

Digitized by the Internet Archive
in 2011 with funding from
Boston Library Consortium Member Libraries

<http://www.archive.org/details/collusionpriceri00athe>

HB 31
.M415
no. 98-23

66

**working paper
department
of economics**

COLLUSION AND PRICE RIGIDITY

**Susan Athey
Kyle Bagwell
Chris Sanchirico**

No. 98-23 November, 1998

**massachusetts
institute of
technology**

**50 memorial drive
cambridge, mass. 02139**



**WORKING PAPER
DEPARTMENT
OF ECONOMICS**

COLLUSION AND PRICE RIGIDITY

**Susan Athey
Kyle Bagwell
Chris Sanchirico**

No. 98-23 November, 1998

**MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY**

**50 MEMORIAL DRIVE
CAMBRIDGE, MASS. 02142**

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

DEC 04 1998

LIBRARIES

Collusion and Price Rigidity

Susan Athey, Kyle Bagwell, and Chris Sanchirico*

Last Revised: November 1998.

Abstract. We consider an infinitely-repeated Bertrand game, in which prices are perfectly observed and each firm receives a privately-observed, i.i.d. cost shock in each period. We focus on symmetric perfect public equilibria (SPPE), 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.

JEL Classification Numbers: C73, L13, L14.

Keywords: Collusion, repeated games, private information, price rigidity, auctions.

*M.I.T. and NBER (Athey); Columbia University Economics Department and Graduate School of Business, and NBER (Bagwell); and Columbia University Economics Department and Law School (Sanchirico). We are grateful to Glenn Ellison, Eric Maskin, David Pearce, Lars Stole, Jean Tirole, and seminar participants at Barcelona (Pompeu Fabra), Boston College, Columbia, Harvard, M.I.T., N.Y.U., Northwestern and Toulouse (I.D.E.I.) for helpful discussions. Athey thanks the National Science Foundation (SBR-9631760) for generous financial support, and the Cowles Foundation at Yale University for hospitality and support.

1. Introduction

In the standard model of collusion, symmetric firms interact in an infinitely-repeated Bertrand game in which past prices are perfectly observed. While the standard model delivers a number of useful insights, a limitation of this model is that it presumes an unchanging market environment. This limitation is important for two reasons. First, the scope for testing a proposed theory of collusion is greater, if the theory offers predictions concerning the manner in which collusive prices may vary with underlying market conditions. Second, a theory of collusion is of potential value for macroeconomists, if the theory determines the response of collusive prices to demand and cost shocks.

These limitations have been partially addressed in two celebrated extensions of the standard collusion model. Rotemberg and Saloner (1986) extend the standard model to include the possibility of publicly-observed demand shocks that are i.i.d. over time. 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, thereby diminishing the incentive to cheat and bringing incentives back in line. In this way, Rotemberg and Saloner (1986) offer a prediction of countercyclical markups for collusive markets that are subject to demand shocks.¹ 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 perfectly monitor the behavior of its rivals. Green and Porter (1984) explore this possibility in an infinitely-repeated Cournot model. They assume that a firm cannot observe 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-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 propose a third extension of the standard collusion model. We consider an

¹ As Haltiwanger and Harrington (1991) and Bagwell and Staiger (1997) explain, the association between collusion and countercyclical markups can be overturned when the assumption of i.i.d. demand shocks is relaxed.

² While the Green-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 Abreu, Pearce and Stacchetti (1986, 1990), Fudenberg, Levine and Maskin (1994) and Porter (1983).

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 each period. We thus represent collusion in the context of a repeated adverse-selection (hidden-information) model with perfectly-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) have concluded that prices are more rigid in concentrated industries. This empirical finding suggests that collusion is associated with a greater tendency toward price rigidity. Further support for this suggestion is offered by Scherer (1980, pp. 184-93), who summarizes a number of studies that find that collusion is often implemented with “rules of thumb,” whereby the collusive price is fixed at some focal level, such as a certain percentage above a commonly-observed wholesale price.⁴ When firms collude in this manner, the collusive price is independent of the private cost positions (beyond the wholesale price) of the member firms.

At the same time, the Industrial Organization literature has as yet failed to provide 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. The theory employs an *ad hoc* behavioral postulate, compresses a dynamic story into a static framework, and does not determine the collusive price.⁵ Given these limitations, Industrial Organization economists have 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:⁶

³ Our modeling framework is related to recent work that extends the Green-Porter (1984) model to allow for privately-observed demand signals. Compte (1998) and Kandori and Matsushima (1998), for example, suppose that firms publicly choose “messages” after privately observing their respective demand signals. In our setting, by contrast, firms are privately informed as to their respective costs and the public action is a (payoff-relevant) price choice.

⁴ Evidence from other sources also indicates that rigid pricing is a ubiquitous feature of collusive endeavors. For example, in a *Business Week* (1975) article, a number of executives in a variety of industries describe informal (and sometimes formal) agreements to maintain a particular price. Hall and Hitch (1939) provide similar accounts. Scherer (1980) describes a collusive scheme whereby General Electric and Westinghouse adopted a pricing formula for turbine generators. Collusion seems especially prevalent in markets with homogeneous goods (see, e.g., Hay and Kelley (1974) and Scherer (1980, p. 203)) and relatively inelastic demands (see, e.g., Eckbo (1976)).

⁵ In the kinked-demand-curve theory, it is assumed that a firm believes that rivals will (not) match a price (increase) decrease. This behavioral postulate implies a kink in the demand curve, from which it follows that price and output are unresponsive to small marginal cost fluctuations. Maskin and Tirole (1988) offer an equilibrium model of the kinked-demand curve theory, based on the hypothesis that firms make alternating price choices in an infinite game. Their theory does not include the possibility of cost shocks, however.

⁶ For a similar view, see Scherer (1980, p. 180), who writes “Most oligopolies...appear willing to forego the modest gains associated with micro-meter like adjustment of prices to fleeting changes in demand and costs in order to avoid the risk of more serious losses from poorly coordinated pricing policies and price warfare.” See also Scherer (1980, pp. 168, 188).

“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. 914-15).

In the present paper, we develop 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 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.

Our formal analysis begins with the static Bertrand game with private cost information. This game, which has been previously analyzed by Spulber (1995), constitutes the stage game of our repeated-game model. In the unique symmetric Nash equilibrium, the pricing strategy is strictly increasing in the firm’s cost level. An appealing feature of this pricing scheme 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. The pricing function is strictly increasing as a consequence of the winner-take-all nature of the game, together with the single-crossing property present in the firm’s expected profit function. The single-crossing property arises because when a firm has lower costs, it values more the increase in expected sales that a lower price implies. At the Nash equilibrium, a higher-cost firm is thus unwilling to “mimic” the lower price that it would charge were costs lower, despite the increase in expected sales that a lower price would confer.

We turn next to the repeated-game model and explore whether firms can support better-than-Nash profits when they interact repeatedly. 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 cooperative or price-war phases. By contrast, in an *asymmetric perfect public equilibrium (APPE)*, firms keep track of past sales and may coordinate current price selections in order to treat more favorably a firm that has experienced low sales in the recent past. Athey and Bagwell (1998) study optimal APPE. After describing our results about SPPE, we will return to compare the two kinds of equilibria.

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, Pearce and Stacchetti (1986, 1990) and Fudenberg, Levine and Maskin (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 of the repeated game 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. For example, a price war is analogous to 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 incentive constraint is met. In this analysis, the on-schedule incentive constraint 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, thereby securing for itself a lower price and 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 under the single-crossing property higher-cost firms find lower prices less appealing. Second, following the selection of lower prices, the collusive scheme may sometimes call for an equilibrium-path price war in future periods. The current-period benefits 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.

A rich array of collusive schemes fit within the SPPE category. At one extreme, firms may 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. At the other extreme, firms may adopt a *monopoly-pricing scheme*, where each firm observes its cost for the period and then selects its monopoly price. If demand slopes down (so that the monopoly price varies with cost), this scheme satisfies the on-schedule incentive constraint only if equilibrium-path wars sometimes follow the selection of lower-cost monopoly prices. More generally, collusive pricing schemes may involve both distorted pricing and equilibrium-path wars. While all SPPE collusive pricing schemes call for prices that are (weakly) increasing in costs, some schemes allow for “rigid regions” (i.e., an interval of cost realizations that result in the same price), while other

schemes are fully sorting.

One scheme of particular interest is the *rigid-pricing scheme*. Under this scheme, each firm selects the same price in each period, regardless of its current cost position, and so the informational costs of collusion are neutralized: a high-cost firm has no incentive to misrepresent its cost, since the same price is required under any realization. Under a rigid-pricing scheme, therefore, any deviation is off-schedule, representing an unambiguous violation of the agreement. The threat of a price war may thus be limited to off-equilibrium path events. The downside of the rigid-pricing scheme, however, is that it does not offer efficiency benefits: one firm may have lower costs than its rivals, and yet the firms share the market. We thus see the outline of a broad tradeoff: if firms seek a collusive scheme that emphasizes efficiency benefits (e.g., a fully-sorting scheme), then the informational costs of collusion may be large; while if firms seek a collusive scheme that moderates the informational costs of collusion (e.g., the rigid-pricing scheme), then they may sacrifice efficiency benefits.

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 in the best such equilibrium firms employ the Nash pricing scheme: it is better to have low (Nash) prices and no actual price wars than higher (e.g., monopoly) prices and actual price wars. Thus, if colluding firms are to improve upon Nash profits, the collusive pricing scheme must exhibit some rigidity.

We next consider the full class of SPPE collusion schemes and report a second main finding: if firms are sufficiently patient and demand is inelastic, then an optimal SPPE collusive scheme can be achieved without recourse to equilibrium-path price wars. We also provide sufficient conditions under which the optimal SPPE does not rely on equilibrium-path price wars, for patient firms when demand is downward sloping; however, we identify a limited set of circumstances where they might be used. The finding that optimal collusion generally does not require price wars stands in sharp contrast to the analysis presented by Green and Porter (1984), suggesting a fundamental difference between repeated games with hidden actions and those with hidden information and publicly-observed actions.

Armed with these findings, we next add some additional structure and provide a characterization of the optimal SPPE collusive scheme. Specifically, when firms are patient and the distribution of cost types is log concave, we establish a third main finding: if demand is inelastic, 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 provide as well sufficient conditions for rigid-pricing behavior in the optimal SPPE when demand is downward sloping. Our results 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, considering impatient firms. We begin by showing that impatience on the part of firms creates an additional disadvantage to using 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 cooperation in the future. Our second (no-wars) finding therefore continues to hold, even when firms are impatient. Next, we observe that the off-schedule incentive constraint is particularly demanding for lower-cost types. Intuitively, when a firm draws a lower-cost type, the temptation to 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 depressed below that of higher-cost types in order to mitigate the former's 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 favorable 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.

We further show that the optimal pricing problem for impatient firms can also be studied using the tools of static mechanism design, although the problem is more difficult than the standard case where the outside option, which affects the participation constraint, is a constant. In the repeated game context, the value of the outside option (that is, openly cheating on the agreement) varies both with the firm's cost type and the collusive agreement itself, through the price specified for each cost type in the present, as well as through the future value of cooperation.

Finally, we consider the consequences of relaxing our restriction to symmetric schemes. Asymmetric schemes allow one firm to enjoy a more profitable continuation value than another, and in this way such schemes facilitate transfers *from* one firm *to* another. While we do not propose a full treatment of APPE (see Athey and Bagwell (1998)), we do provide a simple example that captures the qualitative benefits offered by asymmetric schemes. The scheme we study allows only two prices on the equilibrium path. The low price may be arbitrarily close to the high one; but, over every two-period "cycle," each firm can charge the low price only once. The scheme achieves

some efficiency benefits in the first period of each cycle, since a high-cost firm will hesitate to expend its only opportunity to charge the low price when it knows that tomorrow its expected costs will be lower. In contrast, low-cost firms will wish to increase their expected market share in the first period, since on average, winning will be less valuable in the second period.

Although APPE can provide such efficiency benefits, we focus here on SPPE for three main reasons. First, under many conditions the optimal SPPE is stationary. While we do not propose a theory of how firms learn and coordinate on a particular equilibrium, the rigid-price SPPE (and more generally, stationary symmetric equilibria) are appealingly simple. By contrast, a necessary condition for APPE to improve upon SPPE is that the asymmetric equilibrium is nonstationary. Such schemes are more sophisticated, and may be most plausible when a small number of firms interact frequently and communicate explicitly. Second, most of the informal literature about collusion (discussed above) highlights firms' fear of industry-wide breakdowns in collusion, and it is thus interesting to analyze SPPE, which isolate this consideration. Finally, SPPE are the only alternative if firms can observe the prices offered in the market, but they cannot observe (or infer) the identities of the firms who offer these prices. This situation arises in procurement auctions with more than two bidders, if the winning bid - but not the name of the winning supplier- is announced.

Our findings are related to those developed by McAfee and McMillan (1992), in their analysis of bidding rings. In a static model, they show that a fixed price (i.e., "identical bidding") 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). The analysis that we present for the inelastic-demand case may be understood in terms of a procurement auction. In fact, in our analysis of the optimal SPPE, we generalize the weak-cartel model, since the static mechanism we analyze is directly derived from a repeated game and allows for a restricted class of transfers (corresponding to symmetric price wars).⁷ Our rigid-pricing finding thus provides additional theoretical support for the common practice of identical bidding.⁸ We also extend the previous analysis in other directions, including downward-sloping demand and impatient firms.

The paper proceeds as follows. In Sections 2 and 3, we describe the static and repeated games, respectively. We present our findings for SPPE among patient firms in Section 4. Impatient firms are considered in Section 5. We explore simple asymmetric schemes in Section 6. Concluding thoughts are presented in Section 7.

⁷ McAfee and McMillan (1992) also consider "strong" cartels, in which transfers are unrestricted. (See also Cramton and Palfrey (1990) and Kihlstrom and Vives (1992).) In comparison, our analysis of SPPE restricts transfers to a range of symmetric values (namely, SPPE continuation values). Likewise, in APPE, firms have available an expanded (but still restricted) class of transfers (namely, APPE continuation values).

⁸ See, for example, Comanor and Schankerman (1976), who show that fixed-price rules (or closely related simple rotation schemes) have been used by a substantial fraction of the documented cases of bidding rings in procurement auctions. McAfee and McMillan (1992) summarize a large literature concerning the prevalence of identical bidding.

2. The Static Game

We begin by considering a static game of Bertrand competition when firms possess private information. This game illustrates the immediate tradeoffs that confront firms in determining their pricing policies. In subsequent sections, we examine a dynamic game, so that we may consider as well the long-term implications of different pricing policies. The static game serves as the stage game for the dynamic analysis.

2.1 The Model

We suppose that there are n *ex ante* identical firms that engage in Bertrand competition for sales in a homogenous-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. Specifically, we assume that firm i 's "type" θ_i is drawn in an i.i.d. fashion from the support $[\underline{\theta}, \bar{\theta}]$ according to the commonly known distribution function $F(\theta)$, where the corresponding density $f(\theta) \equiv F'(\theta)$ is everywhere positive. After the firms learn their respective cost types, they simultaneously choose prices. Let $\rho_i \in R_+$ denote the price chosen by firm i . The associated price profile is then represented by the vector, $\rho \equiv (\rho_1, \dots, \rho_n)$. On the demand side of the market, we assume a unit mass of identical consumers, each of whom possesses a demand function $D(\rho)$, where ρ is the price at which the consumer buys. We allow that the demand function is either (i) inelastic, or (ii) positive, twice continuously differentiable and decreasing (i.e., "downward sloping"). In the former case, we suppose that $D(\rho) = 1$ up to some reservation price r , where $r \geq \bar{\theta}$.⁹

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, R_+ . The function p_i is assumed continuously differentiable, except perhaps at a finite number of points (so as to allow for schedules with jumps).¹⁰ A price strategy profile is thus a vector $\mathbf{p}(\theta) \equiv (p_i(\theta_i), \mathbf{p}_{-i}(\theta_{-i}))$, where $\theta \equiv (\theta_i, \theta_{-i})$ is the vector of cost types and $\mathbf{p}_{-i}(\theta_{-i})$ 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 $\pi(\rho, \theta) \equiv [\rho - \theta] D(\rho)$ 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. We assume that $\pi(\rho, \theta)$ has a unique maximizer, $p^m(\theta)$, and is strictly quasi-concave.¹¹ Second, we specify a Bertrand market-share-allocation function, $m_i(\rho)$, that indicates firm i 's market share

⁹ In the case of inelastic demand, we may therefore interpret the static model as a model of a procurement auction. We assume that $r \geq \bar{\theta}$ for simplicity.

¹⁰ This restriction includes a wide range of possible pricing strategies, and it serves to minimize technical complications that are not central to our discussion.

¹¹ Under inelastic demand, $p^m(\theta) \equiv r$. When demand slopes down, $p^m(\theta)$ is increasing in θ .

when the vector of realized prices is ρ . This function allocates consumers evenly among firms that tie for the lowest price in the market; therefore, $m_i(\rho) = 1$ if $\rho_i < \min_{j \neq i} \rho_j$, $m_i(\rho) = 0$ if $\rho_i > \min_{j \neq i} \rho_j$, and $m_i(\rho) = 1/k$ if there are $k - 1$ other firms that tie firm i for the lowest price.

We may now define firm i 's *ex post* profit as $\pi_i(\rho, \theta_i) \equiv \pi(\rho_i, \theta_i) m_i(\rho)$. This is the profit that firm i enjoys after all prices are posted, when it wins a share $m_i(\rho)$ of consumers. It is also convenient to 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, $\mathbf{p}_{-i}(\theta_{-i})$. Firm i 's interim profit function is thus defined as $\Pi_i(\rho_i, \mathbf{p}_{-i}, \theta_i) \equiv E[\pi_i(\rho_i, \mathbf{p}_{-i}(\theta_{-i}), \theta_i)]$, where the expectation is taken over θ_{-i} . Finally, we may also define firm i 's *ex ante* profit. This is the profit that firm i expects before any firm (including i itself) learns its cost type, when it is anticipated that firms' eventual prices are determined by the pricing strategy profile, $\mathbf{p}(\theta)$. Firm i 's ex ante profit is given by $\Pi_i(\mathbf{p}) \equiv E[\Pi_i(\mathbf{p}(\theta), \theta_i)]$, where the remaining expectation is taken over firm i 's own cost level.

2.2 Equilibrium Pricing and Profit

We return now to the interim profit function defined above and describe the essential tradeoff that confronts a firm when it sets its price. Interim profit may be written as $\Pi_i(\rho_i, \mathbf{p}_{-i}, \theta_i) = \pi(\rho_i, \theta_i) E m_i(\rho_i, \mathbf{p}_{-i}(\theta_{-i}))$, indicating that firm i 's price ρ_i affects its profit both by altering its "profit-if-win" (i.e., $\pi(\rho_i, \theta_i)$) and its probability of winning (i.e., $E m_i(\rho_i, \mathbf{p}_{-i}(\theta_{-i}))$). Thus, when a firm of a given type considers whether to lower its price from some status quo level, it must weigh the effect of an increase in the chance of winning against the direct effect of the price reduction on profit-if-win.

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 more appealing than do higher types. In the first place, since lower types have lower total costs, they have higher profit-if-win (i.e., $\pi_\theta(\rho, \theta) < 0$). Therefore, the increased chance of winning (i.e., the increase in $E m_i(\rho, \mathbf{p}_{-i}(\theta_{-i}))$) from lowering price is of greater benefit to a lower-cost firm. Secondly, the reduction in profit-if-win from lowering price is less severe for lower types (i.e., $\pi_{\rho\theta}(\rho, \theta) \geq 0$): revenue changes are the same across all types, but since lower types also have lower marginal cost, their costs rise by less in response to increased demand. This second effect is strict when demand slopes down, as then the lower price does raise demand, and it is neutral in the case of inelastic demand. As is standard, the single crossing property implies that higher-cost firms must select higher prices (i.e., $p_i(\theta_i)$ is non-decreasing).

Employing techniques from the auction literature (Maskin and Riley, 1984), Spulber (1995) analyses this stage game in detail and finds in part:

Proposition 1 (Spulber (1995)) *The static game has a unique Nash equilibrium, which 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: $p^e(\theta) < p^m(\theta), \forall \theta < \bar{\theta}$ 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.*

Notice that the symmetric equilibrium pricing strategy p^e is continuous and strictly increasing. These properties arise as a consequence of the “winner-take-all” nature of the static game. If, for example, the schedule were to jump up discontinuously at $\hat{\theta}$, then a firm with type $\hat{\theta}$ would do better increasing its price into the “gap” (provided $p^e(\hat{\theta}) < p^m(\hat{\theta})$); while if the price schedule were to have a flat section that included type $\hat{\theta}$, then a firm of this type would gain by shading its price downward a slight amount (converting “ties” into “wins”). These restrictions are standard in the literature, but, as we shall see, they disappear when we move to the repeated game.

3. The Repeated Game

In this section, we define the repeated game that is the focus of our study. We also present a “Factored Program” and establish a relationship between solutions to this program and optimal SPPE. We use this relationship in subsequent sections, where we characterize 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

$$v_i(\{\theta^t, \mathbf{p}^t\}) = \sum_{t=1}^{\infty} \delta^{t-1} \pi_i(\rho^t, \theta_i^t),$$

where θ_i^t is the cost level for firm i at date t , $\delta \in [0, 1)$ is the discount factor, and $\pi_i(\rho^t, \theta_i^t)$ is ex post profit as previously defined. At the close of period T , firm i possesses an information set,

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

We would like our dynamic game to have a recursive structure, so that dynamic-programming techniques (as developed by Abreu, Pearce and Stacchetti (1986, 1990) [APS]) may be applied. We do this at two levels. First, we assume that firm types are i.i.d. across time (and firms), so that the game is “structurally recursive.” Second, we follow Fudenberg, Levine and Maskin (1994) [FLM] and restrict attention to a recursive solution concept; namely, we focus upon those sequential equilibria in which players condition only on the history of realized prices and not on their own private history of types and price schedules. Such strategies are called *public strategies* and such sequential equilibria are called *perfect public equilibria* (PPE). Note that PPE are tested against deviations that do condition on private histories; that is, they are true sequential equilibria in the game. Note also that a public strategy may be abbreviated as a map from finite public histories $\{\rho^t\}_{t=1}^T$ to price schedules.

Using familiar arguments (see, e.g., FLM), these two restrictions imply:

Lemma 2 (*Recursion of Perfect Public Equilibria*) *The continuation of any PPE after any history is itself a PPE in the full game and yields payoffs in the continuation equal to what would have been obtained had the strategy profile been used from the start.*

Before moving on to a formal analysis of the perfect public equilibria for the repeated game, we highlight some of the novel features of this game.

First, a PPE in a repeated game with hidden information must be immune to two kinds of current-period deviations. 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_i . When a firm prices in this manner, it has unambiguously deviated; consequently, if the collusive scheme prescribes a punishment (i.e., a reduction in continuation values) following such a deviation, and if the prospect of such a punishment deters a firm from undertaking the off-schedule deviation (as will be true if firms are sufficiently patient), then that punishment will never actually occur (i.e., it is off the equilibrium path). By contrast, a firm deviates “on-schedule” when it chooses a price that is assigned 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, an on-schedule deviation is not perfectly observable, as a deviation, to rival players: a rival can not be sure that the deviating firm was not truly of the cost type that it is imitating. Thus, firms may be tempted to

construct a collusive scheme that imposes a harsh punishment when low prices are chosen, as this would serve to prevent on-schedule deviations by higher-cost types. But such a punishment would be costly, since it would occur along the equilibrium path of play, whenever firms actually realized low costs.

Secondly, we note that the repeated game admits a richer range of equilibrium pricing schedules than did the static game. As discussed in Section 2, the equilibrium pricing schedule for the static game can have neither flat segments nor jumps. This is no longer true in the repeated game. For example, we may specify that undercutting a flat price leads to a harsh punishment, and this may be sufficient to deter the (off-schedule) deviation. More generally, the repeated game adds a third ingredient to the tradeoff that a firm confronts when it considers reducing price. Now lowering price means not just: 1) an increased chance of winning, and 2) lower profit-if-win, but also 3) a different continuation value.

Finally, we observe that the mixture of on- and off-schedule incentive problems is endogenous to the firms' choice of price schedules. A completely flat schedule means a relatively acute off-schedule incentive problem, but no on-schedule problem. A price schedule that spans the full relevant range of prices has no off-schedule problem, but possibly an acute on-schedule problem. With the choice of a price schedule, therefore, firms endogenously determine the support of actual prices that may be observed.

3.2 Optimal Symmetric Collusion: The Program

In a *symmetric perfect public equilibrium (SPPE)*, following every public history, firms adopt symmetric price schedules: $s_i(\rho^1, \dots, \rho^T) = s_j(\rho^1, \dots, \rho^T), \forall i, j, T, \rho^1, \dots, \rho^T$. Symmetry means that all firms suffer future punishments and rewards together. For example, it might be that firm 1 runs the risk of triggering an industry-wide breakdown in collusion should it repeatedly charge implausibly low prices. Symmetric equilibria thus capture the “fear of breakdown” that collusive firms may experience, as discussed in the Introduction.¹²

Our goal is to formulate the problem of characterizing the optimal SPPE in such a way that we can employ the tools of (static) mechanism design. We begin by constructing a programming problem (The “Factored Program”) in terms of current-period price schedules and “continuation

¹² Despite its clear restrictiveness, a rich array of strategy profiles still fit under the symmetric heading. For instance, any amount of history dependence is allowed. Thus, price wars can be triggered by any number of periods of “bad behavior.” This means in particular that the law of large numbers may be utilized: since types are drawn i.i.d. across time and firms, an accurate measure of the average type of each firm becomes available, and this can be compared to the realized average announced type (i.e., the type corresponding to the announced price) to determine whether a given firm has cheated. (See, e.g., Radner (1981) and Rubinstein and Yaari (1983) for a discussion of how the law of large numbers may be used to implement such “review strategies.”) Moreover, gradations of punishment are available: price wars can vary in severity and length and so may be tailored to the “fit the crime.”

payoff functions,” and claim from APS logic that any solution to this problem corresponds to an optimal SPPE. To this end, we first make some notational adjustments that reflect the symmetric setting. Let $\Pi(p)$ denote the ex ante profit to a firm when all firms use the symmetric pricing strategy p , and let $\Pi(p', p)$, represent the ex ante profit to a firm when it adopts the pricing schedule p' while all other firms employ the symmetric pricing strategy p . (The notation for the continuation payoff function developed below is interpreted similarly.) Let $V_s \in \mathfrak{R}$ denote the set of SPPE continuation values and write $\underline{v} \equiv \inf V_s$ and $\bar{v} \equiv \sup V_s$. Note, initially, that with a continuum of possible pricing strategies there is no *a priori* basis from which to argue that either $\bar{v} \in V_s$ or $\underline{v} \in V_s$.

Any symmetric public strategy profile $s = (s, \dots, s)$ can be *factored* into a first-period price schedule p and a *continuation payoff function* $v: \mathfrak{R}_+^n \rightarrow \mathfrak{R}$. 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 $\rho = (\rho_1, \dots, \rho_n) \in \mathfrak{R}_+^n$. Each firm’s payoff from s can then be written as $\Pi(p) + \delta E v(p)$.

Now consider the following problem in which we choose factorizations (p, v) rather than choosing strategies directly:

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

$$\Pi(p) + \delta E v(p)$$

$$\text{subject to: } \forall \rho \in \mathfrak{R}_+^n, v(\rho) \in V_s, \text{ and}$$

$$\forall p', \Pi(p) + \delta E v(p) \geq \Pi(p', p) + \delta E v(p', p).$$

An important ingredient in our argument is contained in the following result:

Lemma 3 *If the symmetric public strategy profile $s^* = (s^*, \dots, s^*)$ factors into (p^*, v^*) , and (p^*, v^*) solves the Factored Program, then s^* is an optimal SPPE.*

The lemma follows from the fact that (p, v) is the factorization of some SPPE if and only if it satisfies the Factored Problem’s two constraints. To see this equivalence, note first that, by Lemma 2, every SPPE must satisfy $v(\rho) \in V_s$. Every SPPE also satisfies the second constraint, because an SPPE must be immune to deviations in the first period. Conversely, one shows that any (p, v) satisfying the constraints is the factorization of an SPPE by appending to p the SPPE continuation strategy profiles underlying $v(\rho)$, and then arguing from the one-stage-deviation-principle that the constructed strategy profile is itself an SPPE. We refer the reader to APS for full details.

To analyze the Factored Program further, we rewrite it in equivalent interim-profit form, parsing the constraints into on- and off-schedule conditions. Using abbreviated notation in light of the symmetric analysis:

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

$$E_i[\Pi(p(\theta_i), p, \theta_i) + \delta E_{-i}v(p(\theta_i), p)]$$

subject to:

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

$$(IC\text{-off1}) \quad \forall \rho_{-i}, \quad v(\rho', \rho_{-i}) \in V_s$$

$$(IC\text{-off2}) \quad \forall \theta_i, \quad \Pi(p(\theta_i), p, \theta_i) + \delta E_{-i}v(p(\theta_i), p) \geq \Pi(\rho', p, \theta_i) + \delta E_{-i}v(\rho', p)$$

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

$$(IC\text{-on1}) \quad \forall \rho_{-i}, \quad v(p(\hat{\theta}_i), \rho_{-i}) \in V_s$$

$$(IC\text{-on2}) \quad \forall \theta_i, \quad \Pi(p(\theta_i), p, \theta_i) + \delta E_{-i}v(p(\theta_i), p) \geq \Pi(p(\hat{\theta}_i), p, \theta_i) + \delta E_{-i}v(p(\hat{\theta}_i), p).$$

Notice how the on-schedule constraints are written in “direct” form: for given p and v , the on-schedule constraint requires that a firm with type θ_i does better by “announcing” that its type is θ_i than by announcing some other type, $\hat{\theta}_i$. This observation suggests that the on-schedule constraint can be understood as corresponding to a “truth-telling” constraint in an appropriate mechanism design formulation.

4. Optimal Symmetric Collusion Among Patient Firms

We break our analysis of symmetric collusion into two parts. In the current section, we suppose that firms are patient, so that the off-schedule incentive constraints are satisfied. The assumption of patient firms is then relaxed in Section 5.

Our strategy is to show that when firms are patient, the Interim Program can be further relaxed to yield what we call the Mechanism Design Program. We then use standard tools from the mechanism design literature to characterize optimal SPPE, so as to offer predictions regarding observed pricing patterns in collusive settings.

4.1 The Mechanism Design Program

To begin, we define a new function that represents interim profit in direct form: $H(\hat{\theta}, \theta; p) \equiv \Pi(p(\hat{\theta}), p, \theta)$. We now state the following program:

The Mechanism Design Program: Choose price schedule p and function $T : [\underline{\theta}, \bar{\theta}] \rightarrow \Re$ to maximize

$$\begin{aligned} & E[H(\theta, \theta; p) - T(\theta)] \\ & \text{subject to : } \forall \theta, T(\theta) \geq 0 \\ & \text{(IC-onM)} \forall \hat{\theta}, \theta, H(\theta, \theta; p) - T(\theta) \geq H(\hat{\theta}, \theta; p) - T(\hat{\theta}). \end{aligned}$$

To relate the Interim and Mechanism Design Programs, we first observe that, while the Mechanism Design Program uses different notation, the objective function is the same. We thus focus on comparing the constraints. Observe that for any (p^I, v^I) which satisfies the constraints of the Interim Program, we can find a corresponding T^I such that (p^I, T^I) satisfies the constraints of the Mechanism Design Program. To see this, recall that $\bar{v} \equiv \sup V_s$, and define $T^I(\theta_i) = \delta(\bar{v} - E_{-i} v^I(p^I(\theta_i), p^I))$. Then, the restriction $v \in V_s$ implies that $v(\rho) \leq \bar{v}$ for all ρ , and thus $T^I(\theta) \geq 0$. Further, (p^I, v^I) satisfies (IC-on2) in the Interim Program if and only if (p^I, T^I) satisfies (IC-onM).

Now consider a (p^D, T^D) which satisfies the constraints of the Mechanism Design Program. As we discuss below, the off-schedule constraints of the Interim Program can be satisfied when firms are sufficiently patient. Turning to the on-schedule constraints, we can define v^D such that $T^D(\theta_i) = \delta(\bar{v} - E_{-i} v^D(p^D(\theta_i), p^D))$, whereby (IC-onM) and (IC-on2) are equivalent. But (IC-on1) requires further that $v^D(p^D(\theta_i), \rho_{-i}) \in V_s$, and so it remains to verify that such a v^D can be found. There are two cases: (a) the “no wars” case, where $T^D \equiv 0$; and (b) the case where T^D is sometimes positive.

Consider first case (a), whereby our construction requires $v^D(p^D(\theta), p^D) = \bar{v}$ for all θ . If V_s is closed, so that $\bar{v} \in V_s$, this also satisfies (IC-on1). We proceed by showing that \bar{v} can indeed be obtained as an SPPE if the solution to the Mechanism Design Program is $(p^*, T^* \equiv 0)$.

Thus, suppose for the moment that $(p^*, T^* \equiv 0)$. If we relax the Interim Program by replacing the (IC-on1) constraint, $v(p(\theta), \rho_{-i}) \in V_s$, with the constraint $v(p(\theta), \rho_{-i}) \in V_s \cup \bar{v}$, two conclusions follow. First, $(p^*, v^*(p^*(\theta), p^*)) \equiv \bar{v}$ solves the relaxed Interim Program. Second, any solution to the relaxed Interim Program yields profits at least as large as the supremum of the profits available in the (unrelaxed) Interim Program, \bar{v} . Formally, $EH(\theta, \theta; p^*) + \delta \bar{v} \geq \bar{v}$, or $EH(\theta, \theta; p^*) / (1 - \delta) \geq \bar{v}$.

But, since $(p^*, T^* \equiv 0)$ satisfies (IC-onM), p^* must satisfy (IC-On2) for any continuation value function that does not vary with the announced cost type. Thus, the collusive scheme where p^* is used in every period is an SPPE so long as the off-schedule constraints are satisfied. If p^* yields per-period profits greater than the static Nash equilibrium (as we will establish in Sections 4.3.4 and

4.4.3), the off-schedule constraints can be satisfied when firms are sufficiently patient (supported by the threat of forever reverting to the static Nash equilibrium). Therefore, since $EH(\theta, \theta; p^*)/(1 - \delta) \in V_s$, the definition of \bar{v} implies $\bar{v} \geq EH(\theta, \theta; p^*)/(1 - \delta)$. As the reverse inequality was established above, we thus conclude that the best SPPE is the stationary equilibrium where each firm uses p^* in every period. Summarizing:

Proposition 4 (Staticizing) *Suppose that the Mechanism Design Program is solved by (p^*, T^*) , where $T^* \equiv 0$. Then $\exists \delta^* \in (0, 1)$ such that, if $\delta \geq \delta^*$,*

- (i) p^* maximizes $EH(\theta, \theta; p)$ subject to $\forall \hat{\theta}, \theta, H(\theta, \theta; p) \geq H(\hat{\theta}, \theta; p)$.
- (ii) In an optimal SPPE, firms adopt p^* after all equilibrium-path histories.
- (iii) $\bar{v} = EH(\theta, \theta; p^*)/(1 - \delta) \in V_s$.

This result establishes conditions under which the problem of finding an optimal SPPE can be reduced to a simple and familiar (static) mechanism design problem. In the remainder of Section 4 we show that, under many conditions, the solution to the Mechanism Design Program (p^*, T^*) in fact entails $T^* \equiv 0$. However, we also describe a limited set of circumstances under which price wars may be used (case (b) from above). It is possible to generalize part (iii) of the staticizing proposition to the case of nonstationary equilibria, showing that when firms are sufficiently patient, $E[H(\theta, \theta; p^*) - T^*(\theta)]/(1 - \delta) = \bar{v}$. However, such a result requires that V_s is convex, so that all positive, finite punishments are available when firms are sufficiently patient.¹³ Because most of our results concern the case where $T^* \equiv 0$, we will not pursue that extension formally here.

4.2 Consequences of Incentive Compatibility

Now that the problem of finding an optimal SPPE has been reduced to a standard programming problem, we proceed to analyze the program. Our first step in this process is to characterize the implications of the on-schedule constraint that is found in the Mechanism Design Program (IC-onM). We do this in the following lemma (where we use the notation $H_\theta(\theta, \theta; p) = \frac{\partial}{\partial \theta} H(\hat{\theta}, \theta; p)|_{\hat{\theta}=\theta}$):

Lemma 5 (Constraint Reduction) (p, T) satisfies (IC-onM) if and only if (p, T) also satisfies:

(i). $p(\theta)$ is weakly increasing, and

(ii).
$$H(\theta, \theta; p) - T(\theta) = H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} H_\theta(\tilde{\theta}, \tilde{\theta}; p) d\tilde{\theta}.$$

This result is standard in the mechanism design literature,¹⁴ and it follows from the single-crossing

¹³ As discussed above, we will establish that there is always variation in the set V_s , which can be made arbitrarily large by taking δ close to 1. One way to guarantee convexity would be to extend the model to allow public randomization. For example, for each price vector, we could specify a probability that the firms revert to the static Nash equilibrium.

¹⁴ For a standard presentation, see Fudenberg and Tirole (1991).

property.¹⁵ We will exploit this result throughout the paper.

We observe that the situation analyzed here contrasts with the usual Principal-Agent formulation, since here the “agents” (i.e., firms) design their *own* schedule, with the goal of generating as much profit as possible, in light of their own incentive-compatibility constraints. The on-schedule incentive constraint thus puts an upper bound on the profits that a type θ can earn: if θ were to earn too much, other types would pretend to be θ . The expression in (ii) thus can be interpreted broadly 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), $H(\theta, \theta; p) - T(\theta)$, 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 higher types.¹⁶

To better understand (ii), we consider the consequences for some type θ of an increase in the profit that type $\bar{\theta}$ earns. Following this increase, type $\bar{\theta}$ becomes less tempted to misrepresent a lower type, which in turn enables the type just below to earn additional profit without inducing such mimicry. And as this slightly lower type earns greater profit, its incentive to misreport a slightly lower type yet is diminished, permitting in turn an increase in the profit that can be allocated to the next type down. In this way, an increase in the profit-at-the-top unleashes a cascade of higher profits to lower types, with the increase at each step governed by the actual efficiency advantage that the associated type has as compared to the next type just above. As (ii) indicates, the profit earned by type θ is thus the profit-at-the-top plus the efficiency rents accumulated across all higher types (note that $H_\theta < 0$).

We next interpret further the efficiency-rent term. To do so, we introduce some new notation. We may define the market share expected by firm i when it announces $\hat{\theta}_i$ as $M_i(\hat{\theta}_i; \mathbf{p}) \equiv Em_i(p_i(\hat{\theta}_i), \mathbf{p}_{-i}(\theta_{-i}))$, where the expectation is over the announcements of other firms (assumed truthful). Imposing symmetry and adjusting notation accordingly, we can write $H(\hat{\theta}, \theta; p) \equiv \pi(p(\hat{\theta}), \theta)M(\hat{\theta}; p)$, where $M(\hat{\theta}; p)$ is the chance of winning for a firm that announces $\hat{\theta}$ and thus sets the price $p(\hat{\theta})$. Using this notation, we may exploit the definition of H and use $\pi_\theta(p(\theta), \theta)$

¹⁵ To understand the result, observe that necessity of property (i) follows directly from the single-crossing property of the interim profit function, $H(\theta, \theta; p)$. Necessity of property (ii) follows since, by the envelope theorem, $dH(\theta, \theta; p)/d\theta = H_\theta(\theta, \theta; p) + T'(\theta)$ almost everywhere if $\hat{\theta}=\theta$ is optimal (see Athey, Milgrom and Roberts (1997) for an appropriately general statement of the envelope theorem). Given (i), it is standard to show that (ii) is sufficient as well, since the single-crossing property implies that if no type finds a local deviation desirable, no type will find a large deviation desirable either.

¹⁶ By contrast, in the standard Principal-Agent problem, the principal designs a schedule that extracts as much rent from the agent as possible. Incentive compatibility then puts a *lower* bound on the profit earned by an agent of type θ : if the principal attempts to extract too much rent from type θ , then that type would misreport and announce a different type. In the optimal schedule, the principal typically extracts all of the rent from the highest-cost agent (corresponding to the profit-at-the-top term in (ii)) and necessarily leaves some efficiency rent for other types. As we explain below, when agents design their own schedule, many features of the standard approach are reversed; for example, agents are attracted to schedules that increase the profit-at-the-top.

$= -D(p(\theta))$ to find that

$$- \int_{\theta}^{\bar{\theta}} H_{\theta}(\bar{\theta}, \bar{\theta}; p) d\bar{\theta} = \int_{\theta}^{\bar{\theta}} D(p(\bar{\theta})) M(\bar{\theta}; p) d\bar{\theta} \quad (1)$$

For a given allocation of market share across types, an important implication is that efficiency rents are higher when demand is *higher* (i.e., prices are *lower*). This is a reflection of the single-crossing property of the model: lower types are more attracted to lower prices/higher demands, and so they satisfy incentive compatibility (i.e., deter mimicry) most profitably when demand is high.

A final theme in our analysis of optimal SPPE is a tradeoff that potentially arises when firms wish to sort between types. In particular, the firms have two instruments with which to distinguish between types: the slope of (or jumps in) the pricing function, and wars (T). Thus, it is useful to address the extent to which incentive compatibility constrains the ability of the firms to choose between these alternatives. To do so, we introduce a simple restriction which simplifies the analysis. Consider a scheme (p, T) and let θ_K denote the lowest θ for which $p(\theta) = p(\bar{\theta})$. We thus represent with θ_K the lowest-cost type that selects the highest-cost price, $p(\bar{\theta})$. Of course, if $p(\theta)$ is strictly increasing, then $\theta_K = \bar{\theta}$. We restrict attention to schemes (p, T) for which $H(\theta_K, \theta_K; p) - T(\theta_K) \geq 0$. This restriction entails no loss of generality, and serves to simplify the exposition of our findings.¹⁷

Under this restriction, we can establish that the use of wars does *not* expand the range of sorting alternatives available to the firm. In particular, the firms are free to first choose the extent of sorting (the market share function, M) and then determine whether or not they want to implement this arrangement with a war. Formally:

Lemma 6 *Given an incentive-compatible scheme (p, T) and associated market-share allocation $M(\theta; p)$, there exists an alternative scheme $(\bar{p}, \bar{T} \equiv 0)$ which is also incentive compatible, and such that $M(\theta; p) = M(\theta; \bar{p})$ and $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = H(\bar{\theta}, \bar{\theta}; \bar{p})$.*

In short, given an original incentive-compatible scheme and market-share allocation, we may construct an alternative incentive-compatible scheme that never uses wars, delivers the same market-share allocation and also provides the same profit-at-the-top. Of course, 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) are better off under the alternative scheme. We explore this issue in subsections below.

¹⁷ If a scheme imposed negative profit (inclusive of wars) for θ_K , then all types higher than θ_K would do even worse, and an alternative scheme (in which θ_K and higher types set “non-serious” higher prices and receive zero market share while also avoiding wars) would yield a higher value for $E[H(\theta, \theta; p) - T(\theta)]$. In the repeated-game context, the superior scheme would make such types inactive in the current period and follow such a realization with a no-war continuation value, \bar{v} .

At this point, we have extracted three lessons that will inform the subsequent analysis of the repeated game. First, after accounting for incentive compatibility, firms may be attracted to pricing schemes that raise the profit-at-the-top. Second, pricing schemes that expand demand levels may also be desirable, as such schemes enhance the available efficiency rents. Finally, for a given amount of sorting, firms are free to choose whether to implement the corresponding market-share allocation with wars.

With these lessons in place, we are prepared now to characterize optimal SPPE for patient firms. We proceed by first analyzing the inelastic-demand case, and then exploring the additional complications which arise with downward-sloping demand.

4.3 Inelastic Demand

We develop our analysis of the inelastic-demand case in four steps. First, we develop further the consequences of incentive compatibility for the case of inelastic demand. Building on the results from this exercise, our second step is to consider symmetric collusive schemes that are *fully sorting* (i.e., $p(\theta)$ is strictly increasing). We then explore whether an optimal SPPE for patient firms requires equilibrium-path wars (i.e., $T(\theta) > 0$ for some θ). Finally, we add further structure and characterize the optimal symmetric pricing scheme for patient firms.

4.3.1 Incentive Compatibility and Revenue Equivalence

When demand is inelastic, Lemma 5 has very powerful consequences for the analysis of optimal collusive schemes. In particular, (IC-onM) implies that type θ 's interim profit is determined as:

$$H(\theta, \theta; p) - T(\theta) = H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} M(\tilde{\theta}; p) d\tilde{\theta}. \quad (2)$$

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

Lemma 7 (Revenue Equivalence Theorem) *Suppose that demand is inelastic. 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 $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = H(\bar{\theta}, \bar{\theta}; \tilde{p}) - \tilde{T}(\bar{\theta})$, also satisfies $H(\theta, \theta; p) - T(\theta) = H(\theta, \theta; \tilde{p}) - \tilde{T}(\theta)$ for all θ .*

Intuitively, suppose that firms start with the scheme (p, T) , and then consider an alternative, (\tilde{p}, \tilde{T}) . As (2) indicates, if the new scheme is on-schedule incentive compatible, gives the same profit-at-the-top, and delivers the same market-share-allocation function, then the interim profit is the same between the two pricing schemes. For example, suppose that $\tilde{T} \equiv 0 < T(\theta)$. Since (2) implies that $H(\theta, \theta; p) - T(\theta) = H(\theta, \theta; \tilde{p})$, we see that the schedule (p, T) must entail a higher price at θ : $p(\theta) > \tilde{p}(\theta)$. In particular, as neither p nor T interacts with the firm's actual cost type θ ,

on-schedule incentive compatibility and full sorting requires that a “higher p ” is balanced against a “higher T ” at a fixed conversion rate $1/M$.

It is important to note that the Revenue Equivalence Theorem fails when $\pi_{\rho\theta}(\rho, \theta) > 0$, as in the case of risk-averse bidders in auctions (Maskin and Riley (1984)) as well as in the case of downward-sloping demand (in the present context).

4.3.2 Fully-Sorting Pricing Schemes

We now consider fully-sorting collusive pricing schemes. The set of such schemes includes a variety of candidates. Firms may employ the static Nash equilibrium in each period. Alternatively, the firms may attempt to sort with higher prices, perhaps at or near the monopoly level. Such schemes satisfy on-schedule incentive constraints only if they include equilibrium-path wars.

To begin, notice that under full sorting, the highest type makes no sales, and the profit-at-the-top is simply $-T(\bar{\theta})$.¹⁸ The full-sorting requirement further implies that a firm wins the market if and only if all other firms announce higher types, and it follows that $M(\theta; p) = [1 - F(\theta)]^{n-1}$. Thus, by Lemma 7, any two fully-sorting schemes which satisfy (IC-onM) will differ only if the profit-at-the-top differs. Further, since Lemma 5 implies that when demand is inelastic and full-sorting is imposed,

$$H(\theta, \theta; p) - T(\theta) = -T(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} [1 - F(\tilde{\theta})]^{n-1} d\tilde{\theta}, \quad (3)$$

the best fully sorting scheme must entail $T(\bar{\theta}) = 0$. But, notice that the static Nash equilibrium pricing scheme, p^e , is strictly increasing and further must satisfy (IC-onM) with $T \equiv 0$, since $H(\theta, \theta; p^e)$ is type θ 's expected payoff in the symmetric Nash equilibrium. Thus, we conclude that under a fully-sorting and on-schedule incentive-compatible pricing scheme, the interim profit (inclusive of wars) available to a type θ firm is at best equal to its Nash profit. It follows that an optimal SPPE under full sorting when demand is inelastic 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.¹⁹ Summarizing:

¹⁸ In this case, if $T(\bar{\theta}) > 0$, then $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) < 0$, so this scheme does not satisfy our earlier restriction about profit-at-the-top which was required for Lemma 6. Nevertheless, a no-war scheme with the same market-share allocation can still be constructed, and it sets $\tilde{T}(\bar{\theta}) = 0$ with \bar{p} strictly increasing at the top. However, in this event the no-war scheme yields a strict improvement. We will allow for this possibility in our discussion, arguing that wars at the top are not optimal.

¹⁹ Suppose that a solution to the Mechanism Design Program among the class of fully-sorting pricing schemes entails $T \equiv 0$ and the Nash pricing scheme, p^e . In this case, we do not require variation in V_s or a high value of δ to establish that this scheme is optimal within the fully-sorting SPPE class. Instead, we can simply observe that repeating the static Nash pricing scheme in each period delivers a fully-sorting SPPE. From here, the logic of Proposition 4 can be applied to show that p^e is the optimal fully-sorting SPPE pricing scheme. Note, in particular, that the set of fully-sorting SPPE forms a recursive class, as per Lemma 1.

Proposition 8 *Among the class of fully-sorting pricing schemes, and for any distribution function F and discount factor δ , if demand is inelastic, then an optimal SPPE is the repeated play of the static Nash equilibrium after all histories.*

4.3.3 No Wars on the Equilibrium Path

The analysis in the previous subsection shows in the class of fully-sorting pricing functions, after on-schedule incentive constraints are considered, there is no benefit in supporting higher prices with on-schedule wars. In this subsection, we show that as long as $\pi_{\rho\theta}(\rho, \theta) = 0$, the logic holds for any market-share allocation.

Formally, consider a particular pricing function p , and let $M(\theta; p)$ be the induced market-share-allocation function. Let us suppose further that (p, T) entails positive wars somewhere. We now refer to Lemma 6, which guarantees the existence of an alternative scheme $(\tilde{p}, \tilde{T} \equiv 0)$ that is incentive compatible, generates the same market-share allocation M , and satisfies $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = H(\bar{\theta}, \bar{\theta}; \tilde{p})$. Then, Lemma 7 implies that the alternative no-war schedule $(\tilde{p}, \tilde{T} \equiv 0)$ gives the same interim profit (inclusive of wars) as did the original schedule (p, T) . Further, since $H(\theta, \theta; p) - T(\theta) = H(\theta, \theta; \tilde{p})$, the alternative schedule achieves this profit by exchanging any war in the original schedule for a lower price. Applying this analysis, together with the staticizing proposition, we conclude that:

Proposition 9 *Consider the inelastic-demand case, and allow for any distribution function F . Suppose that there exists an optimal SPPE which specifies a first-period pricing function p and a continuation value function v such that $v(\mathbf{p}(\theta)) < \bar{v}$ for some θ ($T(\theta) > 0$ for some θ). Then, for δ large enough, there exists an alternative optimal SPPE where the strategy for each firm specifies a pricing function $\tilde{p}(\theta)$ to be used after all histories. Further, this strategy satisfies $\tilde{p}(\theta) \leq p(\theta)$, $H(\theta, \theta; \tilde{p})/(1 - \delta) = \bar{v}$, and $v(\mathbf{p}(\theta)) = \bar{v}$ for all θ ($\tilde{T} \equiv 0$).*

We see from Proposition 9 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. This “no-war” finding serves as a useful benchmark for our subsequent analysis, when revenue equivalence breaks down: impatient firms have preferences about the use of wars, as will firms who face downward-sloping demand curves. Further, we show next that no optimal SPPE uses wars, when the distribution function is log concave.

4.3.4 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 the staticizing proposition. We characterize this pricing scheme, and the

optimal SPPE, in the following proposition:

Proposition 10 *Consider the case of inelastic demand, and allow for any distribution function F that is log-concave. If δ is sufficiently large so that off-schedule incentive constraints are satisfied, then the equilibrium path of the optimal SPPE is characterized by:*

(i). *price rigidity: $p^*(\theta) \equiv r$,*

(ii). *no wars: $T^*(\theta) \equiv 0$.*

In other words, for patient firms, if the distribution function is log concave (i.e., F/f is (strictly) increasing), 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.²⁰

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 $1/n^{\text{th}}$ of the market. The content of the proposition is that the benefits of the rigid-pricing scheme exceed the costs, provided that the distribution function is log concave.

Why do the benefits of the rigid-pricing scheme exceed the costs? We may develop the central intuition by posing and answering two simple questions. First, suppose it is true that the optimal pricing scheme gives the greatest possible profit-at-the-top that is consistent with on-schedule incentive constraints. What pricing schedule accomplishes this? Since the single-crossing property implies that $p(\theta)$ cannot decrease, profit-at-the-top is clearly highest when all firms set the same price, so that the highest type is never underpriced. Thus, we may understand the rigid-pricing scheme as the scheme that offers the highest profit-at-the-top, given that downward-sloping pricing schemes are not on-schedule incentive compatible.

Second, when is it true that expected profit is maximized over on-schedule incentive-compatible schemes when profit-at-the-top is maximized? This is where log concavity comes in. 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 profits), conditional on the “probability” that the given type will actually arise. We may thus think of $F(\theta)/f(\theta)$ as a measure of the contribution of an increase in type θ ’s profit to the firm’s expected

²⁰ As mentioned in the Introduction, this finding is closely related to the “weak-cartel” analysis by McAfee and McMillan (1992). We discuss this relationship in more detail below. Also, we note that 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.

profit. The log-concavity condition simply ensures that this measure is increasing, so that the type that contributes most to expected profit is the highest type. In total, then, we have that under log concavity the optimal collusive scheme maximizes profit-at-the-top, and under the on-schedule incentive constraints this is achieved with a rigid price.

To prove this proposition, we consider any collusive scheme that satisfies the on-schedule incentive constraint. If the scheme is a candidate for the optimal SPPE, then the allocation that it achieves can be realized without recourse to an equilibrium-path war. Using our indirect-utility formulation (2) and employing a standard trick from the literature on optimal auctions (Myerson (1981), Bulow and Roberts (1988)), we integrate by parts and rewrite our objective function as

$$E[H(\theta, \theta; p)] = E[\pi(p(\bar{\theta}), \bar{\theta}) \cdot M(\bar{\theta}; p) + \frac{F}{f}(\theta) \cdot D(p(\theta)) \cdot M(\theta; p)] \quad (4)$$

where of course $D(p(\theta)) = 1$ in the case of inelastic demand. The on-schedule incentive requirement that $p(\theta)$ is non-decreasing then implies that the market-share allocation function $M(\theta; p)$ is non-increasing. Thus, the greatest incentive-compatible market share that can be allocated to the highest-cost type is $M(\bar{\theta}; p) = 1/n$. This allocation is delivered by a rigid-pricing scheme. Further, as (4) suggests, when $F(\theta)/f(\theta)$ is increasing, it always pays to take market share away from lower types and give it to higher types. Finally, we note that the rigid-pricing scheme can be implemented with or without wars, but expected profits are strictly higher when there are no wars.²¹

The assumption that the distribution is log-concave is fairly standard in the contracting literature, and many common distributions satisfy the assumption.²² However, since (4) is a pointwise objective function, we can consider the optimal allocation of market share even if F/f is decreasing on some intervals. It is straightforward to show that the optimal pricing rule will entail sorting on intervals where F/f is decreasing, and rigidity elsewhere. Observe that if $f(\theta) > 0$ (as we have assumed), log-concavity always holds in a neighborhood of θ . Thus, it is always optimal for sufficiently patient firms to use rigid pricing at the bottom of the pricing function. For any distribution function with $f(\theta) > 0$, therefore, sufficiently patient firms can sustain SPPE payoffs strictly greater than the static Nash equilibrium (which entails sorting at the bottom).

We close this subsection by comparing our findings to those established in related papers. We recall first McAfee and McMillan's (1992) study of bidding rings. Working with a static mechanism

²¹ Notice that the profits associated with the no-wars rigid-pricing arrangement cannot be implemented with any other feasible combination of prices and wars, given the restrictions that $p \leq r$ and $T \geq 0$. If we attempt to balance an increased price (p) against an increased war (T), then price exceeds r and consumers refuse to buy; while if we attempt to balance a lower price against a reduced war, then the war assumes a negative value, which is infeasible.

²² For example, uniform, normal, logistic, chi-squared, exponential, and Laplace satisfy the restriction. In the typical principal-agent procurement problem, where the agent is privately informed, the principal wishes to minimize the agent's "virtual cost," $\theta + F(\theta)/f(\theta)$ (See, e.g., Myerson (1981) or Bulow and Roberts (1988)). Log-concavity guarantees that higher-cost agents also have a higher virtual cost, and thus it ensures that the principal's optimal contract is fully-sorting. In our model, since the agents' utilities are maximized, the assumption has the opposite effect: it implies that the "contract" entails pooling.

design formulation, McAfee and McMillan show that bidders collude best (when transfers are not allowed) in procurement auctions if they agree to bid the same price, r . Our Proposition 10 is closely related and relies on a similar proof. 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 sequel, we further generalize the analysis to consider the possibilities of downward-sloping demand and impatient firms.

Our findings are also related to those presented in the hidden-action collusion literature. A central feature of the Green-Porter (1984) and APS (1986, 1990) papers is that collusive conduct involves periodic reversions to price wars. Our model can be placed within their hidden-action modeling 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 . At the same time, Proposition 10 offers the starkly distinct prediction that optimal symmetric collusion takes place without reversion to equilibrium-path wars. What accounts for the different predictions?

The main difference between the two modeling approaches 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-Porter (1984) and APS (1986, 1990) modeling 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. We explore the generality of this “no-war” theme in the sequel, where we consider downward-sloping demand and impatient firms. We begin with the former.

4.4 Downward-Sloping Demand

We turn now to the possibility of downward-sloping demand. While it is often argued that collusive behavior is associated with inelastic-demand markets, most industries are characterized by demand curves that have at least some elasticity, and so it is important to investigate the robustness of our results. When demand is downward-sloping, $\pi_{p\theta} > 0$, and our analysis must confront two new issues. First, the Revenue Equivalence Theorem breaks down, and it is therefore no longer clear that wars are not required for an optimal collusive scheme. Second, the monopoly pricing function is strictly increasing, thus suggesting an additional cost to a rigid-price scheme.

To begin the analysis, consider a mechanism (p, T) and the associated market-share allocation function M . Then Lemma 5 implies that

$$H(\theta, \theta, p) - T(\theta) = H(\bar{\theta}, \bar{\theta}, p) - T(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} D(p(\tilde{\theta}))M(\tilde{\theta}; p)d\tilde{\theta}. \quad (5)$$

Using this expression, we consider in turn the extension of our full-sorting, no-wars and rigid-pricing findings for the downward-sloping demand case.²³

4.4.1 Fully-Sorting Pricing Schemes

Consider first the case of full sorting. As in the inelastic-demand case, the Nash pricing schedule is fully sorting, satisfies on-schedule constraints when joined with $T(\theta) \equiv 0$, and also satisfies off-schedule constraints. Further, $H(\bar{\theta}, \bar{\theta}; p) = 0$ and $M(\theta; p) = (1 - F(\theta))^{n-1}$. Thus, the RHS of (5) when evaluated at the Nash prices, $p^e(\theta)$, and with no-wars-at-the-top, gives the Nash-equilibrium interim profit. Any alternative fully-sorting scheme with uniformly higher-than-Nash prices is sure to do strictly worse, since such a schedule has higher prices and thus results in lower efficiency rents (as D is now strictly decreasing in price). Thus, the alternative scheme is strictly inferior to the repeated play of the static Nash equilibrium.

In the Appendix, we extend the analysis to include any alternative scheme that satisfies full sorting and (5) (i.e., not just those with pointwise higher prices than the Nash pricing scheme). We demonstrate there that our finding for fully-sorting pricing schemes is in fact stronger when demand slopes down, as any alternative fully-sorting scheme is then *strictly* inferior to Nash pricing:

Proposition 11 *Among the class of fully-sorting pricing schemes, and for any distribution function F and discount factor δ , if demand is downward-sloping, then the optimal SPPE is the repeated play of the static Nash equilibrium.*

In short, firms cannot achieve collusive (i.e, better-than-Nash) profits with fully-sorting schedules. The efficiency benefits that such schemes afford are overwhelmed by the informational costs that they demand.

It is interesting to contrast this finding with that which would arise were firms' costs perfectly observed. In this case, if firms were sufficiently patient, then the maximal collusive profits could be achieved with the monopoly-pricing scheme, in which all along the equilibrium path each firm selects its monopoly price (given its current costs). This scheme allocates sales to the lowest-

²³ While our analysis in this subsection remains focused on the Bertrand pricing model, we note that a model of Cournot competition, or imperfect substitutes, would yield similar results. To see this, note that if a firm's expected profits are given by $\tilde{H}(\hat{\theta}, \theta; q) = q(\hat{\theta}) \cdot (E_{-i}[P(q(\hat{\theta}) + \sum_{j \neq i} q(\theta_j))] - \theta)$, then $\tilde{H}_\theta(\theta, \theta; q) = q(\theta)$, and Lemma 5 would imply a very similar expression to (5) (we would simply substitute $q(\theta)$ for $D \cdot M$).

cost firm, thus achieving full efficiency benefits, and it does so with no pricing distortion. By contrast, when firms are privately informed as to their respective cost positions, the monopoly-pricing scheme is incentive compatible only if equilibrium-path wars are experienced. When the informational costs of wars are considered, and as Proposition 11 confirms, the monopoly-pricing scheme is no longer optimal: under full sorting, it is better to have low (i.e., Nash) prices and no equilibrium-path wars than to have high (e.g., monopoly) prices and equilibrium-path wars. Importantly, this finding obtains even if the extent of private information is small: Proposition 11 holds for *any* distribution F . It may be, for instance, that each firm's costs are equal to $E\theta$ with probability $1 - \varepsilon$ and uniformly distributed over $[\underline{\theta}, \bar{\theta}]$ with the complementary probability ε . The repeated game thus exhibits a discontinuity, as the monopoly-pricing scheme is superior to the Nash pricing scheme only in the limit, when costs are *perfectly* observed.²⁴

4.4.2 Equilibrium-Path Wars

Our next step is to relax the fully-sorting requirement and consider the use of wars under other market-share allocations. In contrast to the inelastic-demand analysis, we can no longer rule out wars altogether, but we do show that the scope for wars is limited to “intermediate” types. As we explain below, this finding is driven by a novel feature of the downward-sloping-demand model: incentive constraints can bind in both “upward” and “downward” directions.

To illustrate the main ideas, we refer to the three-step pricing schedule depicted in Figure 1. In this scheme, price ρ_k is used on interval $(\theta_k, \theta_{k+1}]$ for $k = 1, 2, 3$, where $\theta_1 = \underline{\theta}$ and $\theta_4 = \bar{\theta}$. Suppose first that demand is inelastic and that the original schedule has a war on, say, the second step. We may then construct an alternative schedule in which the second-step war is removed, and the second-step price (ρ_2) is reduced to keep θ_3 , the highest type on the second step, indifferent. Given that $\pi_{\rho\theta} = 0$, it follows that all types on the second step are indifferent between the original and the alternative schedules: all types on the same step have the same market share, and thus trade off prices and wars at the same rate. The original market-share allocation therefore remains incentive compatible under the alternative schedule. In this way, and in line with the general finding reported in Proposition 9, we may eliminate wars on a step-by-step basis, and thereby “re-engineer” a payoff-equivalent no-war schedule, when demand is inelastic.

How does this argument change when demand is downward sloping? There are two cases to consider, depending on whether the war occurs on an extreme (i.e., first or third) or intermediate (i.e., second) step. To begin, let us suppose that the original schedule specifies a war for the first step. We can then construct an alternative schedule that removes this war and adjusts the first-step

²⁴ Similar discontinuities have been observed for signaling games, commitment games with imperfect observation (Bagwell, 1995), and repeated games with private signals (Matsushima, 1990).

price downward, until the highest type on the first step, θ_2 , becomes indifferent between the original and alternative schedules. Now, since θ_2 is indifferent, there is no new incentive for second-step types to misreport a lower (first-step) type, and so the alternative schedule may simply leave the second and third steps at their original specifications. Notice, however, that with $\pi_{\rho\theta} > 0$ all types below θ_2 strictly prefer the alternative no-war schedule. The alternative schedule thus improves strictly upon the original schedule. Similarly, we can also strictly improve upon an original schedule that imposes a war at the top step. To do this, we focus on the lowest type on the top step, θ_3 , and then construct an alternative scheme that eliminates the top-step war and specifies for the top step the higher-than-original (and higher-than-monopoly) price that leaves θ_3 indifferent. All types below θ_3 are unaffected by this adjustment, while types above θ_3 strictly prefer the higher price.

In the Appendix, we extend this argument beyond the three-step example. We summarize our finding for “extreme steps” as follows:

Proposition 12 *Suppose that demand is downward sloping, and allow for any distribution function F . If (p, T) is a solution to the Mechanism Design Program and yields the market-share-allocation function $M(\theta; p)$, it must satisfy the following properties.*

- (i) *No wars at the extremes: $T(\bar{\theta})=T(\underline{\theta})=0$;*
- (ii) *If p is constant on $[\underline{\theta}, x]$ and strictly increases at x , $p(x) < p^m(x)$;*
- (iii) *If p is constant on $[y, \bar{\theta}]$ and strictly increases at y , $p^m(y) \leq p(y)$.*

We now turn to the second case mentioned above and explore the possibility of a war on an intermediate step. Referring again to Figure 1, suppose that we eliminate a war on the second step by adjusting ρ_2 downward until θ_3 is indifferent between the original and alternative schedules. With downward-sloping demand, lower types on the second step strictly prefer the alternative schedule, and so expected profit on the second step is now higher under the alternative arrangement. But, given that the lowest type on the second step, θ_2 , strictly prefers the alternative schedule, higher types on the first step will violate incentive compatibility and announce that they have second-step cost levels, unless the alternative schedule also improves the profit that they receive on the first step. Thus, when demand slopes down, intermediate-step wars can no longer be eliminated on a step-by-step basis.

Consider then the first-step adjustment that the alternative schedule must include. If the original schedule has no war on the first step, there is only one means through which θ_2 's first-step profit can be improved: the first-step price under the alternative schedule must rise toward θ_2 's monopoly price. But if we raise the first-step price in this way, then we must consider whether lower types on the first step, who find price increases less appealing than does θ_2 , are better off under the alternative schedule. Indeed, if the first-step price crosses the monopoly pricing function on the interior of $(\underline{\theta}, \theta_2)$ (so that θ_2 prefers a higher first-step price but lower types do not), then raising the

profit for θ_2 may actually *lower* average profits on the first step. This suggests that it sometimes may be best to retain the second-step war, so that the first-step price can be set at a more profitable level.

Implicit in this discussion is a non-standard feature of our problem: the incentive constraints may bind in either direction (i.e., intermediate types may most tempted to announce either a higher or a lower cost). However, types on the lowest and highest steps can only misreport in one direction. Thus, our model delivers an interesting variation on the traditional finding for optimal mechanisms that there is “no distortion for the best type.” In our model, at least as far as wars are concerned, there is both “no distortion for the best type” and “no distortion for the worst type.”

In the Appendix, we explore these tradeoffs more fully. As the three-step discussion illustrates, we are unable to rule out wars altogether. We do, however, establish some rather stringent necessary conditions for wars to be optimal:

Proposition 13 *Suppose that demand is downward sloping, and allow for any distribution function F . Let (p, T) be a solution to the Mechanism Design Program and let $(\tilde{p}, \tilde{T} \equiv 0)$ be the no-war, incentive-compatible, allocation-equivalent scheme of Lemma 6. The following are necessary conditions for $T(\theta) > 0$ on a set of positive measure:*

- (i) p has at least two points of strict increase;
- (ii) There exists an $x > \underline{\theta}$ such that $p(\theta)$ is constant on $[\underline{\theta}, x)$;
- (iii) The no-war scheme is better for all types above the lowest step: $H(\theta, \theta; \tilde{p}) \geq H(\theta, \theta; p) - T(\theta)$ for all $\theta > x$.
- (iv) On the lowest step, the no-war scheme is better for types on the right end of the step, but worse on average over the step:

$$p^m(\underline{\theta}) \leq p(\underline{\theta}) < \tilde{p}(\underline{\theta}) \leq p^m(x), \text{ and}$$

$$E[H(\theta, \theta; \tilde{p}) | \theta < x] < E[H(\theta, \theta; p) - T(\theta) | \theta < x].$$

Thus, a sufficient condition for the solution to entail $(p(\theta), T(\theta) \equiv 0)$ is that demand is sufficiently inelastic that any scheme which satisfies IC-OnM, and further uses three or more prices, a step at the bottom, and no wars satisfies $p(\underline{\theta}) < p^m(\underline{\theta})$.

This proposition sharply limits the scope of wars in an optimal SPPE: wars make matters worse for all types but the lowest step. Thus, we can interpret the use of wars as a way to lower the profits of intermediate types, in order to relax the incentive constraint of type x when the lowest step is $[\underline{\theta}, x)$. By lowering the profit to type x , we are able to lower the price on the bottom step, better matching the *average* monopoly price over the entire lower step.

The necessary conditions for a war are, however, stringent. First, Proposition 12 rules out wars at the top and the bottom of the pricing schedule, and thus the schedule must involve significant

sorting (i.e., multiple steps) before wars can be potentially attractive.²⁵ Second, when the extent of sorting is sufficiently great, the lowest-cost price rests below its monopoly level, and Proposition 13 implies that wars are not optimal.²⁶

4.4.3 Optimal Pricing as an Optimization Program

Our next goal is to characterize the optimal pricing policy for patient firms. We proceed in two steps. First, we reformulate the Mechanism Design Program in terms of an optimization program. Second, we analyze the reformulated program and develop some implications for the optimal pricing schedule.

We begin by highlighting a “no-subsidy” constraint that is important for our analysis. A key feature of our model is that firms can not receive subsidies: $T(\theta) \geq 0$. To illustrate the significance of this constraint, we return to the three-step example described above and reconsider the possibility of an intermediate-step war. If we had been able to transfer a “negative war” (i.e., $-T(\theta) > 0$) to members of the first step, then we would have been able to eliminate the second-step war while maintaining incentive compatibility, without altering the first-step price. Thus, the no-subsidy constraints are behind the fact that incentive constraints can bind in both directions,²⁷ as well as the possibility of intermediate-step price wars. To understand why these constraints did not arise when demand is inelastic (except for $T(\bar{\theta})$), recall that in order to satisfy (IC-onM), any positive subsidy would be offset by a proportional decrease in price; but when demand is inelastic, such an exchange affects all types on the step in the same way and is thus revenue neutral.

Using Lemma 5 to incorporate (IC-onM) into the no-subsidy constraint as well as the objective, the firms’ objective is to maximize

$$E[H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) + \int_{\theta}^{\bar{\theta}} D(p(\tilde{\theta}))M(\tilde{\theta}; p)d\tilde{\theta}]$$

subject to p nondecreasing and

$$H(\theta, \theta; p) - H(\bar{\theta}, \bar{\theta}; p) + T(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} D(p(\tilde{\theta}))M(\tilde{\theta}; p)d\tilde{\theta} \geq 0.$$

²⁵ Thus, in a simpler model with only two cost types, wars would not arise at all. Even with three cost types, wars would arise for the middle type only if that type occurs with sufficiently low probability.

²⁶ The meaning of “sufficiently great” sorting varies with the demand curve under consideration. For example, when demand is fairly inelastic, even a modest degree of sorting requires sub-monopoly pricing for the lowest-cost type. (In the extreme case of inelastic demand, the lowest-cost type selects a sub-monopoly price, if the pricing schedule has even one step.) In addition, if the extent of sorting is perfect (full sorting), then incentive compatibility implies that the lowest-cost price cannot exceed its monopoly level, and this is true regardless of the properties of the demand curve.

²⁷ Jullien (1998) shows that when participation constraints are type-dependent, incentive constraints can bind in both directions. He develops a very similar mathematical approach to that found in this section, but the constraints he considers are participation constraints rather than our “no-subsidy” constraints, which are better interpreted as restrictions on transfers. In our analysis of impatient firms, we will interpret the off-schedule incentive-compatibility constraints as participation constraints.

A difficulty with solving this program is that the market-share-allocation function depends on the entire pricing function. Thus, it will be useful to decompose the choice of the pricing function into two components: first, the choice of whether the pricing function is strictly increasing (“sorting”) or flat (“pooling”) on an interval, and second, the choice of the level of prices. To this end, we develop some additional notation. As mentioned in Section 2, we restrict attention to pricing schedules which can be represented by a finite number of subintervals (indexed by k and with boundaries $(\theta_k, \theta_{k+1}]$), and incentive compatibility implies that in any one such subinterval the pricing function is either strictly increasing or flat. To represent such a function, we choose an even integer K large enough and use a nondecreasing vector $\theta^K = (\theta_1 = \underline{\theta}, \theta_2, \dots, \theta_{K+1} = \bar{\theta})$, which assigns market share as follows. If k is odd, $(\theta_k, \theta_{k+1}]$ represents an interval of sorting and market share is given by $(1 - F(\theta))^{n-1}$. If k is even, $(\theta_k, \theta_{k+1}]$ represents an interval of pooling. We denote by ρ_k and T_k the price and war on that interval, and the market share (denoted $\mu(\theta_k, \theta_{k+1})$) is determined as

$$\mu(\theta_k, \theta_{k+1}) = \sum_{j=0}^{n-1} \binom{n-1}{j} \frac{1}{j+1} (F(\theta_{k+1}) - F(\theta_k))^j (1 - F(\theta_{k+1}))^{n-j-1}.$$

If there are two adjacent intervals of pooling at different prices, we would have $\theta_k < \theta_{k+1} = \theta_{k+2} < \theta_{k+3}$ for some k even.

To begin, we hold the market-share allocation fixed, and consider characterizing the level of the prices and the use of wars. To do so, we introduce multipliers for the no-subsidy constraints. On intervals of full sorting, there is a constraint for each type, and we let $\lambda(\theta)$ be the multiplier. However, on a region of pooling $(\theta_k, \theta_{k+1}]$, the no-subsidy constraint only binds for type θ_{k+1} . Thus, we let λ_k be the multiplier at that point. We further let $\Lambda(\theta)$ represent the sum of the multipliers for types less than θ ; formally, $\Lambda(\underline{\theta}) = 0$, and $\Lambda(\theta)$ has density $\lambda(\theta)$ on regions of sorting, while on a region of pooling, $\lambda_k = \Lambda^+(\theta_{k+1}) - \Lambda^-(\theta_{k+1})$, where the $+$ and $-$ indicate limits from the right and the left.

Recall that in Section 4.3.4, we were able to reformulate the objective as a pointwise maximization problem using integration by parts. Here, we use the same approach, except now we must integrate by parts the constraints as well as the objective. To simplify notation somewhat, and to highlight the fact that we are choosing prices and market share separately, we let $M(\theta; \theta^K)$ represent the expected market share for type θ given the rule θ^K . Then, we may derive the following objective:

$$E \left[[H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta})][1 - \Lambda(\bar{\theta})] + \frac{F(\theta) - \Lambda(\theta)}{f(\theta)} D(p(\theta)) M(\theta; \theta^K) \right] \quad (6)$$

$$+ \int_{\underline{\theta}}^{\bar{\theta}} \pi(p(\theta), \theta) M(\theta; \theta^K) d\Lambda(\theta)$$

To analyze this expression, first consider the optimal choice of $T(\bar{\theta})$. Clearly, we must have $\Lambda(\bar{\theta}) = 1$; otherwise, we would wish to either increase T without bound, or else violate the constraint. Next consider the optimality conditions for prices. For simplicity, we begin by ignoring the constraint that the pricing function should be nondecreasing. In regions of sorting, the optimality conditions for $p(\theta)$ require

$$\lambda(\theta)\pi_\rho(p(\theta), \theta) = -D'(p(\theta))[F(\theta) - \Lambda(\theta)].$$

The right-hand side of this expression describes the effect on all lower types of increasing demand by lowering the price, when (IC-onM) is satisfied. In particular, increasing demand has positive and negative externalities. The term $F(\theta)$ captures the benefit of increasing demand, through the increase in the information rents available for all lower types. The term $\Lambda(\theta)$ captures the negative externality of increasing demand: if the information rents increase, it becomes more difficult for lower types to satisfy the incentive compatibility constraint without subsidies. The left-hand side of the expression describes the effect on type θ of raising the price (while still subject to (IC-onM)). This price change raises the profit of type θ , which relaxes the no-subsidy constraint and thus increases the objective by $\lambda(\theta)$ (the shadow value of the constraint).

Now consider the existence of wars. At this point, we must reconsider the constraint (ignored until now) that the pricing function should be nondecreasing. Observe that if $T(\theta) > 0$, then $\lambda(\theta) = 0$, which implies that either (i) $F(\theta) = \Lambda(\theta)$, or (ii) the monotonicity constraint binds. Consider first case (i). Since we have assumed that F is strictly increasing, but $\Lambda(\theta)$ is constant on a region with wars, we cannot have this equality over an open interval of types. Thus, we conclude that wars on regions of sorting occur only if they serve to sort between types charging essentially the same price, but getting different market shares. To ensure that a solution exists in this case, we assume for simplicity that firms can receive different market shares while charging exactly the same price, so long as this is incentive compatible; we will, however, continue to use the term “rigid price” to describe regions of pooling, noting explicitly if sorting occurs at a fixed price.²⁸

For regions of pooling, since price and market share are fixed for all types on such a region, the derivative of the objective with respect to ρ_k is given by

$$\lambda_k \pi_\rho(\rho_k, \theta_{k+1}) + D'(\rho_k) \left[\int_{\theta_k}^{\theta_{k+1}} F(\theta) d\theta - \Lambda^+(\theta_k) \right].$$

The analysis is qualitatively similar to regions of sorting, except that it is possible that $\int_{\theta_k}^{\theta_{k+1}} F(\theta) d\theta = \Lambda^+(\theta_k)$ and $T_k > 0$.

We turn finally to the joint determination of the pricing function and market-share allocation.

²⁸ Of course, there are other ways to handle this problem; for example, we could assume a large but finite number of prices.

With downward-sloping demand, the efficiency costs of a rigid-pricing scheme are more pronounced, as a scheme with some sorting carries the additional advantage that lower-cost types benefit from the additional demand that a lower price implies. Figure 2 illustrates one possibility. In this figure, the monopoly price locus is initially almost flat, then rises sharply, and then becomes almost flat again. Imposed upon this figure is a two-step pricing schedule that closely tracks the monopoly locus. Assuming that a schedule such as this can be constructed that satisfies the on-schedule incentive constraints, we see that the market share is allocated at close to monopoly prices, and there is thus no hope that a schedule with a single rigid price is optimal.

Despite this, log-concavity of F implies that when demand is not too steep, there is pressure to group types together on a single step. However, the above program highlights the complications which arise. First, monotonicity of F/f creates a force in favor of pooling types together into steps and increasing the market share for high realizations of F/f relative to low ones. When only one price is used, incentive-compatibility for lower types is not a problem and $\Lambda(\theta) = 0$. In contrast, once it becomes optimal to break the pricing function into steps, the no-subsidy constraints start to bind. “Pooling” an interval of types together on a step will be optimal over an interval only if, instead, $\frac{F(\theta) - \Lambda(\theta)}{f(\theta)}$ is nondecreasing. Since the latter expression must be nondecreasing at $\underline{\theta}$ when $f(\underline{\theta}) > 0$ (as we have assumed), and since the objective (6) can be maximized pointwise, it follows immediately that we can improve on the (fully-sorting) static Nash equilibrium. In particular, a scheme $(\tilde{p}, \tilde{T} \equiv 0)$ which is fully-sorting except in a neighborhood of $\underline{\theta}$, will be better than Nash.

In the appendix, we formalize this intuition and further analyze the program. Our conclusions are summarized in the following proposition.

Proposition 14 *Suppose that demand is downward-sloping, and allow for any distribution function F . If (p, T) solves the Mechanism Design Program, then:*

- (i). *There exists an interval $[\underline{\theta}, x)$, $x > \underline{\theta}$, on which no wars are used and the price is rigid.*
- (ii). *For any distribution F , no wars are used on a fully-sorting interval where the pricing function has non-negligible slope. On any interval where F is log-concave, the density f is nondecreasing, and the scheme is fully-sorting, no wars are used.*
- (iii). *If demand is sufficiently inelastic and F log-concave, then a rigid price is used for all types, and there are no wars.*

For the downward-sloping demand case, we see that optimal pricing policies necessarily exhibit some rigidity, and wars are used only in a limited set of circumstances. Moreover, using the staticizing proposition, we see further that optimal symmetric collusion among patient firms is characterized by rigid pricing and the absence of equilibrium-path wars, provided that demand is sufficiently inelastic and the distribution function is log-concave. When these assumptions are relaxed, however, collusion with downward-sloping demand has some relatively subtle features. We have iden-

tified circumstances under which firms may engage in seemingly suboptimal behavior for some cost realizations, in an effort to maximize expected profits over all possible costs.

5. Optimal Symmetric Collusion Among Impatient Firms

Taking as our focus the case of inelastic demand, we now consider collusion among impatient firms. We proceed in three steps. First, we determine the critical discount factor above which the optimal symmetric collusive arrangement can be enforced. Second, when firms are less patient, off-schedule incentive constraints begin to bind, and we show that this provides a separate motivation for the avoidance of price wars. Finally, we examine the collusive pricing schemes that may be used by impatient firms.

5.1 Enforcing Rigidity Off Schedule

We begin with the determination of the critical “patience” level that firms must possess, in order to enforce the rigid-pricing scheme. To accomplish this, we must first specify the punishment that follows an off-schedule deviation. Allowing for a range of possibilities, we suppose that following a deviation firms revert to a “war” continuation value, w , drawn from the SPPE set, V_s . This continuation value may correspond to the worst SPPE continuation value (if a worst value exists), or perhaps to the infinite reversion to the Nash equilibrium (p^e) of the static game. Treating w parametrically, we allow for these cases and others as well. The restriction that we place on w is that the corresponding per-period expected payoff, $\pi^w = (1 - \delta)w$, that is received on average over the long run is such that

$$0 \leq \pi^w < \frac{1}{n}\pi(r, E\theta). \quad (7)$$

Note that $\pi^w \geq 0$ must be true in an SPPE, while the second inequality simply captures the meaning of a punishment. So long as the distribution is not too log-convex, the inequalities will be satisfied when $\pi^w = E[H(\theta, \theta; p^e)]$.

The off-schedule incentive constraint captures the tension that colluding firms experience between the current-period incentive to cheat and the long-term cost of the consequent deterioration in cooperation. Suppose, for example, that the firms attempt to collude by maintaining the 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 $1/n^{\text{th}}$ of the market. The firm’s incentive to cheat is thus represented by the market-share gain times the profit-if-win: $\frac{n-1}{n}\pi(\rho, \theta)$. Note that the incentive to cheat is greatest for a firm with the lowest cost level, $\underline{\theta}$,

since the profit-if-win is then highest, and so the gain in market share is most valuable.²⁹ Balanced against the incentive to cheat is the expected discounted value of future cooperation that would be forfeited if the firm were to cheat. The expected discounted value of future cooperation in this case is given as $(\delta/(1-\delta))[\frac{1}{n}\pi(\rho, E\theta) - \pi^w]$. Thus, the proposed rigid-pricing scheme is off-schedule incentive compatible only if

$$\frac{n-1}{n}\pi(\rho, \underline{\theta}) \leq (\delta/(1-\delta))[\frac{1}{n}\pi(\rho, E\theta) - \pi^w]. \quad (8)$$

Clearly, this condition will hold if δ is sufficiently large.

The situation in which $\rho = r$ is of particular interest. When firms seek at all times to maintain the rigid price r , the critical discount factor implied by (8) takes the form

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

Notice that (7) implies that $\delta^* \in (\frac{n-1}{n}, 1)$. With this calculation at hand, we may conclude that the rigid-pricing scheme identified in Proposition 10 satisfies the off-schedule incentive constraint provide that firms are sufficiently patient; i.e., provided that $\delta \geq \delta^*$.

We consider next the possibility that firms are impatient and thus possess a discount factor δ such that $\delta < \delta^*$. It is clear that such firms cannot maintain a rigid price at r : a firm would be overwhelmed by the incentive to cheat when it draws the lowest cost level. But can they enforce a lower rigid price? Or, can they enforce a non-stationary rigid-price scheme, in which they move over time between different rigid prices? We show in the Appendix that the answer to both questions is, No. We thus summarize our findings as follows:

Proposition 15 *For the inelastic-demand case, and for any distribution function F ,*

- (i). *If $\delta \geq \delta^*$, then there exists an SPPE along whose equilibrium path the price is rigid at the consumer valuation, r .*
- (ii). *If $\delta < \delta^*$, then there does not exist an SPPE along whose equilibrium path the price is rigid in every period.*

We now recall from Proposition 10 that the rigid-pricing scheme in which $p(\theta) \equiv r$ along the equilibrium path characterizes the optimal SPPE when demand is inelastic and the distribution function is log concave, provided that the off-schedule incentive constraint is met. Using Proposition 15(i), we thus have:

²⁹ Observe as well that the incentive to cheat is higher when the price is higher, since then the profit-if-win is greater. This fact plays an important role in the subsections that follow.

Corollary 16 *For the inelastic-demand case, if the distribution function F is log concave and $\delta \geq \delta^*$, then the optimal SPPE entails a rigid price at the consumer valuation r .*

As an example, we may suppose that: (i) there are two firms, (ii) costs are uniformly distributed over $[0, 1]$, (iii) the consumer reservation value r is unity, and (iv) following an off-schedule deviation, the firms revert forever after to the Nash equilibrium of the static game. For these parameter values, we may calculate that $\pi^w = 1/6$, and so we may compute directly that $\delta^* = 6/7$.

We close this subsection with a pair of observations. First, we emphasize that Proposition 15 requires no assumption on the distribution function, since each firm is posited to set the same price regardless of its cost draw. Indeed, holding the expected value of θ fixed, the proposition applies as stated even if costs are correlated across firms (but not time). Second, the proposition also indicates a sense in which intertemporal fluctuations in costs diminish the ability of firms to collude. To see this, we may compute the critical discount factor which would sustain collusion at the price r if firms' costs were symmetric and time invariant. This value can be extracted from the formula for δ^* for the situation in which $\underline{\theta} = E\theta$ and $\pi^w = 0$ (corresponding to infinite reversion to the standard Bertrand equilibrium), which yields the critical discount factor value to be $\frac{n-1}{n}$. Thus, when firms' costs vary through time, collusion indeed requires greater patience (i.e., $\delta^* > \frac{n-1}{n}$). Intuitively, cost fluctuations diminish the potential for collusion, since a collusive scheme then must be robust to the event that a firm draws a current cost level that is low as compared to that which it expects in the future.³⁰ This 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 and over time.

5.2 No Wars on the Equilibrium Path

What is the optimal form of symmetric collusion among impatient firms? The previous proposition provides some guidance: for less patient firms, if demand is inelastic, no rigid-pricing scheme (at any levels) satisfies the off-schedule constraint. The collusive pricing schedule must therefore vary, at least somewhat, with current-period costs. In this subsection, we add a second ingredient toward an answer. We establish that the scope for symmetric collusion cannot be improved (and may be strictly harmed) by the inclusion of equilibrium-path wars, at least when demand is inelastic.

We relegate the formal argument to the Appendix and describe here only the central ideas. Suppose that demand is inelastic and let us start with an original SPPE collusive scheme. Following the factorization logic of APS (1986, 1990), we may break the payoff stream into current-period and continuation-value components. Relying on Proposition 9, if we take an original SPPE in which

³⁰ This observation holds even if firms were able to inflict a zero-profit punishment (i.e., $\pi(w) = 0$) when their cost fluctuations are private. If with private information the punishment threat is less severe (i.e., $\pi(w) > 0$), then $\delta^* - \frac{n-1}{n}$ is higher yet.

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 incentive constraint (given that the original did) and thus constitutes a payoff-equivalent SPPE for patient firms. When firms are impatient, however, the off-schedule incentive 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 expected discounted value of cooperation is enhanced. The off-schedule incentive compatibility constraint is therefore now easier to satisfy than under the original scheme.

We note that this argument is quite general: we place no restrictions on the distribution function F or the level of patience δ . Further, the idea can be applied to characterize the schemes used to achieve any value in the equilibrium set V_s , not just optimal points in the set. However, to apply the ideas to intermediate values in V_s , we need to be more precise about what we mean by a “war” when neither today’s pricing function nor tomorrow’s continuation value function is optimal.

Consider a scheme with pricing schedule p and a continuation value function v , and let $T(\theta) = \delta \cdot (\sup_{\theta} E_{-i} v(p(\theta), p) - E_{-i} v(p(\theta), p))$ denote the difference between the maximal expected continuation value that this scheme allows (for some realization of firm i ’s costs) and the value specified by the scheme when firm i has type θ . Then we can rewrite each type’s interim repeated-game payoff as

$$H(\theta, \theta; p) - T(\theta) + \delta \sup_{\theta} E_{-i} v(p(\theta), p).$$

We interpret $T(\theta) \equiv 0$ as the absence of an equilibrium-path price war for the given collusive scheme: any sorting between types is accomplished with variation in the pricing function, not variation in future continuation values.

Finally, recall that in Section 4.2, we imposed an assumption that the lowest type on the highest step must receive non-negative profits ($H(\theta_K, \theta_K; p) - T(\theta_K) \geq 0$). This assumption was required to apply Lemma 6, which guaranteed the existence of an alternative, incentive-compatible scheme ($\tilde{p}, \tilde{T} \equiv 0$) that delivers $M(\theta; p)$ with the same profit-at-the-top. The analog of that assumption here is given by

$$H(\theta_K, \theta_K; p) + \delta E_{-i} v(p(\theta_K), p) \geq \delta \cdot \sup_{\theta} E_{-i} v(p(\theta), p). \quad (9)$$

This condition simply requires that type θ_K gets non-negative profits relative to the no-war continuation value for the comparison scheme (p, v) . This would be satisfied if there were no wars at the top, that is, if $T(\theta_K) = 0$, and price is above cost for θ_K . With this notation in hand, we can then formalize the idea that any SPPE value that is implemented using wars can also be implemented

without them:

Proposition 17 *Assume that demand is inelastic, and consider any distribution function F and any discount factor δ . Suppose that there exists an SPPE which specifies a pricing strategy p in the current period and a continuation value function v where (9) holds and, for some θ , $E_{-i}v(p(\theta), p) < \sup_{\theta} E_{-i}v(p(\theta), p)$.*

Then, there exists an alternative SPPE which yields the same expected payoff for firms in the repeated game, where the firms use a strategy \tilde{p} in the current period and a continuation value function \tilde{v} such that $E_{-i}\tilde{v}(p(\theta), \tilde{p}) = \sup_{\theta} E_{-i}v(p(\theta), p)$ for all θ .

A consequence of this result is that the staticizing proposition may be generalized to the case of impatient firms. Our problem now becomes to look for a solution to the Mechanism Design Program that satisfies the off-schedule constraints.

We consider next the additional considerations that arise when demand is downward sloping. Let us suppose that the original arrangement specifies p as a step function. Consider now a step k , which lies on the interval $[\theta_k, \theta_{k+1}]$ and allows for an equilibrium-path war (i.e., $T_k > 0$). We then may reengineer an alternative schedule (lowering T_k and lowering p_k) that keeps θ_{k+1} 's expected payoff constant. In line with the discussion above, when demand is inelastic, all types on step k are indifferent to the new schedule, and so the alternative schedule preserves on-schedule incentives while making it easier to satisfy the off-schedule constraint. But what if demand is downward sloping? Consider first the impact of the alternative schedule for other types on step k . Notice that θ_k is now made better off under the alternative schedule. Further, as in the inelastic case, θ_k has less incentive to cheat once the price is reduced, and so the off-schedule incentive constraint is clearly relaxed.³¹ The basic logic that eliminating wars can only help off-schedule incentives therefore extends to the downward-sloping demand case, when we restrict attention to a given step.

However, as we showed in Section 4.4, when demand is downward sloping, it is no longer appropriate to make payoff-equivalent adjustments one step at a time. In particular, if we wish to preserve the original market-share allocation under the alternative schedule, then the gain that θ_k enjoys necessitates a gain as well for a slightly lower type (on step $k - 1$). The alternative schedule can create a gain for the lower type, if it reduces the war magnitude and/or increases the price for step $k - 1$. Of course, when this step originally does not specify a war, only the latter avenue is available: the alternative schedule then must specify a higher price for step $k - 1$ than did the original schedule. This has two potential drawbacks. First, as discussed in Section 4.4, a

³¹ As mentioned above, when demand slopes down, a collusive pricing scheme may require that price exceeds the monopoly price for lower types on a given step. In this case, the best off-schedule deviation for such a type is to undercut all the way down to the monopoly price. This does not change the basic message, however, as an alternative schedule with a lower supra-monopoly equilibrium price for such a type continues to reduce the incentive to cheat: as compared to the original schedule, this type earns greater profit at the equilibrium price and the same profit following its best deviation.

price increase for a lower step may reduce expected profit over that step. Second, when firms are impatient, if a higher price is required on a lower step, then types on that step experience a greater incentive to cheat, and so a violation of the off-schedule incentive constraint may ensue. Thus, we conclude that, while adding a war to a given step worsens the off-schedule incentive constraint for types on that step, such a maneuver may be potentially useful as a means to relax the constraints of lower types.

5.3 Partial Rigidity and Collusion Among Impatient Firms

Returning to the inelastic-demand setting, we consider now collusive arrangements among firms when $\delta < \delta^*$. 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. In this subsection, we first present a sufficient condition under which a *two-step pricing scheme* can be enforced with impatient firms. Second, we argue that optimal SPPE for impatient firms must necessarily be characterized by a pricing schedule that is a step function (partial rigidity), if the collusive scheme is to offer better-than-Nash profits.

5.3.1 Introducing A Second Step: An “Escape Clause”

Our first task is to find a sufficient condition for a two-step pricing scheme to relax the off-schedule constraint and enable a cooperative outcome, even when firms are too impatient to support a rigid-pricing scheme. We develop this condition below, and we interpret the two-step pricing scheme in terms of an “escape-clause” provision, under which the collusive arrangement allows (on the equilibrium path) for a lower price when a low-cost realization is experienced. We also illustrate with a uniform-distribution example that two-step pricing schemes are not always successful: when the condition is not met, two-step pricing schemes may violate the off-schedule incentive constraint for exactly the same discount factors (i.e., $\delta < \delta^*$) as does the rigid-pricing scheme.

To gain some intuition, we recall that the rigid-price scheme fails to be enforceable when $\delta < \delta^*$, because a firm that draws the lowest-cost type finds it too tempting 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 will perceive less immediate benefit from cheating. Firstly, this firm now expects greater than a $1/n^{\text{th}}$ 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

the reduction in expected profit (i.e., the discounted value of cooperation) that a two-step scheme involves: 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 sufficiently small for lower-cost types. For other distributions, however, the resolution may be less clear

Maintaining (7), we now state a proposition that confirms this “small-density” intuition for the case of two firms:

Proposition 18 *For the inelastic-demand case with two firms, if*

$$f(\underline{\theta}) < \frac{(1/2)\pi(r, E\theta) - \pi^w}{r \cdot \pi^w}$$

then there exists $\delta^\circ < \delta^$, such that, for every $\delta \in (\delta^\circ, \delta^*)$, there exists a two-step pricing scheme with $p_2 = r > p_1$ and $\theta_2 \in (\underline{\theta}, \bar{\theta})$ that can be supported as a SPPE.*

This proposition is proved in the Appendix. It reveals two main points. First, if the density on the lowest type is sufficiently small, then the firms can enforce a two-step pricing scheme as a SPPE, even if the firms are not sufficiently patient to support a rigid price. Second, the proposition 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.³² In other words, symmetric collusion among impatient firms may call for an “escape clause” provision, under which a firm is allowed to select a low price in the event that a very favorable cost type is realized. This behavior is reminiscent of the findings of Rotemberg and Saloner (1986), although in their case collusive prices adjust *simultaneously* across firms and the adjustment is in response to a public demand shock. In the present context, by contrast, the collusive arrangement allows “rare exceptions” to rigidity on a firm-by-firm basis.

The small-density condition plays an intuitive role, but the assumption is restrictive. Consider, for example, the 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 the condition is not met when $\alpha = 1$ (corresponding to the uniform distribution), and it also fails when $\alpha < 1$. In addition, when the assumption is

³² The proof constructs a two-step pricing equilibrium for θ_2 close to $\underline{\theta}$. In this equilibrium, the low-step price is approximately $r/2$, indicating a 50% reduction from the high-step price.

violated, it may be that, for all $\delta < \delta^*$, no two-step pricing schedule satisfies on- and off-schedule incentive constraints. In this event, two-step schedules offer no scope for improvement: they fail off-schedule constraints for exactly the same discount factors as did the rigid-price schedule. We illustrate this possibility in the Appendix with a simple example in which the distribution function is uniform and firms revert following an off-schedule deviation to the static Nash equilibrium.

5.3.2 Optimal Pricing for Impatient Firms as an Optimization Program

To close this section, we consider the optimal SPPE for impatient firms facing an inelastic demand when rigidity cannot be supported. The problem faced by the firms in this context is particularly subtle, because, as we have illustrated throughout this section, the incentive for each cost type to deviate from the collusive agreement depends not only on its cost type, but also on its own price as well as the expected payoffs of the entire collusive agreement. Perhaps surprisingly, the set of tools available to deal with this kind of problem is somewhat limited.³³ Nonetheless, we will show that the firms' program can be stated and analyzed to yield a few qualitative predictions about the optimal pricing scheme in the presence of binding off-schedule incentive constraints. A question of particular interest is the following: 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?

To simplify the analysis in this section, we build on Proposition 17, which implies that we may restrict attention to stationary pricing schemes. For such schemes, we may simply consider analyzing the Mechanism Design Program, with the following modifications: we impose $T \equiv 0$, and we incorporate the off-schedule incentive constraints. To see how to do this, note that the only off-schedule incentive constraints that bind are the constraints involving the low-cost type on a flat segment of the pricing function. These constraints may be captured using the notation developed above for pricing functions that can be represented by a finite number of subintervals. In particular, observing that $\rho_k - \theta_k = H(\theta_k, \theta_k; p)/M_k$ and rearranging, the off-schedule constraint for type θ_k , at the left endpoint of step k , can be written:

$$H(\theta_k, \theta_k; p) \frac{(1 - F(\theta_k))^{n-1} - M_k}{M_k} - \frac{\delta}{1 - \delta} (E[H(\theta, \theta; p)] - \pi^w) \leq 0.$$

Now consider the full program for the firms. As in Section 4.4.3, we decompose the firms' choice of a collusive pricing scheme into two parts: that of determining the intervals of pooling and sorting, and that of determining the level of the price. However, unlike the case of downward-sloping demand, once these intervals and the price for the high-cost type are chosen, since we know

³³ A few papers (most recently, Jullien (1997)) have extended the standard adverse selection model to include participation constraints which vary according to a player's type. We discuss Jullien's paper at the end of this section.

no wars are used, the rest of the prices are fully determined by incentive compatibility.

We assign a multiplier ν_k to each off-schedule constraint. Recall the objective (4) we derived in Section 4.3.4, using Lemma 5 and integration by parts; we now incorporate the off-schedule constraints into the objective. We introduce some simplifying notation. Let $\psi(\boldsymbol{\nu}) = 1 + \frac{\delta}{1-\delta} \sum_{k=1}^{K/2} \nu_{2k}$, let $\boldsymbol{\theta}^k = (\theta_1, \dots, \theta_k)$, and let (for $k \leq K$ even) $\Upsilon(k, \boldsymbol{\theta}^k) = \sum_{j=1}^{k/2} \nu_{2j} \frac{(1-F(\theta_{2j}))^{n-1} - \mu(\theta_{2j}, \theta_{2j+1})}{\mu(\theta_{2j}, \theta_{2j+1})}$. Then for a given K large enough, the firms' objective can be rewritten as follows (recalling that odd steps are intervals of sorting and even steps are intervals of pooling):

$$\begin{aligned} & \max_{\theta^K \text{ nondecreasing}, \{\nu_k\}} (p(\bar{\theta}) - \bar{\theta}) \mu(\theta_K, \theta_{K+1}) [\psi(\boldsymbol{\nu}) - \Upsilon(K, \boldsymbol{\theta}^K)] - \sum_{k=1}^K \nu_k \pi^w \\ & + \sum_{k=1}^{K/2-1} \mu(\theta_{2k}, \theta_{2k+1}) \left[\psi(\boldsymbol{\nu}) \int_{\theta_{2k}}^{\theta_{2k+1}} F(\tilde{\theta}) d\tilde{\theta} - (\theta_{2k+1} - \theta_{2k}) \Upsilon(2k, \boldsymbol{\theta}^{2k+1}) \right] \\ & + \sum_{k=1}^{K/2} \int_{\theta_{2k-1}}^{\theta_{2k}} (1 - F(\tilde{\theta}))^{n-1} \left[\psi(\boldsymbol{\nu}) F(\tilde{\theta}) - \Upsilon(2k-1, \boldsymbol{\theta}^{2k+2}) \right] d\tilde{\theta}. \end{aligned}$$

The constraints serve to mitigate the incentive of the firms to use rigid pricing. In the objective above, changing θ_k can potentially worsen not only the off-schedule constraint for step k , but also the constraints for all lower types.

We can use this formulation of the objective to address the question of whether firms will use collusive schemes where the pricing function is strictly increasing on an interval. The following proposition provides (somewhat restrictive) conditions under which intervals of strict increase can be ruled out:

Proposition 19 *Suppose that demand is inelastic, F is log-concave, and further, the density is nondecreasing. If there exists a stationary SPPE that yields greater-than-Nash profits, then in an optimal stationary SPPE there exists no open interval of types, (θ', θ'') , where the pricing function is strictly increasing.*

Proposition 19, proved in the Appendix, establishes that we would do better to introduce tiny regions of pooling rather than strictly separate types throughout an open interval. To relate this to our discussions of log-concavity from Sections 4.3.4 and 4.4.3, notice that a region of strict separation's contribution to the objective is

$$E \left[\mathbf{1}_{[\theta_{k-1}, \theta_k]} (1 - F(\theta))^{n-1} \frac{\psi(\boldsymbol{\nu}) F(\theta) - \Upsilon(k-1, \boldsymbol{\theta}^{k-1})}{f(\theta)} \right].$$

Log-concavity of F implies that F/f is increasing, and there is an incentive to allocate market share away from low types and towards high types, by pooling the types on the same price. When the density is nondecreasing, we can further ensure that the expression $-\Upsilon(k-1, \boldsymbol{\theta}^{k-1})/f(\theta)$ is

also nondecreasing. Thus, the incentive to pool types remains. Put this way, the density condition appears as a sufficient, but not necessary, condition to rule out intervals of strict increase: more generally, the density should not decrease too fast relative to the log-concavity of F .³⁴

Summarizing, we have generalized the analysis to incorporate off-schedule incentive constraints. The problem can be analyzed as a static mechanism design program with two unusual features: the participation constraints vary by type and are endogenously determined. Reformulating the objective to incorporate these constraints, we establish sufficient conditions under which the solution is a step function. Under these conditions, therefore, an optimal SPPE pricing scheme exhibits rigidity over regions of costs.

Finally, we note that our results are related to a developing literature in the theory of static mechanism design, which considers problems where participation constraints may vary with an agent's type.³⁵ As discussed in Fudenberg and Tirole (1991, p. 263), there are as yet few general findings for this class of problems. Our formal approach is most closely related to the methods recently put forth by Jullien (1997), in his study of type-dependent participation constraints for adverse-selection problems with a single agent.³⁶ However, our objective has additional complications, since the "outside option" is itself endogenous. The appearance of the expression $\psi(\nu)$ in the objective may be understood in this light. This expression incorporates the additional role of the overall expected profits of the mechanism in relaxing the off-schedule incentive constraints. More generally, our analysis of collusion adds to the motivation for new research on solving mechanism design problems in which participation constraints are type dependent and endogenous, transfers are restricted, and incentive constraints can bind in multiple directions (see Rochet and Stole (1998) and Chone and Rochet (1998) for some recent advances along these lines, motivated by nonlinear pricing problems).

6. Simple Examples of Asymmetric Collusion

In an SPPE, a firm balances the short-term benefit of a low price against the long-term cost of

³⁴ In contrast, if our assumptions fail (for example, if f is decreasing too quickly and Υ is large), the firms may have an incentive to separate types. Intuitively, it may be useful to sacrifice profits on intermediate types in order to relax the off-schedule incentive constraints of lower types.

³⁵ Participation constraints which vary by type arise in studies of countervailing incentives (Lewis and Sappington (1989); Maggi and Rodriguez (1995)), price discrimination (Chone and Rochet (1998); Rochet and Stole (1998)), and in labor market applications where the outside options of workers are positively related to their ability. They also arise in mechanism design problems with renegotiation (Laffont and Tirole (1990)).

³⁶ Jullien reformulates the objective to reveal an agent's "virtual cost" which, in our problem, would reduce to $\theta + \frac{F(\theta) - \Phi(\theta)}{f(\theta)}$, where $\Phi(\theta)$ is a sum of multipliers on lower types' participation constraints. He shows that if this expression is nonincreasing on an interval, pooling may be optimal for the principal. Since in our model we are concerned with maximizing the agents' utilities, our objective is to maximize the information rent, $\frac{F(\theta) - \Phi(\theta)}{f(\theta)}$. Thus, we would wish to separate types when this expression is nonincreasing, and to pool when it is nondecreasing.

possibly triggering an industry-wide price war. When the symmetry restriction is relaxed, however, a new ingredient can be included in the firm's calculation: a firm may refrain from selecting a low price in the current period, if it anticipates that it will then be rewarded with favorable treatment in the future. For example, if the colluding firms keep records of past sales, then they may adopt a collusive scheme under which current prices are set so that a firm that has not enjoyed significant market share in the recent past is granted higher expected market share. In this section, we show that even simple asymmetric collusive schemes can improve upon the rigid-pricing scheme, when firms are sufficiently patient.

The distinction between symmetric and asymmetric schemes may be understood in the context of the analogy to static mechanism design developed in this paper. Recall that in the Mechanism Design Program, the "transfer" $T(\theta)$ corresponds to a continuation value of the repeated game. Given this interpretation, we can understand the symmetry restriction as (i). allowing for "wasteful" transfers (i.e., industry-wide price wars), but (ii). eliminating the ability of firms to implement transfers *from* one firm *to* another. In contrast, asymmetric perfect public equilibria (APPE) allow for both wasteful and cross-firm transfers, provided that the transfers are drawn from the set of APPE continuation values. Thus, in both SPPE and APPE, transfers are restricted, but the restriction is more severe in SPPE.³⁷

We begin with the simplest asymmetric scheme. We say that a collusive scheme is a *take-turns scheme* if the firms alternate monopolizing the market. This scheme clearly involves asymmetric strategies: for any given current cost draw, the firms select different prices, as the firm whose turn it is to monopolize the market must make the sale. When demand is inelastic, the take-turns scheme offers no benefit beyond the rigid-pricing scheme, because under either scheme the transaction price is r and the identity of the "winning firm" is independent of the firms' current cost positions. This absence of efficiency benefits is maintained when demand is downward sloping; however, in this case the take-turns scheme does offer an advantage: the winning firm can tailor its price ($p^m(\theta)$) to its current cost realization.

When off-schedule incentive constraints are considered, however, the take-turns scheme has a disadvantage. Under this scheme, when it is a firm's turn to "sit out," the firm is especially tempted to cheat on the agreement. In the inelastic-demand case, the firm could select a price just below r , thereby converting the current period from a zero-market-share to a full-market-share outcome. By contrast, under the rigid-pricing scheme, such a firm has less incentive to undercut, because it already expects half of the market share in the current period. Thus, the take-turns scheme requires

³⁷ Thus, SPPE and APPE occupy a middle ground between McAfee and McMillan's (1992) "weak cartel" analysis (in which no transfers are permitted) and their "strong cartel" analysis (in which transfers are unrestricted).

greater patience on the part of collusive firms than does the rigid-pricing scheme.³⁸

The central limitation of the take-turns scheme is that it fails to provide efficiency benefits: the low-cost firm will not in general serve the market. We now show that a very simple asymmetric scheme, closely related to the take-turns scheme, *can* offer some efficiency benefits. To begin, recall that among the class of symmetric collusive schemes, one option available to firms is a two-step scheme. Under such a scheme, lower-cost firms select lower prices, with the downward pricing distortion being sufficient to deter higher-cost firms from pricing at the lower level. This pricing distortion represents an informational cost (which outweighs the efficiency benefit when demand is inelastic and F is log-concave). Here, we return to the two-step approach, but our aim now is to construct an asymmetric collusive scheme with two prices that offers efficiency benefits without requiring a significant pricing distortion.

We illustrate the key ideas with a two-firm, inelastic-demand and uniform-cost-distribution example (recall that the uniform distribution is log-concave, so the best SPPE has a rigid price). The scheme that we analyze operates over a two-period cycle. In particular, we refer to a *two-price scheme* as a price pair, (ρ^l, r) , where over every two-period cycle (periods 1 and 2, periods 3 and 4, etc.) each firm can charge the price ρ^l exactly once and the price r exactly once.³⁹ Any other pattern of pricing constitutes an off-schedule deviation. This scheme is asymmetric, since if the firms make different pricing decisions in the first period, they expect different profits in the second period.

The potential appeal of such a scheme is that it may permit some sorting with only a small pricing distortion. Intuitively, if a firm prices at ρ^l in period one, then it forfeits the right to select ρ^l in period two. This forfeiture is only attractive to a firm whose period-one cost is low, since then the firm values greatly the corresponding increase in expected period-one market share that the low price implies. Furthermore, a higher-cost firm in period one now finds that “mimicking” the low price is especially costly, because it thereby “uses up” its right to select the low price in period two, when its cost may be lower. In other words, the prospect of favorable treatment in period two acts as an additional consideration that dissuades the high-cost firm from selecting the low price in period one. As a consequence, under a two-price scheme, the pricing distortion (i.e., $r - \rho^l$) in the period-one price may be small.

³⁸ This conclusion may be illustrated concretely with the following example. Suppose that the distribution of types is uniform over $[0, 1]$, $r = 1$, and firms revert to the infinite play of the Nash equilibrium when an (off-equilibrium path) deviation is observed. As noted above, when firms use rigid pricing, $\delta^* = 6/7 \approx .857$. Under take-turns, the off-schedule incentive constraint is most demanding when it is a firm’s turn to sit out and the firm realizes the lowest possible cost. For this situation, the critical discount factor above which the off-schedule incentive compatibility constraint holds can be calculated as $.9135 > \delta^* \approx .857$.

³⁹ Note that we might choose a (ρ^l, r) pair where ρ^l is very close to r . Thus, when the two prices are close, if the pair were offered as part of a static mechanism without transfers ($T = 0$), all types would choose ρ^l . The constraint that ρ^l can only be chosen once in the cycle implies that r will indeed be chosen in our proposed asymmetric scheme.

Formally, under a two-price scheme, firms in the first period of the cycle sort between the pairs (ρ^l, r) and (r, ρ^l) , where the first (second) entry in the pair indicates the price selected in the first (second) period of the cycle. The on-schedule incentive constraint thus defines a critical θ^* , where a firm with a period-one cost $\theta < \theta^*$ opts for (ρ^l, r) while a firm with $\theta > \theta^*$ prefers to wait and charge ρ^l in the second period. When it is anticipated that the rival employs the same threshold θ^* , we have that θ^* can be determined as that θ which leaves a firm indifferent between the two price pairs. That is, the on-schedule incentive constraint defines θ^* as the θ for which

$$\begin{aligned} & F(\theta^*)\{\pi(\rho^l, \theta) + \pi(r, E\theta)\}\left(\frac{1}{2}\right) + [1 - F(\theta^*)]\pi(\rho^l, \theta) \\ & = F(\theta^*)\delta\pi(\rho^l, E\theta) + [1 - F(\theta^*)]\{\pi(r, \theta) + \delta\pi(\rho^l, E\theta)\}\left(\frac{1}{2}\right), \end{aligned} \quad (10)$$

where the top expression yields expected two-period profit from (ρ^l, r) , given that the rival selects the same pair with probability $F(\theta^*)$, and the bottom expression similarly captures the expected two-period profit from (r, ρ^l) . Notice that if a firm selects ρ^l in period one and its rival selects r , then the firm effectively “sits out” in period two. As this is true regardless of the firms’ period-two cost realizations, we observe that the two-price scheme fails to deliver sorting in period two.

To see the benefits of the two-price scheme, suppose that $r = 1$ and consider a scheme with a minimal pricing distortion: $p^l \approx r$. In this case, we find that $\theta^* \approx (2 - \delta)/2 \in (1/2, 1]$. Thus, the two-price scheme permits some first-period sorting when $\delta > 0$, since it is then possible that exactly one firm has costs below θ^* , and the two-price scheme ensures that this firm indeed makes the period-one sale. Notice further that this efficiency benefit is increasing in δ ; in fact, as δ approaches 1, we have that θ^* approaches $1/2$, and so first-period sorting is maximal within the two-price class. This points to a useful coincidence: the off-schedule incentive constraint is also sure to be met when δ is close to 1. Thus, when firms are sufficiently patient, they can implement a two-price scheme that generates efficiency benefits, imposes only a negligible pricing distortion and therefore beats the rigid-pricing scheme.

More generally, a two-price scheme with $p^l \approx r$ is sure to improve upon rigid pricing regardless of the value that δ takes, provided that (i). $\theta^* < 1$ (i.e., some first-period sorting occurs)⁴⁰ and (ii). the two-price scheme satisfies the off-schedule incentive constraint. It is important to emphasize, however, that the two-price scheme (like the take-turns scheme) meets its off-schedule constraint only if firms are quite patient. In particular, consider the situation in which in the second period of a cycle a firm is supposed to sit out (because only it selected the low period-one price) and the firm draws the lowest cost type possible, $\underline{\theta}$. The firm then has a high incentive to cheat, as by slightly undercutting the rival’s current price r it converts the current period from a zero-market-share to a

⁴⁰ It is direct to confirm that $\theta^* < 1$ when $\delta > (r - 1)/(r - \frac{1}{2})$. This condition is met, e.g., for any $\delta > 0$, if $r = 1$.

full-market-share outcome. Thus, while the proposed two-price scheme beats rigid pricing, it also requires greater patience on the part of firms (but still less than take turns), if the scheme is to be robust against off-schedule deviations.⁴¹

What lessons can be extracted from this example? In comparing the rigid-pricing and two-price schemes, we note that both schemes emphasize rigid pricing (at or near r), but a novel feature of the two-price scheme is that a firm's market share is negatively correlated over time. Of course, the latter prediction has potential empirical significance only if it is paired with some guidance regarding the likely structural characteristics of industries that are best able to employ asymmetric schemes. We do not claim to resolve fully this issue. We do note, however, that asymmetric schemes require greater coordination on the part of collusive firms. It is thus possible that such schemes have greatest relevance for collusive ventures formed between a small number of firms that track firm-level performance and communicate explicitly.⁴²

7. Conclusion

We propose a model of collusion in which firms are privately informed as to their current cost positions. For the inelastic-demand case, our four main findings are:

1. Firms fare poorly under fully-sorting symmetric collusive schemes, since the efficiency benefits that such schemes afford are small in comparison to the informational costs that they demand.
2. Optimal symmetric collusion can be achieved without recourse to price wars on the equilibrium path.
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, a downward distortion in the price of lower-cost firms may be required, in order to attenuate the incentive that such firms have to cheat.

We note that the first finding underscores the basic tradeoff present in our repeated-adverse-selection model of collusion; the second finding contrasts markedly with the Green-Porter (1984) collusion literature on repeated moral hazard and imperfect observability; the third finding offers an equilibrium interpretation of the empirical association between rigid pricing and industry concentration;

⁴¹ For example, in the uniform-distribution example with $r = 1$, $p^l \approx r$ and $\pi^w = 1/6$ (Nash reversion), we show in the Appendix that the off-schedule incentive constraint requires $\delta \geq \delta^{**} \approx .88 > \delta^* \approx .857$.

⁴² Some support for this conclusion is offered by Comanor and Shankerman (1976), who examine collusive-bidding court cases and find that even simple rotation schemes (corresponding to our take-turns scheme) typically involve a small number of conspirators. More sophisticated rotation schemes, which like the two-price scheme convey efficiency benefits, demand an even greater degree of coordination.

and the fourth finding is reminiscent of the logic developed by Rotemberg and Saloner (1986) for collusion in markets with publicly-observed demand shocks.

We further extend the analysis to the case of downward-sloping demand. We find that the qualitative themes from the analysis of inelastic demand are retained, and when demand is sufficiently inelastic, the firms use rigid pricing and avoid price wars. However, we also identify several novel features for the case of downward-sloping demand. In particular, we identify (somewhat restrictive) conditions under which firms engage in seemingly suboptimal behavior, such as price wars or non-rigid pricing, in order to alleviate the incentives for low-cost types to mimic intermediate cost types. Typically in principal-agent models, incentive constraints bind in only one direction; some ambiguity arises in our model because sidepayments are not available, and “transfers” must correspond to symmetric equilibrium continuation values. We also extend the analysis to consider asymmetric collusive schemes. We show by example that asymmetric schemes can offer additional efficiency benefits, provided that firms are sufficiently sophisticated and patient.

We mention now some broader implications of our analysis. One such implication concerns the proper interpretation of an apparent episode of “cheating” within a collusive industry. Consider the following pricing pattern. Over an initial passage of time, all firms maintain a common fixed price. Then, over a shorter passage of time, one firm significantly lowers its price. The other firms, however, do not respond in any way, and before long the price reduction is eliminated. How should such a pricing pattern be interpreted? A first interpretation might be that one firm briefly cheated on a price-fixing agreement. But this interpretation is problematic. If the firm really did cheat, then why did the other firms not “punish” it? And, given that they did not, why doesn’t the firm continue cheating?⁴³ Our analysis points to an alternative interpretation. If the firms are too impatient to always enforce a rigid price, but still sufficiently patient that some collusion can be achieved, then the collusive scheme may include an “escape clause.” From this perspective, the price reduction is not “cheating” but rather an acceptable response by a single firm to a temporary but quite favorable cost position. The price reduction must be significant, in order to ensure that it would be attractive only to a firm that truly has low current costs. Recognizing that such episodes occur periodically, the other firms continue on with the fixed price, until such time as a (perhaps different) firm again receives an unusually favorable cost shock.⁴⁴

⁴³ A similar dilemma is raised in a recent paper by Genesove and Mullin (1998), who argue that apparent cheating did not lead to retaliation in the collusive sugar industry of the 1930’s (although in their case, rival firms tended to match the initial price reduction). They interpret this pattern as revealing a limitation for the repeated-game approach to collusion.

⁴⁴ Following Rotemberg and Saloner (1986), a third interpretation might be that the industry experienced a public demand shock, requiring some adjustment in the collusive price. Their model, however, predicts that *all* firms adjust prices in response to the shock. A modification of their model to include publicly-observed and firm-specific demand shocks might deliver behavior that is more consistent with the posited pattern.

A second implication is that the nature of collusive conduct may vary interestingly with market-demand (and thus business-cycle) conditions. In particular, it is interesting to consider an extended model, in which firms are privately informed as to their respective current cost positions and publicly informed as to the current state of market demand. Following Hamilton (1987), we may imagine, for example, that market demand alternates stochastically between fast-growth (boom) and slow-growth (recession) phases, with firms aware of the current phase at the time that price selections are made. As Bagwell and Staiger (1997) establish (in a model without private cost information), off-schedule incentive constraints are less restrictive in boom phases, provided that demand growth rates are positively correlated through time. For the extended model, the apparent implication is that symmetric collusion among firms is more easily achieved in booms; consequently, if demand is sufficiently inelastic, collusive pricing may exhibit greater rigidity to private cost fluctuations in booms than in recessions.

A third and related implication of our analysis is that collusive prices are more sensitive to publicly- than privately-observed changes in the market environment. Consider, for example, the response of collusive prices to changes in costs, for a market in which demand is downward sloping. As we have shown, when cost fluctuations are privately observed (as when firm-specific cost components fluctuate), the optimal symmetric collusive scheme calls for some rigidity in the collusive price. By contrast, when the cost shock is public (as when a wholesale price changes), firms can re-tailor their collusive scheme to the new environment; as a consequence, the entire collusive pricing schedule may be adjusted in response to the new situation.

In closing, we mention two methodological contributions of our analysis. First, 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 modeling framework. Second, our work motivates some new questions for static mechanism design, and takes some initial steps towards addressing them. For example, we examine how restrictions on transfers affect optimal mechanisms. Further, in our study of collusion among impatient firms, we observe that the participation constraints vary according to both the firm's cost type as well as the collusive pricing agreement itself. Thus, the participation constraints are endogenous.

8. Appendix

Proof of Lemma 6: We prove the result for step functions; strictly-increasing functions are analogous. Let p_k and M_k denote the price and market share for step k , and let the step cover $(\theta_k, \theta_{k+1}]$. First, notice that for each θ , π is strictly increasing in ρ for $\rho < p^m(\theta)$. Thus, we can define $\phi(\pi; \theta) \equiv \pi^{-1}(x; \theta)$. Now observe that $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) \leq \pi(p^m(\bar{\theta}), \bar{\theta})M(\bar{\theta}; p)$. Then, recall our restriction that the lowest type on the highest step makes non-negative profit: $H(\theta_K, \theta_K; p) - T(\theta_K) \geq 0$. This implies $p(\theta_K) \equiv p(\bar{\theta}) \geq \theta_K$, and thus $D(p(\bar{\theta})) \leq D(\theta_K)$, from which it follows that $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) \equiv \pi(p(\bar{\theta}), \bar{\theta})M(\bar{\theta}; p) - T(\bar{\theta})$

$$\begin{aligned} &= [(\theta_K - \bar{\theta})D(p(\bar{\theta})) + (p(\bar{\theta}) - \theta_K)D(p(\bar{\theta}))]M(\bar{\theta}; p) - T(\bar{\theta}) \\ &\geq [(\theta_K - \bar{\theta})D(\theta_K) + (p(\bar{\theta}) - \theta_K)D(p(\bar{\theta}))]M(\bar{\theta}; p) - T(\bar{\theta}) \\ &= \pi(\theta_K, \bar{\theta})M(\bar{\theta}; p) + H(\theta_K, \theta_K; p) - T(\theta_K) \end{aligned}$$

$\geq \pi(\theta_K, \bar{\theta})M(\bar{\theta}; p)$. Thus, we may find $\tilde{p}(\bar{\theta}) \in [\theta_K, p^m(\bar{\theta})]$ such that $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = \pi(\tilde{p}(\bar{\theta}), \bar{\theta})M(\bar{\theta}; p)$.

From there, no-war prices are determined according to (IC-onM), which specifies that $\pi(\rho_k, \theta_k)M_k = \pi(\rho_{k-1}, \theta_k)M_{k-1}$. Each price is then determined as $\rho_{k-1} = \phi(\pi(\rho_k, \theta_k) \cdot \frac{M_k}{M_{k-1}}; \theta_k)$. Since $\frac{M_k}{M_{k-1}} < 1$ and $\pi(\rho_k, \theta_k) \leq \pi(p^m(\theta_k), \theta_k)$, and since $\pi(\tilde{p}(\bar{\theta}), \theta_K) \geq 0$, this algorithm generates prices $\tilde{p}(\theta)$ that lie between cost and the monopoly level at each step.

Finally, notice that if the pricing function is strictly increasing at the top, the no-war schedule will have to give $\tilde{p}(\bar{\theta}) = \bar{\theta}$, whence $H(\bar{\theta}, \bar{\theta}; \tilde{p}) = 0$. To show that the lemma holds for this case, we confirm that $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) = 0$. This follows since when p is strictly increasing at the top, we have that $M(\bar{\theta}; p) = 0$ and thus $H(\bar{\theta}, \bar{\theta}; p) = 0$. Our assumption that $H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}) \geq 0$ along with the requirement that $T(\bar{\theta}) \geq 0$ then confirms the desired equality. ■

Lemma 20 *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 corresponding no-wars mechanism from Lemma 6. Let $Y(\theta) = H(\theta, \theta; \tilde{p}) - [H(\theta, \theta; p) - T(\theta)]$ be the pointwise net advantage of the constructed no-wars schedule \tilde{p} . Then 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 .*

Proof. Y is differentiable almost everywhere by Lemma 5. By the envelope theorem and the fact that p and \tilde{p} are market share equivalent (i.e., $M(\cdot; \tilde{p}) = M(\cdot; p)$), where it exists,

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

The result follows since demand is downward-sloping. ■

Lemma 21 *Suppose demand is downward-sloping and consider any distribution function F . For any θ' where p is strictly increasing or jumps up, if (p, T) satisfies the on-schedule incentive constraint with no wars in the neighborhood of θ' , then $p(\theta') < p^m(\theta')$.*

Proof. If $p(\theta' + \varepsilon) > p(\theta') \geq p^m(\theta')$, then $M(\theta'; p) > M(\theta' + \varepsilon; p)$ and $\pi(p(\theta'), \theta') \geq \pi(p(\theta' + \varepsilon), \theta')$, by the quasi-concavity of π . But such a pricing function cannot be incentive compatible without a war at θ' , since the lower price offers higher market share and higher profit-if-win. ■

Lemma 22 *Suppose demand is downward-sloping and consider any distribution function F . Let (p, T) be incentive-compatible, and suppose that the constructed no-wars mechanism $(\tilde{p}, \tilde{T} \equiv 0)$ that yields the same market share (from Lemma 6) is sub-monopoly for the lowest type: $\tilde{p}(\underline{\theta}) \leq p^m(\underline{\theta})$. (i) Then the no-wars mechanism is pointwise as good as its with-wars counterpart: $H(\theta, \theta; \tilde{p}) \geq [H(\theta, \theta; p) - T(\theta)]$, for all θ . (ii) Further, the no-wars mechanism is strictly better if $T > 0$ on a set of positive measure.*

Proof. (i) 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 Lemma 20 gives $p(\theta^*) = \tilde{p}(\theta^*)$ and so $H(\theta^*, \theta^*; \tilde{p}) = H(\theta^*, \theta^*; p)$. 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 $H(\underline{\theta}, \underline{\theta}; \tilde{p}) \geq H(\underline{\theta}, \underline{\theta}; p)$, and so $Y(\underline{\theta}) \geq 0$.

(ii) Now suppose 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) $H(\theta, \theta; p) - T(\theta) = H(\theta, \theta; \tilde{p})$ almost everywhere. But this implies $\tilde{p}(\theta) \neq p(\theta)$ almost everywhere that $T(\theta) > 0$. From this and Lemma 20 it follows that (b) $Y'(\theta) \neq 0$ on a set of positive measure. Since by Lemma 5, $Y(\theta) = Y(\bar{\theta}) + \int_{\bar{\theta}}^{\theta} Y'(\bar{\theta}) d\bar{\theta}$, (b) is inconsistent with $Y(\theta) = 0$ almost everywhere. ■

Proof of Proposition 11: By Proposition 1 (3), $p^e(\theta) \leq p^m(\theta)$. Then Lemma 22 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 SPPE follows as in Proposition 4. ■

Proof of Proposition 12: Consider first a step at the bottom on $[\underline{\theta}, x]$. Suppose $p(x) > p^m(x)$. But then, we can find a below-monopoly price which gives the same profits to x , thus preserving (IC-onM), but strictly improving the profit for all types below x by the single-crossing property. Then, consider setting $T(\underline{\theta})=0$. Since x is at or below-monopoly, lowering the price and eliminating the war allows us to maintain (IC-onM). Again, all types below x prefer the exchange. Notice that if $p(x) = p^m(x)$, by Lemma 21, $T(\underline{\theta}) > 0$; but then, our previous arguments would imply that profits are improved by lowering the price and eliminating the war. If p is strictly increasing at the bottom, since by Lemma 21 the no-war prices are below-monopoly, similar arguments imply that there are no wars for θ sufficiently close to $\underline{\theta}$, and thus none at $\underline{\theta}$. Now consider a step at the top. Suppose $p(y) < p^m(y)$ and demand smooth in the relevant range. But then, we can find an above-monopoly price which gives the same profits to y , thus preserving (IC-onM), but strictly improving the profit for all types above y by the single-crossing property. Then, consider setting $T(\bar{\theta}) = 0$. Since y is above-monopoly, eliminating the war requires raising the price to maintain (IC-onM). Again, all types above y prefer the exchange. ■

Proof of Proposition 13: Suppose that (p, T) , where $T > 0$ on a set of positive measure, is a solution. Let $(\tilde{p}, \tilde{T} \equiv 0)$ be the no-war, market-share equivalent scheme, and define Y as in Lemma 20. By Lemma 22, if (p, T) is a positive-war solution, it follows that $\tilde{p}(\underline{\theta}) > p^m(\underline{\theta})$. Further, observe that if $p(\theta)$ is constant everywhere, we can improve profits by eliminating wars without violating incentive compatibility. Thus, $p(\theta)$ must have some points of strict increase.

Let $x \geq \underline{\theta}$ be the first point of strict increase of \tilde{p} . Now, $\tilde{p}(\underline{\theta}) > p^m(\underline{\theta})$ implies that $x > \underline{\theta}$ so that \tilde{p} is flat on $[\underline{\theta}, x)$, and further $\tilde{p}(x) < p^m(x)$. Otherwise, we would have a point of strict increase above monopoly prices, contradicting Lemma 21. Since p is market share equivalent to \tilde{p} , p must be flat on $[\underline{\theta}, x)$ as well.

Now suppose that $\tilde{p}(\underline{\theta}) \leq p(\underline{\theta})$. Since $\tilde{p}(\underline{\theta}) > p^m(\underline{\theta})$, this implies that $\pi(\tilde{p}(\underline{\theta}), \underline{\theta}) \geq \pi(p(\underline{\theta}), \underline{\theta})$ and thus $Y(\underline{\theta}) \geq 0$. But if Y is minimized on $(\underline{\theta}, \bar{\theta}]$, arguments in the proof of Lemma 22 establish that $\min_{\theta} Y(\theta) \geq 0$. Thus, in order for (p, T) to be optimal, we require $Y(\theta) = 0$ almost everywhere. But (as in the proof of Lemma 22), this is ruled out by the fact that $T > 0$ on a set of positive measure.

So, it must follow that $\tilde{p}(\underline{\theta}) > p(\underline{\theta})$. Then, since $\tilde{p}(x) < p^m(x)$, we have $\pi(\tilde{p}(x), x) > \pi(p(x), x)$ and so $Y(x) > 0$.

Observe from Lemma 20 that $\tilde{p}(x) > p(x)$ implies $Y'(x) > 0$. Since $Y(x) > 0$ and $Y(\bar{\theta}) = 0$, either (a) $Y(\theta) \geq 0$ for all $\theta \geq x$, or else (b) there exists some $\bar{\theta} \geq \hat{\theta} \geq x$ s.t. $Y(\hat{\theta}) < 0$ and $Y'(\hat{\theta}) = 0$. But (b) contradicts arguments in the proof of part (i) of Lemma 22. Given $Y(x) > 0$ and $Y(\theta) \geq 0$ for all $\bar{\theta} \geq \theta \geq x$, we conclude that $E[Y | \bar{\theta} \geq \theta \geq x] > 0$. This combined with $EY \leq 0$ implies $E[Y | x \geq \theta \geq \underline{\theta}] < 0$. ■

Lemma 23 *Suppose that demand is downward sloping. Fix a vector θ representing the market share allocation. Then if (p, T) is a solution to the Mechanism Design Program, there must exist a nondecreasing function $\Lambda(\theta)$ such that $\Lambda(\underline{\theta}) = 0$ and $\Lambda(\bar{\theta}) = 1$, and further:*

(i). *On regions with sorting, either (a) $\lambda(\theta) = 0$, $T(\theta) > 0$ and $p(\theta)$ is approximately flat, or else (b) $T(\theta) = 0$ and $\Lambda(\theta)$ has density*

$$\lambda(\theta) = -D'(p(\theta)) [F(\theta) - \Lambda(\theta)] / (\pi_\rho(p(\theta), \theta)).$$

(ii). *On regions of pooling, either (a) $\Lambda(\theta)$ is constant, $T_k > 0$ and $p_k \approx p^+(\theta_{k+1})$, (b) $\Lambda(\theta)$ is constant, $T_k > 0$ and $\int_{\theta_k}^{\theta_{k+1}} F(\theta)d\theta = \Lambda^+(\theta_k)$, or else (c) $T_k = 0$ and*

$$\lambda_k \equiv \Lambda^+(\theta_{k+1}) - \Lambda^-(\theta_{k+1}) = -D'(p_k) \left[\int_{\theta_k}^{\theta_{k+1}} F(\theta)d\theta - \Lambda^+(\theta_k) \right] / (\pi_\rho(p_k, \theta_{k+1})).$$

Proof of Lemma 23: We can incorporate the no-subsidy constraints, as well as the incentive constraints, into the firm's problem. On regions of full sorting, each no-subsidy constraint may be binding, and we introduce a multiplier, $\lambda(\theta)$, for each such constraint. On regions of pooling, only the no-subsidy constraint corresponding to the right endpoint of the step will be binding, and so we introduce a multiplier λ_k for that constraint. Rearranging, integrating by parts, and defining $\Lambda(\theta)$ to be the distribution with density $\lambda(\theta)$ on regions of full sorting and jumps of λ_k at θ_{k+1} for k even, we have the following objective, which can be maximized pointwise in p and in $T(\bar{\theta})$:

$$E \left[\begin{aligned} & (H(\bar{\theta}, \bar{\theta}; p) - T(\bar{\theta}))(1 - \Lambda(\bar{\theta})) \\ & + \sum_{k=1}^{K/2} \mathbf{1}_{(\theta_{2k-1}, \theta_{2k}] }(\theta) \left[\frac{(1-F(\theta))^{n-1}}{f(\theta)} D(p(\theta)) \cdot \left[\begin{aligned} & F(\theta) - \Lambda(\theta) \\ & + \lambda(\theta)(p(\theta) - \theta) \end{aligned} \right] \right] \\ & + \sum_{k=1}^{K/2} \mathbf{1}_{(\theta_{2k}, \theta_{2k+1}] }(\theta) \left[\frac{\mu(\theta_{2k}, \theta_{2k+1})}{F(\theta_{2k+1}) - F(\theta_{2k})} D(p_k) \cdot \left[\begin{aligned} & \int_{\theta_{2k}}^{\theta_{2k+1}} F(\theta)d\theta - \Lambda^+(\theta_{2k}) \\ & + \lambda_{2k}(p_{2k} - \theta_{2k+1}) \end{aligned} \right] \right] \end{aligned} \right]$$

We begin by relaxing the constraint that p is nondecreasing; we will return to verify that the constraint is satisfied. Then, necessary conditions for optimality are given as follows. Differentiating with respect to $T(\bar{\theta})$, we must have $1 = \Lambda(\bar{\theta})$. By definition, $\Lambda(\underline{\theta}) = 0$, since if the lowest interval is fully sorting, the distribution will by definition be zero at the bottom, while if the lowest interval is pooling, the first no-subsidy constraint appears at the right endpoint of the interval. Now considering optimality conditions for p , notice that in regions of full sorting, differentiating with respect to $p(\theta)$ yields $D'(p(\theta))[F(\theta) - \Lambda(\theta) + \lambda(\theta)(p(\theta) - \theta)] + D(p(\theta))\lambda(\theta) = 0$, while on regions of pooling, the analogous condition holds.

Now, anywhere that $T > 0$, the corresponding no-subsidy constraint will be slack. From the preceding equations, that holds only if $F(\theta) = \Lambda(\theta)$ (on regions of sorting) and $\int_{\theta_k}^{\theta_{k+1}} F(\theta)d\theta = \Lambda^+(\theta_k)$ on regions of pooling. In contrast, if $T = 0$, the no-subsidy constraint binds and prices are determined by incentive compatibility. But, notice that we have neglected

the monotonicity constraint until now. Consider regions of full sorting. If $T(\theta) = 0$ and $F(\theta) < \Lambda(\theta)$, in order to make $\lambda(\theta) \geq 0$, it must be that $\pi_\rho(p(\theta), \theta) < 0$. But this implies that $p(\theta)$ is greater than the monopoly price. That violates the monotonicity constraint. However, the monotonicity constraint cannot bind while (IC-onM) is satisfied unless wars are used. Thus, we either have $F(\theta) < \Lambda(\theta), T(\theta) > 0$, and the monotonicity constraint binds, or else $F(\theta) \geq \Lambda(\theta)$, and the previous analysis holds. An analogous argument holds for regions of pooling.

Next, focus on regions of sorting. Suppose that $T(\theta) = 0$ on an open interval, (a, b) , and that the monotonicity constraint does not bind. Our previous arguments imply that $F(\theta) = \Lambda(\theta)$ on that same interval. But, notice that if that is true, our expression for $\lambda(\theta)$ implies that $\Lambda(b) = \Lambda(a)$, clearly contradicting $F(b) = \Lambda(b)$ since $F(b) > F(a)$ by our positive density assumption. ■

Proof of Proposition 14: Lemma 23 formally analyzes the program. (i). Notice that $F(\underline{\theta}) = \Lambda(\underline{\theta}) = \lambda(\underline{\theta}) = 0$, while by assumption, $f(\underline{\theta}) > 0$. Thus, $(F(\theta) - \Lambda(\theta))/f(\theta)$ has first derivative equal to 1 at $\theta = \underline{\theta}$. (ii) First part follows from Lemma 23. If the monotonicity constraint does bind, then $\frac{F(\theta) - \Lambda(\theta)}{f(\theta)} D(\rho)$ will be increasing in θ if F is log-concave and $-\Lambda(\theta)/f(\theta)$ is nondecreasing. Since $\Lambda(\theta)$ is constant on an interval with wars, a sufficient condition is that the density is nondecreasing. But if $\frac{F(\theta) - \Lambda(\theta)}{f(\theta)} D(\rho)$ will be increasing in θ , it will not be optimal to sort. (iii) Consider the incentive of a firm to use two adjacent regions of pooling, on $(\theta_k, \theta_{k+1}]$ and $(\theta_{k+2}, \theta_{k+3}]$, where $\theta_{k+1} = \theta_{k+2}$ (there is no region of sorting in between). In this case, a necessary condition for $\theta_k < \theta_{k+1} = \theta_{k+2} < \theta_{k+3}$ to be optimal is that θ_{k+1} satisfying these inequalities maximizes the expression

$$D(p_k) \cdot \mu(\theta_k, \theta_{k+1}) \cdot \int_{\theta_k}^{\theta_{k+1}} F(\theta) d\theta + D(p_{k+2}) \cdot \mu(\theta_{k+1}, \theta_{k+3}) \int_{\theta_{k+1}}^{\theta_{k+3}} F(\theta) d\theta \\ + D(p_k) \cdot \mu(\theta_k, \theta_{k+1}) \cdot [-\Lambda^+(\theta_k) + \lambda_k(p_k - \theta_{k+1})] \\ + D(p_{k+2}) \cdot \mu(\theta_{k+1}, \theta_{k+3}) [-\Lambda^+(\theta_{k+1}) + \lambda_{k+2}(p_{k+2} - \theta_{k+3})].$$

If $D(p_k)$ is close enough to $D(p_{k+2})$, the first line of this expression is proportional to the expression which arises in the case of inelastic demand. We have already established, however, that when F is log-concave, steps are suboptimal, and the best choice is either $\theta_{k+1} = \theta_k$ or $\theta_{k+1} = \theta_{k+3}$. The second part of the expression will be zero if the no-subsidy constraints do not bind. If demand is sufficiently inelastic, the no-subsidy constraints will not affect revenue very much. Thus, for any distribution F , we can always reduce the slope of D enough so that steps are not optimal.

Proof of Proposition 15: Part (i) is proved in the text. For part (ii), suppose that $\delta < \delta^*$ and consider first any stationary rigid-pricing scheme, with $p(\theta) \equiv \rho \in [\underline{\theta}, r]$ all along the equilibrium path. Consider the off-schedule incentive constraint (8). Observe that the LHS increases in ρ at rate $\frac{n-1}{n}$, while the RHS increases at rate $(\delta/(1-\delta))\frac{1}{n}$. Thus, when $\delta > \frac{n-1}{n}$, the RHS increases faster, and so (8) holds for some ρ only if it holds at the highest ρ , which is $\rho = r$. But this is impossible, given $\delta < \delta^*$. If instead $\delta < \frac{n-1}{n}$, then (8) holds for some ρ only if it holds for the lowest ρ , which is $\rho = \underline{\theta}$. But, at that value, the LHS of (8) is zero, whereas the RHS is negative, and so (8) fails. Therefore, when $\delta < \delta^*$, every stationary rigid-pricing scheme violates (8). Consider second any non-stationary rigid-pricing scheme. Let p^h be the highest rigid price that is specified along the equilibrium path, and consider the date at which that price is specified. The off-schedule constraint at that point would be easier to satisfy, if in all future periods firms were to continue with the rigid price p^h , since then the discounted value of cooperation would be higher. But when $\delta < \delta^*$, firms are unable to enforce even this arrangement. It follows that non-stationary

rigid-pricing schemes also violate the off-schedule incentive constraint, when $\delta < \delta^*$.

Proof of Proposition 17: Define $T(\theta)$ as in the text, and notice that if (p, v) satisfies IC-on2 in the Interim Program, then (p, T) satisfies the constraints of the Mechanism Design Program (since shifting T by a constant does not affect on-schedule incentive compatibility constraints). Recall from Lemma 6 (and exploiting the inelastic-demand assumption) that for any (p, T) , we can find an alternative scheme, $(\tilde{p}, \tilde{T} \equiv 0)$ such that $H(\theta, \theta; p) - T(\theta) = \pi(\tilde{p}(\theta), \theta)M(\theta; p)$ 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 $\tilde{p}(\theta) < p(\theta)$. Adding the constant $z \equiv \sup_{\theta} E_{-i}v(p(\theta), p)$ to both T and \tilde{T} does not affect these conclusions.

Now, we compare the off-schedule incentive constraints across the two schedules. To do so, let $w \in V_s$ 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.⁴⁵ Consider then the incentive that a firm on price step k has to undertake an off-schedule deviation in which it undercuts the step- k price. As the original schedule satisfies the off-schedule incentive constraints, it follows that the lowest-cost type on any step k , which we denote as θ_k , will not choose under the original schedule to slightly undercut the step- k price, $p(\theta_k)$:

$$\pi(p(\theta_k), \theta_k)[(1 - F(\theta_k))^{n-1} - M(\theta_k; p)] \leq \delta\{z - T(\theta) - w\} \quad (11)$$

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

$$\pi(\tilde{p}(\theta_k), \theta_k)[(1 - F(\theta_k))^{n-1} - M(\theta_k; p)] \leq \delta\{z - w\}. \quad (12)$$

In (11) and (12), the LHS'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 RHS's represent the expected discounted values of cooperation in the next and all future periods. The RHS 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 cooperation 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. In a similar way, for any other period in which a positive probability of war exists, we may factor payoffs into current-period and continuation values, and then repeat the same argument. ■

Proof of Proposition 18: Consider a candidate two-price scheme, with a top-step price $\rho_2 \in (\underline{\theta}, \bar{r}]$ 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, $p_1(\theta_2)$, as the ρ_1 that solves

$$\pi(\rho_1, \theta_2)M_1 = \pi(\rho_2, \theta_2)M_2, \quad (13)$$

⁴⁵ 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 $\underline{\theta}$ 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). To see this, note that the relevant case is the highest type on a price step. If this type differs from $\bar{\theta}$, then the best price hike is all the way up to a level just below the next highest on-schedule price. Using the on-schedule incentive constraint, and the fact (which we show in the proof) that the next type up does not desire to charge the same deviant price (as a price-cutting off-schedule deviation), it is straightforward to argue that such a price hike does not violate the off-schedule incentive constraint. Finally, for type $\bar{\theta}$, a price hike implies zero current-period market share. Under either scheme, a price hike for this type sacrifices today's profit and tomorrow's continuation value and replaces it with zero profit today and δw tomorrow. This calculation is invariant across the two schemes, ensuring that if such a deviation were unattractive under the original scheme it remains unattractive under the alternative scheme.

where $M_1 = (2 - F_2)/2$ and $M_2 = (1 - F_2)/2$, with $F_2 \equiv F(\theta_2)$. The expected per-period profit to a firm under the two-price scheme is equal to $\int_{\underline{\theta}}^{\theta_2} \pi(\rho_1, \theta) M_1 dF + \int_{\theta_2}^{\bar{\theta}} \pi(\rho_2, \theta) M_2 dF$. Integrating by parts, using $\pi_{\theta} = -1$ and

imposing (13), expected per-period profit is $M_2 \pi(\rho_2, \bar{\theta}) + \sum_{k=1}^2 M_k \int_{\theta_k}^{\theta_{k+1}} F d\theta$

In the event of an off-schedule deviation, a war follows, with π^w satisfying (7). The per-period value of cooperation under the two-step scheme is given by:

$$\omega_T(\theta_2) \equiv M_2 \pi(\rho_2, \bar{\theta}) + \sum_{k=1}^2 M_k \int_{\theta_k}^{\theta_{k+1}} F d\theta - \pi^w,$$

where $\theta_1 = \underline{\theta} < \bar{\theta} = \theta_3$. The expected discounted value of cooperation is thus $[\delta/(1 - \delta)]\omega_T(\theta_2)$. The following properties are direct to confirm:

$$\omega_T(\underline{\theta}) = \left(\frac{1}{2}\right)\pi(\rho_2, E\theta) - \pi^w \quad (14)$$

$$\omega_T(\bar{\theta}) = \left(\frac{1}{2}\right)\pi(\bar{\theta}, E\theta) - \pi^w \quad (15)$$

$$\text{if } F \text{ is log concave, then } \partial\omega_T(\theta_2)/\partial\theta_2 = 0 \text{ implies } \partial^2\omega(\theta_2)/\partial\theta_2^2 > 0. \quad (16)$$

Figure A1 depicts for the case in which $\rho_2 = \tau$. Then (7) and (14) imply $\omega_T(\underline{\theta}) > 0$. Notice from (16) that $\omega_T(\theta_2)$ is maximized at $\theta_2 = \underline{\theta}$, and the function is either declining or takes a minimum value and then turns partially back up. The two-step scheme duplicates the rigid-price scheme when $\theta_2 = \underline{\theta}$, as then all types are on the high step.

Consider next the incentive-to-cheat functions. The incentive to undercut the top-step price, ρ_2 , is most acute for the lowest type on this step, θ_2 . We therefore define $\Omega_2(\theta_2) \equiv \pi(\rho_2, \theta_2)[1 - F_2]/2$. This function is characterized by: (i) $\Omega_2(\underline{\theta}) = (\frac{1}{2})\pi(\rho_2, \underline{\theta}) > 0$; (ii) $\Omega_2(\bar{\theta}) = 0$; (iii) $\partial\Omega_2(\theta_2)/\partial\theta_2 < 0$, if $\pi(\rho_2, \theta_2) \geq 0$. As Figure A1 illustrates, if $\pi(\rho_2, \theta_2) \geq 0$, then θ_2 's incentive to undercut ρ_2 diminishes as θ_2 rises (since it becomes less likely that the rival is pricing on the top step).

The incentive to undercut the low-step price is greatest for the lowest type: $\theta_2 = \underline{\theta}$. With $\rho_1 = p_1(\theta_2)$ determined by (13), we define the corresponding incentive-to-cheat function as $\Omega_1(\theta_2) \equiv \pi(p_1(\theta_2), \underline{\theta})[F_2/2]$. $\Omega_1(\theta_2)$ may fail to be everywhere increasing, but it is characterized by: (i) $\Omega_1(\underline{\theta}) = 0$; (ii) $\Omega_1(\bar{\theta}) = (\frac{1}{2})\pi(\bar{\theta}, \underline{\theta}) > 0$; (iii) at any common value, $\partial\Omega_1(\theta_2)/\partial\theta_2 > \partial\Omega_2(\theta_2)/\partial\theta_2$. Thus, the incentive-to-cheat functions take the form depicted in Figure A1.

Summarizing, a two-price schedule satisfies the off-schedule incentive constraints if there exists a value $\theta_2 \in (\underline{\theta}, \bar{\theta})$ such that:⁴⁶

$$\Omega_1(\theta_2) \leq [\delta/(1 - \delta)]\omega_T(\theta_2) \quad (17)$$

$$\Omega_2(\theta_2) \leq [\delta/(1 - \delta)]\omega_T(\theta_2). \quad (18)$$

⁴⁶ We have not stated the off-schedule incentive constraints that ensure that a firm will not cheat by raising price (in order to avoid a loss). In the analysis that follows, we set $\rho_2 = \tau$, and the constraints are clearly met.

To prove the proposition, we set $\rho_2 = r$ and $\delta = \delta^*$, and we then show that the off-schedule incentive constraints are satisfied with equality when $\theta_2 = \underline{\theta}$ (the rigid-price case) and with slack when θ_2 is slightly higher (under the small-density condition). This slack ensures that a two-step pricing schedule satisfies the off-schedule incentive constraints even when δ lies just below δ^* , for some $\theta_2 > \underline{\theta}$. Formally, observe that

$$[\delta^*/(1 - \delta^*)]\omega_T(\underline{\theta}) = \left(\frac{1}{2}\right)\pi(r, \underline{\theta}) = \Omega_2(\underline{\theta}) > \Omega_1(\underline{\theta}) = 0.$$

Next, it is straightforward to confirm that the slope of $[\delta^*/(1 - \delta^*)]\omega_T(\theta_2)$ exceeds that of $\Omega_2(\theta_2)$, when both are evaluated at $\theta_2 = \underline{\theta}$ if the small-density condition is satisfied. Clearly, $[\delta^*/(1 - \delta^*)]\omega_T(\theta_2)$ continues to exceed $\Omega_1(\theta_2)$ for θ_2 near $\underline{\theta}$. Thus, for some $\theta_2 > \underline{\theta}$, (17) and (18) are satisfied with slack when $\delta = \delta^*$. It follows that these constraints would continue to be satisfied at such θ_2 if δ were slightly lower. Figure A1 illustrates. ■

Two-Step Collusive Schemes: The Uniform-Distribution Example. Suppose costs are uniformly distributed over $[0, 1]$, $r = 1$ and firms revert forever to the static Nash equilibrium following any off-schedule deviation. Calculations reveal that $\pi^w = 1/6$, $p_1(\theta_2) = 1/(2 - \theta_2)$, $\delta^* = 6/7$, $\omega_T(\theta_2) = [1 - 3\theta_2 + 3(\theta_2)^2]/12$, $\Omega_1(\theta_2) = \theta_2/[2(2 - \theta_2)]$ and $\Omega_2(\theta_2) = [1 - \theta_2]^2/2$. As Figure A.2 illustrates, as θ_2 moves from 0 to 1, Ω_1 climbs (monotonically, in this case) from 0 to $\frac{1}{2}$, and Ω_2 falls from $\frac{1}{2}$ to 0. The uniform distribution function is log concave, so $\omega_T(\theta_2)$ takes a convex shape. This function is minimized at the cost level of $\frac{1}{2}$, and it assumes its maximum value (corresponding to rigid pricing at $r = 1$) at either endpoint, with $\omega_T(\underline{\theta}) = 1/12 = \omega_T(\bar{\theta})$. Noting that $\delta^*/(1 - \delta^*) = 6$, we thus have that $[\delta^*/(1 - \delta^*)]\omega_T(\underline{\theta}) = \frac{1}{2} = [\delta^*/(1 - \delta^*)]\omega_T(\bar{\theta})$, confirming that a rigid price with $r = 1$ can just meet off-schedule incentive constraints at $\delta = \delta^*$. Calculations reveal that $\Omega_2(\theta_2) > [\delta^*/(1 - \delta^*)]\omega_T(\theta_2)$ for all $\theta_2 \in (0, \frac{1}{2})$, and $\Omega_1(\theta_2) > [\delta^*/(1 - \delta^*)]\omega_T(\theta_2)$ for all $\theta_2 \in [\frac{1}{2}, 1)$. Thus, as Figure A.2 depicts, no two-step pricing scheme with $\theta_2 \in (0, 1)$ satisfies both off-schedule incentive compatibility constraints when $\delta = \delta^*$, and therefore no two-step pricing scheme can meet these constraints when $\delta < \delta^*$.

Proof of Proposition 19: Consider a candidate solution which specifies strictly increasing pricing on interval $(\theta_{k-1}, \theta_k]$. Consider introducing a tiny step on the interval $(\theta_k - \varepsilon, \theta_k]$ (for example, by introducing an additional potential step into the firms' objective between θ_{k-1} and θ_k ; since the firms can choose the number of intervals, an additional step is always feasible and thus was in the original choice set). If there is some gain to future cooperation, 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 $(\theta_{k-1}, \theta_k]$ to be a region of separation of types. Then, if we chose ε to maximize the objective when $(\theta_{k-1}, \theta_k - \varepsilon]$ is a region of separation and pooling takes place on $(\theta_k - \varepsilon, \theta_k]$, the solution $\varepsilon = 0$ should be a local maximum. Notice that no new off-schedule incentive constraints will be binding between θ_{k-1} and θ_k by assumption, so that $\Upsilon(k, \theta^{k-1}) = \Upsilon(k-1, \theta^{k-1})$.

Then, the part of the objective which depends on ε can be written as follows:

$$\begin{aligned} & \psi(\nu)\mu(\theta_k - \varepsilon, \theta_k) \left[\int_{\theta_k - \varepsilon}^{\theta_k} F(\tilde{\theta})d\tilde{\theta} - \varepsilon \cdot \Upsilon(k-1, \theta^{k-1}) \right] \\ & - \int_{\theta_{k-1}}^{\theta_k - \varepsilon} (1 - F(\tilde{\theta}))^{n-1} \left[\psi(\nu)F(\tilde{\theta}) - \Upsilon(k-1, \theta^{k-1}) \right] d\tilde{\theta}. \end{aligned}$$

This expression has zero first and second derivatives in ε when $\varepsilon = 0$. Then the question of whether $\varepsilon = 0$ is a local maximum

or minimum is determined by the third derivative, which (evaluated at $\varepsilon = 0$) is given as follows:

$$\frac{1}{2}(n-1)(1-F(\theta_k))^{n-2} \left[\psi(\nu)[f(\theta_k)^2 - F(\theta_k)f'(\theta_k)] + \Upsilon(k-1, \theta^{k-1})f'(\theta_k) \right]. \quad (19)$$

When F is strictly log-concave, $f(\theta_k)^2 - F(\theta_k)f'(\theta_k) > 0$. Further, the multipliers on the off-schedule constraints are all non-negative, so $\Upsilon(k-1, \theta^{k-1}) \geq 0$. Thus, if $f'(\theta_k) \geq 0$, (19) will be strictly positive for $\theta_k < \bar{\theta}$.

Analysis of Two-Price Scheme: We calculate expected profit over the two-period cycle as follows:

$$E\pi(\rho^l, r, \theta^*, \delta) \equiv [1 - F(\theta^*)] \left\{ \begin{array}{l} \int_{\underline{\theta}}^{\theta^*} \pi(\rho^l, \theta) dF(\theta) \\ + \frac{1}{2} \int_{\theta^*}^{\bar{\theta}} \left[\pi(r, \theta) + \delta \int_{\underline{\theta}}^{\bar{\theta}} \pi(\rho^l, \tilde{\theta}) dF(\tilde{\theta}) \right] dF(\theta) \end{array} \right\} \quad (20)$$

$$+ F(\theta^*) \left\{ \begin{array}{l} \frac{1}{2} \int_{\underline{\theta}}^{\theta^*} \left[\pi(\rho^l, \theta) + \delta \int_{\underline{\theta}}^{\bar{\theta}} \pi(r, \tilde{\theta}) dF(\tilde{\theta}) \right] dF(\theta) \\ + [1 - F(\theta^*)] \cdot \delta \int_{\underline{\theta}}^{\bar{\theta}} \pi(\rho^l, \theta) dF(\theta) \end{array} \right\}.$$

Under the uniform-distribution assumption, we derive from (20) that $E\pi(\rho^l, r, \theta^*, \delta)$ is

$$(1/4)[\theta^* - (\theta^*)^2 - (1 + \delta)] + (p^l/2)[2\theta^* + \delta - (1 + \delta)(\theta^*)^2] + (r/2)[(1 + \delta)(\theta^*)^2 + 1 - 2\theta^*], \quad (21)$$

where θ^* is given from (10) as

$$\theta^* = \theta^*(\rho^l, \delta) \equiv \frac{(2 - \delta)\rho^l - r + \delta/2}{1 - (1 + \delta)(r - \rho^l)}.$$

Consider now the off-schedule incentive constraint when the two-price scheme entails $p^l \approx r$, a firm draws $\underline{\theta} = 0$ in the second period of a cycle, and the firm is supposed to sit out in the second period. The immediate gain from jumping into the market and undercutting the rival is then approximately r . The off-schedule incentive constraint is thus

$$r \leq \frac{E\pi}{(1 - \delta)(1 + \delta)} - \frac{\pi^w}{1 - \delta} \quad (22)$$

where $E\pi$ is given via (21). With $r = 1$ and $\pi^w = 1/6$ (Nash reversion), we may use (21) to calculate that (22) requires $\delta \geq \delta^{**} \approx .88$.

9. References

- Abreu, Dilip, David Pearce and Ennio Stacchetti, "Optimal Cartel Equilibria with Imperfect Monitoring", *Journal of Economic Theory*, 39.1, June 1986, pp. 251-69.
- Abreu, Dilip, David Pearce and Ennio Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58.5, September 1990, pp. 1041-63.
- Athey, Susan and Kyle Bagwell, "Optimal Collusion under Private Information," Mimeo, MIT, 1998.
- Athey, Susan, Paul Milgrom, and John Roberts, *Robust Comparative Statics*, Unpublished research monograph.
- Bagwell, Kyle, "Commitment and Observability in Games," *Games and Economic Behavior*, 8, 1995, pp.271-80.
- Bagwell, Kyle and Robert W. Staiger, "Collusion over the Business Cycle," *RAND Journal of Economics*, 28.1, Spring 1997, pp.82-106.
- Bulow, Jeremy and John Roberts, "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 97.5, October 1989, pp. 1060-90.
- Business Week*, "Price Fixing: Crackdown Underway," June 2, 1975, 47-48.
- Carlton, Dennis W., "The Rigidity of Prices," *American Economic Review*, 76.4, September 1986, pp. 637-58.
- Carlton, Dennis W., "The Theory and the Facts of How Markets Clear: Is Industrial Organization Valuable for Understanding Macroeconomics?" in Schmalensee, Richard and Robert D. Willig, eds.. *Handbook of Industrial Organization. Volume 1*. Handbooks in Economics, no. 10 Oxford: North-Holland, 1989, pp. 909-46.
- Chone, Philippe and Jean-Charles Rochet, "Ironing, Sweeping, and Multidimensional Screening," forthcoming, *Econometrica*.
- Comanor, W.S. and M.A. Schankerman, "Identical Bids and Cartel Behavior," *Bell Journal of Economics*, 7, 1976, 281-6.
- Compte, Olivier, "Communication in Repeated Games with Private Monitoring," *Econometrica*, 1998, 66: 597-626.
- Cramton, Peter C. and Thomas R. Palfrey, "Cartel Enforcement with Uncertainty about Costs," *International Economic Review*, 31.1, February 1990, pp. 17-47.
- Eckbo, Paul L. 1976. *The Future of World Oil*. Cambridge, Mass: Ballinger.
- Fudenberg, Drew, David J. Levine and Eric Maskin, "The Folk Theorem with Imperfect Public

- Information,” *Econometrica*, 62.5 September 1994, pp. 997-1039.
- Fudenberg, Drew and Jean Tirole, *Game Theory*, Cambridge, Mass. and London: MIT Press, 1991.
- Genesove, David and Wallace Mullin, “Narrative Evidence on the Dynamics of Collusion: The Sugar Institute Case,” Mimeo, 1998, Michigan State University.
- Green, Edward J. and Robert H. Porter, “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, 52.1 January 1984, pp. 87-100.
- Hall, R.L. and C.J. Hitch, “Price Theory and Business Behavior,” *Oxford Economic Papers*, 2, May 1939, pp. 12-45.
- Haltiwanger, John and Joseph Harrington, “The Impact of Cyclical Demand Movements on Collusive Behavior,” *RAND Journal of Economics*, 22, 1991, pp. 89-106.
- Hamilton, James, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, 57, 1987, pp. 357-84.
- Hay, George A. and Daniel Kelley, “An Empirical Survey of Price-Fixing Conspiracies,” *Journal of Law and Economics*, 17, 1974, pp. 13-38.
- Jullien, Bruno, “Participation Constraints in Adverse Selection Models,” Institut d’Economie Industrielle, Universite des Sciences Sociales de Toulouse, Document de Travail 67.
- Kandori, Michihiro and Hitoshi Matshushima, “Private Observation, Communication and Collusion,” *Econometrica*, 1998, 66: 627-652.
- Kihlstrom, Richard and Xavier Vives, “Collusion by Asymmetrically Informed Firms,” *Journal of Economics and Management Strategy*, 1.2, Summer 1992, pp. 371-96.
- Lewis, T. and D. Sappington, “Counterveiling Incentives in Agency Problems,” *Journal of Economic Theory*, 1989, 49: 294-313.
- Maggi, G. and A. Rodriguez, “On Counterveiling Incentives,” *Journal of Economic Theory*, 1995, 66: 238-263.
- Mas-Colell, A., M. Whinston and J. Green (1995), *Microeconomic Theory*. Oxford University Press.
- Maskin, Eric and Jean Tirole, “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica*, 56.3, May 1988, pp. 571-99.
- Matsushima, H., “On the Theory of Repeated Games with non-Observable Actions, part I,” *Economic Letters*, 35, 1990.
- Maskin, Eric and John G. Riley, “Optimal Auctions with Risk Averse Buyers,” *Econometrica*, 52.6, November 1984, pp. 1473-1518.

- McAfee, Preston and John McMillan, "Bidding Rings," *American Economic Review*, 82.3, June 1992, pp. 579-599.
- Means, Gardiner. "Industrial Prices and Their Relative Inflexibility," Senate Document 13, 74th Congress, 1st Session, Jan. 17, 1935. Washington D.C.
- Mills Frederick C. *The Behavior of Prices*. New York: National Bureau of Economic Research, 1927.
- Myerson, R., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, 6, 58-73.
- Porter, Robert H., "Optimal Cartel Trigger Price Strategies," *Journal of Economic Theory*, 29.2, April 1983, pp. 313-38.
- Radner, Roy, "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," *Econometrica*, 49, 1981, pp. 1127-1148.
- Rochet, Jean-Charles and Lars Stole, "Competition in Nonlinear Prices," mimeo, University of Chicago.
- Rotemberg, Julio J. and Garth Saloner, "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review*; 76.3, June 1986, pp. 390-407.
- Rubenstein, Ariel and M. Yaari, "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory*, 30, 1983, pp. 406-13.
- Scherer, F.M., *Industrial Market Structure and Economic Performance*, 2nd ed., Chicago: Rand McNally, 1980.
- Spulber, Daniel F., "Bertrand Competition When Rivals' Costs Are Unknown," *Journal of Industrial Economics*, 43.1, March 1995, pp. 1- 11.
- Stigler, George, "A Theory of Oligopoly," *Journal of Political Economy*, 72.1, 1964, pp. 44-61.
- Sweezy, Paul M, "Demand Under Conditions of Oligopoly," *Journal of Political Economy*, 47.4, August 1939, pp. 568-73.
- Tirole, Jean. *The Theory of Industrial Organization*, Cambridge, Mass. and London: MIT Press, 1988.

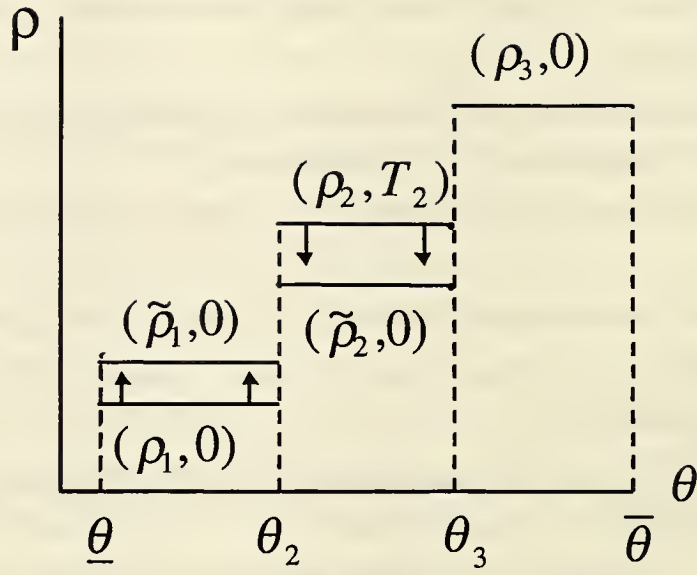


Figure 1

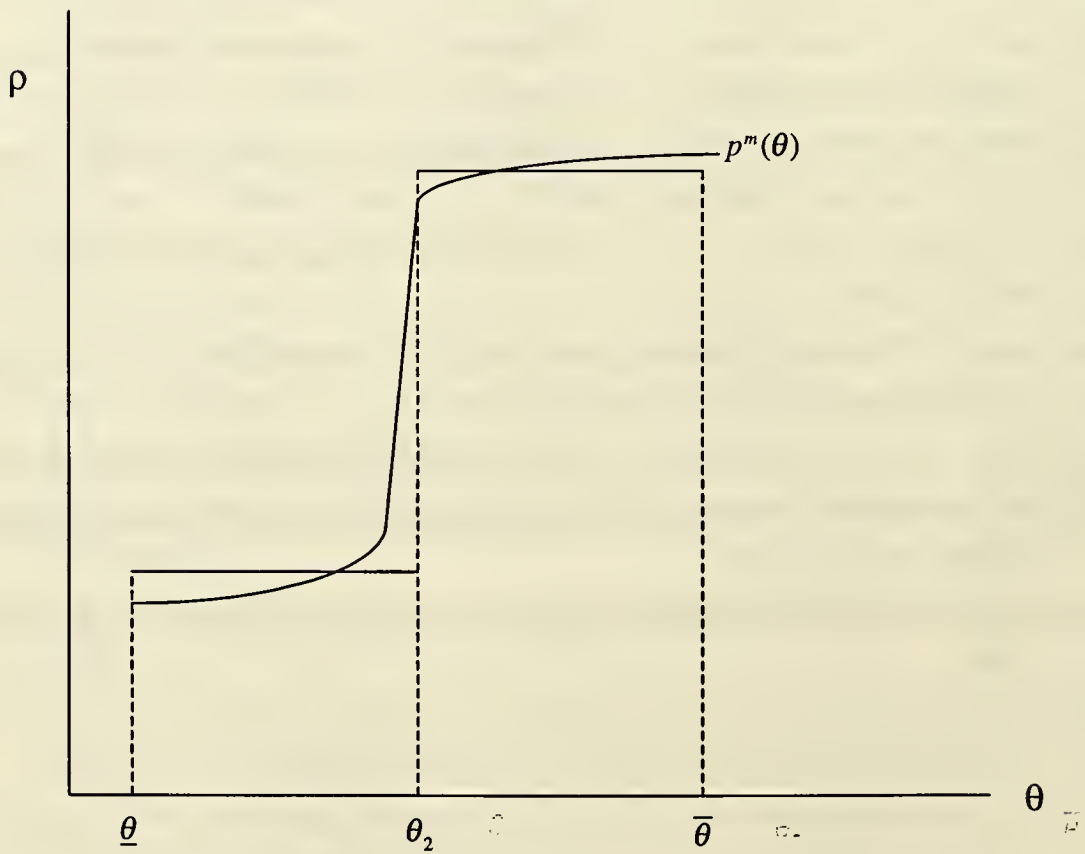


Figure 2

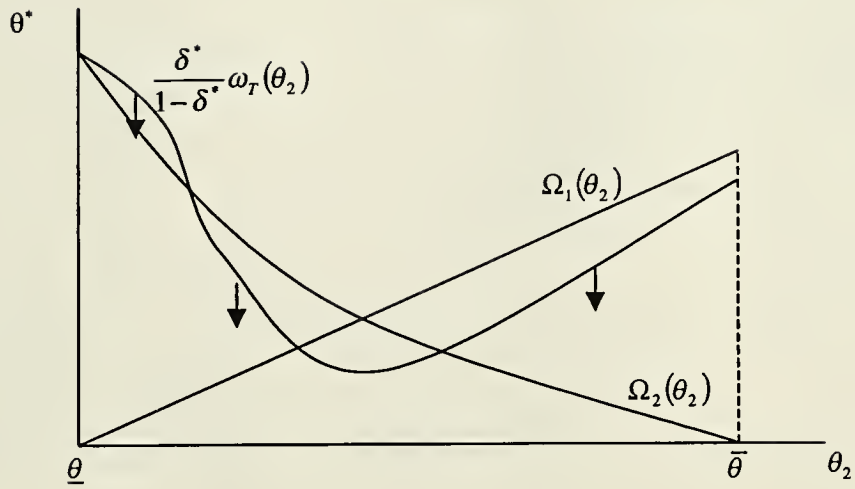


Figure A1: Arrow Represent Shift as δ Decreases.

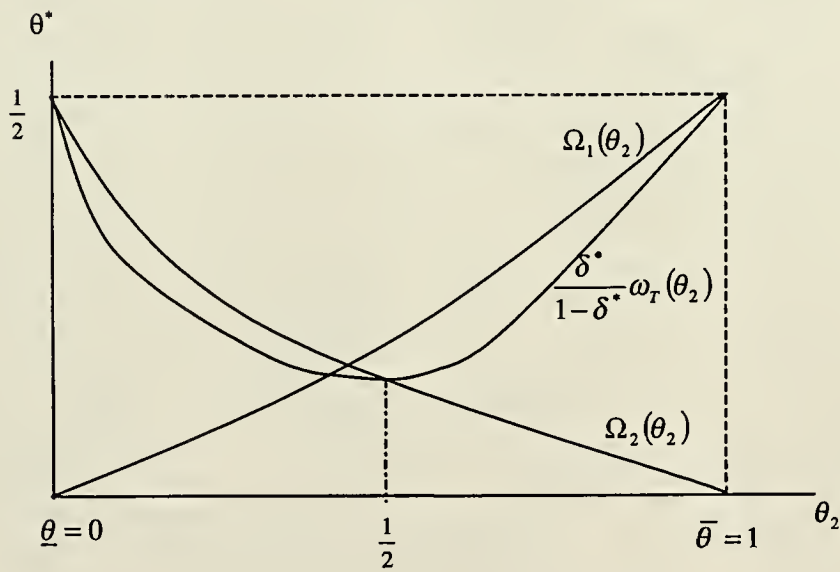


Figure A2

4873 011

Date Due

SEP 30 1999

MAY 1 1999

