*Econometrica Supplementary Material*

# SUPPLEMENT TO "COLLUSION WITH PERSISTENT COST SHOCKS": GENERALIZATIONS AND EXTENSIONS

(*Econometrica*, Vol. 76, No. 3, May 2008, 493–540)

BY SUSAN ATHEY AND KYLE BAGWELL

This document has three parts. The first part analyzes generalizations of our model to downward-sloping demand, imperfect substitutes, Cournot competition, and nonlinear cost functions. The second part describes a general dynamic programming approach to games with serially correlated private information. The third part establishes conditions for the existence of an equilibrium with productive efficiency for perfectly persistent types and also includes an analysis of severe "belief threat" punishments.

## S1. MODEL GENERALIZATIONS AND ROBUSTNESS OF RESULTS

THIS SECTION DISCUSSES the robustness of our results on rigid pricing to several modifications of the basic model, allowing for alternative specifications of demand, quantity competition, and nonlinear cost functions. In the last subsection, we briefly discuss how our results about first-best equilibria generalize.

### S1.1. *Downward-Sloping Demand, Perfect Substitutes*

Modify Model 2 of the main paper as follows. For simplicity of notation, eliminate announcements and quantity restrictions from the model, so that the only choice in the stage game is the price. (This does not affect our analysis.) Maintaining the assumption that goods are perfect substitutes, define the *profit-if-lowest-price* function as $\pi(p_{i,t}, \theta_{i,t}) \equiv (p_{i,t} - \theta_{i,t})D(p_{i,t})$, where $D$ is a differentiable market demand function that satisfies $D > 0 \geq D'$ over the relevant range. We assume that $\pi$ is strictly quasiconcave in $p_{i,t}$, with a unique maximizer, $\rho^m(\theta_{i,t})$, where $\rho^m(\overline{\theta}) \geq \overline{\theta}$. The monopoly price, $\rho^m(\theta_{i,t})$, is nondecreasing in $\theta_{i,t}$.

With these modifications, given period strategies $\boldsymbol{\rho}_t$ and beliefs $\boldsymbol{\nu}_{-i,t}$, the expected market share for firm $i$ in period $t$ can be written

$$\overline{m}_i(\hat{\theta}_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i,t}) = \mathbb{E}_{\boldsymbol{\theta}_{-i,t}}\big[\varphi_i(\boldsymbol{\rho}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}))|\boldsymbol{\nu}_{-i,t}\big]$$

and the interim profit function can be written

$$\overline{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i,t}) = \pi(\rho_{i,t}(\hat{\theta}_{i,t}), \theta_{i,t})\overline{m}_i(\hat{\theta}_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i,t}).$$

We next define $Q(p_{i,t}; \boldsymbol{\rho}_{-i,t}, \boldsymbol{\nu}_{-i,t})$ as the quantity a firm expects to sell when it sets price $p_{i,t}$ and opponents use pricing function $\boldsymbol{\rho}_{-i,t}$. In the present context, $Q_i(p_{i,t}; \boldsymbol{\rho}_{-i,t}, \boldsymbol{\nu}_{-i,t}) = D(p_{i,t}) \cdot \mathbb{E}_{\boldsymbol{\theta}_{-i,t}}[\varphi_i(p_{i,t}, \boldsymbol{\rho}_{-i,t}(\boldsymbol{\theta}_{-i,t}))|\boldsymbol{\nu}_{-i,t}]$, and so the function $\overline{\pi}_i$ may be alternatively expressed as

$$\overline{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i,t}) = (\rho_{i,t}(\hat{\theta}_{i,t}) - \theta_{i,t})Q_i(\rho_{i,t}(\hat{\theta}_{i,t}); \boldsymbol{\rho}_{-i,t}, \boldsymbol{\nu}_{-i,t}).$$

We ask three related questions about the optimality of rigid pricing in this model. First, can we bound the profits from an equilibrium with productive efficiency (if it is an equilibrium at all)? Second, is price rigidity optimal when demand is sufficiently inelastic? (That is, is the price rigidity result of Proposition 2 in the main paper "knife edge"?) Third, what can we say about optimal equilibria more generally?

Following the proof of Proposition 2, let $\breve{R}_i(\hat{\theta}_{i,1})$ and $\breve{Q}_i(\hat{\theta}_{i,1})$ denote the expected future discounted revenues and demand that firm $i$ anticipates if it mimics type $\hat{\theta}_{i,1}$ throughout the game, and let $\breve{M}_i(\hat{\theta}_{i,1})$ be the associated market share. If firm $i$'s type is $\theta_{i,1}$, then the present discounted value of profits for firm $i$ can be represented as $U_i(\hat{\theta}_{i,1}, \theta_{i,1}) \equiv \breve{R}_i(\hat{\theta}_{i,1}) - \theta_{i,1}\breve{Q}_i(\hat{\theta}_{i,1})$. Incentive compatibility implies the monotonicity constraint that $\breve{Q}_i(\theta_{i,1})$ is nonincreasing. Further, we note that local incentive compatibility together with the envelope theorem imply that

$$U_i(\theta_{i,1}, \theta_{i,1}) = U_i(\overline{\theta}, \overline{\theta}) + \int_{\theta_{i,1}}^{\overline{\theta}} \breve{Q}_i(\tilde{\theta})\, d\tilde{\theta}. \tag{S1.1}$$

Thus, using integration by parts, given prior $F_0$, ex ante profits are

$$\mathbb{E}_{\theta_{i,1}}[U_i(\theta_{i,1}, \theta_{i,1})] = U_i(\overline{\theta}, \overline{\theta}) + E_{\theta_{i,1}}\left[\breve{Q}_i(\theta_{i,1})\frac{F_0}{f_0}(\theta_{i,1})\right]. \tag{S1.2}$$

This expression shows that there are important similarities and differences between the downward-sloping demand case and the inelastic demand case. In both cases, monotonicity of $F_0/f_0$ implies that, all else equal, we would like to have $\breve{Q}_i$ nondecreasing as well, allocating more expected demand to higher-cost types. However, the force in favor of $\breve{Q}_i$ nondecreasing conflicts with incentive compatibility and creates a force for pooling.

In addition, in the case of inelastic demand, a rigid price at $r$ maximizes both terms of expression (S1.2). This is no longer true when demand is downward-sloping. Instead, the two terms are in conflict about the level of the price. To maximize the profit to the high-cost type, it would be optimal to have a rigid price, as the market share for type $\overline{\theta}$ is thereby made as large as possible; furthermore, the best rigid price would be the monopoly price for the high-cost type. To maximize the second term (subject to monotonicity constraints), it would also be optimal to have a rigid price, following an argument analogous to the one we used for inelastic demand; however, the best rigid price would now be as low as possible, so as to make the quantity produced as large as possible and maximize the second term (recall that with inelastic demand, the expected market share was fixed at 1). The conflict between the two terms about the optimal price level implies that price rigidity is not necessarily optimal.

At the same time, (S1.2) illustrates forces in favor of at least partial price rigidity, and for certain classes of demand functions it can be used to provide a very simple proof that partial price rigidity must be optimal. In addition, when demand is sufficiently inelastic, full price rigidity will be optimal.

PROPOSITION S1: *Suppose that demand is downward-sloping, firms sell perfect substitutes, cost types are perfectly persistent, and $F_0$ is log-concave. Then there exists $\delta$ sufficiently large such that the following hold.* (i) *Suppose that there exists a price level $p' \geq \overline{\theta}$ such that for $p \leq p'$, demand is inelastic: $D'(p) = 0$ for $p \leq p'$. Then the optimal perfect public Bayesian equilibrium (PPBE) entails a positive probability (ex ante) of production for type $\overline{\theta}$, and thus some productive inefficiency.* (ii) *If demand is sufficiently inelastic, then the best rigid-pricing equilibrium is the optimal PPBE.*

PROOF: (i) Suppose that the equilibrium entails zero probability of production for type $\overline{\theta}$. Let the ex ante expected production in this equilibrium as a function of type be $\breve{Q}_i^S(\cdot)$, where $\breve{Q}_i^S(\overline{\theta}) = 0$.

Observe that even if deviations from a collusive agreement can be punished by giving a firm zero profits forever (as in the belief threat punishment discussed below), any such equilibrium must have prices less than or equal to $\overline{\theta}$. To see why, note that with probability 1, the high-cost type gets zero market share and zero profit in every period, so there can be no future reward to the high-cost type for pricing above cost. Thus, in any period where the market price was greater than $\overline{\theta}$, the high-cost type would have an incentive to deviate and undercut the market price. Thus, prices must be less than or equal to $\overline{\theta}$ in the proposed equilibrium.

We next argue that there exists a rigid- or partially rigid-pricing scheme that dominates the proposed one (recall that we have already proved that the best rigid-pricing scheme is an equilibrium for firms sufficiently patient). Let $K = \mathbb{E}_{\theta_{i,1}}[\breve{Q}_i^S(\theta_{i,1})]$. Let $p^K$ be the maximum price such that $D(p^K)/I \geq K$; note that given our specification of demand, $p^K \geq \overline{\theta}$. Then if the firms use a rigid-pricing scheme with prices equal to $p^K$, the difference between the ex ante expected profits with the rigid scheme and the original scheme can be written (using (S1.2))

$$(p^K - \overline{\theta})K + K \cdot \mathbb{E}_{\theta_{i,1}}\left[\left(1 - \frac{\breve{Q}_i^S(\theta_{i,1})}{K}\right)\frac{F_0}{f_0}(\theta_{i,1})\right]. \tag{S1.3}$$

Since $\mathbb{E}_{\theta_{i,1}}[\breve{Q}_i^S(\theta_{i,1})] = K$, $(\breve{Q}_i^S(\theta_{i,1})/K) \cdot f(\theta_{i,1})$ is a probability density. Then, since $\breve{Q}_i$ is nonincreasing, the probability distribution associated with $f(\theta_{i,1})$ dominates the distribution associated with $(\breve{Q}_i^S(\theta_{i,1})/K) \cdot f(\theta_{i,1})$ by first-order stochastic dominance (FOSD). Since $F_0/f_0$ is nondecreasing and is strictly increasing at $\underline{\theta}$ (using our assumption that $f(\underline{\theta}) > 0$), the second term is positive. If $p^K \geq \overline{\theta}$, then the whole expression (S1.3) is positive and the rigid scheme is better. When $\delta$ is sufficiently high, a (pooling) carrot–stick scheme will be an equilibrium.

(ii) Normalize a family of demand functions so that $D(0) = 1$. Then expected profit is given by

$$\begin{aligned}\mathbb{E}_{\theta_{i,1}}[U_i(\theta_{i,1}, \theta_{i,1})] &= U_i(\overline{\theta}, \overline{\theta}) + \mathbb{E}_{\theta_{i,1}}\left[\breve{Q}_i(\theta_{i,1})\frac{F_0}{f_0}(\theta_{i,1})\right] \\ &= U_i(\overline{\theta}, \overline{\theta}) + \mathbb{E}_{\theta_{i,1}}\left[\breve{M}_i(\theta_{i,1})\frac{F_0}{f_0}(\theta_{i,1})\right] \\ &\quad + \mathbb{E}_{\theta_{i,1}}\left[[\breve{Q}_i(\theta_{i,1}) - \breve{M}_i(\theta_{i,1})]\frac{F_0}{f_0}(\theta_{i,1})\right]\end{aligned}$$

for each member of the family. Assuming log-concavity and inelastic demand, we showed in Proposition 2 that a rigid price at the reservation value uniquely maximizes the first two terms. As demand becomes more inelastic, $\breve{Q}_i(\theta_{i,1})$ approaches $\breve{M}_i(\theta_{i,1})$ for all prices below the reservation value, and so the first two terms dominate. The level of the optimal rigid price approaches the reservation value as demand becomes inelastic; further, as in Proposition 2, this scheme is used in a PPBE for sufficiently patient firms. *Q.E.D.*

The only place we used the restriction on the class of demand curves in part (i) was to guarantee that $p^K \geq \overline{\theta}$, where $p^K$ is the price that generates output equal to average output in the posited equilibrium with separation. The result could be extended in a number of ways; for example, we would have $p^K \geq \overline{\theta}$ as long as demand was sufficiently inelastic for prices below $\overline{\theta}$. More generally, any demand curve and cost distribution that jointly lead to monopoly prices that are much higher than costs will have the feature that the optimal rigid-pricing scheme dominates schemes with productive efficiency for an interval of high-cost types, since such schemes necessarily entail prices below $\overline{\theta}$.

Although we have now argued that neither rigid pricing nor full productive efficiency will typically be optimal, it is useful to illustrate forces in favor of partial pooling (and in particular, intervals of pooling). Suppose that $\breve{Q}_i$ is strictly decreasing on $[x, y] \subset [\underline{\theta}, \overline{\theta}]$. Now define the series of pricing strategies implicitly (where pricing strategies need only be modified on $[x, y]$) so that the associated quantity scheme $\breve{Q}_i'$ is equal to $\breve{Q}_i$ outside of $[x, y]$ but is constant on $[x, y]$, and

$$\mathbb{E}_{\theta_{i,t}}[\breve{Q}_i(\theta_{i,t})|\theta_{i,t} \in [x, y]] = \mathbb{E}_{\theta_{i,t}}[\breve{Q}_i'(\theta_{i,t})|\theta_{i,t} \in [x, y]] = \breve{Q}_i'(y). \tag{S1.4}$$

Then we can define a probability distribution

$$G(\theta_{i,1}; \breve{Q}_i) = \frac{1}{\breve{Q}_i'(y)} \int_x^{\theta_{i,1}} \breve{Q}_i(\theta_{i,1})\, dF_0(\theta_{i,1}|\theta_{i,1} \in [x, y])$$

with $G(\theta_{i,1}; \breve{Q}_i')$ defined in an analogous way. Since $G(\cdot; \breve{Q}_i')$ dominates $G(\cdot; \breve{Q}_i)$ by FOSD, if $F_0$ is log-concave, then

$$E_{\theta_{i,1}}\left[\breve{Q}_i'(\theta_{i,1})\frac{F_0}{f_0}(\theta_{i,1})\right] \geq E_{\theta_{i,1}}\left[\breve{Q}_i(\theta_{i,1})\frac{F_0}{f_0}(\theta_{i,1})\right]$$

and a force in favor of rigidity is thus illustrated. However, the new scheme does not, in general, satisfy local on-schedule incentive constraints at $x$ and $y$, and so this argument does not establish that a step function is optimal (as opposed to a scheme with intervals of pooling and intervals of separation).

A full analysis would incorporate the additional modifications to the scheme that would restore on-schedule incentive compatibility, and also satisfy off-schedule incentive compatibility, and establish conditions under which the modification increases expected profits. For example, we could consider a new scheme $\breve{Q}_i''$ that is equal to $\breve{Q}_i'$ for $\theta_{i,1} \geq y$, is constant on $[x, y]$, and has the same set of intervals as the original pricing function on which the pricing function is strictly increasing; but the new scheme is modified to satisfy on-schedule and off-schedule incentive compatibility constraints for $\theta_{i,1} < x$ (it may or may not be possible to satisfy both on- and off-schedule constraints; let us focus on the case where it is). Then we could compare the overall profits from $\breve{Q}_i''$ to profits from the original scheme. The difference can be expressed as

$$\begin{aligned}&\left(E_{\theta_{i,1}}\left[(\breve{Q}_i'(x) - \breve{Q}_i(\theta_{i,1}))\frac{F_0}{f_0}(\theta_{i,1})\middle|\theta_{i,1} \in [x, y]\right]\right)\Pr(\theta_{i,1} \in [x, y])\\ &\quad + (\breve{Q}_i''(x) - \breve{Q}_i'(x))\left(E_{\theta_{i,1}}\left[\frac{F_0}{f_0}(\theta_{i,1})\middle|\theta_{i,1} \in [x, y]\right]\right)\Pr(\theta_{i,1} \in [x, y])\\ &\quad + \left(E_{\theta_{i,1}}\left[[\breve{Q}_i''(\theta_{i,1}) - \breve{Q}_i(\theta_{i,1})]\frac{F_0}{f_0}(\theta_{i,1})\middle|\theta_{i,1} < x\right]\right)\Pr(\theta_{i,1} < x).\end{aligned}$$

The expression on the first line is positive by our previous analysis, but the second and third expressions are ambiguous. The sign and magnitude depend on how the pricing scheme needs to be adjusted to maintain incentive compatibility, and on the shape of the demand curve. In general, there may be other ways to improve upon a particular pricing scheme with intervals of strict monotonicity (for example, modifying the pricing function above $y$ or modifying the shape of the pricing below $x$ in other ways), so even if the latter expression is negative, it still may be possible to improve on a scheme with intervals where the pricing function is strictly monotone. See Athey, Atkeson, and Kehoe (2005) for an approach to analyzing the optimality of partial pooling in a different but related model.

This analysis suggests that, in general, partial pooling will be optimal, taking a form similar to the market-sharing two-step scheme analyzed in Section 3.2.3 of the main paper (though there may be more than two steps).

### S1.2. *Imperfect Substitutes*

We now consider briefly the possibility that firms sell imperfect substitutes. Note that there are several differences from the perfect substitutes case: with imperfect substitutes, the first-best allocation typically has all firms producing a positive amount, and typically demand is continuous in prices. However, these differences do not affect the analysis of the optimal collusive scheme very much. In this case, given strategies, we can represent demand in period $t$ by

$$Q_i(p_{i,t}; \boldsymbol{\rho}_{-i,t}, \boldsymbol{\nu}_{-i,t}) = \mathbb{E}_{\boldsymbol{\theta}_{-i,t}}\big[D_i(p_{i,t}, \boldsymbol{\rho}_{-i,t}(\boldsymbol{\theta}_{-i,t}))|\boldsymbol{\nu}_{-i,t}\big]$$

and let $\breve{Q}_i(\theta_{i,1})$ be the expected discounted demand a firm expects over the course of the game from mimicking $\theta_{i,1}$. Then (S1.2) still characterizes profits.

In this model, expected demand for each cost type depends on the entire pricing function of opponents. If the demand function is linear in prices, however, players care only about the average price of opponents. Then it is possible to modify pricing functions for each player only on the interval of types $[x, y]$, to change a pricing function from being strictly increasing to being constant on that interval, while leaving the expected price and the expected demand for opponents unchanged outside that interval. Then the arguments of the last section can be applied: in particular, if $\breve{Q}_i$ is strictly decreasing on $[x, y] \subset [\underline{\theta}, \overline{\theta}]$, we can find new pricing strategies such that $\breve{Q}_i'$ is equal to $\breve{Q}_i$ outside of $[x, y]$ but is constant on $[x, y]$, and expected demand on that region is unchanged. Then, as above, there will be a force in favor of pooling. From a formal perspective, this model is closely related to one analyzed by Athey, Atkeson, and Kehoe (2005). We conjecture that their arguments can be modified to show that if $F$ and $1 - F$ are log-concave, and demand is linear, then an optimal PPBE is characterized by intervals of pooling.

### S1.3. *Quantity Competition*

The analysis is similar when firms compete in quantities rather than prices. To see this, let $P$ be the inverse demand function and let firm $i$'s expected profit in period $t$ be given by (where firm $i$'s quantity strategy in period $t$ is $\psi_{i,t}$)

$$\begin{aligned}&\overline{\pi}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\nu}_{-i})\\&\quad = \big(\mathbb{E}_{\boldsymbol{\theta}_{-i,1}}\big[P(\psi_{i,t}(\hat{\theta}_{i,t}), \boldsymbol{\psi}_{i,t}(\boldsymbol{\theta}_{-i,t}))|\boldsymbol{\nu}_{-i,t}\big] - \theta_{i,t}\big)\psi_{i,t}(\hat{\theta}_{i,t}).\end{aligned}$$

Again, we can let $\breve{Q}_i(\theta_{i,1})$ denote the expected discounted demand firm $i$ expects from mimicking $\theta_{i,1}$ throughout the game. Then (S1.2) characterizes profits, and a force in favor of pooling remains. Again, we expect optimal PPBE to be characterized by partial pooling.

### S1.4. *Nonlinear Costs*

So far we have assumed constant marginal costs. Suppose instead that total costs in period $t$ are given by $h(q_{i,t}, \theta_{i,t}) = h^q(q_{i,t})h^\theta(\theta_{i,t})$ when a firm with type $\theta_{i,t}$ produces $q_{i,t}$. Then we provide sufficient conditions for rigid pricing to dominate alternative schemes that have the property that the highest-cost type serves less than $1/I$ of the market in each period. Any scheme with greater period-by-period productive efficiency than rigid pricing will have this feature. At first it might seem impossible to find an incentive-compatible scheme where the highest-cost type serves more than $1/I$; however, even though it might seem pathological, nonlinearities in cost do make it possible. Thus, we stop short of a full proof of the optimality of rigid pricing. Despite this, our analysis makes clear that nonlinear costs do not remove the incentive for pooling nor do they invalidate our overall approach. In addition, our analysis establishes that rigid pricing dominates a wide range of schemes with greater productive efficiency.

PROPOSITION S2: *Suppose that* (i) $h^q$ *and* $h^\theta$ *are differentiable*, *with* $h^\theta > 0$, $h^q > 0$, $h^q_q > 0$, $h^\theta_\theta > 0$, $h^q_{qq} \leq 0$, *and* $h^\theta_{\theta\theta} \geq 0$, *and* (ii)

$$r > h_q(0, \overline{\theta}). \tag{S1.5}$$

*If* $F_0$ *is log-concave*, *then rigid pricing at* $r$ *dominates any scheme satisfying the market share restriction that* $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) \leq 1/I$ *for each* $i$, $t$, *and* $\boldsymbol{\theta}_{-i,1}$.

Before proving the result, we pause to interpret the sufficient conditions. They require economies of scale: marginal costs are nonincreasing. It may seem somewhat surprising that with economies of scale, pooling could be optimal, since cost considerations favor shifting production to one firm even more than in the case of constant marginal costs. However, as in the case of constant marginal cost, allocating more market share to high-cost types relaxes incentive constraints for low-cost types. Thus, expected "information rents" to the firms are higher when more market share goes to high-cost firms. When $r$ is higher than marginal cost for the high-cost type in the relevant range and if rigid pricing increases the market share to the high-cost type, then rigid pricing increases profits to the high-cost type as well.

When the sufficient conditions of the proposition fail, there are generally competing effects and it may be necessary to consider a parameterized model to fully characterize optimal PPBE.

PROOF OF PROPOSITION S2: Following our notation above, let $\breve{R}_i(\hat{\theta}_{i,1})$ denote the expected future discounted revenue that firm $i$ anticipates if it mimics type $\hat{\theta}_{i,1}$ throughout the game, and let $\breve{M}_i(\hat{\theta}_{i,1})$ be the associated market share. Let $\breve{H}_i(\hat{\theta}_{i,1}, \theta_{i,1})$ be the expected discounted total cost when type $\theta_{i,1}$

mimics $\hat{\theta}_{i,1}$ throughout the game, decomposed into two components, so that $\breve{H}_i(\hat{\theta}_{i,1}, \theta_{i,1}) = h^\theta(\theta_{i,1}) \cdot \breve{H}^q(\hat{\theta}_{i,1})$.

Let

$$U_i(\hat{\theta}_{i,1}, \theta_{i,1}) = \breve{R}_i(\hat{\theta}_{i,1}) - \breve{H}_i(\hat{\theta}_{i,1}, \theta_{i,1}).$$

Using the envelope theorem and on-schedule incentive compatibility,

$$U_i(\theta_{i,1}, \theta_{i,1}) = \breve{R}_i(\overline{\theta}) - \breve{H}_i(\overline{\theta}, \overline{\theta}) + \int_{\theta_{i,1}}^{\overline{\theta}} h^\theta_\theta(\tilde{\theta}_{i,1}) \cdot \breve{H}^q(\tilde{\theta}_{i,1})\, d\tilde{\theta}_{i,1}.$$

So, using integration by parts,

$$\mathbb{E}_{\theta_{i,1}}[U_i(\theta_{i,1}, \theta_{i,1})] = \mathbb{E}_{\theta_{i,1}}\left[\breve{R}_i(\overline{\theta}) - \breve{H}_i(\overline{\theta}, \overline{\theta}) + h^\theta_\theta(\theta_{i,1})\breve{H}^q(\theta_{i,1})\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}\right]. \tag{S1.6}$$

The on-schedule incentive constraints also imply that

$$\begin{aligned} U_i(\theta_{i,1}, \theta_{i,1}) &\geq U_i(\hat{\theta}_{i,1}, \theta_{i,1}) \\ &= U_i(\hat{\theta}_{i,1}, \hat{\theta}_{i,1}) + \breve{H}_i(\hat{\theta}_{i,1}, \hat{\theta}_{i,1}) - \breve{H}_i(\hat{\theta}_{i,1}, \theta_{i,1}), \\ U_i(\hat{\theta}_{i,1}, \hat{\theta}_{i,1}) &\geq U_i(\theta_{i,1}, \hat{\theta}_{i,1}) \\ &= U_i(\theta_{i,1}, \theta_{i,1}) + \breve{H}_i(\theta_{i,1}, \theta_{i,1}) - \breve{H}_i(\theta_{i,1}, \hat{\theta}_{i,1}), \end{aligned}$$

so combining,

$$\breve{H}_i(\hat{\theta}_{i,1}, \hat{\theta}_{i,1}) + \breve{H}_i(\theta_{i,1}, \theta_{i,1}) \leq \breve{H}_i(\hat{\theta}_{i,1}, \theta_{i,1}) + \breve{H}_i(\theta_{i,1}, \hat{\theta}_{i,1}).$$

In words, on-schedule incentive compatibility implies that $\breve{H}_i$ is submodular. Using our definitions and monotonicity restrictions, that in turn implies that $\breve{H}^q(\hat{\theta}_{i,1})$ is nonincreasing in $\hat{\theta}_{i,1}$. This does not, however, imply that expected market shares are globally decreasing in $\hat{\theta}_{i,1}$, since $h^q$ is nonlinear.

*Profit-at-the-Top*

Fix a PPBE. Let $m_{i,t}(\boldsymbol{\theta}_1)$ be the market-share allocation to firm $i$ in period $t$ as a function of firm types. For every $i, t$, and $\boldsymbol{\theta}_{-i,1}$, we assume that $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) \leq 1/I$. The ex ante expected value of "profit-at-the-top" to firm $i$ in period $t$ is then

$$\mathbb{E}_{\boldsymbol{\theta}_{-i,1}}\big[\rho_{i,t}(\overline{\theta})m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) - h^\theta(\overline{\theta})h^q(m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}))\big].$$

We claim that this profit-at-the-top cannot exceed that which is achieved when a best rigid-pricing scheme is used in period $t$.

To establish this claim, we fix $i$ and $t$, and suppose that $\rho_{i,t}(\overline{\theta}) < r$ and/or $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) < 1/I$ for a positive measure of values for $\boldsymbol{\theta}_{-i,1}$. (Recall that a best rigid-pricing scheme has $\rho_{i,t}(\overline{\theta}) = r$ and $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) \equiv 1/I$.) Clearly, if $\rho_{i,t}(\overline{\theta}) < r$, then profit-at-the-top would be increased if the price were raised to $r$. Suppose then that $\rho_{i,t}(\overline{\theta}) = r$ and $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) < 1/I$ for a positive measure of values for $\boldsymbol{\theta}_{-i,1}$. For each such $\boldsymbol{\theta}_{-i,1}$, if we were to increase firm $i$'s market share from $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1})$ to $1/I$, then firm $i$ would enjoy a strict increase in profit when state $(\overline{\theta}, \boldsymbol{\theta}_{-i,1})$ occurs. This follows since $r - h^{\theta}(\overline{\theta})h^{q}_{q}(q_{i,t}) \geq r - h^{\theta}(\overline{\theta})h^{q}_{q}(0) = r - h_{q}(0, \overline{\theta}) > 0$ for $q_{i,t} \in [0, 1/I]$, where the first inequality uses $h^{q}_{qq}(q_{i,t}) \leq 0$ and the second inequality uses (S1.5). Over a positive measure of values for $\boldsymbol{\theta}_{-i,1}$, these pointwise improvements imply a strict increase of profit-at-the-top to firm $i$ in period $t$.

Thus, when a best rigid-pricing scheme is used, the effect on the profit-at-the-top in each period is positive, and so the aggregate effect must be positive. With regard to the profit-at-the-top, we have therefore established that any PPBE with the feature that $m_{i,t}(\overline{\theta}, \boldsymbol{\theta}_{-i,1}) \leq 1/I$ for each $i$, $t$, and $\boldsymbol{\theta}_{-i,1}$ is dominated by using the best rigid-pricing scheme in each period. Since rigid pricing at $r$ is a PPBE for sufficiently large $\delta$, it provides greater expected profit-at-the-top than any other PPBE that satisfies the market-share restriction.

*Information Rents*

The last part of the ex ante expected profits expression (S1.6), referred to as the information rents, is

$$\mathbb{E}_{\theta_{i,1}}\left[h^{\theta}_{\theta}(\theta_{i,1})\breve{H}^{q}(\theta_{i,1})\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}\right].$$

For any $\breve{M}_i$ associated with a PPBE, consider an alternative market-share allocation function derived from rigid pricing, so that market shares are equal for all types and all firms in each period, and $\breve{H}^{q}(\theta_i)$ is constant in $\theta_i$. Rigid pricing dominates from the perspective of information rents if

$$\mathbb{E}_{\theta_{i,1}}\left[\left(h^{\theta}_{\theta}(\theta_{i,1})\frac{h^{q}(1/I)}{1-\delta} - h^{\theta}_{\theta}(\theta_{i,1})\breve{H}^{q}(\theta_{i,1})\right)\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}\right] \geq 0.$$

The expression can be rewritten as

$$\text{(S1.7)} \quad \operatorname{cov}\left(\left(\frac{h^{q}(1/I)}{1-\delta} - \breve{H}^{q}(\theta_{i,1})\right), h^{\theta}_{\theta}(\theta_{i,1})\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}\right) + \mathbb{E}_{\theta_{i,1}}\left[\left(\frac{h^{q}(1/I)}{1-\delta} - \breve{H}^{q}(\theta_{i,1})\right)\right] \cdot \mathbb{E}\left[h^{\theta}_{\theta}(\theta_{i,1})\frac{F_0(\theta_{i,1})}{f_0(\theta_{i,1})}\right].$$

Note that $\mathbb{E}_{\theta_{i,1}}[h^\theta_\theta(\theta_{i,1})F_0(\theta_{i,1})/f_0(\theta_{i,1})] \geq 0$ since the integrand is positive everywhere. In addition, concavity of $h^q$ and Jensen's inequality imply that for each $t$,

$$\begin{aligned} h^q(1/I) &= h^q\big(\mathbb{E}_{\boldsymbol{\theta}_1}[m_{i,t}(\boldsymbol{\theta}_1)]\big) \geq \mathbb{E}_{\boldsymbol{\theta}_1}\big[h^q(m_{i,t}(\boldsymbol{\theta}_1))\big] \\ &= \mathbb{E}_{\theta_{i,1}}\big[\mathbb{E}_{\boldsymbol{\theta}_{-i,1}} h^q(m_{i,t}(\boldsymbol{\theta}_1))\big], \end{aligned}$$

so that the second term of (S1.7) is positive, using the definition of $\breve{H}^q(\theta_{i,1})$. On the first term, given log-concavity, $h^\theta_\theta > 0$, and $h^\theta_{\theta\theta} \geq 0$, $h^\theta_\theta(\theta_{i,1})F_0(\theta_{i,1})/f_0(\theta_{i,1})$ is nondecreasing. By incentive-compatibility, $(h^q(1/I))/(1-\delta) - \breve{H}^q(\theta_{i,1})$ is nondecreasing. The covariance of two nondecreasing functions of a single random variable must be positive. Therefore, (S1.7) must be positive and the alternative with equal market shares must be preferred. *Q.E.D.*

We conclude by noting that the restriction on market shares was used only in the analysis of profit-at-the-top, so that there will still be a strong force for at least partial pooling (from the information rents term) even when the market share-restriction fails.

### S1.5. *Robustness of Results About Efficient Collusion*

Although we have not conducted a full numerical analysis of alternative models, the general approach we take to constructing first-best equilibria can be generalized to alternative models. The first step would be to construct a punishment equilibrium analogous to the carrot–stick pooling equilibrium (or some alternative). Although critical discount factors would clearly differ and some details of the construction would differ, there is no reason to expect difficulties generalizing the carrot–stick equilibrium to alternative models. The dynamic programming approach we employ for analyzing first-best equilibria generalizes directly to other models of product market competition. Since first-best equilibrium payoffs are higher than payoffs from pooling equilibria, for sufficiently patient firms, off-schedule incentive constraints should not bind.

The main challenge in eliciting truthful revelation alongside efficient market-share allocation is to provide future rewards and punishments that provide sufficient incentives for firms to give up market share, while still implementing efficient allocation in the reward and punishment continuation equilibria. With discrete types and perfect substitutes (either Bertrand or Cournot), states of the world arise with positive probability where market share can be shifted among firms. For particular functional forms, it is straightforward to write the system of equations and incentive constraints that must be satisfied to implement a first-best equilibrium, analogous to the system described in the paper. One could then describe the parameter values for which a first-best equilibrium exists. We note that if firms compete in quantities, communication plays a more substantive role in coordinating firms on the desired quantities as a function of cost.

With imperfect substitutes, there is typically a uniquely optimal market-share allocation for any cost vector, even when costs are identical. Thus, some inefficiency should be expected for any discount factor strictly less than 1. Similarly, with nonlinear costs, diseconomies of scale imply that firms should share the market in a particular way. Thus, distorting production to provide rewards or punishments for past revelation will lead to some inefficiency. In both cases, the requisite inefficiency should decrease as patience increases.

## S2. USING DYNAMIC PROGRAMMING TO ANALYZE THE FIRST-BEST SCHEME

This section provides more details about applying dynamic programming approaches in the spirit of Abreu, Pearce, and Stacchetti (1986, 1990) and Cole and Kocherlakota (2001) to our Model 1. Although this is not necessary to establish that the first-best scheme is a PPBE, the techniques are useful more generally for analyzing self-generating sets of PPBE values either numerically or analytically.

We retain the notation from the main paper, but also introduce some new notation. Let $\mathcal{V}$ be the set of functions $\mathbf{v}=(v_1, \dots, v_I)$ such that $v_i : \Theta_i \to \mathbb{R}$. This is the set of possible type-contingent payoff functions. Let $\mathcal{W} = \mathcal{V} \times \Delta\Theta$. The set of PPBE then corresponds to a subset of $\mathcal{W}$. Each equilibrium is described by a set of initial beliefs about opponents, $\boldsymbol{\mu} \in \Delta\Theta$, and a function $\mathbf{v} \in \mathcal{V}$ that specifies the payoff each player expects to attain, conditional on the player's true type.

Consider continuation value and belief functions $(\mathbf{V}, \mathbf{M})$ mapping $\mathcal{Z}$ to $\mathcal{W}$. For every possible publicly observed outcome $\mathbf{z}_t$ from period $t$, these functions specify an associated belief and a type-contingent continuation payoff function. That is, $V_i(\mathbf{z}_t)$ is the type-contingent payoff function that will be realized following observed actions $\mathbf{z}_t$, and $V_i(\mathbf{z}_t)(\theta_{i,t+1})$ is the payoff firm $i$ expects starting in period $t+1$ if $\mathbf{z}_t$ was the vector of observed actions in period $t$ and its true type in period $t+1$ is $\theta_{i,t+1}$. Note that this structure makes it possible to compute firm $i$'s expected future payoffs if firm $i$ mimics another type in period $t$. The chosen actions affect which continuation payoff function is used through $\mathbf{z}_t$, but the firm's true type determines the firm's beliefs about $\theta_{i,t+1}$ and thus the firm's continuation value.

Let $\eta_j(a_{j,t}, s_{j,t}, \mu_{j,t})$ represent the belief that firm $i \neq j$ has about player $j$ in period $t$ after firm $j$ has made announcement $a_{j,t}$, given that firm $i \neq j$ began the period with beliefs $\mu_{j,t}$ and that it posits that player $j$ uses period strategy $s_{j,t}$. Then, for a continuation value function $\mathbf{V} = (V_1, \dots, V_I)$, define expected discounted payoffs for firm $i$ in period $t$, when firm $i$'s type is $\theta_{i,t}$, after announcements have been observed to be $\mathbf{a}_t$ and firm $i$ chooses actions $p_{i,t}, q_{i,t}$ following these announcements:

$$\begin{aligned}
&u_i(\mathbf{a}_t, p_{i,t}, q_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}, V_i) \\
&\quad = (p_{i,t} - \theta_{i,t}) \cdot \mathbb{E}_{\boldsymbol{\theta}_{-i,t}}\big[\varphi_i\big((p_{i,t}, \boldsymbol{\rho}_{-i,t}(\mathbf{a}_t, \boldsymbol{\theta}_{-i,t})); (q_{i,t}, \boldsymbol{\psi}_{-i,t}(\mathbf{a}_t, \boldsymbol{\theta}_{-i,t}))\big) \\
&\qquad\qquad |\boldsymbol{\eta}_{-i}(\mathbf{a}_{-i,t}, \mathbf{s}_{-i,t}, \boldsymbol{\mu}_{-i,t})\big] \\
&\qquad + \delta\mathbb{E}_{\theta_{i,t+1}, \boldsymbol{\theta}_{-i,t}}\big[V_i((a_{i,t}, p_{i,t}, q_{i,t}), \mathbf{s}_{-i,t}(\mathbf{a}_t, \boldsymbol{\theta}_{-i,t}))(\theta_{i,t+1}) \\
&\qquad\qquad |\boldsymbol{\eta}_{-i}(\mathbf{a}_{-i,t}, \mathbf{s}_{-i,t}, \boldsymbol{\mu}_{-i,t}), \theta_{i,t}\big].
\end{aligned}$$

As in Section 4.2 of the main paper, the following equality represents firm $i$'s expected payoffs in period $t$, before announcements are made, when firm $i$ has type $\theta_{i,t}$ and mimics type $\hat{\theta}_{i,t}$:

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

Then, following Cole and Kocherlakota (2001), we define a mapping $B : \mathcal{W} \to \mathcal{W}$ such that

$$\begin{aligned}
B(W) = \Big\{&(\mathbf{v}, \boldsymbol{\mu}) \in \mathcal{W} : \forall i, \exists s_i^* \in \mathcal{S}_i \text{ and } (\mathbf{V}, \mathbf{M}) : \mathcal{Z} \to W, \\
&\text{s.t. } \mathbf{M}(\mathbf{z}) \in \tilde{\mathbf{T}}(\boldsymbol{\mu}, \mathbf{s}^*, \mathbf{z}) \ \forall \mathbf{z} \in \mathcal{Z}, v_i(\theta_i) = \overline{u}_i(\theta_i, \theta_i, \mathbf{s}^*, \boldsymbol{\mu}_{-i}, V_i) \\
&\forall \theta_i \in \Theta_i, \text{ and (IC) } s_i^* \in \arg\max_{s_i \in \mathcal{S}_i} \overline{u}_i(\theta_i, \theta_i, (s_i, \mathbf{s}^*_{-i}), \boldsymbol{\mu}_{-i}, V_i)\Big\}.
\end{aligned} \tag{S2.1}$$

Note that in this definition, the requirement that $(\mathbf{V}, \mathbf{M}) : \mathcal{Z} \to W$ is restrictive, since $W$ is not a product set. The requirement ensures that the continuation value and belief functions are compatible: given the beliefs that arise given posited strategies and observations, $\mathbf{M}(\mathbf{z})$, only a subset of potential continuation value functions satisfy $(\mathbf{V}(\mathbf{z}), \mathbf{M}(\mathbf{z})) \in W$. Intuitively, today's actions reveal information about cost types that restrict the expected value of costs, and thus feasible payoffs, tomorrow.

Following Cole and Kocherlakota (2001), standard arguments can be adapted to show that the operator $B$ is monotone (where a set $A$ is larger than $B$ if $A \supseteq B$). Showing that the set of PPBE is the largest fixed point of the operator $B$ involves more work, because our model differs in a few respects from that of Cole and Kocherlakota. Cole and Kocherlakota's (2001) assumptions about the monitoring technology imply that $\tilde{\mathbf{T}}$ is single-valued, since it is impossible to observe outcomes $\mathbf{z}$ that are inconsistent with strategies $\mathbf{s}^*$. Our definition of $B$ has an additional degree of freedom, since $B(W)$ may include different elements that are supported using different off-equilibrium-path beliefs. Although this in itself does not pose a difficulty, the fact that we consider a hidden-information game does raise additional issues. In particular, beliefs are not continuous in strategies: in the limit as a separating strategy approaches pooling (for example, announcements are uninformative and the prices chosen by different types approach one another), the sequence of separating strategies induces very different beliefs than those induced by the limiting pooling strategy.

Two additional technical differences from the literature arise. First, the strategy space in each period is a compact set but is not finite. Second, there is a discontinuity in payoffs due to the Bertrand stage game: a firm can receive discretely higher payoffs by selecting a slightly different price. We believe that these differences can be addressed, but a full treatment is beyond the scope of this paper and is not necessary for our purposes.

The most useful insight for analyzing a particular class of PPBE is that if we can explicitly construct a set $W$ such that $W \subseteq B(W)$, this set must be a PPBE. Lemma S1 states this formally.

LEMMA S1: *Let $W^*$ be the set of PPBE type-contingent payoff functions and beliefs. For any compact set $W \subseteq \mathcal{W}$ such that $W \subseteq B(W)$, $W$ is a self-generating equilibrium set*: $W \subseteq W^*$.

This lemma follows by adapting the findings of Abreu, Pearce, and Stacchetti (1986) and Cole and Kocherlakota (2001) to our game; the extension is straightforward given the definitions.

Each PPBE is described by $\mathbf{w} = (\mathbf{v}, \boldsymbol{\mu}) \in W^*$, which is the type-contingent payoff function and the belief. Since each $\mathbf{w} \in W^*$ corresponds to a PPBE outcome, we will simply refer to $\mathbf{w} \in W^*$ as a PPBE. Further, since $\mathbf{w} \in W^*$ implies $\mathbf{w} \in B(W^*)$, each such equilibrium can be "decomposed" into the period strategies, $\mathbf{s}^*$, and the continuation belief and payoff mappings $(\mathbf{V}, \mathbf{M}) : \mathcal{Z} \to W$ that are guaranteed to exist by the definition of $B$. We rely heavily on this way of describing and analyzing equilibria below.

The incentive constraint in the definition of $B$ includes all deviations. We wish to separate the types of deviations into *on-schedule* deviations, whereby one type mimics another, and *off-schedule* deviations, where a type chooses an action that was not assigned to any type. Unlike the case of perfect persistence, here there is always a chance that types change from period to period, and so any kind of mimicking behavior constitutes an on-schedule deviation, even if the firm mimics different types at different points in time. The on-schedule constraint for firm $i$ can be written as

$$\text{(S2.2)} \qquad \overline{u}_i(\theta_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_t, V_i) \geq \overline{u}_i(\hat{\theta}_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}, V_i) \quad \forall \hat{\theta}_{i,t}, \theta_{i,t}.$$

The off-schedule constraint is written

$$\text{(S2.3)}\quad u_i\big(\boldsymbol{\alpha}_t(\boldsymbol{\theta}_t), \rho_{i,t}(\boldsymbol{\alpha}_t(\boldsymbol{\theta}_t), \theta_{i,t}), \psi_{i,t}(\boldsymbol{\alpha}_t(\boldsymbol{\theta}_t), \theta_{i,t}), \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}, V_i\big)$$
$$\geq u_i\big((a_{i,t}, \boldsymbol{\alpha}_{-i,t}(\boldsymbol{\theta}_{-i,t})), p_{i,t}, q_{i,t}, \theta_{i,t}; \mathbf{s}_t, \boldsymbol{\mu}_{-i,t}, V_i\big) \quad \text{for all } \theta_{i,t}, \boldsymbol{\theta}_{-i,t},$$
$$\text{and all}\quad (a_{i,t}, p_{i,t}, q_{i,t}) \notin \big\{(a'_{i,t}, p'_{i,t}, q'_{i,t}) : \exists \hat{\theta}_{i,t} \in \Theta_i \text{ s.t. } a'_{i,t} = \alpha_{i,t}(\hat{\theta}_{i,t}),$$
$$p'_{i,t} = \rho_{i,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), \hat{\theta}_{i,t}),$$
$$q'_{i,t} = \psi_{i,t}(\boldsymbol{\alpha}_t(\hat{\theta}_{i,t}, \boldsymbol{\theta}_{-i,t}), \hat{\theta}_{i,t})\big\}.$$

If, for all $i$, both of these constraints are satisfied for $\mathbf{s} = \mathbf{s}^*$, then $s_i^* \in \arg\max_{s_i \in \mathcal{S}_i} \overline{u}_i(\theta_i, \theta_i, (s_i, \mathbf{s}^*_{-i}), \boldsymbol{\mu}_{-i}, V_i)$.

### S2.1. *Applying Dynamic Programming Tools to the First-Best Scheme*

To verify that a particular first-best scheme is a PPBE, there are two steps. First, we construct a set $W^{fb} \subset \mathcal{W}$, which requires constructing the (state-contingent payoff functions, belief) pairs induced by the first-best scheme, as well as associated continuation value functions. Using this development, the second step is to define a set of incentive constraints and verify that they are satisfied, thus ensuring that strategies of the first-best scheme are indeed a PPBE.

The paper defines period strategies $\tilde{s}_i$ and continuation value functions $\tilde{\mathbf{V}}$ for each state, and it states a system of equations that yields the state-contingent payoff functions $\tilde{\mathbf{v}}(j, \boldsymbol{\theta}_{t-1}) \in \mathcal{V}$.

Recall that $v^{cs}$ and $v^r$ are the type-contingent payoff functions from the worst carrot–stick and best rigid-pricing equilibria, respectively. We consider a modified version of a carrot–stick scheme where players announce their types in each period. Since behavior does not depend on beliefs about opponents on or off of the equilibrium path in this scheme, firms are indifferent about their reports, and so the carrot–stick scheme with truthful reporting is a PPBE as long as the original carrot–stick scheme was a PPBE. For the case of $I = 2$, let the set of possible beliefs in a fully separating equilibrium be $\mathcal{M}^{FS} = \{\boldsymbol{\mu} \in \Delta\Theta^2 : \boldsymbol{\mu} = F(\cdot; \boldsymbol{\theta}_{t-1}) \text{ for some } \boldsymbol{\theta}_{t-1} \in \Theta^2\}$. Then define the following sets of (state-contingent payoff function–belief) pairs:

$$W^e = \{(\mathbf{v}, \boldsymbol{\mu}) \in \mathcal{W} : \exists \omega_j(\boldsymbol{\theta}_{t-1}) \in \Omega^e \text{ s.t. } \boldsymbol{\mu} = \mathbf{F}(\cdot; \boldsymbol{\theta}_{t-1})$$
$$\text{and } \mathbf{v} = \tilde{\mathbf{v}}(j, \boldsymbol{\theta}_{t-1})\},$$
$$W^p = \{(\mathbf{v}, \boldsymbol{\mu}) \in \mathcal{W} : \mathbf{v}(\boldsymbol{\theta}_{t-1}) = (v^{cs}, v^{cs})\ \forall \boldsymbol{\theta}_{t-1}, \boldsymbol{\mu} \in \mathcal{M}^{FS}\}$$
$$\cup \{(\mathbf{v}, \boldsymbol{\mu}) \in \mathcal{W} : \mathbf{v}(\boldsymbol{\theta}_{t-1}) = (v^r, v^r)\ \forall \boldsymbol{\theta}_{t-1}, \boldsymbol{\mu} \in \mathcal{M}^{FS}\}.$$

We wish to show that $W^{fb} = W^e \cup W^p$ is a self-generating set in the sense of Abreu, Pearce, and Stacchetti (1986), and as applied to our problem in

Lemma S1. Informally, we require that for each element $w$ of $W^{fb}$, we can find strategies $\mathbf{s}^*$ and a continuation value function $\mathbf{V}: \mathcal{Z} \to \mathcal{V}$, such that three conditions hold. First, we require feasibility: $(\mathbf{V}(\mathbf{z}_t), \mathbf{F}(\cdot; \mathbf{a}_t)) \in W^{fb}$, so that the continuation valuation function gives a state-contingent payoff function that is in $W^{fb}$ given the beliefs induced by the period's public outcomes. Second, we require *promise-keeping*: the strategies and continuation valuation functions deliver the promised state-contingent payoffs. Third, the strategies $\mathbf{s}^*$ must be best responses.

The following result formalizes the sufficient conditions required to verify that a first-best scheme is a PPBE. It provides an alternative to the approach outlined in the paper.

PROPOSITION S3: *Fix $I = 2$ and consider the two-type model with imperfect persistence, with primitives $\delta, r, L, H$, and $\mathbf{F}$, with $\delta \geq \delta_c$. Fix the specification of a first-best scheme $\tilde{g}, \tilde{q}_1$, and $\tilde{T}$, and consider the corresponding $\tilde{\mathbf{s}}$ and $\tilde{\mathbf{V}}$. If for each $(i, j, \boldsymbol{\theta}_{t-1}) \in \{1, 2\}^2 \times \Theta^2$, the on-schedule and off-schedule constraints, (S2.2) and (S2.3), hold when $\boldsymbol{\mu}_t = \mathbf{F}(\cdot; \boldsymbol{\theta}_{t-1})$, $\mathbf{V} = \tilde{\mathbf{V}}(j, \boldsymbol{\theta}_{t-1})$, and $\mathbf{s}^* = \tilde{\mathbf{s}}(j, \boldsymbol{\theta}_{t-1})$, then $W^{fb} \subseteq B(W^{fb})$ and $W^{fb}$ is a self-generating PPBE set that yields first-best profits in every period.*

PROOF: We established in the paper that $W^p$ is a self-generating equilibrium set. By Lemma S1, it remains to show that $\mathbf{w} \in W^e$ implies $\mathbf{w} \in B(W^{fb})$. By construction, each $\mathbf{w} \in W^e$ is associated with a $\omega_j(\boldsymbol{\theta}_{t-1}) \in \Omega^e$. Let $\mathbf{s}^* = \tilde{\mathbf{s}}(j, \boldsymbol{\theta}_{t-1})$ and $\mathbf{V} = \tilde{\mathbf{V}}(j, \boldsymbol{\theta}_{t-1})$, and note that $\mathbf{s}^* \in \mathcal{S}$ and, for all $\mathbf{z}$, $(\mathbf{V}(\mathbf{z}), \mathbf{F}(\cdot; \mathbf{a})) \in W^{fb}$, so that feasibility is satisfied. Further, letting $v_i = \tilde{v}_i(j, \boldsymbol{\theta}_{t-1})$ for each $i$, it follows by definition of $\tilde{v}_i$ that promise-keeping holds: $v_i(\theta_i) = \overline{u}_i(\theta_i, \theta_i, \mathbf{s}, \boldsymbol{\mu}_{-i}, V_i)$. Finally, if (S2.2) and (S2.3) hold with these definitions, $\mathbf{s}_i^*$ is a best response to $\mathbf{s}_{-i}^*$ for each $i$. Thus, $\mathbf{w} \in B(W^{fb})$, as desired. *Q.E.D.*

## S3. PRODUCTIVE EFFICIENCY WITH PERFECT PERSISTENCE

This section establishes conditions under which equilibria with productive efficiency exist when cost types are perfectly persistent. We focus on the case of two firms and then discuss the multi-firm extension. The equilibrium we construct requires a severe form of punishment in the case of an off-schedule deviation. We also discuss the use of severe punishments as a means of supporting the best rigid-pricing equilibrium.

### S3.1. *Separating Equilibria With Productive Efficiency*

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

$$U_i(\theta_{i,1}, \theta_{i,1}) = U_i(\overline{\theta}, \overline{\theta}) + \int_{\tilde{\theta}_i=\theta_{i,1}}^{\overline{\theta}} \breve{M}_i(\tilde{\theta}_i)\, d\tilde{\theta}_i = \frac{1}{1-\delta} \int_{\tilde{\theta}_i=\theta_{i,1}}^{\overline{\theta}} (1 - F_0(\tilde{\theta}_i))^{I-1}\, d\tilde{\theta}_i, \tag{S3.1}$$

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

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

PRODUCTIVE EFFICIENCY SCHEME: This is a set of strategies such that, in each period, announcements are uninformative and market-share proposals are not binding ($q_{i,t} \geq 1$). The first-period pricing strategy of firm $i$ is denoted $\rho_{i,1}(\theta_{i,1})$, and is strictly increasing and symmetric across firms. Each firm infers the other firm's cost once first-period prices are observed. Let $\theta_w$ and $\theta_l$ denote, respectively, the inferred cost of the "winner" (the lower-cost firm) and "loser" (the higher-cost firm) in the period-1 pricing contest. Each firm adopts a stationary price along the equilibrium path in periods $t > 1$. Let $\beta(\theta_w, \theta_l)$ denote the price selected by the winner in periods $t > 1$ and suppose that the loser charges $\varepsilon$ more. We restrict attention to $\beta(\theta_w, \theta_l) \in [\theta_w, \theta_l]$.[2] In all periods,

[1] To establish these relationships, simply compare the profits of the two equilibria for each type and note that the difference is strictly convex for all $\theta \in [\underline{\theta}, \overline{\theta})$. The distributional condition holds (with equality) for the uniform distribution, for example.

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

any off-schedule deviation induces the carrot–stick belief threat punishment, as described below.

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

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

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

We next establish that a productive efficiency equilibrium exists if two conditions hold. The first condition is that

$$\text{(S3.2)} \qquad \inf_{\theta'_{i,1} > \theta''_{i,1}} \frac{f_0(\theta'_{i,1})}{f_0(\theta''_{i,1})} > \frac{\delta}{1-\delta(1-\delta)} \quad \text{and} \quad \inf_{\theta'_{i,1} < \theta''_{i,1}} \frac{f_0(\theta'_{i,1})}{f_0(\theta''_{i,1})} > \delta;$$

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

$$\text{(S3.3)} \qquad \frac{2f_0(\theta_{i,1})}{(1-F_0(\theta_{i,1}))^2} \int_{\tilde{\theta}_i=\theta_{i,1}}^{\overline{\theta}} (1-F_0(\tilde{\theta}_i))\, d\tilde{\theta}_i > \delta.$$

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

PROPOSITION S4: *Consider Model* 2 *and suppose* $I = 2$. *If* (S3.2) *and* (S3.3) *are satisfied*, *then there exists a productive efficiency equilibrium*. *Specifically*, *in the first period*, *each firm i uses the strategy*

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

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

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

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

PROOF: In a separating equilibrium with productive efficiency, $\rho_{i,1}(\theta_{i,1})$ is strictly increasing and symmetric across firms, $\beta(\theta_w, \theta_l) \in [\theta_w, \theta_l]$, and thus $\beta(\theta_w, \theta_w) = \theta_w$. Let $\rho(\theta_{i,1})$ denote the symmetric first-period pricing function. Suppose further that $\beta(\theta_w, \theta_l)$ is monotonic, in that it is strictly increasing in each argument.

Fix $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$. First, suppose firm $i$ engages in a downward deviation by mimicking the first-period price of $\hat{\theta}_{i,1}$ slightly below $\theta_{i,1}$. Consider the types $\tilde{\theta}_{j,1}$ for the rival firm $j$ such that the rival loses (i.e., $\tilde{\theta}_{j,1} > \hat{\theta}_{i,1}$) and the winner chooses a future price that exceeds $\theta_{i,1}$ (i.e., $\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1}) > \theta_{i,1}$). Observe that $\beta(\hat{\theta}_{i,1}, \theta_{i,1}) < \theta_{i,1}$; further, $\beta(\theta_{i,1}, \overline{\theta}) > \beta(\theta_{i,1}, \theta_{i,1}) = \theta_{i,1}$, and so for $\hat{\theta}_{i,1}$ slightly below $\theta_{i,1}$, we have that $\beta(\hat{\theta}_{i,1}, \overline{\theta}) > \theta_{i,1}$. We conclude that there exists a unique value $\theta_c(\hat{\theta}_{i,1}, \theta_{i,1}) \in (\theta_{i,1}, \overline{\theta})$ that satisfies $\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1}) = \theta_{i,1}$. Second, suppose firm $i$ engages in an upward deviation by mimicking the first-period price of $\hat{\theta}_{i,1}$ slightly above $\theta_{i,1}$. Consider the types $\tilde{\theta}_{j,1}$ for the rival such that the rival wins (i.e., $\tilde{\theta}_{j,1} < \hat{\theta}_{i,1}$) and chooses a future price that exceeds $\theta_{i,1}$ (i.e., $\beta(\tilde{\theta}_{j,1}, \hat{\theta}_{i,1}) > \theta_{i,1}$). Observe that $\beta(\theta_{i,1}, \hat{\theta}_{i,1}) > \theta_{i,1}$; further, $\beta(\underline{\theta}, \hat{\theta}_{i,1}) < \theta_{i,1}$ for $\hat{\theta}_{i,1}$ sufficiently little above $\theta_{i,1}$. We conclude that there exists a unique value $\theta_b(\hat{\theta}_{i,1}, \theta_{i,1}) \in (\underline{\theta}, \theta_{i,1})$ that satisfies $\beta(\tilde{\theta}_{j,1}, \hat{\theta}_{i,1}) = \theta_{i,1}$.

Consider the following downward deviation: Firm $i$ with type $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$ mimics $\hat{\theta}_{i,1}$ slightly below $\theta_{i,1}$ (i.e., chooses $\rho(\hat{\theta}_{i,1}) < \rho(\theta_{i,1})$) and then (i) if $\tilde{\theta}_{j,1} < \hat{\theta}_{i,1}$, firm $i$ makes no first-period sale and exits (e.g., prices above $r$) in all future periods; (ii) if $\tilde{\theta}_{j,1} \in (\hat{\theta}_{i,1}, \theta_c)$, firm $i$ makes the first-period sale and exits (e.g., prices above $r$) in all future periods; and (iii) if $\tilde{\theta}_{j,1} > \theta_c$, firm $i$ makes the first-period sale and mimics thereafter the equilibrium pricing behavior of type

$\hat{\theta}_{i,1}$ (i.e., sets $\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1})$ in all future periods). As $\hat{\theta}_{i,1} \uparrow \theta_{i,1}$, the deviating firm $i$'s payoff approaches that which it earns in the putative equilibrium. Thus, a necessary feature of a productive efficient equilibrium is that firm $i$ does better by announcing $\hat{\theta}_{i,1} = \theta_{i,1}$ than any other $\hat{\theta}_{i,1} < \theta_{i,1}$, given the associated strategies described in (i)–(iii) above.

Under (i)–(iii), the profit from a downward deviation is defined as

$$\text{(S3.4)} \qquad \pi^D(\hat{\theta}_{i,1}, \theta_{i,1}) = [\rho(\hat{\theta}_{i,1}) - \theta_{i,1}][1 - F_0(\hat{\theta}_{i,1})] + \frac{\delta}{1-\delta}\int_{\theta_c}^{\overline{\theta}} [\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1}) - \theta_{i,1}]\, dF_0(\tilde{\theta}_{j,1}).$$

LEMMA S2: *For any $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$, if derivatives are evaluated as $\hat{\theta}_{i,1} \uparrow \theta_{i,1}$,*

$$\text{(S3.5)} \qquad \pi^D_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) = [\rho(\theta_{i,1}) - \theta_{i,1}][-F_0'(\theta_{i,1})] + [1 - F_0(\theta_{i,1})]\rho'(\theta_{i,1}) + \frac{\delta}{1-\delta}\int_{\theta_{i,1}}^{\overline{\theta}} \beta_{\theta_w}(\theta_{i,1}, \tilde{\theta}_{j,1})\, dF_0(\tilde{\theta}_{j,1})$$

$$\text{(S3.6)} \qquad \pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1}) = F_0'(\theta_{i,1})\left[1 - \frac{\delta}{1-\delta}\frac{\beta_{\theta_w}(\theta_{i,1}, \theta_{i,1})}{\beta_{\theta_l}(\theta_{i,1}, \theta_{i,1})}\right].$$

PROOF OF LEMMA S2: Using (S3.4) and the definition of $\theta_c$, we find that

$$\pi^D_{\hat{\theta}_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) = [\rho(\hat{\theta}_{i,1}) - \theta_{i,1}][-F_0'(\hat{\theta}_{i,1})] + [1 - F_0(\hat{\theta}_{i,1})]\rho'(\hat{\theta}_{i,1}) + \frac{\delta}{1-\delta}\int_{\theta_c}^{\overline{\theta}} \beta_{\theta_w}(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1})\, dF_0(\tilde{\theta}_{j,1}).$$

Differentiating with respect to $\theta_{i,1}$ and using $\partial\theta_c/\partial\theta_{i,1} = 1/\beta_{\theta_l}(\hat{\theta}_{i,1}, \theta_c)$, we obtain

$$\text{(S3.7)} \qquad \pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) = F_0'(\hat{\theta}_{i,1}) - \frac{\delta}{1-\delta}\frac{\beta_{\theta_w}(\hat{\theta}_{i,1}, \theta_c)}{\beta_{\theta_l}(\hat{\theta}_{i,1}, \theta_c)}F_0'(\theta_c).$$

Finally, as $\hat{\theta}_{i,1} \uparrow \theta_{i,1}$, we observe that $\theta_c \downarrow \theta_{i,1}$, and so we obtain the desired expressions. *Q.E.D.*

Consider now the following upward deviation: Firm $i$ with type $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$ mimics $\hat{\theta}_{i,1}$ slightly above $\theta_{i,1}$ (i.e., chooses $\rho(\hat{\theta}_{i,1}) > \rho(\theta_{i,1})$) and then (i) if $\tilde{\theta}_{j,1} < \theta_b$, firm $i$ makes no first-period sale and exits (e.g., prices above $r$) in all future periods; (ii) if $\tilde{\theta}_{j,1} \in (\theta_b, \hat{\theta}_{i,1})$, firm $i$ makes no first-period sale, undercuts the rival's price $\beta(\tilde{\theta}_{j,1}, \hat{\theta}_{i,1})$ in the second period, and then exits (e.g., prices above $r$) in all future periods; and (iii) if $\tilde{\theta}_{j,1} > \hat{\theta}_{i,1}$, firm $i$ makes the first-period sale and mimics thereafter the equilibrium pricing behavior of type $\hat{\theta}_{i,1}$ (i.e., sets $\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1})$ in all future periods). As $\hat{\theta}_{i,1} \downarrow \theta_{i,1}$, the deviating firm $i$'s payoff approaches that which it earns in the putative equilibrium. Thus, a necessary feature of a productive efficient equilibrium is that firm $i$ does better by announcing $\hat{\theta}_{i,1} = \theta_{i,1}$ than any other $\hat{\theta}_{i,1} > \theta_{i,1}$, given the associated strategies described in (i)–(iii) above.

Under (i)–(iii), the profit from an upward deviation is defined as

$$\begin{aligned} \text{(S3.8)} \qquad \pi^U(\hat{\theta}_{i,1}, \theta_{i,1}) &= [\rho(\hat{\theta}_{i,1}) - \theta_{i,1}][1 - F_0(\hat{\theta}_{i,1})] \\ &\quad + \delta \int_{\theta_b}^{\hat{\theta}_{i,1}} [\beta(\tilde{\theta}_{j,1}, \hat{\theta}_{i,1}) - \theta_{i,1}]\, dF_0(\tilde{\theta}_{j,1}) \\ &\quad + \frac{\delta}{1-\delta} \int_{\hat{\theta}_{i,1}}^{\overline{\theta}} [\beta(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1}) - \theta_{i,1}]\, dF_0(\tilde{\theta}_{j,1}). \end{aligned}$$

LEMMA S3: *For any $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$, if derivatives are evaluated as $\hat{\theta}_{i,1} \downarrow \theta_{i,1}$,*

$$\begin{aligned} \text{(S3.9)} \qquad \pi^U_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) &= [\rho(\theta_{i,1}) - \theta_{i,1}][-F_0'(\theta_{i,1})] + [1 - F_0(\theta_{i,1})]\rho'(\theta_{i,1}) \\ &\quad + \frac{\delta}{1-\delta} \int_{\theta_{i,1}}^{\overline{\theta}} \beta_{\theta_w}(\theta_{i,1}, \tilde{\theta}_{j,1})\, dF_0(\tilde{\theta}_{j,1}), \end{aligned}$$

$$\text{(S3.10)} \qquad \pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1}) = F_0'(\theta_{i,1})\left[1 + \delta\left(\frac{\delta}{1-\delta} - \frac{\beta_{\theta_l}(\theta_{i,1}, \theta_{i,1})}{\beta_{\theta_w}(\theta_{i,1}, \theta_{i,1})}\right)\right].$$

PROOF OF LEMMA S3: Using (S3.8), the definition of $\theta_b$, and $\beta(\hat{\theta}_{i,1}, \hat{\theta}_{i,1}) = \hat{\theta}_{i,1}$, we find that

$$\begin{aligned} &\pi^U_{\hat{\theta}_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) \\ &\quad = [\rho(\hat{\theta}_{i,1}) - \theta_{i,1}][-F_0'(\hat{\theta}_{i,1})] + [1 - F_0(\hat{\theta}_{i,1})]\rho'(\hat{\theta}_{i,1}) \\ &\qquad - [\hat{\theta}_{i,1} - \theta_{i,1}]F_0'(\hat{\theta}_{i,1})\frac{\delta^2}{1-\delta} + \delta \int_{\theta_b}^{\hat{\theta}_{i,1}} \beta_{\theta_l}(\tilde{\theta}_{j,1}, \hat{\theta}_{i,1})\, dF_0(\tilde{\theta}_{j,1}) \\ &\qquad + \frac{\delta}{1-\delta} \int_{\hat{\theta}_{i,1}}^{\overline{\theta}} \beta_{\theta_w}(\hat{\theta}_{i,1}, \tilde{\theta}_{j,1})\, dF_0(\tilde{\theta}_{j,1}). \end{aligned}$$

Differentiating with respect to $\theta_{i,1}$ and using $\partial\theta_b/\partial\theta_{i,1} = 1/\beta_{\theta_w}(\theta_b, \hat{\theta}_{i,1})$, we obtain

$$\text{(S3.11)} \qquad \pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) = F_0'(\hat{\theta}_{i,1})\left[1 + \frac{\delta^2}{1-\delta}\right] - \delta\frac{\beta_{\theta_l}(\theta_b, \hat{\theta}_{i,1})}{\beta_{\theta_w}(\theta_b, \hat{\theta}_{i,1})}F_0'(\theta_b).$$

Finally, as $\hat{\theta}_{i,1} \downarrow \theta_{i,1}$, we observe that $\theta_b \uparrow \theta_{i,1}$ and so we obtain the desired expressions. *Q.E.D.*

We now report two corollaries:

COROLLARY S1: *For any* $\theta_{i,1} \in (\underline{\theta}, \overline{\theta})$,

$$\begin{aligned}
\pi^D_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) &= \pi^U_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) \\
&= [\rho(\theta_{i,1}) - \theta_{i,1}][-F_0'(\theta_{i,1})] + [1 - F_0(\theta_{i,1})]\rho'(\theta_{i,1}) \\
&\quad + \frac{\delta}{1-\delta}\int_{\theta_{i,1}}^{\overline{\theta}} \beta_{\theta_w}(\theta_{i,1}, \tilde{\theta}_{j,1})\, dF_0(\tilde{\theta}_{j,1}).
\end{aligned}$$

COROLLARY S2: *Suppose that*

$$\frac{\beta_{\theta_w}(\theta_w, \theta_l)}{\beta_{\theta_l}(\theta_w, \theta_l)} = 1 - \delta. \tag{S3.12}$$

*Then*

$$\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1}) = \pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1}) = F_0'(\theta_{i,1})[1-\delta] > 0, \tag{S3.13}$$

$$\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) = F_0'(\hat{\theta}_{i,1}) - \delta F_0'(\theta_c), \tag{S3.14}$$

$$\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) = F_0'(\hat{\theta}_{i,1})\left[1 + \frac{\delta^2}{1-\delta}\right] - \frac{\delta}{1-\delta}F_0'(\theta_b). \tag{S3.15}$$

The corollaries follow directly from Lemmas S2 and S3 and expressions (S3.7) and (S3.11). The latter corollary motivates the specification for $\beta(\theta_w, \theta_l)$ in Proposition S4, which satisfies monotonicity and (S3.12).

We now confirm that the pricing functions specified in Proposition S4 satisfy local incentive compatibility with respect to our two deviation candidates. Define

$$\pi(\hat{\theta}_{i,1}, \theta_{i,1}) = 1_{\{\hat{\theta}_{i,1} \le \theta_{i,1}\}}\pi^D(\hat{\theta}_{i,1}, \theta_{i,1}) + 1_{\{\hat{\theta}_{i,1} > \theta_{i,1}\}}\pi^U(\hat{\theta}_{i,1}, \theta_{i,1}). \tag{S3.16}$$

Since $\pi^D_{\hat{\theta}_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ and $\pi^U_{\hat{\theta}_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ exist everywhere, and $\pi^D_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) = \pi^U_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1})$ (as shown in Corollary S1), it follows that $\pi_{\hat{\theta}_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ exists everywhere. Imposing the specification for $\beta(\theta_w, \theta_l)$ in Proposition S4, we may use Corollary S1 to find that local incentive compatibility holds if and only if $\pi^D_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) = \pi^U_{\hat{\theta}_{i,1}}(\theta_{i,1}, \theta_{i,1}) = 0$ or, equivalently,

$$\begin{aligned}
&[\rho(\theta_{i,1}) - \theta_{i,1}][-F_0'(\theta_{i,1})] \\
&\quad + [1 - F_0(\theta_{i,1})]\rho'(\theta_{i,1}) + \frac{\delta}{2-\delta}[1 - F_0(\theta_{i,1})] = 0.
\end{aligned} \tag{S3.17}$$

Thus, we can characterize the first-period pricing function that achieves local incentive compatibility by

$$\text{(S3.18)}\quad \rho(\overline{\theta}) = \overline{\theta},$$

$$\text{(S3.19)}\quad \rho'(\theta_{i,1}) = \frac{F_0'(\theta_{i,1})}{1 - F_0(\theta_{i,1})}(\rho(\theta_{i,1}) - \theta_{i,1}) - \frac{\delta}{2-\delta}.$$

It is now straightforward to verify that the first-period pricing function specified in Proposition S4 solves (S3.18) and (S3.19).

We next confirm that the specified pricing functions satisfy global incentive compatibility with respect to our two deviation candidates. As established in Corollary S2, $\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ and $\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ exist everywhere and $\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1}) = \pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\theta_{i,1}, \theta_{i,1})$ for the $\beta(\theta_w, \theta_l)$ function that we specify. It follows that $\pi_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ exists everywhere as well. Now consider the sign of $\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ for $\hat{\theta}_{i,1} < \theta_{i,1}$. Using (S3.14), we see that $\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ is positive if $f_0(\hat{\theta}_{i,1})/f_0(\theta_c) > \delta$. Since $\theta_c > \hat{\theta}_{i,1}$, we may draw the following conclusion: given $\beta(\theta_w, \theta_l)$ is specified as in Proposition S4, for every $\hat{\theta}_{i,1} < \theta_{i,1}$, $\pi^D_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) > 0$ if the second inequality in (S3.2) holds. Next, consider the sign of $\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ for $\hat{\theta}_{i,1} > \theta_{i,1}$. Using (S3.15), we see that $\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ is positive if $f_0(\hat{\theta}_{i,1})/f_0(\theta_b) > \delta/[1-\delta(1-\delta)]$. Since $\theta_b < \hat{\theta}_{i,1}$, we may draw the following conclusion: given $\beta(\theta_w, \theta_l)$ is specified as in Proposition S4, for every $\hat{\theta}_{i,1} > \theta_{i,1}$, $\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) > 0$ if the first inequality in (S3.2) holds. Thus, under (S3.2), $\pi_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1})$ is positive everywhere. Then standard arguments can be used to show that local incentive compatibility implies global incentive compatibility.[3]

Next, we determine conditions under which the first-period pricing function is strictly increasing. Differentiating the first-period pricing function specified in Proposition S4, we can confirm that $\rho'(\theta_{i,1}) > 0$ if (S3.3) holds.

Guided by the foregoing analysis, we can specify a separating equilibrium with productive efficiency when (S3.2) and (S3.3) hold. Along the equilibrium path, firms use the pricing strategies specified in Proposition S4. Following any history where an off-schedule deviation has been observed, the carrot–stick belief threat punishment is induced. This punishment is characterized in Proposition S5, and it ensures that a firm that undertakes an off-schedule deviation makes approximately zero expected profit over the subsequent periods. In the

[3]For $\hat{\theta}_{i,1} < \theta_{i,1}$, observe that $\pi(\hat{\theta}_{i,1}, \theta_{i,1}) - \pi(\theta_{i,1}, \theta_{i,1}) = \pi^D(\hat{\theta}_{i,1}, \theta_{i,1}) - \pi^D(\theta_{i,1}, \theta_{i,1}) < 0$, where the inequality follows from standard arguments, given that $\pi^D(\hat{\theta}_{i,1}, \theta_{i,1})$ satisfies local incentive compatibility and positive cross partials. For $\hat{\theta}_{i,1} > \theta_{i,1}$, the same argument applies with $\pi^U$ replacing $\pi^D$.

event that firm $i$ undertakes an on-schedule deviation in period 1, we specify that firm $i$'s subsequent behavior is determined as specified in the downward and upward deviation candidates discussed above.

To complete the proof of Proposition S4, we now confirm that no deviation is attractive. Clearly, no firm would gain by taking an off-schedule deviation in the first period (i.e., by deviating outside of the range of the first-period pricing function). Likewise, if a firm did not deviate in the first period, then it would not gain by taking an off-schedule deviation in a later period. A losing firm would clearly not gain from undercutting $\beta(\theta_w, \theta_l)$; and a winning firm would not gain from raising price above $\beta(\theta_w, \theta_l)$, since the immediate gain is approximately zero (the future price of the losing firm is $\beta(\theta_w, \theta_l) + \epsilon$) and the induced subsequent profits are also approximately zero. Next, suppose that firm $i$ took an on-schedule deviation in the first period and consider its optimal play in subsequent periods. Under our specification, if firm $i$ takes an off-schedule deviation in a later period, then firm $j$ is induced to follow the carrot–stick belief punishment thereafter. Thus, if firm $i$ takes an on-schedule deviation in period 1, then it can do no better than to follow the behavior prescribed by the downward and upward deviation candidates discussed above in periods 2 and later. This observation, combined with our work above, ensures as well that firm $i$ does not gain from taking an on-schedule deviation in period 1. *Q.E.D.*

Conditions (S3.2) and (S3.3) are satisfied in a rich parameter space; however, when they are not satisfied, a productive efficiency equilibrium may fail to exist. Intuitively, the highest-cost type ($\overline{\theta}$) gets no future profit and thus prices at cost in the first period. All other types, however, distort their first-period prices upward, in an attempt to signal higher costs and thereby secure a higher future price. If firms are very patient, the benefit of a higher future price is significant, and greater distortions in the first-period price are incurred. It is then possible that higher-cost types may price above $\overline{\theta}$ and thus above the first-period price of the highest-cost type. This implies a nonmonotonicity in the first-period pricing function, in contradiction to the hypothesis of a separating equilibrium.

We conclude that separating equilibria with productive efficiency exist under certain conditions. Such equilibria are characterized by strategic signaling in the first period. They thus represent the Bertrand counterpart to the separating equilibria constructed by Mailath (1989) for a two-period model with differentiated products and perfectly persistent cost types.

### S3.2. *More Than Two Firms*

It is straightforward to generalize the description of the productive efficiency scheme to more than two firms. However, stating sufficient conditions for the scheme to be a PPBE becomes more complex. The pricing strategies can be written

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

Let $\theta_w = \min_{i\in I}\theta_{i,1}$, while $\theta_l = \max_{A\in 2^I,|A|=I-1}\min_{i\in A}\theta_{i,1}$. If firm $i$ is the low-cost firm in period 1, then for all $t>1$, firm $i$ sets price

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

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

The sufficient conditions for global incentive compatibility from the $I=2$ case generalize as

$$\text{(S3.20)}\qquad \inf_{\theta'_{i,1}>\theta''_{i,1}} \frac{(1-F_0(\theta'_{i,1}))^{I-2}f_0(\theta'_{i,1})}{(1-F_0(\theta''_{i,1}))^{I-2}f_0(\theta''_{i,1})} > \frac{\delta}{1-\delta(1-\delta)} \quad\text{and}$$

$$\inf_{\theta'_{i,1}<\theta''_{i,1}} \frac{(1-F_0(\theta'_{i,1}))^{I-2}f_0(\theta'_{i,1})}{(1-F_0(\theta''_{i,1}))^{I-2}f_0(\theta''_{i,1})} > \delta.$$

The second condition is that $\delta$ is small enough that, for all $\theta_{i,1}$,

$$\text{(S3.21)}\qquad \frac{2(I-1)f_0(\theta_{i,1})}{(1-F_0(\theta_{i,1}))^{I}}\int_{\tilde{\theta}_i=\theta_{i,1}}^{\overline{\theta}} (1-F_0(\tilde{\theta}_i))^{I-1}\,d\tilde{\theta}_i > \delta.$$

However, the first condition in (S3.20) will fail when $I>2$ (let $\theta'_{i,1} = \overline{\theta}$). In the proof of Proposition S4, the condition was used to establish global incentive compatibility for upward deviations: that is, we used it to show that type $\theta_{i,1}$ does not want to mimic type $\hat{\theta}_{i,1} > \theta_{i,1}$. However, the condition is stronger than what is needed: if $\pi^U(\hat{\theta}_{i,1}\theta_{i,1})$ is the payoff from such mimicry, the condition guarantees that $\pi^U_{\hat{\theta}_{i,1}\theta_{i,1}}(\hat{\theta}_{i,1}, \theta_{i,1}) \geq 0$ for all $\hat{\theta}_{i,1} > \theta_{i,1}$. What is necessary is that $\pi^U(\hat{\theta}_{i,1}, \theta_{i,1}) \leq \pi^U(\theta_{i,1}, \theta_{i,1})$ for all $\hat{\theta}_{i,1} > \theta_{i,1}$. This can be verified directly for a particular distribution of types; numerical calculations indicate that for $F_0$ uniform and $I>2$, the critical discount factor *below* which global incentive compatibility holds is greater than zero, but diminishes rapidly with $I$.

### S3.3. *Belief Threat Punishment*

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

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

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

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

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

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

The (carrot–stick) belief threat punishment implies a new critical discount factor for the best rigid-pricing equilibrium, as is stated formally in the following proposition.

PROPOSITION S5: *Consider Model* 2 *and suppose* $\delta > (I-1)/I$. *Then there exists a best rigid-pricing equilibrium. If firm i deviates, the continuation entails a carrot–stick belief threat punishment, and so firms* $j \neq i$ *price at* $\underline{\theta} + 2\varepsilon$ *in subsequent periods and firm i prices above* $\underline{\theta} + 2\varepsilon$ *unless its cost type is less than* $\underline{\theta} + 2\varepsilon$.

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

$$\text{(S3.22)} \quad \frac{r-\theta_{i,1}}{I}\frac{1}{1-\delta} \geq r - \theta_{i,1} + \delta\xi(\theta_{i,1}).$$

Rewriting, we obtain $((I-1)(r-\theta_{i,1}) + I\delta\xi(\theta_{i,1}))/(I(r-\theta_{i,1}) + I\delta\xi(\theta_{i,1})) \leq \delta$. The left-hand side is increasing in $\xi(\theta_{i,1})$. Thus, for $\chi$ sufficiently small, $\xi(\theta_{i,1})$ is arbitrarily close to zero for all $\theta_{i,1}$ and we are sure to satisfy (S3.22) if $\delta > (I-1)/I$. *Q.E.D.*

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

While the equilibrium of Proposition S5 entails undominated strategies, one may object that the nondeviating firms might relinquish their worst-case beliefs after a deviation if the deviant firm consistently did not price at $\underline{\theta} + \varepsilon$. Our specification requires a dogged pessimism: even if the deviant firm $i$ has not priced at $\underline{\theta} + \varepsilon$ yet, each firm $j \neq i$ remains sure that firm $i$ will do so tomorrow. Standard refinements also do not eliminate this equilibrium. The belief threat punishment as stated, however, is not robust to the possibility of imperfect persistence.

*Dept. of Economics, Harvard University, 1875 Cambridge Street, Cambridge, MA 02138, U.S.A. and NBER; athey@fas.harvard.edu*

*and*

*Dept. of Economics and School of Business, Columbia University, 420 West 118th Street, New York, NY 10027, U.S.A. and NBER; kwb8@columbia.edu.*

*Manuscript received August, 2004; final revision received November, 2007.*

## REFERENCES

ABREU, D., D. PEARCE, AND E. STACCHETTI (1986): "Optimal Cartel Equilibria With Imperfect Monitoring," *Journal of Economic Theory*, 39, 251–269. [11,13,14]

——— (1990): "Toward a Theory of Discounted Repeated Games With Imperfect Monitoring," *Econometrica*, 58, 1041–1063. [11]

ATHEY, S., A. ATKESON, AND P. KEHOE (2005): "The Optimal Degree of Monetary Policy Discretion," *Econometrica*, 73, 1431–1476. [5,6]

COLE, H., AND N. KOCHERLAKOTA (2001): "Dynamic Games With Hidden Actions and Hidden States," *Journal of Economic Theory*, 98, 114–126. [11-13]

MAILATH, G. (1989): "Simultaneous Signaling in an Oligopoly Model," *Quarterly Journal of Economics*, 104, 417–427. [23]