition, impatient firms may use pricing schemes that entail an “escape clause”, whereby a firm is allowed to depart from the rigid price and set a lower price when it experiences a favourable cost shock. In an extended model, we find that such a departure is especially likely when demand is temporarily high. Finally, we establish conditions under which the optimal collusive pricing scheme for less patient firms must be a step function (partial rigidity).

6.1. Enforcing rigidity off schedule

We begin with the determination of the critical discount factor. The rigid-pricing scheme satisfies off-schedule constraints if a firm always regards the current-period benefit from undercutting the rigid price as small in comparison to the discounted value of future co-operation. In turn, future co-operation is more valuable when firms are more patient and the punishment that would follow a deviation is more severe. The critical discount factor is therefore determined as the lowest discount factor at which firms can enforce the rigid-pricing scheme, when a deviation leads to the most severe punishment, \underline{V}_s .

Formally, let us suppose that the firms attempt to maintain a rigid price $\rho \geq \underline{\theta}$ in all periods for all cost realizations. If a firm of type θ were to cheat and undercut (by ϵ) this price, then the firm would win the entire market, as opposed to just $1/n$ -th of the market, and so the firm’s incentive to cheat is $\frac{n-1}{n}\pi(\rho, \theta)$. Importantly, this incentive is greatest for a firm with

18. To see this, note that if a firm’s expected profit is given by $\tilde{\Pi}(\hat{\theta}, \theta; q) = q(\hat{\theta}) \cdot (E_{-i}[P(q(\hat{\theta}) + \sum_{j \neq i} q(\theta_j))] - \theta)$, then $\tilde{\Pi}_{\theta}(\hat{\theta}, \theta; q) = -q(\theta)$, and so the expression (5.3) for expected profit, as well as the forces in favour of pooling, would be the same.

the lowest cost level, $\underline{\theta}$, since the profit-if-win is then highest and the gain in market share is thus most valuable. If a firm were to cheat, however, it would forfeit the discounted value of future co-operation. This value is measured in relation to the cost that a firm expects in the future, $E\theta$. For example, if a deviation is punished by an infinite reversion to the static Nash equilibrium, then the proposed rigid-pricing scheme is off-schedule incentive compatible if and only if

$$\frac{n-1}{n}\pi(\rho, \underline{\theta}) \leq (\delta/(1-\delta)) \left[\frac{1}{n}\pi(\rho, E\theta) - \pi^{NE} \right]. \tag{6.1}$$

If the proposed scheme yields greater-than-Nash profit, (6.1) holds when δ is sufficiently large.

We note as well that both the incentive to cheat and the future value of co-operation increase with the rigid price, ρ . As long as $\delta > \frac{n-1}{n}$, however, the latter effect dominates, so that the rigid price that is easiest to support has $p(\theta) \equiv r$ in each period. Using (6.1), we find that the critical discount factor δ^* above which firms can use the Nash punishment threat to enforce the rigid-pricing scheme is

$$\delta^* \equiv \frac{(n-1)\pi(r, \underline{\theta})}{(n-1)\pi(r, \underline{\theta}) + \pi(r, E\theta) - n\pi^{NE}}. \tag{6.2}$$

It is straightforward to verify that $\delta^* \in (\frac{n-1}{n}, 1)$, if $\frac{1}{n}\pi(r, E\theta) > \pi^{NE}$. We may now state the following proposition.

Proposition 7.

- (i) *If F is log-concave, then for all discount factors δ , $\underline{v}_s = \pi^{NE}/(1-\delta)$.*
- (ii) *If F is log-concave and $\delta < \delta^*$, then there does not exist an SPPE with rigid pricing.*
- (iii) *If F is log-concave or $r - \bar{\theta}$ is large enough, then $\delta^* \in (\frac{n-1}{n}, 1)$ and for all $\delta > \delta^*$ any optimal SPPE is characterized by rigid pricing at r in every period.*

Thus, when F is log-concave, Nash reversion is in fact the worst punishment, and the rigid-pricing scheme can be enforced if and only if $\delta > \delta^*$, where $\delta^* \in (\frac{n-1}{n}, 1)$. To see an example, suppose there are two firms, costs are uniformly distributed over $[0, 1]$ and $r = 1$. Then, $\pi^{NE} = 1/6$, and an SPPE with rigid pricing exists and is optimal if and only if $\delta \geq \delta^* = 6/7$.

It is striking that the lowest SPPE continuation value, \underline{v}_s , corresponds to Nash play when F is log-concave. This is true despite the fact that SPPE may exist in which some firm types price below cost. For example, there may exist non-stationary SPPE, in which higher-cost types price below cost in the first period, sustained by the promise of a better future equilibrium. Of course, SPPE continuation values cannot be driven too low: the scheme must offer the highest-cost type overall expected payoffs greater than $\delta \underline{v}_s$ (or else the firm will deviate off-schedule), and lower-cost types cannot be deprived of the available efficiency rents. In searching for the lowest SPPE continuation value, we thus consider pricing schemes that minimize efficiency rents. Following the logic of Section 5.3, the minimum efficiency rent is attained using a strictly increasing pricing scheme when F is log-concave, and with this we can establish that it is not possible to sustain punishments worse than Nash. As we confirm in Lemma 6 in the Appendix, however, when the log-concavity assumption is relaxed, on-schedule incentive constraints may be compatible with below-Nash efficiency rents. If firms are sufficiently patient, non-stationary SPPE with below-cost pricing can then be constructed that yield below-Nash payoffs.¹⁹

19. The scheme used to generate below-Nash payoffs requires some firms to price below cost, and such a firm must be dissuaded from deviating to a higher price. Indeed, a firm would undertake just such a deviation, if the market-clearing price were public but individual prices were otherwise not. In this case, the worst SPPE involves the repeated play of the static Nash equilibrium (for any F and δ). When individual prices are public, however, a firm can be induced to price below cost, and when F is not log-concave this may describe the worst SPPE.

Recall that when firms have access to a public randomization device, the set of equilibrium values is convex. In that case, Proposition 7 provides a complete characterization of the SPPE set when F is log-concave and $\delta > \delta^*$.

An interesting implication of Proposition 7 is that intertemporal fluctuations in costs diminish the ability of firms to collude. Intuitively, when costs fluctuate through time, collusion requires greater patience, as the scheme must withstand the incentive imbalance that occurs when a firm draws a low current cost level (θ), and thus faces a great incentive to cheat, while assessing the long-term value of co-operation with reference to an average cost level ($E\theta$). Formally, $\delta^* > \frac{n-1}{n}$, where $\frac{n-1}{n}$ is the critical discount factor for the standard Bertrand supergame, in which firms' costs are time-invariant. This implication is broadly consistent with the common assessment (see, *e.g.* Scherer, 1980, p. 205) that collusion is more difficult when costs are variable across firms.

6.2. *No wars on the equilibrium path*

How do firms best collude when they are unable to enforce the rigid-pricing scheme? In this section, we take a first step toward answering this question. Allowing for impatient firms, we establish that the scope for symmetric collusion cannot be enhanced (and may be strictly reduced) by equilibrium-path wars.

The central idea is simple. Let us start with an original SPPE collusive scheme. Relying on Proposition 4, if there is a positive probability of an equilibrium-path war associated with some cost type, then we can re-engineer an alternative collusive scheme—by eliminating the war and reducing the price for that type a corresponding amount—that yields for this type the same expected payoff. The alternative scheme satisfies the on-schedule constraint (given that the original did) and thus constitutes a payoff-equivalent SPPE for patient firms. When firms are impatient, however, the off-schedule constraint is also a concern, and it is here that the alternative schedule offers an actual advantage: by shifting profit from the current period (price is reduced) to the future (wars are eliminated), the incentive to cheat is reduced while the future value of co-operation is enhanced. The off-schedule constraint is therefore now easier to satisfy than under the original scheme.

As the following proposition confirms, this argument is quite general:

Proposition 8. *Allow for any distribution function F and any discount factor δ . If an SPPE exists with the optimal payoff \bar{V}_s , then there exists a stationary SPPE, where the same pricing strategy is used following every equilibrium-path history, with the optimal payoff \bar{V}_s .*

This argument further indicates that “revenue equivalence” does not extend to impatient firms. A scheme that uses price wars may violate the off-schedule constraint when a payoff-equivalent scheme without price wars does not.

6.3. *Partial rigidity and collusion among impatient firms*

The propositions developed above suggest that our search for collusive schemes among impatient firms should emphasize two ingredients: the absence of rigid pricing and no equilibrium-path wars. But exactly how do impatient firms price in an optimal SPPE? In this section, we first present sufficient conditions under which a *two-step pricing scheme* can be enforced and is optimal for impatient firms. Second, we consider an extended model that includes publicly observed fluctuations in industry demand. Finally, we argue that optimal SPPE for impatient

firms is characterized by a pricing schedule that is a step function (partial rigidity), if the collusive scheme is to offer better-than-Nash profits.

6.3.1. Introducing a second step: an “Escape Clause”. Recall that the rigid-price scheme fails to be enforceable when $\delta < \delta^*$, because a firm that draws the lowest-cost type is too tempted to undercut the rigid price r and increase its market share. A natural conjecture is that this problem may be overcome when a two-step pricing scheme is employed, with prices ρ_1 and ρ_2 , where $\rho_1 < \rho_2$, and a break-point θ_2 . In this case, the lowest-cost firm has less incentive to cheat. Firstly, this firm now expects greater than a $1/n$ -th share of the market, and so the gain in market share that accompanies a price cut is diminished. Secondly, any given gain in market share is now less profitable, since the lower-cost firm has a lower price, and thus the profit-if-win it experiences on the market share it enjoys is now lower.

This, however, is not the whole story. Balanced against this diminished incentive to cheat is a reduction in expected long-term profit: if the distribution function is log-concave, a two-step scheme yields lower expected profit than does a rigid scheme, and so the firm also now has less to lose in the future if it cheats today. Complicating matters further, the net resolution of these conflicting effects for the off-schedule incentive constraint may hinge upon the nature of the distribution function. A two-step scheme will satisfy the off-schedule incentive constraint if it lowers the incentive that the lowest-cost firm has to cheat without substantially altering the expected profit that firms anticipate in the future. Intuitively, this will be the case if the density is small for lower-cost types, so that these types occur infrequently in the future.

Proposition 9. *If F is log-concave and*

$$f(\underline{\theta}) < \frac{\frac{1}{n}\pi(r, E\theta) - \pi^{NE}}{(n-1)\pi(r, \underline{\theta})\pi^{NE}} \tag{6.3}$$

then there exists $\delta^o < \delta^$, such that, for every $\delta \in (\delta^o, \delta^*)$, there exists an optimal SPPE that is stationary and uses a two-step pricing scheme, with $p_2 = r > p_1$ and $\theta_2 \in (\underline{\theta}, \bar{\theta})$.*

When F is log-concave, the two-price scheme is optimal for δ just below δ^* , since then the two-price scheme departs from the desired rigid-pricing scheme only at the lowest-cost types.

Proposition 9 describes a situation in which the realization of an unlikely and low cost type results in a marked reduction in the firm’s price, suggesting that rare but pronounced price cuts may occur under symmetric collusion schemes when firms are impatient.²⁰ In other words, symmetric collusion among impatient firms may call for an “escape clause” provision, under which a firm is allowed to select a lower price in the event that a very favourable cost type is realized. The price reduction must be substantial, in order to ensure that the low price is attractive only when a firm’s cost type is low.

This behaviour is reminiscent of the findings of Rotemberg and Saloner (1986), but there are important differences: in their case, collusive prices adjust across *all* firms in response to a *public* demand shock. Proposition 9 also may be useful when interpreting an apparent episode of “cheating” in a collusive industry. Imagine, for example, a situation in which a single firm charges a low price and yet faces no retaliation. It is difficult to reconcile such an observation with standard collusion models. In our private-information setting, however, optimal collusion among impatient firms may allow for “rare exceptions” to rigidity, in which a firm substantially cuts its price and faces no retaliation.

20. The proof constructs a two-step pricing equilibrium for θ_2 close to $\underline{\theta}$. In this equilibrium, the low-step price is approximately $[r + \underline{\theta}(n-1)]/n$ indicating a discrete reduction of amount $(r - \underline{\theta})(n-1)/n$ from the high-step price of r . The two-step pricing scheme thus calls for a greater price reduction when markets are less concentrated.

The “small-density condition” (6.3) plays an intuitive role, but the assumption is restrictive. Obviously, it is satisfied if $f(\underline{\theta})$ is close enough to zero. To see a more subtle example, consider the (log-concave) distribution function family $F(\theta) = \theta^\alpha$, with $\underline{\theta} = 0 < 1 = \bar{\theta}$. The small-density condition is satisfied for any $\alpha > 1$, but it fails when $\alpha < 1$. The condition also fails when $\alpha = 1$ (corresponding to the uniform distribution). When (6.3) is violated, it may be that, for all $\delta < \delta^*$, no two-step pricing schedule satisfies on- and off-schedule constraints; this is the case for the uniform distribution (see Athey *et al.*, 1998).

A second example highlights an interesting prediction: if the support of the distribution increases, the optimal collusive scheme may switch from a rigid-pricing to a two-step pricing scheme. Consider a distribution $F(\theta; \mu, z)$, where the mean is constant at μ , but the support is parametrized by z , so that $\underline{\theta} = \mu - z$ and $\bar{\theta} = \mu + z$. Suppose $F(\theta; \mu, z)$ is log-concave and satisfies the small-density condition. An example is the “triangle” distribution, where the density $f(\theta; z)$ is symmetric about μ , and $f(\theta; z) = \frac{1}{z^2}(\theta - (\mu - z))$ on $[\mu - z, z]$.²¹ We make two observations. First, while increasing z leaves the per-period profit from rigid pricing unchanged, it introduces lower-cost types that are especially tempted to cheat and thereby increases the critical discount factor for rigid pricing, δ^* . Second, consider increasing z while holding the discount factor fixed at $\frac{n-1}{n} < \delta < 1$. When z is small, information is approximately complete and rigid pricing can be supported. At a critical z , however, the rigid-pricing scheme breaks down, and Proposition 9 implies that a two-price scheme is then optimal. Thus, increased “spread” in the cost distribution leads to increased price variability. In an application of this framework, Simon (2000) argues that inflation can lead to an increase in the spread of costs and establishes that prices are more variable when inflation is high.

6.3.2. Observable fluctuations in demand. The intuition underlying Proposition 9 suggests that any exogenous variation in the economic environment that heightens the short-term incentive to cheat and/or reduces the long-term value of co-operation may result in lower and more variable prices. A variation of particular empirical relevance occurs when industry demand fluctuates over time. Following Rotemberg and Saloner (1986), we now extend our model to an environment in which industry demand fluctuates in an i.i.d. fashion between low and high states, $\phi \in \{\phi_L, \phi_H\}$, where $\phi_H > \phi_L$. Profit is proportional to the demand state, which is publicly observed at the beginning of each period, before cost shocks are realized.

In this model, the long-term value of co-operation is proportional to the demand that is expected in future periods, $E\phi$, which is independent of the current demand state. By contrast, the incentive to cheat is greatest when current demand is high. The off-schedule constraint therefore binds first for the high-demand state. Formally, we may modify (6.1) to calculate for the rigid-pricing scheme a critical discount factor,

$$\delta_H^* \equiv \frac{(n-1)\pi(r, \underline{\theta})\phi_H}{(n-1)\pi(r, \underline{\theta})\phi_H + [\pi(r, E\theta) - n\pi^{NE}]E\phi},$$

where $\delta_H^* > \delta^*$, at which the off-schedule constraint binds when current demand is high. Similarly, we may define $\delta_L^* < \delta^*$ as the critical discount factor for the rigid-pricing scheme in the low-demand state. When the discount factor falls slightly below δ_H^* , it is no longer possible to enforce a rigid price for all cost levels in the high-demand state; however, it remains possible to do so when the demand state is low. Assuming that $\frac{\phi_H}{E\phi}$ is not too large, so that $1 > \delta_H^*$, a modification of Proposition 9 implies the following:

21. Note that $f(\underline{\theta})$ and $f(\bar{\theta})$ are equal to 0, which strictly speaking violates our maintained assumption; but it is straightforward to show that all of our results extend as long as $f'(\underline{\theta}) > 0$ and $f'(\bar{\theta}) < 0$.

Proposition 10. *For two firms, when the market size is i.i.d. with $\phi \in \{\phi_L, \phi_H\}$, if F is log-concave and (6.3) holds, then there exists $\delta^o \in [\delta_L^*, \delta_H^*]$, such that, for every $\delta \in (\delta^o, \delta_H^*)$, there exists an optimal SPPE that is stationary and satisfies:*

- (i) *in the low demand state ($\phi = \phi_L$), firms use the rigid-pricing scheme of Proposition 7;*
- (ii) *in the high demand state ($\phi = \phi_H$), firms use the two-price scheme of Proposition 9.*

This proposition extends a theme of the previous subsection: symmetric collusion between impatient firms may be marked by occasional (and perhaps substantial) price reductions by individual firms. We learn here that these departures are most likely to occur when one firm receives a favourable cost shock *and* current demand is high.

One implication of the model is that the countercyclical-pricing finding of Rotemberg and Saloner (1986) is robust to the presence of private cost fluctuations. This model can be generalized in a number of directions. For example, following Bagwell and Staiger (1997), we may consider an alternative stochastic process, in which the demand growth rate follows a Markov process, so that recessions are characterized by slow growth and booms are characterized by fast growth. In such a model, recessions are the time when collusion is most difficult. Thus, given the tradeoffs we outlined above, we would expect rigid prices in booms and variable prices in recessions. This is a striking and testable prediction. Indeed, Reynolds and Wilson (1998) use this stochastic process in their empirical work, and they find that in 14 out of 15 industries prices are more variable in recessions than in booms.²²

6.3.3. Optimal pricing for impatient firms. Next, we consider the general features of optimal SPPE pricing schemes for impatient firms. This problem is subtle, because a firm's incentive to deviate depends on its own cost type, its own price and the expected payoffs of the entire collusive agreement. The off-schedule constraint is thus analogous to a "participation constraint" in a static mechanism design model, except that the constraint depends on the type and the outside option is endogenous.

Our analysis is simplified by two observations. First, as established in Proposition 8, we may restrict attention to stationary pricing schemes. Second, we observe that all off-schedule incentive constraints are satisfied if they hold for the lowest-cost type on any step (*i.e.* for type θ_k on any step k defined by endpoints (θ_k, θ_{k+1}) over which the pricing schedule is flat). Clearly, this constraint is more difficult to satisfy for a given step as the step gets larger, since then the market-share gain from an off-schedule price cut is larger.

To understand the main tradeoffs, suppose that the distribution function is log-concave (at least in the relevant region) and consider whether a decrease in θ_k , and thus an increase in the length of step k , might be optimal. There are three effects. First, the off-schedule constraint for step k is exacerbated, since a deviation results in a larger increase in market share. Second, expected collusive profit increases, and this relaxes the off-schedule constraints. As in our two-step analysis, which of these two effects dominates depends on the shape of the distribution function. Finally, a multi-step scheme introduces a third effect as well: when an intermediate step is lengthened, the on-schedule constraints may require adjustments in prices on other steps, and these adjustments may in turn tighten the off-schedule constraints at these steps.

22. As a further extension, suppose the reservation value fluctuates in an i.i.d. fashion over time between $r_H = r + z$ and $r_L = r - z$, where each state occurs with probability 1/2 and $z \in [0, r - \bar{\theta}]$. A rigid-pricing scheme then entails always setting price equal to the current reservation value. While the per-period expected profit from this scheme is insensitive to z , the off-schedule constraint would bind for lower-cost types in state r_H if z were sufficiently large. The collusive price would then be reduced by individual firms in the "high-demand" state. The average collusive price and profit thus may decrease as the spread in reservation values increases.

The resolution of these tradeoffs depends on the shape of F . In our discussion paper (Athey *et al.*, 1998), we specify an optimization program which can be solved, either numerically or analytically, given specific functional forms and parameters. Here, we consider a more qualitative question: Are the off-schedule incentive constraints ever so severe that the firms are induced to use an interval of strictly increasing prices, even when the distribution is log-concave? The following proposition summarizes our findings.

Proposition 11. *Suppose $\delta > \frac{n-1}{n}$.*

- (i) *If $r > \bar{\theta}$, then $\bar{V}_s > \pi^{NE}/(1-\delta)$ and $p(\bar{\theta}) = r$.*
- (ii) *If $\bar{V}_s > \pi^{NE}/(1-\delta)$, and both F and $1-F$ are strictly log-concave, then in an optimal stationary SPPE there exists no open interval of types, (θ', θ'') , throughout which the pricing function is strictly increasing.*

Proposition 11(i) establishes that, when $\delta > \frac{n-1}{n}$ and $r > \bar{\theta}$, firms can achieve above-Nash payoffs and therefore enjoy at least partial collusion.²³ Intuitively, when $r > \bar{\theta}$, the introduction of a small interval of pooling for the highest-cost types at the price $p(\bar{\theta}) = r$ improves expected future profit; furthermore, when $\delta > \frac{n-1}{n}$, this improvement overwhelms the higher incentive to cheat that higher types then face. In part (ii), we assume directly that some collusion is attainable. We then show that, if both F and $1-F$ are strictly log-concave, then firms would do better to introduce tiny regions of pooling rather than strictly separate types throughout an open interval. The introduction of a small region of pooling has a first-order benefit for expected future profit when F is log-concave, and for a small step the gain in market share from undercutting the collusive price is small. There remains, however, the third effect mentioned above, associated with cross-step externalities and off-schedule constraints. For any particular type θ , an off-schedule constraint might bind above or below θ , and our assumption that both F and $1-F$ are strictly log-concave ensures that pooling is optimal, whether the cross-step externality extends to the fraction F of lower types or the fraction $1-F$ of higher types.²⁴

Under the conditions of Proposition 11, then, in an optimal stationary collusive scheme, there is no open interval of cost types where efficiency benefits are attained. An optimal SPPE pricing scheme thus exhibits rigidity over regions of costs, and the observed distribution of prices will have mass points. Such an observation may offer guidance in interpreting allegations of collusion. For example, in the NASDAQ stock exchange (see, *e.g.* Christie and Schultz, 1999), dealers systematically restricted their price quotes to multiples of \$0.25, and such behaviour was associated with higher average bid-ask spreads (and thus, presumably, higher profits).

6.4. Downward-sloping demand

In our discussion paper (Athey *et al.*, 1998), we also analyse the case of downward-sloping demand. Here, we simply make two points. First, the results derived in this section (Propositions 7–11) all generalize to the case in which demand is sufficiently inelastic that rigid pricing

23. Notice that this result complements Proposition 5(ii) by offering a specific lower bound for δ , under the further assumption that $r > \bar{\theta}$.

24. More precisely, for $\theta > \theta_j$, the on-schedule constraint implies $\Pi(\theta, \theta; p) = \Pi(\theta_j, \theta_j; p) - \int_{\theta_j}^{\theta} M(\tilde{\theta}; p) d\tilde{\theta}$, from which we may derive $E_{\theta}[\Pi(\theta, \theta; p) | \theta > \theta_j] = \Pi(\theta_j, \theta_j; p) - E_{\theta}[\frac{1-F(\theta)}{f(\theta)} M(\theta; p) | \theta > \theta_j]$. If $1-F$ is strictly log-concave, so that $-\frac{1-F(\theta)}{f(\theta)}$ is strictly increasing, the expected profit for $\theta > \theta_j$ is maximized when these types are pooled.

is optimal for patient firms.²⁵ Second, while our no-wars finding (Proposition 8) holds even for impatient firms when demand is sufficiently inelastic, we cannot rule out the possibility of a war in an optimal SPPE for impatient firms and general demand functions.

7. CONCLUSION

We propose a model of collusion in which firms are privately informed as to their current cost positions. Under the assumption of inelastic demand, we establish five main findings:

- (1) Firms fare poorly under fully sorting symmetric collusive schemes, since the efficiency benefits that such schemes afford are small relative to the informational costs.
- (2) Optimal symmetric collusion can be achieved without equilibrium-path price wars.
- (3) If firms are sufficiently patient and the distribution of costs is log-concave, optimal symmetric collusion is characterized by price rigidity and the absence of price wars on the equilibrium path.
- (4) If firms are less patient, optimal symmetric collusion may be characterized by price rigidity over intervals of costs (a step function), where the price of a lower-cost firm is distorted downward to diminish the incentive that such a firm has to cheat.
- (5) If firms are less patient and the model is modified to include i.i.d. public demand shocks, under optimal symmetric collusion, the downward pricing distortion that accompanies a firm's lower-cost realization may occur only when current demand is high.

We note that the first finding underscores the basic tradeoff present in our repeated-adverse-selection model of collusion; the second finding contrasts with the Green and Porter (1984) collusion literature on repeated moral hazard; the third finding offers an equilibrium interpretation of the empirical association between rigid pricing and industry concentration, as well as the commonly observed collusive practices of identical bidding and price fixing with stable market shares; and the fourth and fifth findings are reminiscent of the logic developed by Rotemberg and Saloner (1986) for collusion in markets with publicly observed demand shocks, but associate low collusive prices with individual firm behaviour in high-public-demand and low-private-cost states.

We also discuss an extended model with downward-sloping demand. When demand is sufficiently inelastic, the firms use rigid pricing and avoid price wars. In addition, we highlight some novel features that arise when demand slopes down.

Our analysis contributes as well at a methodological level. We develop the precise connections between static and dynamic analyses, making clear the similarities and differences, and laying the groundwork for treating other repeated-game problems within the mechanism design framework. Our work also motivates some new questions for static mechanism design, and takes some steps towards addressing them. For example, we examine how restrictions on transfers affect optimal mechanisms.

While we study collusion, our analysis provides a foundation for other applications. In one family of applications, players interact repeatedly but a single player observes private information in each period (*e.g.* a government in a "policy game"). The question arises as to whether the informed player should follow a "rule" (*i.e.* adopt behaviour that is never responsive to private information) or be granted "discretion" (*i.e.* adopt behaviour that is sometimes responsive to private information). In the latter case, incentive compatibility may require that the informed

25. Showing that repeated play of the static Nash equilibrium yields the worst punishment when F is log-concave and demand is sufficiently inelastic entails additional work. Details are available from the authors.

player is sometimes punished. Our analysis of SPPE provides a foundation for such applications, since in each case all informed players bear a symmetric punishment.

APPENDIX

Proof of Lemma 3. We prove this result allowing for the possibility of downward-sloping demand, considered in a subsequent extension of the model. Letting p^m be the monopoly pricing function, we assume that $\pi(\rho, \theta) = (\rho - \theta)D(\rho)$ is strictly quasiconcave and differentiable in ρ , where $D(\rho) \equiv 1$ for $\rho \leq r \equiv p^m(\theta)$ when demand is inelastic. Then, for any $x \leq \pi(p^m(\theta), \theta)$, we can find the lowest price ρ that satisfies $\pi(\rho, \theta) = x$.

We begin with the case where p is a step function. Let p_k and M_k denote the price and market share on step k , where $k = 1, \dots, K$ and step k covers $(\theta_k, \theta_{k+1}]$. We seek a new pricing function, \tilde{p} , such that $(\tilde{p}, \tilde{T} \equiv 0)$ satisfies (IC-onM) and $M(\theta; p) = M(\theta; \tilde{p})$, while preserving the profit-at-the-top. The function \tilde{p} is described by a K -dimensional price vector, (ρ_1, \dots, ρ_K) , which we now construct. For the highest step K , observe that $\Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) \leq \pi(p^m(\bar{\theta}), \bar{\theta})M_K$. Recall our restriction that the lowest type on the highest step makes non-negative profit: $\Pi(\theta_K, \theta_K; p) - T(\theta_K) \geq 0$. This implies $p_K \geq \theta_K$, and thus $D(p_K) \leq D(\theta_K)$, from which it follows that

$$\begin{aligned} \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) &\equiv \pi(p_K, \bar{\theta})M_K - T(\bar{\theta}) = [(\theta_K - \bar{\theta})D(p_K) + (p_K - \theta_K)D(p_K)]M_K - T(\bar{\theta}) \\ &\geq [(\theta_K - \bar{\theta})D(\theta_K) + (p_K - \theta_K)D(p_K)]M_K - T(\bar{\theta}) \\ &= \pi(\theta_K, \bar{\theta})M_K + \Pi(\theta_K, \theta_K; p) - T(\theta_K) \geq \pi(\theta_K, \bar{\theta})M_K. \end{aligned}$$

Thus, we can find $\rho_K \in [\theta_K, p^m(\bar{\theta})]$ such that $\Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = \pi(\rho_K, \bar{\theta})M_K$. For steps $k = K - 1, \dots, 1$, we define ρ_k inductively as the lowest price satisfying

$$\pi(\rho, \theta_{k+1})M_k = \pi(\rho_{k+1}, \theta_{k+1})M_{k+1}.$$

Since $M_k > M_{k+1}$ and π is strictly increasing in ρ for $\rho < p^m(\theta)$, it follows that $\rho_k < \rho_{k+1}$.

When demand is inelastic, inspection of (IC-onM) implies that the constructed price sequence satisfies $\rho_k = p_k - T_k/M_k$, where T_k denotes the war on step k . Note that $\text{sign}\{p(\theta) - \tilde{p}(\theta)\} = \text{sign}\{T(\theta)\}$.

General functions p can now be treated using standard limiting arguments.²⁶ ||

Lemma 5. *Suppose demand is downward-sloping and consider any distribution function F . Let (p, T) satisfy the constraints of the Mechanism Design Program and let $(\tilde{p}, \tilde{T} \equiv 0)$ be the constructed no-wars mechanism from Lemma 3. Let $Y(\theta) = \Pi(\theta, \theta; \tilde{p}) - [\Pi(\theta, \theta; p) - T(\theta)]$ be the pointwise net advantage of the constructed no-wars schedule \tilde{p} .*

- (i) *Y is differentiable almost everywhere, and where it exists, $\text{sign}[Y'(\theta)] = \text{sign}[\tilde{p}(\theta) - p(\theta)]$: the net advantage increases where and only where \tilde{p} exceeds p .*
- (ii) *For any θ' where p is strictly increasing or jumps up, if T does not change at θ' , then $p(\theta') < p^m(\theta')$.*
- (iii) *Suppose that $\tilde{p}(\underline{\theta}) \leq p^m(\underline{\theta})$. Then the no-wars mechanism is pointwise as good as its with-wars counterpart: $\Pi(\theta, \theta; \tilde{p}) \geq [\Pi(\theta, \theta; p) - T(\theta)]$, for all θ , and further, the no-wars mechanism is strictly better if $T > 0$ on a set of positive measure.*

Proof. (i) By Milgrom and Segal's (2002) envelope theorem, Y is differentiable almost everywhere. Using Lemma 2 and the fact that p and \tilde{p} are market share equivalent (i.e. $M(\cdot; \tilde{p}) = M(\cdot; p)$), where it exists,

$$Y'(\theta) = \Pi_{\theta}(\theta, \theta; \tilde{p}) - \Pi_{\theta}(\theta, \theta; p) = M(\theta; p)(D(p(\theta)) - D(\tilde{p}(\theta))).$$

The result follows since demand is downward-sloping. (ii) If $p(\theta' + \varepsilon) > p(\theta') \geq p^m(\theta')$, then $M(\theta'; p) > M(\theta' + \varepsilon; p)$ and $\pi(p(\theta'), \theta') > \pi(p(\theta' + \varepsilon), \theta')$, by the strict quasiconcavity of π . But such a pricing function cannot be incentive compatible unless T is decreasing at θ' , since the lower price offers higher market share and higher profit-if-win. (iii) It suffices to show $\min_{\theta} Y(\theta) \geq 0$. If $\bar{\theta}$ minimizes Y , then we are done since \tilde{p} was constructed so that $Y(\bar{\theta}) = 0$. A necessary condition for a minimum at θ^* in the interior $(\underline{\theta}, \bar{\theta})$ is $Y'(\theta^*) = 0$. But then part (i) gives $p(\theta^*) = \tilde{p}(\theta^*)$ and so $\Pi(\theta^*, \theta^*; \tilde{p}) = \Pi(\theta^*, \theta^*; p)$. Given $T(\theta^*) \geq 0$, we conclude $Y(\theta^*) \geq 0$. Lastly, if $\underline{\theta}$ minimizes Y , then $Y'(\underline{\theta}) \geq 0$, implying $p(\underline{\theta}) \leq \tilde{p}(\underline{\theta})$. Since $\tilde{p}(\underline{\theta}) \leq p^m(\underline{\theta})$, we conclude $\Pi(\underline{\theta}, \underline{\theta}; \tilde{p}) \geq \Pi(\underline{\theta}, \underline{\theta}; p)$, and so $Y(\underline{\theta}) \geq 0$. Now suppose

26. For any non-decreasing function, find a sequence of step functions that converges to p pointwise, such that each element of the sequence is equal to p on intervals where p is constant. Following the construction above, then find the corresponding sequence of functions such that $(\tilde{p}, \tilde{T} \equiv 0)$ satisfies (IC-onM), and let \tilde{p}^* be the limit of a convergent subsequence. (IC-onM) is satisfied for \tilde{p}^* , since (IC-onM) is a weak inequality restriction, and since limiting payoffs are continuous in prices where p is strictly increasing.

that $T(\theta) > 0$ on a set of positive measure. If (p, T) is optimal, since the no-war scheme is weakly better pointwise, it follows that (a) $\Pi(\theta, \theta; p) - T(\theta) = \Pi(\theta, \theta; \bar{p})$ almost everywhere. But this implies $\bar{p}(\theta) \neq p(\theta)$ almost everywhere that $T(\theta) > 0$. From this and part (i) it follows that (b) $Y'(\theta) \neq 0$ on a set of positive measure. Since by Lemma 2, $Y(\theta) = Y(\bar{\theta}) + \int_{\bar{\theta}}^{\theta} Y'(\tilde{\theta})d\tilde{\theta}$, (b) is inconsistent with $Y(\theta) = 0$ almost everywhere. \parallel

Proof of Proposition 6. (i) By Spulber (1995), $p^e(\theta) \leq p^m(\theta)$. Then Lemma 5 (iii) applies to any fully sorting (p, T) and $(p^e, T^e \equiv 0)$ and so the latter is the unique solution to the Mechanism Design Program among fully sorting price schedules. The argument that p^e must then be the optimal fully sorting SPPE follows as in Proposition 2. (ii) Consider a family of demand functions normalized so that $D(0) = 1$. Then, expected profit is given by

$$E_{\theta}[\Pi(\theta, \theta; p) - T(\theta)] = \Pi(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + E_{\theta} \left[M(\theta; p) \frac{F}{f}(\theta) \right] + E_{\theta} \left[M(\theta; p)[D(p(\theta)) - 1] \frac{F}{f}(\theta) \right]$$

for each member of the family. Assuming log-concavity and inelastic demand, we show in Proposition 5 that a rigid price at the reservation value uniquely maximizes the first two terms. As demand becomes more inelastic, $D(p(\theta))$ approaches 1 for all prices below the reservation value, and so the first two terms dominate. The level of the optimal rigid price approaches the reservation value as demand becomes inelastic; further, as in Proposition 5, this scheme can be enforced by sufficiently patient firms. \parallel

Lemma 6. Take $p^* \in \arg \min_p E_{\theta} \left[\frac{F}{f}(\theta) M(\theta; p) \right]$ s.t. $\Pi(\theta, \theta; p) \geq \Pi(\hat{\theta}, \theta; p)$ for all $\hat{\theta}, \theta$, and let $\hat{R} = E_{\theta} \left[\frac{F}{f}(\theta) M(\theta; p^*) \right]$. Suppose that there exists an $x \leq \bar{\theta}$ such that p^* is increasing on $[\underline{\theta}, x]$ and constant on $(x, \bar{\theta}]$.²⁷

- (i) If $r > \bar{\theta}$, $\delta > \frac{n-1}{n}$ and $x < \bar{\theta}$, $\hat{R}/(1-\delta) \leq \underline{V}_s < \pi^{NE}/(1-\delta)$.
- (ii) If firms are sufficiently patient, $\underline{V}_s = \hat{R}/(1-\delta) \in \mathcal{V}_s$.
- (iii) If F is log-concave ($x = \bar{\theta}$), then for all discount factors δ , $\underline{V}_s = \pi^{NE}/(1-\delta) = \hat{R}/(1-\delta)$.

Proof. We begin with part (ii). Suppose that (p, v) implements \underline{V}_s . The off-schedule constraint for $\bar{\theta}$ requires that $-(\bar{\theta} - p(\bar{\theta}))M(\bar{\theta}; p) + \delta v(p(\bar{\theta}); p) \geq \delta \underline{V}_s$. Using (5.1), the on-schedule constraints imply that

$$-(\bar{\theta} - p(\bar{\theta}))M(\bar{\theta}; p) + \delta v(p(\bar{\theta}); p) + E_{\theta} \left[\frac{F}{f}(\theta) M(\theta) \right] = \underline{V}_s. \quad (\text{A.1})$$

Substituting yields

$$E_{\theta} \left[\frac{F}{f}(\theta) M(\theta; p) \right] \leq (1-\delta)\underline{V}_s. \quad (\text{A.2})$$

Thus, no scheme can yield a lower per-period continuation value than \hat{R} .

Consider the following scheme. In the first period, firms use the pricing scheme $p^L(\theta)$, where $p^L(\bar{\theta}) = \frac{1}{n}\bar{\theta} + \frac{n-1}{n}x$ and the pricing scheme is rigid on $[x, \bar{\theta}]$ and strictly increasing elsewhere. Under the assumptions of the lemma, the pricing scheme p^L minimizes expected informational rents. The continuation value function v^L specifies $v^L(\mathbf{p}(\theta)) = \underline{V}_s$ if $\theta_j < x$ for any firm j (i.e. if the market price is below $p^L(\bar{\theta})$), while $v^L(\mathbf{p}(\theta)) = v_r$ otherwise (i.e. if the market price is $p^L(\bar{\theta})$). Off-schedule deviations are punished by returning to \underline{V}_s . We choose v_r to satisfy with equality the off-schedule constraint for $\bar{\theta}$, so that profit-at-the-top is $\delta \underline{V}_s$. The on-schedule constraints then imply that the payoff from this collusive scheme is $\delta \underline{V}_s + \hat{R}$. Now, we know from (A.2) that $\delta \underline{V}_s + \hat{R} \leq \underline{V}_s$. Further, if this collusive scheme is an SPPE scheme, then the definition of \underline{V}_s implies that $\delta \underline{V}_s + \hat{R} \geq \underline{V}_s$. Hence, if this collusive scheme satisfies all other off-schedule incentive constraints, and if firms are sufficiently patient that $v_r \in \mathcal{V}_s$, then $\underline{V}_s = \hat{R}/(1-\delta) \in \mathcal{V}_s$.

To see that the proposed scheme satisfies all other constraints, notice that since $p^L(\bar{\theta}) > x$, on-schedule incentive compatibility determines p^L for $\theta < x$; in particular, it requires that $p^L(\theta) > \theta$ for all θ . Since the pricing function is strictly increasing on this interval, there are no additional off-schedule constraints. Further, as may be confirmed, at the price $p^L(\bar{\theta}) = \frac{1}{n}\bar{\theta} + \frac{n-1}{n}x$, the off-schedule constraint for x (who is tempted to under-cut) is satisfied exactly when it is for $\bar{\theta}$ (who is tempted to raise price).²⁸

27. This is true, for example, if F/f is non-decreasing on $[\underline{\theta}, x]$ and non-increasing on $(x, \bar{\theta}]$, and further, $\sup_{\theta \in [\underline{\theta}, x]} F(\theta)/f(\theta) \leq \inf_{\theta \in [x, \bar{\theta}]} F(\theta)/f(\theta)$. Then, a pricing scheme of the shape described puts the maximum market share on the lowest realizations of $F(\theta)/f(\theta)$.

28. We note that the price described requires the least patience for implementation. If we were to raise $p^L(\bar{\theta})$, this would increase the incentive of type x to undercut the collusive price, which would tighten the off-schedule constraint for type x and thereby require a greater future reward v_r . Thus, overall firm profits go up. Lowering $p^L(\bar{\theta})$ relaxes the

Part (iii) follows from part (ii), together with the fact that repeated play of the static Nash equilibrium is always an SPPE. For part (i), observe that firms can use a flat scheme on an interval $[y, \bar{\theta}]$. Setting y close to $\bar{\theta}$, the incentive to deviate off-schedule can be made arbitrarily small. Proposition 11 establishes that some future reward greater than $\pi^{NE}/(1-\delta)$ is available when $\delta > \frac{n-1}{n}$ and $r > \bar{\theta}$. \parallel

Proof of Proposition 7. Part (i) follows directly from Lemma 6. Parts (ii) and (iii) follow from (i) and the analysis provided in the text.

Proof of Proposition 8. Let (p, v) be the factorization of an original scheme that constitutes an optimal SPPE. As before, let $T(\theta) = \delta[\bar{V}_s - \bar{v}(p(\theta); p)]$ and note that (p, T) must satisfy the constraints of the Mechanism Design Program. Using Lemma 3 and exploiting the inelastic-demand assumption, we can find an alternative scheme, $(\tilde{p}, \tilde{T} \equiv 0)$, such that $\Pi(\theta, \theta; p) - T(\theta) = \pi(\tilde{p}(\theta), \theta)M(\theta; p)$ for all θ . Clearly, this implies that $\tilde{p}(\theta) = p(\theta)$ for any θ such that $T(\theta) = 0$ while $T(\theta) > 0$ implies that $\tilde{p}(\theta) < p(\theta)$. The factorization associated with the alternative scheme is $(\tilde{p}, \tilde{v} \equiv \bar{V}_s)$.

Now, we compare the off-schedule incentive constraints across the two schemes. To do so, let $w \in \mathcal{V}_s$ be the punishment used following an off-schedule deviation in the original scheme. When the on-schedule incentive constraint is met, the off-schedule incentive constraint is satisfied if the lowest type on any interval of costs over which prices are rigid does not prefer to charge a slightly lower price.²⁹ As the original schedule satisfies the off-schedule incentive constraints, it follows that the lowest-cost type θ_k on any step k will not choose to slightly undercut the step- k price, $p(\theta_k)$:

$$\pi(p(\theta_k), \theta_k)[(1 - F(\theta_k))^{n-1} - M(\theta_k; p)] \leq \delta\{\bar{V}_s - T(\theta)/\delta - w\}. \quad (\text{A.3})$$

Now consider the analogous off-schedule incentive constraint for the alternative schedule:

$$\pi(\tilde{p}(\theta_k), \theta_k)[(1 - F(\theta_k))^{n-1} - M(\theta_k; p)] \leq \delta\{\bar{V}_s - w\}. \quad (\text{A.4})$$

In (A.3) and (A.4), the L.H.S.'s represent the current-period incentive to cheat. This incentive is either the same under the two schedules (when $T(\theta_k) = 0$) or strictly lower under the alternative schedule (when $T(\theta_k) > 0$, since then $\tilde{p}(\theta_k) < p(\theta_k)$). The R.H.S.'s represent the expected discounted values of co-operation in the next and all future periods. The R.H.S. is also either the same (when $T(\theta_k) = 0$) or strictly higher under the alternative schedule (when $T(\theta_k) > 0$). Thus, eliminating the war in this way simultaneously raises the expected discounted value of co-operation and lowers the current incentive to cheat; as a consequence, if there is no incentive to undercut in the original collusive arrangement, then there will certainly be no such incentive under the alternative arrangement. We then conclude that $(\tilde{p}, \tilde{v} \equiv \bar{V}_s)$ satisfies all of the constraints of the Interim Program. This in turn implies that $\bar{V}_s = \Pi(\theta, \theta; \tilde{p})/(1-\delta)$, corresponding to an optimal, stationary SPPE where \tilde{p} is used in every period. \parallel

Proof of Proposition 9. Consider a candidate two-price scheme, denoted $\check{p}(\theta; \theta_2)$, with a top-step price $\rho_2 \in (\underline{\theta}, r]$ and a breakpoint $\theta_2 \in (\underline{\theta}, \bar{\theta})$. Given ρ_2 and θ_2 , the on-schedule incentive constraint determines the low-step price, ρ_1 , as the ρ_1 that solves

$$\pi(\rho_1, \theta_2)\mu(\underline{\theta}, \theta_2) = \pi(\rho_2, \theta_2)\mu(\theta_2, \bar{\theta}), \quad (\text{A.5})$$

where $\mu(\theta_k, \theta_{k+1})$ is the market share for cost types on a step on $[\theta_k, \theta_{k+1}]$. If the lowest step is small enough, the gain in market share from undercutting ρ_1 is small, and the binding off-schedule constraint is the constraint for θ_2 , the lowest type on the top step. The off-schedule constraint is written:

$$(\rho_2 - \theta_2)[\mu(\theta_2, \bar{\theta}) - (1 - F(\theta_2))^{n-1}] + \frac{\delta}{1-\delta}[E_\theta \Pi(\theta, \theta; \check{p}(\cdot; \theta_2)) - \pi^{NE}] \geq 0 \quad (\text{A.6})$$

off-schedule constraint for type x ; but it tightens the off-schedule constraint for type $\bar{\theta}$, again requiring a greater future reward v_r . However, the required increase in the reward is exactly equal to the reduction in firm profits, so that the overall scheme implements the same value. Yet, a reward of greater magnitude is required, and for firms of moderate patience, a greater reward may not always be available.

29. This statement follows from four observations. First, over a segment for which price is strictly increasing, the lowest type clearly has nothing to gain from a small price cut. Second, over a segment for which price is flat, the on-schedule incentive constraint requires that either the lowest type on this segment is θ or that the price schedule jumps discontinuously down for lower types. Third, over a segment for which price is flat, the incentive to undercut is greatest for the lowest-cost type. Together, these observations imply that the collusive scheme is robust against off-schedule price-cutting deviations, so long as the lowest type on a flat segment does not choose to cut price slightly. Fourth, off-schedule price-increasing deviations are unattractive under the alternative schedule (given that this is true for the original schedule and given that on-schedule constraints are satisfied). Details associated with this final observation are in our working paper.

At $\theta_2 = \underline{\theta}$, $\rho_2 = r$ and $\delta = \delta^*$, (A.6) holds with equality. Notice that increasing θ_2 affects both the incentive to cheat as well as the future value of co-operation. Since introducing a small lower step results in a low price for only a correspondingly small region of types, $E_{\theta} \Pi(\theta, \theta; \check{p}(\cdot; \theta_2))$ decreases smoothly in θ_2 , and thus (A.6) is differentiable in θ_2 at $\theta_2 = \underline{\theta}$. Taking the derivative of the L.H.S. of (A.6) with respect to θ_2 , evaluated at $\theta_2 = \underline{\theta}$ and $\rho_2 = r$, and using (A.5) yields

$$\frac{n-1}{n} \left[f(\underline{\theta})(n-1)(r-\underline{\theta}) + 1 - \frac{\delta}{1-\delta}(r-E\underline{\theta})f(\underline{\theta}) \right].$$

This expression is non-negative at $\delta = \delta^*$ if and only (6.3) holds. Thus, we conclude that increasing θ_2 at $\theta_2 = \underline{\theta}$ relaxes the off-schedule incentive constraint.

When the distribution function is log-concave, rigid pricing at r is optimal so long as $\delta = \delta^*$. This implies that $E_{\theta} \Pi(\theta, \theta; \check{p}(\cdot; \theta_2))$ is maximized at $\theta_2 = \underline{\theta}$ and/or $\bar{\theta}$. Consider first $\theta_2 = \bar{\theta} + \varepsilon$ for ε small. When introducing a small step at or near $\bar{\theta}$, the low-step price must be set to satisfy (A.5). As $\mu(\theta_2, \bar{\theta}) \rightarrow 0$ when $\varepsilon \rightarrow 0$, (A.5) implies that $\rho_1 \rightarrow \bar{\theta}$ when $\varepsilon \rightarrow 0$. Since this low price is used by all types on $[\underline{\theta}, \theta_2)$, expected per-period profits are approximately $\pi(\bar{\theta}, E\underline{\theta})/n$ for a small step. Consider second $\theta_2 = \underline{\theta} - \varepsilon$ for ε small. We may then set $\rho_2 = r$, with ρ_1 then determined by (A.5). Since the high price is used by all types on $(\theta_2, \bar{\theta}]$, expected per-period profits are approximately $\pi(r, E\underline{\theta})/n$, as in the rigid-pricing scheme, for a small step. Recalling that $r \geq \bar{\theta}$, we conclude that $E_{\theta} \Pi(\theta, \theta; \check{p}(\cdot; \theta_2))$ is maximized at $\theta_2 = \underline{\theta}$.

Finally, we must show that when δ is close to δ^* , the optimal two-step scheme dominates any other scheme with other market-share allocation schemes. Consider an alternative scheme that has more than one point of strict increase; for simplicity, suppose that this pricing scheme is a step function. Consider the bottom two steps of such a scheme. If we condition on types on the bottom two steps of this scheme, Proposition 5 can be applied without further modifications. When the distribution is log-concave, it is always better two combine the two bottom steps into one (thus reallocating market share to types with a higher $F(\theta)/f(\theta)$); or to turn a region of strict increase into one of pooling. Proceeding in this manner, a multi-step scheme is successively dominated by schemes with fewer steps. Thus, the best two-step scheme, if it clears the off-schedule constraints, is better than any other scheme with a larger number of steps. Since a two-step scheme with a small lower step is better than any other two-step scheme, and (6.3) and δ close to δ^* imply that this scheme can be supported, it must also dominate any schemes with a larger number of steps. \parallel

Proof of Proposition 11. We treat part (ii) first. Recall the notation developed above for pricing functions that can be represented by a finite number of subintervals. Let $M_k \equiv \mu(\theta_k, \theta_{k+1})$ be the market share for cost types on a step on $[\theta_k, \theta_{k+1}]$. Consider a candidate solution which specifies strictly increasing pricing on interval $(\theta_{j-1}, \theta_j]$. Consider introducing a tiny step on the interval $(\theta_j - \varepsilon, \theta_j]$. If there is some gain to future co-operation, then for ε small enough, introducing a tiny step does not introduce a new off-schedule incentive constraint, at least not one which is binding. Suppose it were indeed optimal for $(\theta_{j-1}, \theta_j]$ to be a region of separation of types. Then, if we chose ε to maximize the objective when $(\theta_{j-1}, \theta_j - \varepsilon]$ is a region of separation and pooling takes place on $(\theta_j - \varepsilon, \theta_j]$, the solution $\varepsilon = 0$ should be a local maximum.

Note that $\rho_k - \theta_k = \Pi(\theta_k, \theta_k; p)/M_k$. Rearranging, we thus observe that the off-schedule constraint for type θ_k , at the left endpoint of step k , can be written as follows:

$$\frac{\delta}{1-\delta} (E_{\theta} [\Pi(\theta, \theta; p)] - \mathcal{V}_s) - \Pi(\theta_k, \theta_k; p) \frac{(1-F(\theta_k))^{n-1} - M_k}{M_k} \geq 0. \tag{A.7}$$

We can then express the off-schedule constraint for step $k < j$ as a function of ε . Tedious calculations (using in particular the fact that $\mu_1(\theta_j, \theta_j) = \frac{1}{2} \frac{\partial}{\partial \theta_j} (1-F(\theta_j))^{n-1}$, and $\mu_{11}(\theta_j, \theta_j) = \frac{1}{3} \frac{\partial^2}{\partial \theta_j^2} (1-F(\theta_j))^{n-1} - \frac{1}{6} f'(\theta_j)(n-1)(1-F(\theta_j))^{n-2}$) establish that the first and second derivatives of (A.7) with respect to ε are equal to 0 when $\varepsilon = 0$, while the third derivative is given by

$$\frac{\delta}{1-\delta} [f(\theta_j)^2 - F(\theta_j)f'(\theta_j)] + f'(\theta_j) \frac{(1-F(\theta_k))^{n-1} - M_k}{M_k}. \tag{A.8}$$

If (A.8) is positive, $\varepsilon = 0$ is a local minimum; in that case as ε increases, the off-schedule constraint is relaxed, so that the strictly increasing scheme is dominated. If the density is non-decreasing, (A.8) is positive by log-concavity of the distribution. Suppose now that $f'(\theta_j) < 0$. When $\delta > \frac{n-1}{n}$, since $\frac{(1-F(\theta_k))^{n-1} - M_k}{M_k} < n-1$, (A.8) is positive if $f(\theta_j)^2 - F(\theta_j)f'(\theta_j) \geq -f'(\theta_j)$. But this is true whenever $1-F$ is log-concave.

Now return to part (i). Using a similar logic, we consider starting from a strictly increasing pricing schedule (*i.e.* the static Nash equilibrium) and introducing a small step on $[\bar{\theta} - \varepsilon, \bar{\theta}]$ at $p(\bar{\theta}) = r$. This will improve per-period profits if $r > \bar{\theta}$. We must then verify that it improves per-period profits faster than it tightens the off-schedule constraint for this upper step. Tedious calculations show that the first $(n-2)$ derivatives of the off-schedule constraint with respect to ε are

zero, while the $(n-1)$ -th derivative is equal to $\frac{n-1}{n} \left(\frac{\delta}{1-\delta} - (n-1) \right) (r - \bar{\theta}) f(\bar{\theta})^2 > 0$. Thus, $\varepsilon = 0$ is a local minimum. Since the off-schedule constraint is satisfied at $\varepsilon = 0$, introducing a small step relaxes it. Finally, for a given market-share allocation that has a step at the top (of arbitrary size), consider the optimal $p(\bar{\theta})$. Recalling (4.2), increasing $p(\bar{\theta})$ increases $E_{\theta}[\Pi(\theta, \theta; p)]$ at the same rate that it increases $\Pi(\theta_k, \theta_k; p)$. Since $\frac{(1-F(\theta_k))^{n-1} - M_k}{M_k} < n-1$, inspection of (A.7) implies that increasing $p(\bar{\theta})$ increases the future value of co-operation, $\frac{\delta}{1-\delta} E_{\theta} \Pi(\theta, \theta; p)$, faster than the incentive to deviate off-schedule whenever $\delta > \frac{n-1}{n}$. \parallel

Acknowledgements. We are grateful to Mark Armstrong, Doug Bernheim, Dennis Carlton, Glenn Ellison, Eric Maskin, Wally Mullin, David Pearce, Michael Riordan, Jesse Schwartz, Scott Stern, Lars Stole, Jean Tirole, three anonymous referees and seminar participants at Berkeley, Bilkent, Boston College, Canada Competition Bureau, Columbia, Georgetown, Harvard, ITAM, MIT, NYU, Northwestern, Oxford, Penn, Pompeu Fabra, SMU, Stanford, Texas, Toronto, Toulouse, UCL, UWO and Vanderbilt for helpful discussions. We thank the National Science Foundation (Athey: SBR-9631760 and SES-9983820; Bagwell: SES-9905460) for generous financial support. Athey thanks the Cowles Foundation at Yale University for hospitality and support.

REFERENCES

- ABREU, D., PEARCE, D. and STACCHETTI, E. (1986), "Optimal Cartel Equilibria with Imperfect Monitoring", *Journal of Economic Theory*, **39** (1), 251–269.
- ABREU, D., PEARCE, D. and STACCHETTI, E. (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", *Econometrica*, **58** (5), 1041–1063.
- ATHEY, S. and BAGWELL, K. (2001), "Optimal Collusion with Private Information", *RAND Journal of Economics*, **32** (3), 428–465.
- ATHEY, S. and BAGWELL, K. (2002), "Optimal Collusion with Persistent Cost Shocks" (Mimeo, Stanford University).
- ATHEY, S., BAGWELL, K. and SANCHIRICO, C. (1998), "Collusion and Price Rigidity" (Working Paper No. 98-23, MIT) available at www.ssrn.com.
- BAGWELL, K. and STAIGER, R. W. (1997), "Collusion Over the Business Cycle", *RAND Journal of Economics*, **28** (1), 82–106.
- BULOW, J. and ROBERTS, J. (1989), "The Simple Economics of Optimal Auctions", *Journal of Political Economy*, **97** (5), 1060–1090.
- BUSINESS WEEK (1975), "Price Fixing: Crackdown Underway", 2 June, 47–48.
- BUSINESS WEEK (1998), "Justice's Cartel Crackdown", 27 July, 50–51.
- CARLTON, D. W. (1986), "The Rigidity of Prices", *American Economic Review*, **76** (4), 637–658.
- CARLTON, D. W. (1989), "The Theory and the Facts of How Markets Clear: Is Industrial Organization Valuable for Understanding Macroeconomics?", in R. Schmalensee and R. D. Willig (eds.) *Handbook of Industrial Organization, Volume 1*, Handbooks in Economics, No. 10 (Oxford: North-Holland) 909–946.
- CHAKRABARTI, S. (2001), "Information Revelation and Collusion in Oligopolies with Unknown Costs" (Working Paper, Cornell University).
- CHRISTIE, W. and SCHULTZ, P. (1999), "The Initiation and Withdrawal of Odd-Eighth Quotes Among NASDAQ Stocks: An Empirical Analysis", *Journal of Financial Economics*, **52** (3), 409–442.
- COMPTE, O. (1998), "Communication in Repeated Games with Private Monitoring", *Econometrica*, **66** (3), 597–626.
- CRAMTON, P. C. and PALFREY, T. R. (1990), "Cartel Enforcement with Uncertainty About Costs", *International Economic Review*, **31** (1), 17–47.
- ECKBO, P. L. (1976) *The Future of World Oil* (Cambridge, MA: Ballinger).
- EUROPEAN COMMISSION (1994), "The Commission Condemns a Cartel of European Cartonboard Producers and Imposes Substantial Fines", Press Release, 13 July.
- FUDENBERG, D., LEVINE, D. J. and MASKIN, E. (1994), "The Folk Theorem with Imperfect Public Information", *Econometrica*, **62** (5), 997–1039.
- GREEN, E. J. and PORTER, R. H. (1984), "Noncooperative Collusion Under Imperfect Price Information", *Econometrica*, **52** (1), 87–100.
- HALL, R. L. and HITCH, C. J. (1939), "Price Theory and Business Behavior", *Oxford Economic Papers*, **2**, 12–45.
- HALTIWANGER, J. and HARRINGTON, J. (1991), "The Impact of Cyclical Demand Movements on Collusive Behavior", *RAND Journal of Economics*, **22** (1), 89–106.
- HAY, G. A. and KELLEY, D. (1974), "An Empirical Survey of Price-Fixing Conspiracies", *Journal of Law and Economics*, **17** (1), 13–38.
- KANDORI, M. and MATSHUSHIMA, H. (1998), "Private Observation, Communication and Collusion", *Econometrica*, **66** (3), 627–652.
- KIHLSTROM, R. and VIVES, X. (1992), "Collusion by Asymmetrically Informed Firms", *Journal of Economics and Management Strategy*, **1** (2), 371–396.
- LACASSE, C. (1999), "Price-Fixing with Incomplete Information" (Mimeo, University of Alberta).
- MCAFEE, P. and MCMILLAN, J. (1992), "Bidding Rings", *American Economic Review*, **82** (3), 579–599.
- MEANS, G. (1935), "Industrial Prices and Their Relative Inflexibility", *Senate Document 13, 74th Congress, 1st Session* (Washington, D.C.).
- MILGROM, P. and SEGAL, I. (2002), "Envelope Theorems for Arbitrary Choice Sets", *Econometrica*, **70** (2), 583–601.

- MILLS, F. C. (1927) *The Behavior of Prices* (New York: NBER).
- MYERSON, R. (1981), "Optimal Auction Design", *Mathematics of Operations Research*, **6**, 58–73.
- PORTER, R. H. (1983), "Optimal Cartel Trigger Price Strategies", *Journal of Economic Theory*, **29** (2), 313–338.
- REYNOLDS, S. S. and WILSON, B. J. (1998), "Price Movements over the Business Cycle in U.S. Manufacturing Industries" (Mimeo, University of Arizona).
- ROBERTS, K. (1985), "Cartel Behavior and Adverse Selection", *Journal of Industrial Economics*, **33**, 401–413.
- ROTEMBERG, J. J. and SALONER, G. (1986), "A Supergame-Theoretic Model of Price Wars during Booms", *American Economic Review*, **76** (3), 390–407.
- SCHERER, F. M. (1980) *Industrial Market Structure and Economic Performance*, 2nd edn (Chicago: Rand McNally).
- SIMON, J. (2000), "Essays in Macroeconomics" (Ph.D. Dissertation, MIT).
- SPULBER, D. F. (1995), "Bertrand Competition when Rivals' Costs are Unknown", *Journal of Industrial Economics*, **43** (1), 1–11.
- STIGLER, G. (1964), "A Theory of Oligopoly", *Journal of Political Economy*, **72** (1), 44–61.
- SWEEZY, P. M. (1939), "Demand Under Conditions of Oligopoly", *Journal of Political Economy*, **47** (4), 568–573.
- TIROLE, J. (1988) *The Theory of Industrial Organization* (Cambridge, MA and London: MIT Press).