

Wolitzky, Alexander

Article

Dynamic monopoly with relational incentives

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Wolitzky, Alexander (2010) : Dynamic monopoly with relational incentives, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 5, Iss. 3, pp. 479-518,
<https://doi.org/10.3982/TE701>

This Version is available at:

<https://hdl.handle.net/10419/150144>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>

Dynamic monopoly with relational incentives

ALEXANDER WOLITZKY

Department of Economics, Massachusetts Institute of Technology

This paper studies the price-setting problem of a monopoly that in each time period has the option of failing to deliver its good after receiving payment. The monopoly may be induced to deliver the good if consumers expect that the monopoly will not deliver in the future if it does not deliver today. If the good is nondurable and consumers are anonymous, the monopoly's optimal strategy is to set a price equal to the static monopoly price each period if the discount factor is high enough, and otherwise to set the lowest price at which it can credibly promise to deliver the good. If the good is durable, we derive an intuitive lower bound on the monopoly's optimal profit for any discount factor and show that it converges to the optimal static monopoly profit as the discount factor converges to 1, in contrast to the Coase conjecture. We also show that rationing the good is never optimal for the monopoly if there is an efficient resale market and that the best equilibrium in which the monopoly always delivers involves a strictly decreasing price path that asymptotes to a level strictly above the ratio of the monopoly's marginal cost to the discount factor.

KEYWORDS. Coase conjecture, durable goods, monopoly pricing, nondurable goods, rationing, relational incentives.

JEL CLASSIFICATION. C70, D42, L12.

1. INTRODUCTION

The possibility of trade is often threatened by the possibility of opportunism. For example, a consumer who purchases a good from an online retailer must trust that the good will actually be delivered—as taking legal action in the case of nondelivery would be very costly—and must also believe that the retailer is not about to cut its price dramatically. Fortunately, long term incentives can mitigate the risk of opportunistic behavior: in the above example, the retailer may both deliver the good and keep prices high to preserve its standing with its consumers, even if it has no fear of the legal consequences of nondelivery. In particular, either failing to deliver the good or cutting prices may lead consumers to believe that the firm will not deliver the good in the future, as either of these actions could be interpreted as an indication that the firm is trying to

Alexander Wolitzky: wolitzky@mit.edu

I am grateful to Glenn Ellison for numerous helpful discussions and comments on several drafts, as well as continual advice and support; to Daron Acemoglu, Bengt Holmström, Parag Pathak, Michael Powell, Iván Werning, and Muhamet Yildiz for helpful discussions and comments on earlier drafts; to the co-editor and to two anonymous referees for helpful comments; and to the NSF for financial support.

Copyright © 2010 Alexander Wolitzky. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](http://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: 10.3982/TE701

maximize its short run profits and then quit the market.¹ This reasoning suggests that a seller who is tempted to fail to deliver her product may still do quite well if the future is sufficiently important. This paper studies this idea in the context of both nondurable- and durable-goods monopoly, focusing primarily on the more involved durable-goods case.

The above intuition contrasts starkly with the Coase conjecture (Coase 1972) that a patient durable-goods seller who cannot commit to future prices earns little profit. As we will see, the Coase conjecture relies on the assumption that the seller is committed to delivering the good at her quoted price. In particular, the Coasian temptation to cut prices is absent when a price cut leads to a continuation equilibrium in which no consumers make purchases (expecting nondelivery) and the seller never delivers (expecting no future purchases).² Thus, even if the seller cannot commit to a price path, she can still earn high profits if she is not committed to delivering the good either.³ This suggests that the Coase conjecture may not apply to any institutional setting: If the seller can legally commit herself to both a price path and delivery of the good, she should do so. If she can legally commit herself to delivery, but not to a price path, she should not.^{4,5}

Throughout, we consider an infinitely repeated interaction between a monopoly seller and a continuum of buyers, where, in every period, the seller first sets a price, consumers then choose whether to pay, and finally the seller chooses whether to deliver the good to each consumer. All actions are perfectly observable. If the good is nondurable and consumers are anonymous, we completely characterize the optimal perfect Bayesian equilibrium of this game for the seller: if the seller is sufficiently patient, she sets the static monopoly price each period and delivers the good to all consumers who purchase, while if she is less patient, she charges a higher price so as to reduce the quantity demanded and thereby reduce her temptation to fail to deliver.⁶

¹An alternative story, which we discuss below, is that the retailer is contractually obligated to deliver something, but that the quality of the good it delivers is unverifiable. In this case, it is natural to think that a price cut may suggest to consumers that the retailer intends to deliver a low-quality good. For an example of an online market in which lower-priced goods seem to be of extremely low quality, see Ellison and Ellison (2009).

²Of course, the seller now has an incentive to fail to deliver the good, so the result that the seller can earn high profits is not trivial.

³This reasoning is similar to Bernheim and Whinston's (1998) point that if some aspects of behavior are noncontractible, it is often optimal to fail to contract on other aspects as well.

⁴This is a slight oversimplification, as there will be many equilibria in our model, not all of which yield high profits. For example, if consumers believe that the monopoly will never deliver the good unless it legally commits itself to do so, then of course so committing is the right move. Alternatively, the dynamic contracting literature often uses profit maximization as an equilibrium refinement and it does not seem more unreasonable than usual to do so here.

⁵In some environments, the seller may be automatically committed to delivering the good, for example, if nondelivery is viewed by courts as breaching an implicit contract. To address this issue, in Section 7 we show that our results extend to a setting, where, in each period, the seller has an exogenous chance of being unable to deliver the good. We feel that in such a setting the issue that nondelivery may be viewed as breaching an implicit contract does not arise, since nondelivery always occurs occasionally.

⁶The first part of this statement also holds when consumers are nonanonymous, in contrast to the results of Hart and Tirole (1988). See the discussion following Proposition 1.

When the good is durable, the structure of any equilibrium in which the seller delivers the good is complicated: sales must continue forever, since the seller would never deliver the good to the last consumer, and the price path must fall slowly enough that consumers do not always wait for lower prices but quickly enough that sales do not occur so rapidly that the seller gives in to her temptation to fail to deliver. Indeed, with a general distribution of consumer valuations, it is very difficult to construct any equilibria in which the seller always delivers the good.⁷ We, therefore, take an indirect approach to analyzing this model by first considering an auxiliary model where the seller has the ability to set a maximum sales quantity each period in addition to the price, thereby rationing the good. Our main result in this model with rationing, which we see as being of some independent interest, is that using rationing is never optimal for the seller. We then show that the seller's optimal profit in the original model must exceed her profit in any equilibrium involving rationing.

This observation allows us to derive a lower bound on the seller's profit in the original model—where constructing equilibria is very difficult—by constructing simple equilibria in the model with rationing. In particular, we construct equilibria in which price is constant over time but quantity sold every period is restricted via rationing. These quantity restrictions lead to positive residual demand, which gives the seller a reason to deliver the good. We show that a patient seller can approximate her static optimal profit level by setting price equal to the static monopoly price every period and selling to those consumers who are willing to buy at this price at a constant rate. Furthermore, for any discount factor δ , the seller's optimal profit is at least as high as the static monopoly profit of a seller with cost of delivering the good equal to c/δ , where c is the cost of delivering the good in the dynamic model, as this is precisely the profit level that can be attained by setting the price equal to the static monopoly price of a seller with cost c/δ and then selling (at cost c) at the fastest rate at which the seller is willing to deliver in the dynamic model. We also use the relationship between our model and the model with rationing to show that the best equilibria for the seller, in which she delivers the good to all consumers who purchase, involve a strictly declining price path that asymptotes to a price no lower than c/δ .

We proceed as follows: [Section 2](#) relates this paper to the literatures on the Coase conjecture, strategic rationing, and relational contracting. [Section 3](#) introduces our general model of both durable- and nondurable-goods monopoly with relational incentives. [Section 4](#) analyzes the model in the simpler case of a nondurable-goods monopoly. It is included both for completeness and because of connections between it and the subsequent analysis of the durable-goods model. [Section 5](#) introduces the model with a durable-goods monopoly, as well as the model with rationing, and studies the connection between the two, ultimately showing that the best equilibrium without on-path nondelivery for the seller in the model without rationing yields profit at least as high as that in any equilibrium without on-path nondelivery in the model with rationing. Building off this insight, [Section 6](#) presents our main results on the durable-goods model:

⁷As discussed below, it is much easier to construct equilibria in which the seller sometimes fails to deliver the good, but these equilibria may be unappealing for other reasons.

profits are bounded from below by those of a static monopoly with cost c/δ , and the best equilibrium price path along which the seller always delivers strictly declines over time and *asymptotes* to at least c/δ . Section 7 extends our analysis to a setting in which the seller is sometimes (exogenously) unable to deliver the good, where our assumption that the seller has the option of nondelivery seems particularly appropriate. Section 8 concludes and discusses some applications and empirical predictions of our model. Several proofs are deferred to Appendix A, and Appendix B discusses equilibria in which the seller does not always deliver the good along the equilibrium path.

2. RELATION TO THE LITERATURE

As indicated above, our results stand in stark contrast to the Coase conjecture (Coase 1972), which was formalized and explored by Stokey (1982), Bulow (1982), Fudenberg et al. (1985), and Gul et al. (1986).^{8,9} Our model would coincide with the standard “no-commitment” durable-goods monopoly model if the seller, while still lacking commitment power over prices, was committed to delivering the good to all consumers who purchase. In this sense, our model has less commitment than this standard no-commitment case, although of course the reason the seller does better in our model is not that it has less commitment power, but rather that committing to delivering the good to all consumers who purchase may not be wise, as after making such a commitment, the seller is tempted to cut prices.

The literature on the Coase conjecture draws a sharp distinction between the “gap case” in which the lowest consumer valuation is strictly greater than the seller’s marginal cost and the alternative “no-gap case.” In the gap case, Fudenberg et al. (1985) and Gul et al. (1986) show that there is generically a unique perfect Bayesian equilibrium, which is Markovian and satisfies the Coase conjecture. In the no-gap case, a seminal paper by Ausubel and Deneckere (1989) constructs non-Markovian equilibria that yield static monopoly profits as the discount factor approaches 1. The reason for the difference between the cases is that in the gap case, the seller is always tempted to cut prices to the lowest consumer valuation, which allows the problem to be solved by backward induction, while in the no-gap case, the possibility that the price may fall to marginal cost very quickly if the seller deviates from a prescribed price path allows the seller to

⁸Some of the many influential papers in the subsequent literature, in addition to those discussed in the text, are Ausubel and Deneckere (1987) and Gul (1987) on durable-goods oligopoly; Sobel (1991) on the entry of new consumers; Bagnoli et al. (1989) on finite populations; Bond and Samuelson (1987), Karp (1996), and Deneckere and Liang (2008) on depreciation; Kahn (1986) and McAfee and Wiseman (2008) on capacity constraints; Olsen (1992) on learning-by-doing in production; Cabral et al. (1999) and Mason (2000) on network externalities; Dudine et al. (2006) on storable goods; Deneckere and Liang (2006) and Hörner and Vieille (2009) on interdependent values; Biehl (2001) on changing consumer valuations; and Board (2008) on time-varying demand. There is also a large literature on durable-goods monopoly with bilateral offers, the early part of which is surveyed in Section 10.4 of Fudenberg and Tirole (1991).

⁹In traditional Coase conjecture papers, like Fudenberg et al. (1985), Gul et al. (1986), and Ausubel and Deneckere (1989), the model may be interpreted as a monopoly selling to either a continuum of consumers with a known distribution of valuations or to a single consumer with unknown valuation. In the current paper, only the first interpretation is applicable, as in the single-buyer case, the monopoly would never deliver the good after the buyer purchased, so there would be no equilibrium in which trade occurs.

maintain high prices in equilibrium. This distinction between the gap and no-gap cases does not arise in our model, since in our model the off-path expectation that prevents the seller from cutting prices is that the seller will not deliver the good, not that the seller will rapidly cut prices. Our analysis of durable-goods monopoly does more than show that the possibility of nondelivery allows Ausubel and Deneckere-style equilibria to be constructed in the gap case; however, as indicated above, we also provide a natural lower bound on seller payoffs for a fixed discount factor δ and prove that, for any δ , the best equilibrium for the seller in which there is no nondelivery has declining prices converging to a price no lower than c/δ . Results for fixed δ and characterizations of optimal equilibria are rare in the durable-goods monopoly literature. For example, for δ bounded away from 1, none of the early papers on the Coase conjecture cited above contains results about optimal seller profits or the asymptotic behavior of the optimal price path.

Because our approach relies on comparing our model to an auxiliary model in which the seller is able to ration the good, our paper connects to the literature on strategic rationing. One lesson from this literature is that rationing in the absence of an efficient resale market, i.e., when the highest-valuation consumers do not always receive the good when there is a shortage, can help the seller both when she can commit to a price path (Van Cayseele 1991) and when she cannot (Denicolò and Garella 1999). Both Van Cayseele and Denicolò and Garella consider short finite horizons and state that rationing in the presence of an efficient resale market is never optimal. As part of our analysis of the durable-goods model, we show that this result holds in an infinite-horizon setting.¹⁰ Our focus is very different from that of Van Cayseele and Denicolò and Garella, as they are interested primarily in cases where allowing rationing can increase profits, while we are interested precisely in cases where allowing rationing cannot increase profits, so that we can use the model with rationing to derive results about the model without rationing.

Finally, our paper is related to the literature on relational contracting, particularly that part of the relational contracting literature that studies durable goods with hidden quality, which originated with the famous papers of Klein and Leffler (1981) and Shapiro (1982, 1983).¹¹ While traditional models of durable-goods monopoly can be thought of as “relational,” in that they study the effect of dynamic incentives on a seller’s decision to cut prices, we go further and assume that dynamic incentives also govern the seller’s decision to deliver the good. Thus, the difference between our model and the existing literature on the dynamic seller is that we move a decision—delivery—from formal to relational enforcement. Also, the equilibria we construct induce cooperation through the Nash threat of breaking off trade, as in many relational contracting models (e.g., Bull 1987, Levin 2003). Indeed, a key difference between our model and traditional models of dynamic monopoly is that our model admits a Nash equilibrium in which the seller receives her minmax value.

¹⁰The relevant result (Proposition 4) assumes that the seller has the option of failing to deliver the good, but the proof shows that the result continues to hold when the seller does not have this option.

¹¹For an up-to-date survey of this rapidly expanding literature, see Malcomson (forthcoming). For a recent contribution with some similarities to the current paper, see Masten and Kosová (2009).

3. MODEL

Throughout, we consider a seller who can provide a good at marginal cost $c > 0$ facing a continuum of consumers of mass 1 with valuations (per period in the case of non-durables; net present value in the case of durables) $v \sim F(v)$ with bounded support $[\underline{v}, \bar{v}]$ with $\underline{v} \geq 0$, $\bar{v} > c$, and F continuously differentiable with strictly positive density f . There is a continuum of consumers with each valuation in $[\underline{v}, \bar{v}]$, so that if a random fraction x of consumers receive the good in some period, then that fraction x of consumers with every valuation receive the good. We do not make any assumptions as to whether \underline{v} is greater than or less than c , i.e., as to whether we are in the gap or no-gap case. Let p^m be the static monopoly price of a seller facing consumers with valuations $v \sim F(v)$ and marginal cost c .

The traditional no-commitment model of dynamic monopoly is the following infinitely repeated game:

1. At time $t \in \{0, 1, \dots\}$, the seller chooses a menu of price–delivery probability pairs $\{(p_{t,n}, x_{t,n})\}_n$.
2. Every consumer either selects a price–delivery probability pair $(p_{t,n}, x_{t,n}) \in \{(p_{t,n}, x_{t,n})\}_n$ or rejects. Consumers who select $(p_{t,n}, x_{t,n})$ pay $p_{t,n}$ and receive the good with probability $x_{t,n}$. The seller gets payoff $p_{t,n} - c$ from each consumer who pays $p_{t,n}$ and receives the good, and gets $p_{t,n}$ from each consumer who pays $p_{t,n}$ and does not receive the good. A consumer with valuation v who pays $p_{t,n}$ gets payoff $v - p_{t,n}$ if he receives the good and gets payoff $-p_{t,n}$ if he does not receive the good.
3. Repeat 1 and 2, discounting by (common) discount factor δ .

In our model, the seller has the option of nondelivery. The game becomes the following:

1. At time $t \in \{0, 1, \dots\}$, the seller chooses a menu of prices $\{p_{t,n}\}_n$.
2. Every consumer either selects a price $p_{t,n} \in \{p_{t,n}\}_n$ or rejects. Consumers who select $p_{t,n}$ pay $p_{t,n}$. Let $Q_{t,n}$ be the mass of consumers who pay $p_{t,n}$.
3. For each $p_{t,n}$, the seller chooses what fraction $x_{t,n} \in [0, 1]$ of those $Q_{t,n}$ consumers who pay $p_{t,n}$ receive the good. Each consumer who pays $p_{t,n}$ receives the good with probability $x_{t,n}$. Payoffs are as above.
4. Repeat 1 and 3, discounting by δ .

We assume that players use strategies that depend on consumers' decisions at time t only through $Q_{t,n}$. This entails assuming that the seller does not condition her strategy on play by measure-zero sets of consumers, as is standard in the durable-goods monopoly literature, as well as that consumers are anonymous.¹² In particular, the seller

¹²See the discussion following [Proposition 1](#) for more on this point.

cannot discriminate among consumers on the basis of their past play in either her pricing or delivery decisions.

Crucially, we assume that all decisions of the seller are publicly observed. Formally, let the history h^t at the start of period t be

$$(\{p_{0,n}\}, \{Q_{0,n}\}, \{x_{0,n}\}, \dots, \{p_{t-1,n}\}, \{Q_{t-1,n}\}, \{x_{t-1,n}\}).$$

Each of the seller's (pure) strategies is a pair of maps from histories h^t to $\{p_{t,n}\}$, where $p_{t,n} \in [0, \infty)$ for all t, n , and from histories $(h^t, \{p_{t,n}\}, \{Q_{t,n}\})$ to $x_{t,n} \in [0, 1]$ for all $Q_{t,n}$, while a consumer's (pure) strategy is a map from histories $(h^t, \{p_{t,n}\})$ to $\{\{p_{t,n}\}, \emptyset\}$, corresponding to accepting a price $p_{t,n}$ or rejecting. Note that, for any strategy profile, changing the strategy of a single consumer does not affect the probability distribution over histories h^t for any t ; that is, a deviation by a single consumer does not affect the path of play.

Throughout, our solution concept is pure strategy perfect Bayesian equilibrium, which we simply abbreviate PBE. Of course, the assumption that the seller uses a pure strategy does *not* imply that she chooses $x_{t,n} \in \{0, 1\}$, but rather than she does not randomize over different choices of $\{p_{t,n}\}$ or $\{x_{t,n}\}$. We have not explored whether mixed strategy equilibria can differ substantially from pure strategy equilibria; however, our main results that the seller can earn high profits in equilibrium can only be strengthened by considering mixed strategy equilibria.

We observe immediately that in either the nondurable- or the durable-goods version of our model there is a Nash equilibrium in which consumers reject all price offers and the seller sets $x_{t,n} = 0$ for all t, n . The threat of reversion to this equilibrium following any deviation may induce the seller to conform to a prescribed price path as well as to deliver the good to those consumers who purchase. No such Nash equilibrium exists in the traditional no-commitment model.

We make frequent use of the following definition.

DEFINITION 1. A PBE is *optimal* if there is no other PBE that yields strictly higher payoff for the seller.

Finally, we briefly note an alternative interpretation of our model in terms of product quality. Suppose that the seller is (for whatever reason) contractually obligated to deliver at least a low-quality good (at cost normalized to zero) to any consumer who purchases and is able to deliver a high-quality good at additional cost c , and that quality is non-contractable. If every consumer has valuation zero for the low-quality good, our model is unchanged, with "low-quality delivery" substituted for "nondelivery." This interpretation depends on every consumer having valuation zero for the low-quality good, and thus may be most attractive when quality is extremely difficult to verify. For example, the good may be a complicated, high-tech upgrade of an existing piece of hardware or software that has no value at all for consumers if it is not superior to the original product

and that outside observers are unable to verify whether the upgrade is, in fact, better than the original.¹³

4. NONDURABLE-GOODS MONOPOLY

In this section, each consumer demands one unit of the good each period, and v is a consumer's per-period valuation. We also assume, for this section only, that $v - (1 - F(v))/f(v)$ is weakly increasing, so that in the static monopoly allocation, every consumer with positive virtual surplus receives the good.¹⁴

Our main result in this section is that, in the optimal equilibrium,¹⁵ the seller sets the (single) price equal to the static monopoly price if she is sufficiently patient, and otherwise sets the lowest price at which she is willing to deliver the good. The intuition is that the seller's incentive to fail to deliver the good is increasing in quantity, so if the seller is impatient, she must restrict quantity to credibly commit to delivery, and the most profitable way to do this is to increase price. In particular, the seller sets $p = \max\{p^m, c/\delta\}$ every period. To see why c/δ is the lowest price at which the seller is willing to deliver the good, let $D(p) \equiv 1 - F(p)$ be demand at price p , and note that in every period the seller gains $cD(p)$ from failing to deliver and gains $(\delta/(1 - \delta))(p - c)D(p)$ from delivering. The latter is weakly greater than the former if and only if $p \geq c/\delta$. The idea of the proof is to first note that the seller can in effect commit to any price path, since deviations in price-setting may lead consumers to believe that the seller will not deliver the good and thus lead to zero sales. Next observe that the best dynamic sales mechanism for the seller is stationary, as increasing one period's profits also relaxes the seller's incentive compatibility (willingness to deliver) constraints from earlier periods. Finally, use standard static mechanism techniques to characterize the optimal stationary mechanism that is incentive compatible for the seller. The proof is deferred to [Appendix A](#).

PROPOSITION 1. *If $\bar{v} \geq c/\delta$, the equilibrium path of the optimal PBE of the nondurable-goods model is given by $p_{t,n} = \max\{p^m, c/\delta\}$ for all t, n , buyers accept if and only if $v \geq p_{t,n}$, and the seller delivers the good with probability 1 to all buyers who accept each period. That is, the seller offers only a posted price p in every period, $p = p^m$ if $\delta \geq c/p^m$ and $p = c/\delta > p^m$ if $\delta < c/p^m$. If $\bar{v} < c/\delta$, there is no PBE in which the seller ever delivers the good or receives positive payments.*

¹³Our results do not apply if consumers have positive valuations for the low-quality good, since in this case the model need not have a Nash equilibrium that yields zero profit. However, two recent papers illustrate interesting phenomena that may occur in such settings. [Inderst \(2008\)](#) shows that a durable-goods monopoly that sells low- and high-quality goods may serve the entire market in the first period, selling the low-quality good to low-valuation consumers as a means of committing itself not to subsequently offer the high-quality good at a lower price. [Hahn \(2006\)](#) shows that this logic may provide an incentive for a durable-goods monopoly to introduce a damaged version of its good and argues that this often has negative welfare consequences.

¹⁴This assumption is for technical convenience only.

¹⁵The proof of [Proposition 1](#) shows existence and uniqueness of an optimal equilibrium.

Recall that we have assumed that buyers are anonymous. Nonetheless, it is not hard to construct equilibria that yield static monopoly profits even if buyers are nonanonymous, provided that $\delta \geq c/p^m$. For example, let the seller set $p = p^m$ in every period, and deliver the good if and only if she has both always delivered the good to all consumers who have purchased and set $p = p^m$ in the past, and let each consumer purchase the good every period if and only if her valuation exceeds p^m and the seller has always delivered the good to all consumers who have purchased and has always set $p = p^m$. In every period, the seller gains $cD(p^m)$ from failing to deliver and gains $(\delta/(1-\delta))(p^m - c)D(p^m)$ from delivering, so the seller will deliver if $\delta \geq c/p^m$. This result differs dramatically from the classic analysis of nondurable-goods monopoly with nonanonymous consumers provided by Hart and Tirole (1988). Hart and Tirole show that, in a finite-horizon model with nondurable goods and nonanonymous consumers, equilibrium is governed by the ratchet effect: in every PBE, if $\underline{v} > c$, then $p_t = \underline{v}$ for all but the last few periods. Technically, the difference between our result and theirs comes from the fact that the stage game in our model has a bad Nash equilibrium (reject any offer, never deliver), which can be used as an off-equilibrium threat to prevent the seller from using information revealed early on against high-valuation buyers.¹⁶ The key economic point is that the usual repeated game trade-off between a short term gain from cheating and a long term gain from cooperation on the part of the seller is absent in the Hart–Tirole model: in their model, the seller is free to “cheat” by raising the price she charges to buyers that reveal themselves to have high valuations, but buyers cannot credibly retaliate by refusing to buy at the higher price. In our model, the option of the seller to fail to deliver the good lets the buyer credibly punish the seller for raising the price, allowing the seller to commit to keeping the price constant. Alternatively, we must now keep track of the seller’s incentive to deliver the good. If $\delta \geq c/p^m$, this incentive constraint is slack, so the seller can attain her full-commitment optimum.

5. DURABLE-GOODS MONOPOLY AND RATIONING

5.1 Preliminaries

For the remainder of the paper, each consumer demands only one unit of the (durable) good, and v is a consumer’s net present value of receiving the good. In the traditional model of this situation (see Section 3), Fudenberg et al. (1985) and Gul et al. (1986) show that the Coase conjecture applies if the lowest valuation \underline{v} is greater than c : for generic parameters, there is a unique PBE, and as δ goes to 1, the seller’s profit goes to $\underline{v} - c$ and the price drops to \underline{v} very quickly.

Our main result implies that the Coase conjecture does not apply to this model when the seller has the option of nondelivery (see Section 3), which we call the *relational contracting model* or Γ . Much of our analysis focuses on a particular class of PBE, which we call full-delivery PBE. A *full-delivery PBE* is a PBE in which the seller sets $x_{t,n} = 1$

¹⁶The infinite-horizon version of the Hart–Tirole model has equilibria that yield seller profits above $\underline{v} - c$, though how much above $\underline{v} - c$ has to our knowledge not been studied in the literature. Thus, it is possible that some of the difference in results is due to the difference in time horizons.

for all n at all histories on the equilibrium path. It is important to note that the seller may set $x_{t,n} < 1$ off the equilibrium path in a full-delivery PBE. A full-delivery PBE is a *best* full-delivery PBE if there is no other full-delivery PBE that yields strictly higher payoff for the seller—we use the word “optimal” for the best PBE overall and “best” for the best full-delivery PBE to help avoid confusion. Note that on the equilibrium path of a full-delivery PBE there is no reason for the seller to offer a menu of prices, as each consumer will either accept the lowest offered price or reject, so we simplify notation by writing p_t for the lowest price offered by the seller at time t on the equilibrium path. Furthermore, a consumer who pays p_t always receives the good at time t ; we say that a consumer who pays p_t at time t on the equilibrium path of a full-delivery PBE *purchases* the good at time t . Since we have restricted attention to pure strategy equilibria, every consumer purchases at exactly one time in every full-delivery PBE, with the convention that a consumer who never receives the good purchases at $t = \infty$.

Clearly, an optimal PBE of the relational contracting model can yield no higher payoff to the seller than an optimal PBE of the full-commitment model in which the requirement that the seller’s strategy is sequentially rational is relaxed, and it follows from standard results that an optimal PBE of this full-commitment model yields profits equal to optimal static monopoly profits. Our main result is the following theorem, which implies that the Coase conjecture does not hold in this game regardless of the relationship between \underline{v} and c , and also provides a lower bound on the seller’s profit for any fixed δ .

THEOREM 1. *In the relational contracting model, the following statements hold.*

- (i) *An optimal PBE exists.*
- (ii) *A best full-delivery PBE exists.*
- (iii) *As δ approaches 1, profit in a best full-delivery PBE approaches static monopoly profit.*
- (iv) *If $\bar{v} > c/\delta$ and cost equals c , there exists a full-delivery PBE in which profit is strictly greater than static monopoly profit when cost equals c/δ .*
- (v) *If $\bar{v} > c/\delta$, any best full-delivery PBE has a strictly decreasing price path and involves positive sales in every period.*
- (vi) *If $\bar{v} > c/\delta$, $p_t \geq \max\{\underline{v}, c/\delta\}$ for all t in any best full-delivery PBE.*
- (vii) *If $\bar{v} \leq c/\delta$, there is no PBE in which the seller ever delivers the good or receives positive payments.*

Sections 5 and 6 were devoted to establishing **Theorem 1**: parts (i) and (ii) are proved in this section (in Propositions 2 and 5) and parts (iii)–(vii) are proved in **Section 6** (in Propositions 6–9). We therefore take a moment to motivate devoting so much attention to full-delivery PBE. Full-delivery PBE are those equilibria in which on-path delivery is as in both the full-commitment model (in which the seller commits to both a price path $(p)_t$ and a delivery path $(x)_t$) and the traditional no-commitment model described

in Section 3, which makes them a natural class of equilibria to study. Indeed, on-path nondelivery—the equivalent of the seller selling lottery tickets that entitle consumers to receive the good with some probability less than 1—may be unappealing in some settings, for example, if consumers can tell whether the seller has failed to deliver the good to anyone but not whether the seller has delivered to some exact fraction of consumers. Furthermore, Theorem 1 implies that the profit lost by the seller in a best full-delivery PBE as opposed to an optimal PBE is bounded from above by the difference between static monopoly profit when cost equals c and when cost equals c/δ , which is small for δ close to 1. Nonetheless, we conjecture that, in general, the optimal PBE is not full delivery, for reasons we discuss in Appendix B. Appendix B proves the analogues of parts (iii) and (iv) of Theorem 1 for non-full-delivery equilibria directly, i.e., without relying on the connection between the relational contracting model and the related model with rationing introduced below. The approach of Appendix B also has the advantage of explicitly constructing equilibria in the relational contracting model, while the approach taken in the body of the paper is nonconstructive. Thus, there are at least two very different kinds of PBE that yield high seller profits: full-delivery PBE with declining price paths, whose existence is proven nonconstructively in the text, and non-full-delivery equilibria with constant price paths, which are constructed in Appendix B.

We adopt a novel approach to proving Theorem 1. We first introduce the following variant of the relational contracting model, in which the seller can artificially restrict the quantity of the good supplied each period.

1. The seller chooses a price p_t and a maximum quantity to supply $q_t \in [0, 1]$.¹⁷
2. Every consumer chooses whether or not to accept p_t . If less than q_t consumers accept, all consumers who accept pay p_t . Otherwise, the q_t consumer with the highest valuation among those who accept pay p_t . Formally, a consumer with valuation v who accepts pays if and only if the mass of consumers with valuation strictly greater than v who accept is strictly less than q_t .
3. If measure Q_t of consumers pay p_t (which we call the period t quantity), the seller chooses what fraction $x_t \in [0, 1]$ of these consumers receive the good. Each consumer who pays p_t receives the good with probability x_t .
4. Repeat 1 and 3, discounting by δ .

We have not allowed the seller to offer menus of prices, as this would only complicate notation, since we restrict attention to full-delivery PBE in what follows.

We call this game the relational contracting model with rationing, or simply the model with rationing, or Γ_R .¹⁸ Optimal, full-delivery, and best full-delivery PBE in Γ_R

¹⁷For the remainder of the paper, q_t refers to the quantity cap in period t and Q_t refers to the number of consumers who pay in period t (i.e., the period t quantity). By construction of the model with rationing, $Q_t \leq q_t$.

¹⁸In defining Γ_R we have made two assumptions on the rationing technology: that types on the boundary between receiving the good and not do not receive the good, and that any rationing that occurs is “efficient” in that the highest-valuation consumers are eligible to receive the good. The first assumption is

are defined as in Γ . The main reason we introduce Γ_R is that full-delivery equilibria in Γ_R may have flat price paths, while every full-delivery equilibrium in Γ must involve price cuts, as otherwise there would be no way to delay sales and thereby induce delivery.¹⁹ Full-delivery equilibria with flat price paths are easy to analyze, as consumers' incentives in such equilibria are trivial: if the price is fixed at p in a full-delivery equilibrium, a consumer with valuation $v \geq p$ wants to purchase as soon as possible, while a consumer with $v < p$ will never purchase. We will show that full-delivery equilibria with flat price paths exist in Γ_R that approximate static monopoly profits for high δ . Furthermore, we will show that a price–quantity path $(p, Q)_t$ is a best full-delivery PBE price–quantity path in Γ if and only if it is a best full-delivery PBE price–quantity path in Γ_R (Corollary 1 in Section 5.4). Therefore, the best full-delivery PBE profit attainable by the seller is the same in Γ and Γ_R , so the above observation that simple full-delivery PBE exist in Γ_R in which profits approximate static optimal profits immediately yields part (iii) of Theorem 1, even though no such simple full-delivery PBE exist in Γ . The proofs of parts (ii) and (iv)–(vii) of Theorem 1 also rely on Corollary 1, as we will see; thus, Corollary 1 is the key to our approach to proving Theorem 1.

To summarize the above roadmap, Sections 5 and 6 establish the chain of inequalities

$$\begin{aligned}
 & \text{optimal PBE profit in } \Gamma \\
 & \geq \text{best full-delivery PBE profit in } \Gamma \quad (\text{by Definition 1}) \\
 & = \text{best full-delivery PBE profit in } \Gamma^R \quad (\text{by Corollary 1}) \\
 & > \text{best full-delivery, constant-price PBE profit in } \Gamma^R \quad (\text{by Proposition 7}) \\
 & = \text{static monopoly profit with cost } c/\delta \quad (\text{by Corollary 2}).
 \end{aligned}$$

Before beginning our analysis of Γ_R , we first prove part (i) of Theorem 1 directly. The proof proceeds by first showing that the seller's profit is continuous in price–delivery paths $(p, x)_t$ and then showing that any price–delivery path can be supported in PBE by endowing consumers with the belief that the seller will never deliver the good if she ever deviates from her prescribed price–delivery path. The details are deferred to Appendix A.

PROPOSITION 2 (Theorem 1(i)). *An optimal PBE exists in Γ .*

5.2 Existence of best full-delivery PBE in the model with rationing

We now begin our analysis of the full-delivery PBE of Γ and Γ_R , and the relationship between them. The goal of this subsection is to show that a best full-delivery PBE exists in Γ_R . We start with a definition.

only for technical convenience and simplifies the proof of Lemma 3. The second assumption is substantive, as Van Cayseele (1991) shows that under full commitment a monopoly can achieve profits above static monopoly profits by using “inefficient” rationing. The second assumption is descriptive in the presence of a frictionless resale market. Alternatively, one could view the model with rationing entirely as a technical aid in analyzing the model without rationing.

¹⁹Our results about Γ_R , especially Proposition 4, may also be of independent value to readers interested in strategic rationing.

DEFINITION 2. Given a price path $(p)_t$, a valuation v is *generic with respect to* $(p)_t$ if

$$\delta^t(v - p_t) \neq \delta^{t'}(v - p_{t'})$$

for all $t \neq t'$. If not, v is *nongeneric with respect to* $(p)_t$.

That is, a valuation v is generic with respect to $(p)_t$ if a consumer with valuation v is not indifferent between purchasing at any two times t and t' when prices are given by $(p)_t$. For any price path $(p)_t$, there are only countably many valuations that are nongeneric with respect to $(p)_t$, so the assumption that F admits a strictly positive density immediately yields the following observation.

LEMMA 1. *For any price path $(p)_t$, the set of valuations $v \in [\underline{v}, \bar{v}]$ that are generic with respect to $(p)_t$ has measure 1.*

We now present a series of lemmas that are needed to prove existence of a best full-delivery PBE in Γ_R . The longer proofs are deferred to [Appendix A](#).

[Lemma 2](#) simply states that any two consumers with the same valuation receive the same payoff in any PBE, and consumers with higher valuations receive higher payoffs.

LEMMA 2. *In any PBE of Γ or Γ_R , any two consumers with the same valuation, v , receive the same PBE payoff, V_v . If $v \geq v'$, then $V_v \geq V_{v'}$.*

PROOF. The first part follows because at any PBE, a consumer with valuation v can deviate to the strategy of another consumer with valuation v and receive the same payoff as him, because the actions of a single consumer do not affect the path of play (in either Γ or Γ_R). The second part follows because at any PBE, a consumer with valuation $v \geq v'$ can deviate to the strategy of a consumer with valuation v' and receive a weakly higher payoff than him (in Γ_R , this relies on the fact that a consumer with higher valuation can purchase whenever a consumer with lower valuation can do so), again because the actions of a single consumer do not affect the path of play. \square

The next two lemmas show that, across all full-delivery PBE, the price–rationing path $(p, q)_t$ uniquely determines the quantity path $(Q)_t$. [Lemma 3](#) is not trivial because the set of times at which a consumer is able to purchase under price–rationing path $(p, q)_t$ depends on the times at which higher-valuation consumers are purchasing. The intuition for the result is that if a consumer with valuation v cannot purchase at the same set of times under two PBE, then there must be a nontrivial mass of higher-valuation consumers who cannot purchase at the same set of times under the two PBE either, because otherwise almost all higher-valuation consumers would purchase at the same times under both PBE and the original consumer would not have been “rationed out” of purchasing at his preferred time. Therefore, there can be no valuation v that is approximately the highest valuation that gets rationed out, which implies that no valuation can be rationed out.

LEMMA 3. *Given a price–rationing path $(p, q)_t$ in Γ_R and a valuation v that is generic with respect to $(p)_t$, there exists a time τ_v such that every consumer with valuation v purchases at τ_v in any full-delivery PBE in Γ_R with price–rationing path $(p, q)_t$.*

Combining Lemma 1 and Lemma 3 immediately yields the following lemma.

LEMMA 4. *Given price–rationing path $(p, q)_t$, every full-delivery PBE in Γ_R with price–rationing path $(p, q)_t$ has the same quantity path $(Q)_t$.*

In fact, this quantity path $(Q)_t$ can be viewed as a continuous function of the price–rationing path $(p, q)_t$.

LEMMA 5. *The unique quantity path $(Q)_t$ that may occur in a full-delivery PBE in Γ_R with price–rationing path $(p, q)_t$ is continuous in $(p, q)_t$ in the product topology.*

We now show that a best full-delivery PBE exists in the model with rationing (Proposition 3). This holds because the set of full-delivery PBE price–rationing–quantity paths can be shown to be compact in the product topology²⁰ and the seller's profit is continuous in price–rationing–quantity paths. It is straightforward to show that the set of full-delivery PBE price–rationing paths is compact: the seller can be induced to set any price–rationing path if consumers believe that she will never deliver the good if she sets the wrong path, and the seller is willing to deliver Q units of the good if she is willing to deliver $Q - \varepsilon$ for all small ε . The difficulty is showing that small changes in the price–rationing path induce small changes in the quantity path. This is taken care of by Lemmas 4 and 5, which are both proved in Appendix A.

PROPOSITION 3. *A best full-delivery PBE exists in Γ_R .*

PROOF. Let \mathcal{F} be the set of full-delivery PBE price–rationing–quantity paths $(p, q, Q)_t$ in Γ_R satisfying $p_t \in [\underline{v}, \bar{v}]$ for all t . Note that if a PBE is best in the set of PBE with price–rationing paths in \mathcal{F} , then it is best overall, as any PBE with $p_t > \bar{v}$ for some t yields no more profit than a PBE with an identical price–rationing path but with $p_t = \bar{v}$ for all such t instead, and similarly for $p_t < \underline{v}$. Given a price–rationing–quantity path $(p, q, Q)_t$, the associated profit for the seller is

$$\sum_{t=0}^{\infty} \delta^t (p_t - c) \min\{q_t, Q_t\},$$

which is obviously continuous in $(p, q, Q)_t$ in the product topology. We will show that \mathcal{F} is compact in the product topology, and then apply Weierstrass' Theorem to complete the proof.

Observe that $\mathcal{F} \subseteq \prod_{t=0}^{\infty} ([\underline{v}, \bar{v}], [0, 1], [0, 1])_t$, which is compact by Tychonoff's Theorem. Therefore, to show that \mathcal{F} is compact in the product topology, it suffices to show

²⁰Technically, this holds for price paths with $p_t \in [\underline{v}, \bar{v}]$ for all t , to which we can restrict attention without loss of generality.

that \mathcal{F} is closed in the product topology. To see that it is, consider a sequence of paths $\{(p, q, Q)_t\}_n \in \mathcal{F}$ converging pointwise to $(p^*, q^*, Q^*)_t$. We must show that there exists a full-delivery PBE with price-rationing-quantity path $(p^*, q^*, Q^*)_t$. Consider the following strategy profile:

1. The seller sets price-rationing path $(p^*, q^*)_t$ and $x_t = 1$ as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$, $q_t = 1$, and $x_t = 0$ for all future periods. In particular, the seller sets $x_t = 0$ in any period in which she has set $p_t \neq p_t^*$.
2. A consumer with valuation v who has not yet received the good at t accepts at t if and only if the seller has never deviated from her prescribed strategy and $\delta^t(v - p_t^*) \geq \delta^\tau(v - p_\tau^*)$ for all $\tau \geq t$.

To establish that this profile is a PBE, we first observe that if the seller ever sets $p_\tau \neq p_\tau^*$, she receives zero continuation payoff. Since this is her minmax value, she cannot receive continuation payoff strictly less than this in any PBE; in particular, her on-path continuation value after τ along $(p, q, Q)_{t,n}$ is weakly positive for every n , so by continuity of profits in $(p, q, Q)_t$, we see that her on path continuation value after τ along $(p^*, q^*, Q^*)_t$ is also weakly positive. This implies that setting $p_\tau \neq p_\tau^*$ on path is not a profitable deviation. Similarly, the fact that setting $x_t = 1$ is optimal on path along $(p, q, Q)_{t,n}$ for all n implies that setting $x_t = 1$ is optimal on path in this strategy profile, because the cost of delivery and on-path continuation values are continuous in $(p, q, Q)_{t,n}$, while the payoff of zero that results from deviating from the equilibrium path in this profile is at least as bad as the payoff from deviating in any PBE. Also, the seller's off-path play is optimal because off-path price-setting does not affect her payoffs and off-path delivery imposes a positive cost at no benefit.

We next check that each consumer's play is optimal. It is again obvious that his off-path play is optimal, as paying is costly and yields no benefit when the seller sets $x_t = 0$. To see that his on-path play is optimal given $(p^*, q^*, Q^*)_t$, note that accepting at t yields $\delta^t(v - p_t^*)$ if he pays (i.e., if he is allowed to purchase the good) and yields his continuation payoff otherwise, while rejecting always yields his continuation payoff and $\delta^t(v - p_t^*)$ is weakly greater than his continuation payoff if $\delta^t(v - p_t^*) \geq \delta^\tau(v - p_\tau^*)$ for all $\tau \geq t$.

Finally, we must check that the prescribed consumer behavior actually induces quantity path $(Q^*)_t$. By Lemma 4, for any price-rationing path $(p, q)_t$, there is a unique quantity path $(Q)_t$ that occurs in a full-delivery PBE with price-rationing path $(p, q)_t$, and $(Q)_t$ is continuous in $(p, q)_t$ by Lemma 5. Therefore, the fact that $(p, q)_{t,n}$ converges to $(p^*, q^*)_t$ implies that $(Q)_{t,n}$ converges to $(Q^*)_t$. Thus, there exists a full-delivery PBE with price-rationing-quantity path $(p^*, q^*, Q^*)_t$.

We have shown that \mathcal{F} is closed, and therefore compact, in the product topology. Weierstrass' Theorem now implies that there is a point in \mathcal{F} that maximizes profits, which completes the proof. \square

5.3 Nonoptimality of rationing in the model with rationing

We now show that any best full-delivery PBE in Γ_R involves no rationing on the equilibrium path. This is the central step in showing equivalence of best full-delivery PBE in Γ and Γ_R (Corollary 1), which is in turn our main tool in proving Theorem 1.

By Lemma 3, the path of play of a full-delivery PBE is given by a price–rationing path $(p, q)_t$, up to differences in the play of the measure-zero set of consumers with nongeneric valuations with respect to $(p)_t$. Let us write $D_\tau((p, q)_t)$ for the quantity demanded at time τ given price–rationing path $(p, q)_t$, i.e., the measure of consumers who would prefer to receive the good at time τ at price p_τ than to receive their PBE payoff.²¹ Similarly, we say that a consumer demands the good at τ if she prefers receiving the good at time τ at price p_τ to receiving her PBE payoff. Finally, we say that rationing occurs along a price–quantity–rationing path $(p, q)_t$ if there exists a time τ such that $D_\tau((p, q)_t) > q_\tau > 0$.²² Note that in a full-delivery PBE in which $D_\tau((p, q)_t) \leq q_\tau$, a consumer with nongeneric valuation who demands the good at τ must purchase at τ .²³

We show that every best full-delivery PBE in Γ_R involves no rationing by arguing that any full-delivery PBE involving rationing can be strictly improved upon by another full-delivery PBE. The basic idea is that if rationing occurs at time t^* , modifying the equilibrium by slightly increasing price at t^* , such that quantity sold at t^* remains constant, and using additional rationing to ensure that quantity sold in every other period does not increase, leads the timing of all sales to remain constant and therefore yields an increase in profits. However, the proof is complicated by the fact that, without first ruling out rationing, we cannot ensure that the price path is decreasing and cannot establish the usual skimming property that higher-valuation consumers purchase earlier. The heart of the proof involves showing that slightly increasing price at t^* and using additional rationing to ensure that sales do not increase elsewhere cannot lead to a decrease in sales at some other time τ . If it did, then those consumers who used to purchase at τ must now purchase at some other time that is better for them than τ , as they still have the option of earning surplus by purchasing at τ . Additionally, the fact that they have this new opportunity means that some other higher-valuation consumers must also be purchasing at a different time. Since higher-valuation consumers must purchase at some point rather than never purchasing if lower-valuation consumers do so, following this “trail” of consumers who purchase at different times ultimately shows that every consumer (with generic valuation) who purchased before the price increase still purchases after the price increase. The details of the proof are deferred to Appendix A.

PROPOSITION 4. *In Γ_R , no rationing occurs along a best full-delivery PBE price–quantity–rationing path.*

²¹Throughout the paper, $D(p) \equiv 1 - F(p)$ is the static demand at price p , while $D_\tau((p, q)_t)$ is the time- τ demand in the dynamic model under price–rationing path $(p, q)_t$.

²²If $q_t = 0$, it is irrelevant whether we consider the monopoly to be rationing at t or to be setting price equal to infinity. We do not refer to this case as rationing for technical convenience.

²³If $D_\tau((p, q)_t) = q_\tau$, this may fail for a measure-zero set of consumers who demand the good at τ but are unable to purchase at τ due to rationing. Since measure-zero sets of consumers are irrelevant for our analysis, we ignore this case in the discussion.

5.4 Equivalence of best full-delivery PBE in the model with and without rationing

We are finally ready to prove [Corollary 1](#), which establishes a very close relationship between best full-delivery PBE in the relational contracting model with and without rationing. The intuition for [Corollary 1](#) is simple: by [Proposition 4](#), no rationing occurs on the equilibrium path in a best full-delivery PBE of Γ_R , and the worst possible off-path punishment (breaking off trade) does not require rationing, so a best full-delivery PBE of Γ_R can be no better than a best full-delivery PBE of Γ . The details of the proof, which involves constructing a PBE in Γ corresponding to a given price–quantity path in Γ_R and vice versa, is deferred to [Appendix A](#). The constructed PBE have the same grim-trigger structure as the PBE described in the proof of [Proposition 3](#) and in [Section 6.1](#).

COROLLARY 1. *A price–quantity path $(p, Q)_t$ is a best full-delivery PBE price–quantity path in Γ_R if and only if it is a best full-delivery PBE price–quantity path in Γ .*

[Corollary 1](#) combined with [Proposition 3](#) immediately yields part (ii) of [Theorem 1](#).

PROPOSITION 5 ([Theorem 1\(ii\)](#)). *A best full-delivery PBE exists in Γ .*

6. PROPERTIES OF BEST FULL-DELIVERY EQUILIBRIA

6.1 High profits and supermonopoly pricing

In this subsection, we use the facts about Γ_R and its relationship to Γ established in [Section 5](#) to prove parts (iii) and (iv) of [Theorem 1](#).

We first show that profits in a best full-delivery PBE in Γ_R (which exists, by [Proposition 3](#)) converge to the static monopoly profit as δ approaches 1, which is not difficult. [Corollary 1](#) then implies that the same is true in Γ . To see why payoffs in the best full-delivery PBE in Γ_R converge to static monopoly profits as δ approaches 1, let $D(p) \equiv 1 - F(p)$ —the static demand at price p —and consider the following strategy profile, where γ is a constant in $(0, (p^m - c)/p^m)$.

1. The seller sets price–rationing–delivery path $p_t = p^m$, $q_t = \gamma(1 - \gamma)^t D(p^m)$, and $x_t = 1$ as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$, $q_t = 1$, and $x_t = 0$ for all future periods. In particular, the seller sets $x_t = 0$ in any period in which she has set $p_t \neq p_t^m$.
2. A consumer with valuation v who has not yet received the good accepts if and only if the seller has never deviated from her prescribed strategy and $v \geq p^m$.

That is, the seller keeps price fixed at the static monopoly price, p^m , and sells to fraction γ of those consumers who demand the good each period, while consumers accept if and only if $v \geq p^m$ and the seller has never deviated. It is clear that consumers' play is optimal, and that the seller can never benefit from setting a different value of p_t or q_t , so checking that this profile is an equilibrium reduces to checking that the seller prefers to deliver the good. The proof of [Proposition 6](#) shows that the seller does, in fact, prefer to deliver the good if $\gamma \leq (\delta p^m - c)/(\delta p^m)$, and if δ is close to 1, then this strategy profile

yields approximately static monopoly profits, as the cost of delay involved in selling to only fraction γ of the consumers who demand the good each period is small. Therefore, profits in a best full-delivery PBE in Γ_R must approximate static monopoly profits for δ close to 1 as well.

PROPOSITION 6 (Theorem 1 (iii)). *For both Γ and Γ_R , for all $\varepsilon > 0$, there exists $\bar{\delta} < 1$ such that, for all $\delta > \bar{\delta}$, there exists a full-delivery PBE under which the seller's payoff is within ε of the static monopoly payoff.*

PROOF. We prove the result for Γ_R below. **Proposition 3** then implies that, for every $\delta > \bar{\delta}$, there exists a best full-delivery PBE in Γ_R under which the seller's payoff is within ε of the static monopoly payoff. **Corollary 1** in turn implies that the same is true in Γ .

Recall that p^m is the static monopoly price, so the static monopoly payoff is $(p^m - c)D(p^m)$. Suppose that $p^m > c$, i.e., that positive profits are possible: the case where this fails is trivial.

Consider the strategy profile described above for γ some constant in $(0, (p^m - c)/p^m)$. It is clear that each consumer's strategy is a best reply. Note also that $q_t = Q_t$ for all t along the equilibrium path. To check only that this profile describes a PBE, we must check only that the seller has an incentive to deliver the good along the equilibrium path, since any other deviation yields continuation payoff zero against positive continuation payoff from conforming. This condition is

$$\sum_{\tau=1}^{\infty} \delta^\tau q_{t+\tau} (p_{t+\tau} - c) \geq q_t c \quad \text{for all } t \geq 0.$$

For any t , this can be rewritten as

$$\gamma(1 - \gamma)^t \left(\frac{\delta(1 - \gamma)}{1 - \delta(1 - \gamma)} \right) D(p^m)(p^m - c) \geq \gamma(1 - \gamma)^t D(p^m)c$$

or

$$\left(\frac{\delta(1 - \gamma)}{1 - \delta(1 - \gamma)} \right) (p^m - c) \geq c.$$

Rearranging this inequality gives

$$\gamma \leq \frac{\delta p^m - c}{\delta p^m}. \tag{1}$$

Thus, the strategy profile above is a PBE for any γ satisfying (1). Since $p^m > c$, there exists $\gamma > 0$ such that the strategy profile above is a PBE for high enough δ , in particular for $\delta > c/p^m$.

Suppose that $\delta > c/p^m$ and fix any positive γ satisfying (1). Note that this strategy profile yields profit

$$\left(\frac{\gamma}{1 - \delta(1 - \gamma)} \right) D(p^m)(p^m - c)$$

for the seller. As δ approaches 1, this converges to $D(p^m)(p^m - c)$, completing the proof. \square

The intuition for this result is that, for δ high enough ($\delta > c/p^m$), the seller can credibly deliver the good to those consumers willing to pay the monopoly price at a fixed positive rate γ , and taking δ to 1 means that the loss from delay involved in this strategy is insignificant. Observe that, while the proof of [Proposition 6](#) shows that, in Γ_R , there exists a single strategy profile that is a PBE for all sufficiently high δ and that yields profits converging to static monopoly profits as δ converges to 1, such a strategy profile need not exist in Γ .

Note that the strategy profile described in the proof of [Proposition 6](#), with $p = p^m$, is not a best full-delivery PBE in Γ_R for fixed $\delta < 1$. Indeed, there exist full-delivery PBE in Γ_R with constant-price paths (i.e., $p_t = p_{t'}$ for all t, t') that yield higher profits. To see this, consider the strategy profile in the proof of [Proposition 6](#) with p^m replaced by some price p . Let

$$\gamma^*(p) \equiv \frac{\delta p - c}{\delta p}.$$

The argument in the proof of [Proposition 6](#) that led to equation (1) shows that $\gamma^*(p)$ is the fastest rate at which the seller can sell in a full-delivery PBE in which price is fixed at p . This implies that the seller's profit in the best full-delivery PBE with a constant price path at p and a constant sales rate γ is

$$\left(\frac{\gamma^*(p)}{1 - \delta(1 - \gamma^*(p))} \right) D(p)(p - c),$$

which equals

$$\left(\frac{p - \frac{c}{\delta}}{p - c} \right) D(p)(p - c). \quad (2)$$

Note that the first term of (2) represents the cost of the delay in sales required to induce the seller to deliver, while the second term is simply the static profit at price p . Raising p above p^m yields a first-order increase in the first term in (2) and a second-order decrease in the product of the second and third terms, so the seller does better to sell at price above p^m . The intuition is similar to that of [Section 4](#): raising price reduces quantity, which reduces the seller's temptation to fail to deliver, and, with durable goods, this allows the seller to sell at a faster rate. More specifically, the required delay in sales forces a seller who would receive $p - c$ per unit sold under full commitment to receive only $p - c/\delta$ per unit sold, so, with a constant-price path, a seller with cost c can do no better than imitating the pricing of a static monopoly with cost c/δ . That is, (2) equals

$$\left(p - \frac{c}{\delta} \right) D(p),$$

from which it is clear that the best full-delivery, fixed-price PBE in which the seller sells at a constant rate is given by price $p^m(c/\delta)$, the monopoly price when cost equals c/δ ,

and sales rate $\gamma = \gamma^*(p^m(c/\delta))$. In fact, it is not hard to show that this is the best full-delivery, fixed-price PBE overall: all that remains to show this is to establish that selling at the constant rate $\gamma^*(p)$ is optimal given that prices are fixed at any given p , which follows from a standard dynamic programming argument.²⁴

COROLLARY 2. *If $\bar{v} > c/\delta$, the best full-delivery, constant-price PBE in Γ_R is given by $p_t = p^m(c/\delta)$ and $q_t = \gamma^*(p^m(c/\delta))(1 - \gamma^*(p^m(c/\delta)))^t D(p^m(c/\delta))$. Furthermore, $(p^m(c/\delta) - c/\delta)D(p^m(c/\delta))$ is a lower bound on the best full-delivery PBE profit in both Γ and Γ_R .*

PROOF. Given the first part of the result, the second part follows immediately from Corollary 1.

Suppose $p_t = p$ for all t . Let Q be the static demand for price p . The problem of finding the best full-delivery PBE with a constant price p in Γ^R reduces to finding the best number of consumers to sell to in every period while maintaining the seller's incentive to deliver the good; i.e., to solving the functional equation

$$V(Q) = \max_{q \leq Q \text{ such that } \delta V(Q-q) \geq qc} (p - c)q + \delta V(Q - q). \tag{3}$$

Standard dynamic programming results imply that there is at most one solution to this equation with a nontrivial set satisfying the constraints. Conjecture that $V(Q) = ((\delta p - c)/\delta)Q$. The right-hand side of (3) then becomes

$$\begin{aligned} \max_{q \leq ((\delta p - c)/(\delta p))Q} (p - c)q + (\delta p - c)(Q - q) &= (p - c) \left(\frac{\delta p - c}{\delta p} \right) Q + (\delta p - c) \left(\frac{c}{\delta p} \right) Q \\ &= \left(p - \frac{c}{\delta} \right) Q, \end{aligned}$$

where the constraint set is nontrivial if $p > c/\delta$. Therefore, $(p - c/\delta)Q$ is the highest profit attainable by a price path fixed at $p > c/\delta$ when there are Q remaining consumers with valuations greater than p , and 0 is the highest such profit if $p \leq c/\delta$ (as the solution to (3) must be nonincreasing in p). Setting $Q = D(p)$ and maximizing over p completes the proof. \square

Finally, we note that (nonconstant-price) full-delivery PBE of Γ_R exist that yield profits strictly above static monopoly profits with cost equal to c/δ if $\bar{v} > c/\delta$. For example, consider modifying the best full-delivery, constant-price path by increasing p_0 from $p^m(c/\delta)$ to $p^m(c/\delta) + \varepsilon$ for ε small. We claim that for small ε , q_0 consumers will still pay p_0 . This follows because a consumer with valuation v demands the good at time 0 and price p_0 if $v - p^m(c/\delta) - \varepsilon \geq \delta(v - p^m(c/\delta))$ or $\varepsilon \leq (1 - \delta)(v - p^m(c/\delta))$. This holds for all consumers with $v > p^m(c/\delta)$ in the limit as ε goes to 0, and $q_0 = ((\delta p^m(c/\delta) - c)/(\delta p^m(c/\delta)))D(p^m(c/\delta))$, which is strictly less than $1 - F(p^m(c/\delta))$. In addition, there exists $\varepsilon > 0$ such that more than q_0 consumers demand the good at

²⁴Corollary 2 applies only to the case $\bar{v} > c/\delta$. Proposition 9 shows that if $\bar{v} \leq c/\delta$, there is no full-delivery PBE in Γ or Γ_R in which the seller ever delivers the good or receives positive payments.

time 0 when $p_0 = p^m(c/\delta) + \varepsilon$. And the continuation path of play from $t = 1$ onward is the same under the modified strategy profile as under the best constant-price PBE, so the modified profile yields strictly higher profits overall. This yields part (iv) of [Theorem 1](#).

PROPOSITION 7 (Theorem 1(iv)). *If $\bar{v} > c/\delta$, there exists a full-delivery PBE of Γ_R (when cost equals c) yielding profits strictly greater than static monopoly profits when cost equals c/δ . By [Corollary 1](#), the same is true of full-delivery PBE of Γ .*

Before leaving this subsection, note that [Corollary 2](#) suggests that the best full-delivery PBE of the relational contracting model may involve pricing above the static monopoly level. We demonstrate this here in a simple, two-type example.²⁵

EXAMPLE 1. Suppose that half the consumers have valuation 2.36, while the other half have valuation 2.12. Let $c = .38$ and $\delta = .4$. Note that the static monopoly price is 2.12, as this yields profit 1.74, while setting price equal to 2.36 yields profit .99. In the dynamic model, the discussion preceding [Corollary 2](#) implies that the best full-delivery PBE with price fixed at 2.36 yields profit $(2.36 - \frac{.38}{.4}) \cdot .5 = .71$, while the best PBE with price fixed at 2.12 yields profit $(2.12 - \frac{.38}{.4}) \cdot 1 = 1.17$. Alternatively, one can check that setting $p_0 = 2.26$ and $p_t = 2.12$ for all $t \geq 1$, and selling to all high-valuation consumers in period zero and then selling to the low-valuation consumers at the fastest possible rate yields profit $(2.26 - .38) \cdot .5 + .4(2.12 - \frac{.38}{.4}) \cdot .5 = 1.174$. Furthermore, this is a PBE price–quantity path, as high-valuation consumers receive $2.36 - 2.26 = .1$ from purchasing in period zero and at most $.4(2.36 - 2.12) = .096$ from purchasing at a later date, while the seller gains $.4(2.12 - \frac{.38}{.4}) \cdot .5 = .234$ from delivering the good at time zero and gains $.38 * .5 = .19$ from failing to deliver.²⁶ Since this full-delivery PBE yields higher profit than the best PBE that fixes price at the monopoly price of 2.12, which is clearly the best PBE in which all prices are weakly below the monopoly price, the best full-delivery PBE in this example must have $p_t > p^m$ for some time t . \diamond

6.2 Declining prices

Finally, we establish three additional important properties of best full-delivery PBE of Γ and Γ_R , which hold for any fixed discount factor (parts (v)–(vii) of [Theorem 1](#)). We first use the possibility of rationing to ensure that best full-delivery PBE involve strictly decreasing price paths and positive sales each period. The idea is that delaying sales is wasteful and rationing can be used to ensure that speeding up sales does not violate the seller’s incentive compatibility constraint, which might otherwise be a concern.

²⁵This example does not exactly fit our model, as we have assumed a continuous distribution of valuations. However, the example can be slightly perturbed to yield a distribution that satisfies our assumptions, and noting that every best full-delivery PBE price path is decreasing (by [Proposition 8](#)), we conjecture that the best full-delivery PBE in the perturbed example will have $p_0 > p^m$.

²⁶Corresponding off-path play may be taken to be as in the strategy profile in the proof of [Proposition 6](#), for example.

PROPOSITION 8 (Theorem 1(v)). *If $\bar{v} > c/\delta$, any best full-delivery PBE of Γ or Γ_R has a strictly decreasing price path and strictly positive sales each period.*

PROOF. We prove the result for Γ_R , whence the result for Γ follows by Corollary 1. If $\bar{v} > c/\delta$, full-delivery PBE exist in which the seller makes positive profits (by Proposition 7), so any best full-delivery PBE of Γ_R yields positive profits.²⁷ Suppose that $(p, q)_t$ is such a best full-delivery PBE price path (which exists by Proposition 3). By Proposition 4, $D_\tau((p, q)_t) \leq q_\tau$ for all τ , so $Q_\tau = D_\tau((p, q)_t)$ for all τ . Suppose that there exists some time τ such that $D_\tau((p, q)_t) = 0$. Let t^* be the first such time. If $t^* = 0$, then define a new path by letting $p'_t = p_{t+1}$ and $q'_t = q_{t+1}$, i.e., shifting the original price-rationing path forward one period, which implies that $Q'_t = Q_{t+1}$, so profits under the new path are $1/\delta$ times profits under the original path, contradicting the optimality of the original path. If $t^* > 0$, let v_{t^*-1} be the lowest valuation such that a consumer with valuation v_{t^*-1} demands the good at $t^* - 1$, which is well defined because a positive measure of consumers demand the good at $t^* - 1$, by definition of t^* . We first claim that $v_{t^*-1} > p_{t^*-1}$. To see this, first note that a consumer with valuation v_{t^*-1} can demand the good at $t^* - 1$ only if $v_{t^*-1} \geq p_{t^*-1}$. If $v_{t^*-1} = p_{t^*-1}$, then it must be true that $p_\tau = p_{t^*-1}$ for all $\tau > t^* - 1$, since the price path is weakly decreasing by assumption, and if the price ever falls strictly below p_{t^*-1} , then all consumers with valuations sufficiently close to p_{t^*-1} prefer to wait until this time to purchase and all but at most a set of measure 0 of these consumers have the option of doing so since $D_\tau((p, q)_t) \leq q_\tau$ for all τ . The fact that $D_\tau((p, q)_t) \leq q_\tau$ for all τ then implies that $Q_\tau = 0$ for all $\tau > t^* - 1$, as all consumers prefer to purchase at $t^* - 1$ than at any later time. Therefore, continuation profits from time $t^* - 1$ onward equal 0, which implies that the seller does not deliver at $t^* - 1$. This in turn implies that no consumers pay at $t^* - 1$, so that continuation profits from time $t^* - 2$ onward equal 0 as well. By induction, continuation profits from time 0 onward are 0, contradicting the fact that any best full-delivery PBE yields positive profits if $\bar{v} > c/\delta$.

Now consider modifying $(p, q)_t$ by changing p_{t^*} to $(p_{t^*-1} - (1 - \delta)v_{t^*-1})/\delta$. Since $v_{t^*-1} > p_{t^*-1}$, we have $p_{t^*} < p_{t^*-1}$, and it is easy to check that all consumers with valuation weakly greater than v_{t^*-1} continue to demand the good at $t^* - 1$. By the skimming property (which is easily seen to hold due to declining prices and no rationing), the seller can sell a positive quantity at date $t^* + \tau$ only if $p_{t^*+\tau} < (p_{t^*-1} - (1 - \delta^{\tau+1})v_{t^*-1})/\delta^{\tau+1}$, so the seller strictly prefers selling to some mass of consumers at t^* at the new price to selling to them at any point in the future. Next observe that under the new price there is strictly positive demand at t^* , since at the new price a consumer with valuation v_{t^*-1} strictly prefers to purchase at $t^* - 1$ than to purchase at any other time except t^* , and is indifferent between purchasing at $t^* - 1$ and purchasing at t^* , so a consumer with valuation slightly below v_{t^*-1} strictly prefers purchasing at t^* to purchasing at any other time. Furthermore, the total sales at all future dates to consumers who do not buy at t^* is left unchanged, so total profits are strictly higher under the new path. Finally, the potential complication that the seller's incentive compatibility constraint may be violated at t^* can be addressed by rationing at t^* , since the necessity of positive continuation profits

²⁷See Proposition 9 for why this is not true if $\bar{v} < c/\delta$.

from t^* on implies that the seller can credibly sell a strictly positive quantity at t^* . So the modified path (possibly with rationing at t^*) strictly improves on the original path, contradicting the assumption that $D_\tau((p, q)_t) = 0$ for some τ .

We have shown that every best full-delivery PBE induces strictly positive sales at every date. Since every best full-delivery PBE involves no rationing, this is possible only if every best full-delivery PBE has a strictly declining price path. \square

We are now ready to complete the proof of [Theorem 1](#) by proving parts (vi) and (vii), which show that every best full-delivery PBE of Γ (or Γ_R) has an equilibrium price path $(p)_t$ that asymptotes to a price at least as high as $\max\{\underline{v}, c/\delta\}$ as t goes to infinity. The intuition is that a best full-delivery PBE has a declining price path by [Proposition 8](#), there is no reason to price below \underline{v} , and prices must be at least c/δ in any full-delivery PBE with a declining price path in which the seller ever delivers, in analogy with [Proposition 1](#). The following lemma formalizes the last part of this intuition.

LEMMA 6. *In any full-delivery PBE of Γ or Γ_R with price–quantity path $(p, q)_t$ in which $p_t \geq p_{t+1}$ for all t and a strictly positive quantity of the good is delivered along the equilibrium path, $p_t > c/\delta$ for all t .*

PROOF. Consider Γ_R first. Suppose that Q consumers have not yet received the good at time t^* . We first note that the seller's continuation profit from time t^* onward is bounded from above by her continuation profit from time t^* onward in a best full-delivery PBE of the modified continuation game where she is constrained to price weakly below p_{t^*} and all remaining consumers' valuations are set to p_{t^*} . This follows because in the modified game, the seller can set the original continuation price path $(p)_{t \geq t^*}$ and use rationing to sell according to the original price–quantity path.

The seller's continuation value at t^* in a full-delivery PBE of the modified game is therefore bounded from above by the solution to equation (3) with $p = p_{t^*}$. As shown in the proof of [Corollary 2](#), equation (3) has a solution with $V(Q) > 0$ if and only if $\delta p_{t^*} > c$. So if $p_{t^*} \leq c/\delta$, the seller's continuation value at t^* equals 0 in any full-delivery PBE in the modified game and, therefore, equals 0 in any full-delivery PBE of the unmodified game as well. This implies the seller delivers 0 units of the good at time t^* , which then implies that no buyers pay anything to the seller at time t^* , so the seller's continuation value at $t^* - 1$ equals 0 as well. By induction, the seller's continuation value equals 0 at all periods, and the seller never delivers a positive quantity of the good.

By [Proposition 3](#) and [Corollary 1](#), the above argument shows that in any full-delivery PBE of Γ with a declining price, the seller's continuation value starting from any t^* satisfying $p_{t^*} \leq c/\delta$ is 0. As above, this implies that the seller never delivers any positive quantity of the good. \square

PROPOSITION 9 ([Theorem 1](#)(vi) and (vii)). *Any best full-delivery PBE of Γ or Γ_R has $p_t > c/\delta$ and $p_t \geq \underline{v}$ for all t if $\bar{v} > c/\delta$. If $\bar{v} \leq c/\delta$, there is no PBE in Γ or Γ_R in which the seller ever delivers the good or receives positive payments.*

PROOF. If $\bar{v} > c/\delta$, the price path of any best full-delivery PBE of Γ or Γ_R is declining, by Proposition 8, and any full-delivery PBE with a declining price path has $p_t > c/\delta$ for all t , by Lemma 6. Finally, modifying any declining price path in Γ_R by replacing all $p_t < \underline{v}$ with \underline{v} and using rationing to ensure delivery yields a strict increase in profits if $p_t < \underline{v}$ for any t (as sales occur in every period in a best full-delivery PBE, by Proposition 8), so the result for $\bar{v} > c/\delta$ holds for Γ_R . Corollary 1 then implies that it also holds for Γ .

Suppose that $\bar{v} \leq c/\delta$ and that mass Q consumers have not yet received the good at some time t in Γ or Γ_R . If the seller delivers q units of the good at time t , she cannot receive total payments of more than $\bar{v}q$ and must of course be willing to deliver the q units. Therefore, her continuation payoff from time t onward is bounded from above by the solution to equation (3) with $p = \bar{v}$. As we have seen, the only solution to equation (3) when $\bar{v} \leq c/\delta$ is $V(Q) = 0$ for all Q . So no PBE in Γ or Γ_R yields positive profits if $\bar{v} \leq c/\delta$, which, as in the proof of Lemma 6, implies that no PBE involves delivery or positive payments. \square

7. AN EXTENSION: EXOGENOUS CHANCE OF NONDELIVERY

Our analysis is based on the assumption that the seller has the option of failing to deliver the good after receiving payment. We have argued that the presence of equilibria that yield high profits for the seller under this assumption suggests that sellers may try to avoid committing themselves to delivering the good. However, in some environments sellers may be automatically committed to delivery; for example, taking payment for a good and then failing to provide it may be viewed by courts as breaching an implicit contract, particularly if the seller has always provided the good to paying customers in the past (as is the case in full-delivery PBE). In this section, we show that our model can easily be extended to an environment in which this concern that the seller may be involuntarily committed to delivery does not apply. In particular, we assume that in every period there is an exogenous, independent probability $\eta > 0$ that the seller privately learns that she is unable to deliver the good after receiving payment.²⁸ For example, the seller may require certain specialized inputs to produce the final good, and these inputs may not always be available (and consumers and courts may be unable to observe whether the inputs are available). In this model, the seller periodically fails to deliver the good even if she wishes to deliver in every period, and since courts cannot tell whether failure to deliver results from lack of inputs or opportunistic behavior by the seller, there is no possibility that the seller can be involuntarily committed to trying to deliver the good in every period.

The equilibria we have constructed for both the nondurable- and durable-goods models can easily be adapted to this environment by specifying that no purchases or delivery occur after any nondelivery by the seller (so that trade eventually breaks down on the equilibrium path) and that prior to the breakdown of trade, consumers take into account that they receive the good only with probability $1 - \eta$ even if they pay (since

²⁸I thank the editor for suggesting I pursue this analysis.

consumers are risk-neutral, this implies that the mass of consumers who wish to purchase at price p is now $D(p/(1 - \eta))$ rather than $D(p)$. That is, our results are continuous in η . Rather than formally stating this rather natural finding, we instead focus on characterizing the best full-delivery, constant-price PBE in Γ_R , in analogy to [Corollary 2](#), which provides an intuitive lower bound on the best full-delivery PBE profit in both Γ and Γ_R . It turns out that the analysis of [Section 6.1](#) carries through with the sole modification that $D(p)$ is replaced by $D(p/(1 - \eta))$: the intuition for this result is that, in the best full-delivery PBE, the seller is indifferent between delivering the good and breaking off trade, so she is not made worse off by the possibility that trade may break off exogenously (except insofar as this causes consumers with valuations $v \in [p, p/(1 - \eta)]$ to reject her price offer). Finally, we remark that our original definition of a full-delivery PBE does not allow for the possibility that trade breaks down in equilibrium, which leads us to use the following, somewhat ad hoc, definition in the statement of the result.

DEFINITION 3. A *modified full-delivery PBE* is a PBE in which the seller sets $x_t = 1$ at all on-path histories at which $Q_t > 0$ and sets $x_t = 0$ at all on-path histories at which $Q_t = 0$.

Our earlier results pertaining to full-delivery PBE (in particular, [Corollary 1](#)) also apply to modified full-delivery PBE.

PROPOSITION 10. *If $\bar{v} > c/\delta$, the best modified full-delivery, constant-price PBE in Γ_R is given by $p_t = \arg \max_p (p - c/\delta)D(p/(1 - \eta)) \equiv p^*(\eta)$ and $q_t = \gamma^*(p^*(\eta)) \times (1 - \gamma^*(p^*(\eta)))^t D(p^*(\eta)/(1 - \eta))$, where $\gamma^*(p) = (\delta p - c)/(\delta p)$ as in [Section 6.1](#).*

PROOF. A consumer who demands the good at price p receives it with probability at most $1 - \eta$, so at most $D(p/(1 - \eta))$ consumers ever purchase in a full-delivery PBE with constant price p . The argument in the proof of [Corollary 2](#) shows that if the seller faces this demand curve and can freely choose what quantity to deliver in every period, her best (modified) full-delivery PBE profit with constant price p equals $(p - c/\delta)D(p/(1 - \eta))$. Therefore, $(p^*(\eta) - c/\delta)D(p^*(\eta)/(1 - \eta))$ is an upper bound on the seller's best modified full-delivery, constant-price PBE profit when, in each period, she may be unable to deliver the good with probability η .

We claim that the following strategy profile attains this upper bound: the seller sets (p_t, q_t) as in the statement of the proposition, sets $x_t = 1$ until the first time that delivery is impossible, and subsequently sets $x_t = 0$; a consumer with valuation v demands the good if and only if $v \geq p/(1 - \eta)$ and the seller has always set $p_t = p^*(\eta)$ and delivered the good in the past. The only nontrivial part of verifying that this profile is a PBE is checking that it is optimal for the seller to deliver the good when prescribed. Nondelivery leads to continuation payoff 0 and in every period prior to the first nondelivery, the seller fails to deliver with probability η . Therefore, the condition that it is optimal for the seller to deliver the good when prescribed at time t is

$$\sum_{\tau=1}^{\infty} (1 - \eta)^{\tau-1} \delta^{\tau} q_{t+\tau} (p_{t+\tau} - c(1 - \eta)) \geq q_t c.$$

Substituting in the specified (p_t, q_t) yields

$$\begin{aligned} \gamma^*(p^*(\eta))(1 - \gamma^*(p^*(\eta)))^t & \left(\frac{\delta(1 - \gamma^*(p^*(\eta)))}{1 - \delta(1 - \gamma^*(p^*(\eta)))(1 - \eta)} \right) D\left(\frac{p^*(\eta)}{1 - \eta}\right) \\ & \times (p^*(\eta) - c(1 - \eta)) \\ & \geq \gamma^*(p^*(\eta))(1 - \gamma^*(p^*(\eta)))^t D\left(\frac{p^*(\eta)}{1 - \eta}\right) c \end{aligned}$$

or

$$\left(\frac{\delta(1 - \gamma^*(p^*(\eta)))}{1 - \delta(1 - \gamma^*(p^*(\eta)))(1 - \eta)} \right) (p^*(\eta) - c(1 - \eta)) \geq c.$$

This can be rewritten as

$$\gamma^*(p^*(\eta)) \leq \frac{\delta p^*(\eta) - c}{\delta p^*(\eta)},$$

exactly as in (1), which holds by definition of $\gamma^*(p^*(\eta))$. This verifies that the above strategy profile is a modified full-delivery, constant-price PBE, and it is straightforward to check that it yields expected profit $(p^*(\eta) - c/\delta)D(p^*(\eta)/(1 - \eta))$. \square

Thus, [Proposition 10](#) shows that the lower bound on optimal monopoly profits derived in [Section 6.1](#) extends naturally to environments with an exogenous change of nondelivery, where it may be more realistic to view the seller as having the option of nondelivery.

8. CONCLUSION

The main insight of this paper is that the optimal pricing strategy of a dynamic monopoly may be very different from that in traditional models when the relationship between the seller and consumers is regulated by relational incentives. Unlike in [Hart and Tirole \(1988\)](#), a nondurable-goods monopoly in our model can earn high profits even if consumers are nonanonymous, provided the discount factor is sufficiently high. Unlike in [Coase \(1972\)](#), a durable-goods monopoly can earn approximately static monopoly profits in the limit as the discount factor approaches 1, even if the lowest consumer valuation is above the marginal cost of production. A durable-goods monopoly can also earn high profits when the discount factor is bounded away from one.

While our model has many equilibria, restricting attention to the best equilibria for the seller brings out some novel economic intuitions and empirical predictions. First, for both nondurable- and durable-goods monopolies, the temptation to fail to deliver provides an incentive for pricing above the static monopoly level.²⁹ The intuition is the same in both cases: The larger the quantity of the good a monopoly is supposed to deliver, the greater is its incentive to renege. So the monopoly benefits from restricting quantity, and the most profitable way for it to restrict quantity is to raise price. Second, in the durable-goods case, the monopoly has an incentive to gradually cut prices

²⁹This possibility that dynamic monopolies may price higher than static monopolies is a prediction of our model that differs from standard models of dynamic monopoly pricing.

over time, using high prices rather than rationing to restrict sales early on. These new effects have potentially interesting applications for regulation. In traditional models, observing a monopoly cutting its price is a sign that consumers are doing better than they would be if the monopoly had full commitment power, since they are paying lower prices and (if the discount factor is high) are not facing costly delays in purchasing. In our model, however, consumers may be better off when the monopoly has full commitment power for two reasons: they may face lower prices (since without commitment, the monopoly may price above the static monopoly price) and they may receive the good significantly faster. This also points to an important empirical prediction of our model: in contrast to the standard full-commitment and no-commitment models of durable-goods monopoly, our model predicts that a monopoly will cut prices over time, but will do so slowly enough that the costs from delay are significant.

We also introduce two methodological innovations. First, we use an augmented model with rationing to help analyze the durable-goods seller problem. This greatly simplifies the analysis by allowing us to construct simple equilibria with flat price paths in the model with rationing, and then use the relationship between the model with and without rationing to draw conclusions about best full-delivery equilibria in the model without rationing. Second, and more generally, we use relational incentives to replace the temptation to deviate at the contract offer stage (price offers in our model) with the temptation to deviate at the contract execution stage (delivery of the good in our model), which may have applications to other areas where studying dynamics in the presence of adverse selection has proved difficult. For example, recall that in our model of non-durable goods and nonanonymous consumers, the dynamic enforcement constraint that the seller delivers the good replaced the ratchet effect in price-setting. Perhaps further insights may be gained from applying this idea to dynamic principal-agent problems with adverse selection, where characterizing dynamics in models with no commitment is difficult due to the ratchet effect (see, e.g., [Laffont and Tirole 1988](#)).

APPENDIX A: OMITTED PROOFS

PROOF OF PROPOSITION 1. We first observe that the problem of finding the best PBE for the seller is equivalent to finding the best PBE for the seller when she can fully commit to her sequence of prices $(\{p_{t,n}\}_n)_t$. To see this, note that we can specify off-path beliefs for buyers such that each buyer expects the seller to never deliver the good following any deviation in price-setting by the seller. Given these beliefs, no buyer will ever accept a strictly positive price in any period following a deviation in price-setting by the seller, so the seller always receives continuation payoff zero, equal to her minmax payoff, after any such deviation.

Using this observation and applying the revelation principle to each period, we can write the problem in a standard mechanism design notation, writing T for transfers:

$$\max_{\{T_t(\cdot), x_t(\cdot)\}_t} \sum_{t=0}^{\infty} \delta^t \int_{\underline{v}}^{\bar{v}} (T_t(v) - cx_t(v))f(v) dv$$

subject to

$$vx_t(v) - T_t(v) \in \arg \max_{v'} vx_t(v') - T_t(v') \quad \text{for all } v \text{ and } t \tag{IC}$$

$$vx_t(v) - T_t(v) \geq 0 \quad \text{for all } v \text{ and } t \tag{IR}$$

and

$$\sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} (T_{t+\tau}(v) - cx_{t+\tau}(v))f(v) dv \geq c \int_{\underline{v}}^{\bar{v}} x_t(v)f(v) dv \quad \text{for all } t. \tag{DE}$$

Note that the third constraint is the seller’s incentive compatibility (IC) constraint, which we also refer to as the dynamic enforcement (DE) constraint. Substituting for $T_t(v)$ using the (IR) and (IC) constraints in the usual way, and temporarily ignoring the resulting monotonicity constraint lets us rewrite the problem as

$$\max_{\{x_t(\cdot)\}_t} \sum_{t=0}^{\infty} \delta^t \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x_t(v) dv$$

subject to the (DE) constraint

$$\sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x_{t+\tau}(v) dv \geq c \int_{\underline{v}}^{\bar{v}} x_t(v)f(v) dv \quad \text{for all } t.$$

Let $\{x_t^*(v)\}_t$ be a solution to this problem. Note that, for all t , $x_t^*(v)$ must solve

$$\max_{x_t(\cdot)} \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x_t(v) dv$$

subject to

$$\sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x_{t+\tau}^*(v) dv \geq c \int_{\underline{v}}^{\bar{v}} x_t(v)f(v) dv,$$

since the solution to this program maximizes both the original objective and the left-hand side of each original constraint over all $x_t(\cdot)$ that satisfy the original time t constraint. This implies that, for all t, t' , if $(v - c)f(v) - (1 - F(v)) > 0$, then $x_t^*(v) > x_{t'}^*(v)$ if $\sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} x_{t+\tau}^*(v) dv > \sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} x_{t'+\tau}^*(v) dv$, while if $(v - c)f(v) - (1 - F(v)) < 0$, then $x_t^*(v) = x_{t'}^*(v) = 0$. Since $\sum_{\tau=1}^{\infty} \delta^{t+\tau} \int_{\underline{v}}^{\bar{v}} x_{t+\tau}^*(v) dv$ is bounded from above, there exists a finite $x^*(\cdot)$ such that $x^*(v) = \sup_t x_t^*(v)$ if $(v - c)f(v) - (1 - F(v)) \geq 0$ and $x^*(v) = 0$ otherwise.

We claim that $x_t(v) = x^*(v)$ for all t and v in any solution to this problem. Clearly, the profit corresponding to this allocation is an upper bound on the profit in any solution. Furthermore,

$$\begin{aligned} \sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x^*(v) dv \\ = \sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v))) \sup_t x_t^*(v) dv \end{aligned}$$

$$\begin{aligned} &\geq \sup_t \sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x_{t+\tau}^*(v) dv \\ &\geq \sup_t c \int_{\underline{v}}^{\bar{v}} x_t^*(v)f(v) dv \\ &= c \int_{\underline{v}}^{\bar{v}} x^*(v)f(v) dv, \end{aligned}$$

where the first equality is by the definition of $x^*(v)$, the second inequality is immediate, the third inequality follows because $\{x_t^*(v)\}_t$ satisfies the (DE) constraint for all t , and the fourth equality follows because $x_t^*(v) \geq x_{t'}^*(v)$ if and only if $x_t^*(v') \geq x_{t'}^*(v')$ for any t, t', v , and v' , so the sup may be moved inside the integral. The above chain of inequalities implies that repeating $x^*(v)$ satisfies the seller's incentive compatibility constraint. Finally, if there exists t such that $x_t^*(v) \neq x^*(v)$, then the allocation $\{x_t^*(v)\}_t$ yields strictly lower profit than repeating $x^*(v)$ in period t and yields weakly lower profit in all other periods, so every solution to the original problem has the same allocation rule in every period.

We have shown that the optimal allocation rule is stationary, so the problem becomes

$$\max_{x(\cdot)} \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x(v) dv$$

subject to the (DE) constraint

$$\sum_{\tau=1}^{\infty} \delta^\tau \int_{\underline{v}}^{\bar{v}} ((v - c)f(v) - (1 - F(v)))x(v) dv \geq c \int_{\underline{v}}^{\bar{v}} x(v)f(v) dv \quad \text{for all } t.$$

The (DE) constraint may be rewritten as

$$\int_{\underline{v}}^{\bar{v}} (vf(v) - (1 - F(v)))x(v) dv \geq \left(\frac{c}{\delta}\right) \int_{\underline{v}}^{\bar{v}} x(v)f(v) dv.$$

If the constraint is slack, we have standard monopoly pricing. If the constraint is binding, noting that our assumptions on $F(v)$ imply that $x(v)$ continues to take a cutoff form whereby $x(v) = 0$ if $v < v^*$ and $x(v) = 1$ if $v \geq v^*$ for some v^* yields that, for any $v \geq v^*$, price equals $v - \int_{v^*}^v x(s) ds = v^*$. Additionally, complementary slackness implies that the constraint is binding if and only if $p^m \leq c/\delta$. Finally, note that in any case these solutions satisfy the monotonicity constraint.

If $\bar{v} < c/\delta$, then $v^* > \bar{v}$, so $x(v) = 0$ for all v , which implies that the seller never delivers the good or receives positive payments in any optimal PBE. Since the seller's minmax payoff is zero, every PBE is optimal if $\bar{v} < c/\delta$, which proves the result in the $\bar{v} < c/\delta$ case. \square

PROOF OF PROPOSITION 2. The proof is similar to the proof of Proposition 3, so we omit some details. Let \mathcal{F} be the set of PBE price–quantity–delivery paths $(p, Q, x)_t$ satisfying $p_t \in [\underline{v}, \bar{v}]$ for all t . If a PBE is optimal in the set of PBE with price–demand–delivery

paths in \mathcal{F} , then it is optimal overall. Furthermore, it is clear that the seller's PBE payoff is continuous in price–quantity–delivery paths $(p, Q, x)_t$ in the product topology.

Next, we note that the continuation value of a consumer with valuation v facing price–quantity–delivery path $(p, Q, x)_t$ at time t is continuous in $(p, Q, x)_t$ in the product topology.³⁰ To see this, observe that the maximum gain in continuation value over an ε ball about $(p, Q, x)_t \in \mathcal{F}$ is no more than $(1 + \bar{v})\varepsilon/(1 - \delta)$, corresponding to receiving the good, valued at \bar{v} , with additional probability ε in each period, and paying ε less in each period. This converges to 0 as ε does.

We now show that \mathcal{F} is compact in the product topology. Observe that $\mathcal{F} \subseteq \prod_{t=0}^{\infty}([\underline{v}, \bar{v}], [0, 1], [0, 1])_t$, which is compact by Tychonoff's Theorem. Therefore, it suffices to show that \mathcal{F} is closed in the product topology. To see that it is, consider a sequence of paths $\{(p, Q, x)_t\}_n \in \mathcal{F}$ converging pointwise to $(p^*, Q^*, x^*)_t$. We must show that there exists a PBE with price–demand–delivery path $(p^*, Q^*, x^*)_t$. Consider the following strategy profile.

1. The seller sets price–delivery path $(p^*, x^*)_t$ as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$ and $x_t = 0$ for all future periods. In particular, the seller sets $x_t = 0$ in any period in which she has set $p_t \neq p_t^*$.
2. A consumer with valuation v who has not yet received the good at t pays at t if and only if the seller has never deviated from her prescribed strategy and $x_t^*v - p_t^* \geq \delta x_t^*V_{t+1}^v$, where V_{t+1}^v is the continuation value of such a consumer facing $(p^*, Q^*, x^*)_t$.

The proof that the seller's play is optimal is as in the proof of Proposition 3. To see that each consumer's play is optimal, first note that it is obvious that her off-path play is optimal, as paying is costly and yields no benefit when the seller sets $x_t = 0$. To see that her on-path play is optimal given $(p^*, Q^*, x^*)_t$, note that paying at t gives expected payoff $x_t^*v - p_t^* + \delta(1 - x_t^*)V_{t+1}^v$, while not paying gives δV_{t+1}^v , so paying is optimal if and only if $x_t^*v - p_t^* \geq \delta x_t^*V_{t+1}^v$.

That the prescribed consumer behavior induces quantity path $(Q^*)_t$ follows from the observation that each consumer's payoff is continuous in $(p, Q, x)_t$ and that each consumer plays a best response to each $(p, Q, x)_{t,n}$ in equilibrium. This completes the argument that \mathcal{F} is closed, and therefore compact, in the product topology. Weierstrass' Theorem then implies that there is a point in \mathcal{F} that maximizes profits, completing the proof. \square

PROOF OF LEMMA 3. Fix a price–rationing path $(p, q)_t$ and two full-delivery PBE σ and σ' . Let \mathcal{V} be the set of generic valuations v such that there exists a consumer with valuation v who purchases at different times under σ and σ' . Suppose, toward a contradiction, that \mathcal{V} is nonempty. Then \mathcal{V} has a supremum, which we denote by v^* . Let V_{v^*} be the payoff of a consumer with valuation v^* under σ , let V'_{v^*} be the payoff of a consumer with valuation v^* under σ' , and, without loss of generality, assume that $V_{v^*} \geq V'_{v^*}$.

³⁰This continuation value is well defined here by standard dynamic programming arguments, because, unlike in the model with rationing, each consumer faces the same optimization problem regardless of the behavior of the other consumers.

We first claim that $V_{v^*} = V'_{v^*}$. To see this, suppose that there exists a consumer with valuation v^* who purchases at time τ_{v^*} under σ and purchases at time $\tau'_{v^*} \neq \tau_{v^*}$ under σ' , with $\delta^{\tau_{v^*}}(v^* - p_{\tau_{v^*}}) > \delta^{\tau'_{v^*}}(v^* - p_{\tau'_{v^*}})$, so that the consumer receives a higher payoff under σ . This is possible only if the consumer is unable to purchase at time τ_{v^*} under σ' , which in turn is possible only if strictly more than $q_{\tau_{v^*}}$ consumers accept price $p_{\tau_{v^*}}$ at time τ_{v^*} under σ' . Since the consumer is able to purchase at time τ_{v^*} under σ , which is possible only if no more than $q_{\tau_{v^*}}$ consumers accept price $p_{\tau_{v^*}}$ at time τ_{v^*} under σ , this implies that there is a positive measure μ of consumers with valuations greater than v^* who purchase at τ_{v^*} under σ' but not under σ . By Lemma 1, this implies that there exists a consumer with valuation $v' > v^*$ and v' generic with respect to $(p)_t$ who purchases at different times under σ and σ' , which contradicts the fact that $v^* = \sup\{v : v \in \mathcal{V}\}$. This implies that $V_{v^*} = V'_{v^*}$, which also implies that $v^* \notin \mathcal{V}$, because if $V_{v^*} = V'_{v^*}$, then either every consumer with valuation v^* purchases at the same time under σ and σ' or v^* is nongeneric with respect to $(p)_t$.

If $V_{v^*} = V'_{v^*} = 0$, then there is no time t at which $v^* > p_t$ and a consumer with valuation v^* is able to purchase under either σ or σ' . This implies that there is no time t at which $v \geq p_t$ and a consumer with valuation v is able to purchase under either σ or σ' , for any $v \in \mathcal{V}$, because $v < v^*$ for all $v \in \mathcal{V}$ and a consumer with a lower valuation is able to purchase at a weakly smaller set of times. Therefore, a consumer with valuation v never purchases under either σ or σ' , for all $v \in \mathcal{V}$, which implies that \mathcal{V} is empty, a contradiction.

If $V_{v^*} = V'_{v^*} > 0$, then for any $\eta \in (0, V_{v^*})$ there exist at most finitely many times t such that there exists $v \in [\underline{v}, \bar{v}]$ such that $\delta^t(v - p_t) \geq V_{v^*} - \eta$ and $q_t > \tilde{Q}_t$, where \tilde{Q}_t is the measure of consumers who purchase at time t under σ and have valuations greater than v^* (as $p_t \geq 0$ for all t); call the set of such times \mathcal{T} . Let $\varepsilon_t \equiv q_t - \tilde{Q}_t$ and let $\varepsilon \equiv \min\{\varepsilon_t : t \in \mathcal{T}\}/2 > 0$. Since every consumer with generic valuation greater than v^* purchases at the same time under σ and σ' , by definition of \mathcal{V} , and the set of consumers with nongeneric valuations is of measure 0, by Lemma 1, the measure of consumers with valuations greater than $v^* - \varepsilon$ who purchase at any t under σ' is less than $\tilde{Q}_t + \varepsilon$. By definition of ε , this implies that the measure of consumers with valuations greater than $v^* - \varepsilon$ who purchase at any $t \in \mathcal{T}$ under σ' is less than q_t . So any consumer with valuation $v > v^* - \varepsilon$ can purchase at any time t with $\delta^t(v - p_t) \geq V_{v^*} - \eta$ under σ' at which she can purchase under σ . By the same argument, there exists $\varepsilon' > 0$ such that a consumer with valuation $v > v^* - \varepsilon'$ can purchase at any time t with $\delta^t(v - p_t) \geq V_{v^*} - \eta$ under σ at which she can purchase under σ' . Therefore, letting $\varepsilon'' \equiv \min\{\varepsilon, \varepsilon'\}$, we see that a consumer with valuation $v > v^* - \varepsilon''$ can purchase at the same set of times t with $\delta^t(v - p_t) \geq V_{v^*} - \eta$ under σ and σ' . Furthermore, a consumer with valuation close enough to v^* can purchase at any time at which a consumer with valuation v^* can purchase (by our specification of rationing), so there exists ε^* such that a consumer with valuation $v > v^* - \varepsilon^*$ receives a payoff of at least $V_{v^*} - \eta$ under both σ and σ' . Finally, by definition of v^* , there exists $v \in \mathcal{V}$ such that $v > v^* - \min\{\varepsilon'', \varepsilon^*\}$. A consumer with valuation v receives a payoff of at least $V_{v^*} - \eta$ under both σ and σ' , which implies that he purchases at a time t with $\delta^t(v - p_t) \geq V_{v^*} - \eta$ under both σ and σ' . The set of such times at which the consumer can purchase is the same under σ and σ' . Since v is generic

with respect to $(p)_t$, the consumer has a strict preference ordering over purchase times, which implies that he purchases at the same time under σ and σ' , which contradicts the assumption that $v \in \mathcal{V}$. \square

PROOF OF LEMMA 5. Consider the problem of maximizing Q_t over price-rationing paths $(p', q')_t$ in an ε ball about $(p, q)_t$. As $\varepsilon \rightarrow 0$, the measure of consumers who have different preference orderings over purchase times (i.e., over the $\{\delta^t(v - p_t)\}_t$) under $(p', q')_t$ and $(p, q)_t$ converges to 0. Furthermore, the maximum difference between Q_τ and a Q'_τ corresponding to $(p, q')_t$ in an ε ball about $(p, q)_t$ (holding $(p)_t$ fixed) is no more than $\sum_{t=0}^\infty \max\{\varepsilon, Q_t\}$, the maximum measure of consumers whose purchasing times can be affected by decreasing q_t by ε for all t and holding other consumers' purchasing times fixed; this follows because if rationing prevents measure μ consumers from purchasing at some time t , each of these consumers cannot alter his play in a way that leads more than one total consumer to purchase at time τ (i.e., he can purchase at time τ himself or he can displace one other consumer through rationing at some other time).³¹ Thus, the maximum variation in Q_t over an ε ball about $(p, q)_t$ converges to $\lim_{\varepsilon \rightarrow 0} \sum_{t=0}^\infty \max\{\varepsilon, Q_t\}$ as $\varepsilon \rightarrow 0$, so the following technical lemma completes the proof.

LEMMA 7. *Given any quantity path $(Q)_t$, $\lim_{\varepsilon \rightarrow 0} \sum_{t=0}^\infty \max\{\varepsilon, Q_t\} = 0$.*

PROOF. First, note that

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} \sum_{t=0}^\infty \max\{\varepsilon, Q_t\} &= \left(\lim_{\varepsilon \rightarrow 0} \varepsilon \#\{t : Q_t > \varepsilon\} \right) + \left(\lim_{\varepsilon \rightarrow 0} \sum_{t: Q_t < \varepsilon} Q_t \right) \\ &= \lim_{\varepsilon \rightarrow 0} \varepsilon \#\{t : Q_t > \varepsilon\}. \end{aligned}$$

Let $N_\varepsilon \equiv \#\{t : Q_t > \varepsilon\}$ to simplify notation. Assume, toward a contradiction, that the lemma is false, i.e., that there exists $\delta > 0$ such that for all $\bar{\varepsilon} > 0$ there exists $\varepsilon < \bar{\varepsilon}$ satisfying $\varepsilon N_\varepsilon > \delta$. Fix such a $\delta > 0$ and let $\varepsilon_0 > 0$ satisfy $\varepsilon_0 N_0 > \delta$. Now for all $n \geq 1$, let $\bar{\varepsilon}_n = \varepsilon_{n-1}/2^n$ and let ε_n be a strictly positive number strictly less than $\bar{\varepsilon}_n$ satisfying $\varepsilon_n N_{\varepsilon_n} > \delta$. Note that $\varepsilon_n/\varepsilon_{n-1} < 1/2^n$.

Observe that, for any n , $N_{\varepsilon_n} < 1/\varepsilon_n$, for otherwise the total quantity of sales made in the N_{ε_n} periods in which $Q_t > \varepsilon_n$ would exceed 1. Since $N_{\varepsilon_n} < 1/\varepsilon_n$ and $\varepsilon_{n+1} N_{\varepsilon_{n+1}} > \delta$, we have that $N_{\varepsilon_{n+1}} - N_{\varepsilon_n} > \delta/\varepsilon_{n+1} - 1/\varepsilon_n$. Now $N_{\varepsilon_{n+1}} - N_{\varepsilon_n}$ is the number of periods in which Q_t is between ε_{n+1} and ε_n , so total sales made in all periods is at least

$$\begin{aligned} \sum_{n \geq 0} (N_{\varepsilon_{n+1}} - N_{\varepsilon_n}) \varepsilon_{n+1} &> \sum_{n \geq 0} \left(\delta - \frac{\varepsilon_{n+1}}{\varepsilon_n} \right) \\ &> \sum_{n \geq 0} \left(\delta - \frac{1}{2^{n+1}} \right) \\ &= \infty. \end{aligned}$$

³¹We omit the measure-theoretic details of this argument, which are similar to those in the proof of Proposition 4.

This contradicts the assumption that the population of consumers is of measure 1. \square

PROOF OF PROPOSITION 4. Suppose that rationing occurs at time t^* along a full-delivery PBE path $(p, q)_t$. We show that $(p, q)_t$ cannot be a best full-delivery PBE path.

First, consider the path $(p', q')_t$ given by $p'_t = p_t$ for all t and $q'_t = Q_t$ for all t , where $(Q)_t$ is the unique (by Lemma 4) quantity path corresponding to $(p, q)_t$. All consumers are best-responding if they purchase at the same time under $(p', q')_t$ as they did under $(p, q)_t$, and by Lemma 3, this purchasing schedule is unique up to the measure-zero set of consumers who are indifferent between purchasing at different times, so the seller's profit is the same in any full-delivery PBE corresponding to $(p', q')_t$ and in any full-delivery PBE corresponding to $(p, q)_t$. Furthermore, $D_{t^*}((p', q')_t) > q'_{t^*}$. Since F admits a strictly positive density, there is a small enough strict increase in p_{t^*} , Δ_p , such that demand at t^* still exceeds q_{t^*} when price at t^* is increased by Δ_p . So consider the path $(p^*, q^*)_t$ given by $p^*_{t^*} = p_{t^*} + \Delta_p$, $p^*_t = p_t$ for all $t \neq t^*$, and $q^*_t = Q_t$ for all t . We claim that $Q^*_t = Q_t$ for all t , which then implies that profit is higher under $(p^*, q^*)_t$ than under $(p', q')_t$ (and therefore $(p, q)_t$), since $Q_{t^*} > 0$ (by the definition of rationing occurring at t^*).

Since $q^*_t = Q_t$ for all t , we have $Q^*_t \leq Q_t$ for all t , so since $Q^*_{t^*} = Q_{t^*}$ by definition of Δ_p , it suffices to show that $\sum_{t \neq t^*} Q^*_t \geq \sum_{t \neq t^*} Q_t$. Suppose, toward a contradiction, that $\sum_{t \neq t^*} Q_t - \sum_{t \neq t^*} Q^*_t \equiv \mu > 0$. For any $\tau \neq t^*$, if $Q_\tau - Q^*_\tau \equiv \mu_\tau > 0$, then $D_\tau((p^*, q^*)_t) = q^*_\tau - \mu_\tau$. Since the price at τ is the same under $(p^*, q^*)_t$ and $(p, q)_t$, this is possible only if there are measure μ_τ consumers who demanded the good at τ under $(p, q)_t$ and have higher PBE payoffs under $(p^*, q^*)_t$. Since prices are weakly higher in each period under $(p^*, q^*)_t$, this implies that at least μ_τ consumers who purchase at τ under $(p, q)_t$ must purchase at times under $(p^*, q^*)_t$ at which they could not purchase under $(p, q)_t$. This argument applies to all τ such that $\mu_\tau > 0$, so at least $\mu = \sum_t \mu_t$ consumers purchase at times under $(p^*, q^*)_t$ at which they could not purchase under $(p, q)_t$, and they receive higher payoffs under $(p^*, q^*)_t$. Let \mathcal{D} be the set of consumers who purchase at times under $(p^*, q^*)_t$ at which they could not purchase under $(p, q)_t$ and receive higher payoffs under $(p^*, q^*)_t$. Now measure μ of consumers can purchase at times under $(p^*, q^*)_t$ at which none of them can purchase under $(p, q)_t$ only if there exists a measure-preserving injection $\psi: \mathcal{D} \rightarrow [\underline{v}, \bar{v}]$ (mapping consumers who do better under $(p^*, q^*)_t$ to consumers they displace) from these consumers to another set of consumers of mass μ satisfying the following conditions.

- (i) For all $v \in \mathcal{D}$, $\psi(v) > v$.
- (ii) If a consumer with generic (with respect to $(p^*)_t$) valuation v purchases at time t under $(p^*, q^*)_t$, then every consumer with valuation $\psi(v)$ purchases at time t under $(p, q)_t$, and (since ψ is measure-preserving) for every t the measure of consumers in the preimage who purchase at time t under $(p^*, q^*)_t$ equals the measure of consumers in the image who purchase at time t under $(p, q)_t$.
- (iii) A consumer in the image of ψ who purchases at time t under $(p, q)_t$ purchases at some time $t' \neq t$ under $(p^*, q^*)_t$.

Note that each of the consumers in the image of ψ retains, under $(p^*, q^*)_t$, the option of purchasing at the same time at which he purchased under $(p, q)_t$, because his valuation is higher than that of the corresponding consumer in the preimage, so since he does not do so, it must be either that he purchases at a time t' at which he could not purchase under $(p^*, q^*)_t$ and receives a higher payoff under $(p^*, q^*)_t$ or that $t = t^*$, in which case purchasing at t has become less attractive. That is, if a consumer is in the image of ψ , then either he is also in \mathcal{D} (the preimage of ψ) or he purchases at t^* under $(p, q)_t$ but not under $(p^*, q^*)_t$. Iterating the procedure of constructing such a measure-preserving injection from consumers who purchase at different times under $(p^*, q^*)_t$ and (p, q) , and receive higher payoffs under $(p^*, q^*)_t$ to the consumers they displace implies that there are μ consumers who did not purchase at t^* under $(p, q)_t$ who do purchase at t^* under $(p^*, q^*)_t$, that all of them receive higher payoffs under $(p^*, q^*)_t$ than under $(p, q)_t$, and that a measure-preserving bijection satisfying conditions (i)–(iii) exists between the set of consumers who receive a higher payoff under $(p^*, q^*)_t$ than under $(p, q)_t$ and the set of consumers who purchase at t^* under $(p, q)_t$ who do not purchase at t^* under $(p^*, q^*)_t$.

By the preceding paragraph, the measure of consumers who purchase at t^* under $(p, q)_t$ who do not purchase at t^* under $(p^*, q^*)_t$ is at least μ . Since all consumers who purchase at t^* under $(p^*, q^*)_t$ but not under $(p, q)_t$ receive a higher payoff under $(p^*, q^*)_t$, it follows that every consumer who purchases at t^* under $(p, q)_t$ has a higher valuation than any of these consumers, and, therefore, has a higher valuation than any consumer who receives a higher payoff under $(p^*, q^*)_t$ than under $(p, q)_t$. Therefore, every consumer who purchases at t^* under $(p, q)_t$ but not under $(p^*, q^*)_t$ prefers to purchase at any τ satisfying $\mu_\tau > 0$ to never purchasing. Furthermore, suppose that mass ε of such consumers, with valuations with infimum v , purchase at time $\hat{\tau}$ satisfying $\mu_{\hat{\tau}} = 0$ under $(p^*, q^*)_t$. Then there must exist mass ε of consumers each with valuation strictly less than v who purchase at $\hat{\tau}$ under $(p, q)_t$ but not under $(p^*, q^*)_t$. Consider such a consumer with valuation $v' < v$, fix any τ satisfying $\mu_\tau > 0$, and suppose toward a contradiction that $v' < p_\tau$. We have that $v > p_\tau$, a consumer with valuation v prefers purchasing at time $\hat{\tau}$ and price $p_{\hat{\tau}}$ to purchasing at time τ and price p_τ (by revealed preference at $(p^*, q^*)_t$, since there is no rationing at τ under $(p^*, q^*)_t$), and a consumer with valuation v' also prefers purchasing at time $\hat{\tau}$ and price $p_{\hat{\tau}}$ to purchasing at time τ and price p_τ (since $v' > p_{\hat{\tau}}$ by revealed preference at $(p, q)_t$ and $v' < p_\tau$ by assumption). Now there also exists a consumer who purchases at time τ and price p_τ under $(p, q)_t$ and obtains a higher payoff under $(p^*, q^*)_t$, since $\mu_\tau > 0$. Such a consumer must have valuation $v'' \in [p_\tau, v)$, so $v'' > v'$, which implies that such a consumer has the option of purchasing at $\hat{\tau}$ under $(p, q)_t$. Therefore, such a consumer must prefer purchasing at time τ and price p_τ to purchasing at time $\hat{\tau}$ and price $p_{\hat{\tau}}$. Thus, the assumption that $v' < p_\tau \leq v'' < v$ yields a violation of single crossing. Therefore, each of the μ consumers who purchases at t^* under $(p, q)_t$ but not under $(p^*, q^*)_t$ either purchases at a τ such that $\mu_\tau > 0$ under $(p^*, q^*)_t$ or else displaces another consumer who prefers to purchase at any such τ to never purchasing. So the measure of consumers who purchase at some (finite) time under $(p^*, q^*)_t$ must weakly exceed the measure of consumers who purchase at some time under $(p, q)_t$. Since $Q_{t^*}^* = Q_{t^*}$, this implies that $\sum_{t \neq t^*} Q_t^* \geq \sum_{t \neq t^*} Q_t$, completing the proof that profit is higher under $(p^*, q^*)_t$ than under $(p, q)_t$.

It remains only to check that there exists a full-delivery PBE with price-rationing path $(p^*, q^*)_t$. This follows from the fact that there exists a full-delivery PBE with price-rationing path $(p, q)_t$, because, since $Q_t^* = Q_t$ for all t and $p_t^* \geq p_t$ for all t , the seller's gain from nondelivery is the same in every period under $(p^*, q^*)_t$ as under $(p, q)_t$ and her gain from delivery is weakly higher in every period under $(p^*, q^*)_t$ in a strategy profile in which consumers expect the seller to never deliver in the future if she does not deliver in the current period. \square

PROOF OF COROLLARY 1. Suppose that $(p^*, q^*)_t$ is a best full-delivery PBE price-rationing path in Γ_R . Consider the following strategy profile in Γ , which we denote by σ .

1. The seller sets price path $(p^*)_t$ and $x_t = 1$ as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$, $x_t = 0$ for all future periods. In particular, the seller sets $x_t = 0$ in any period in which she has set $p_t \neq p_t^*$.
2. A consumer with valuation v who has not yet received the good at τ accepts at τ if and only if the seller has never deviated from her prescribed strategy and $\tau \in \arg \max_t \delta^t (v - p_t^*)$.³²

To establish that σ is a PBE, we first observe that a consumer with valuation v receives the same payoff V_v under σ as under any full-delivery PBE with price-rationing path $(p^*, q^*)_t$ in Γ_R . This follows because, since no rationing occurs along $(p^*, q^*)_t$ in Γ_R (by Proposition 4) and the path of play does not depend on an individual consumer's actions, a consumer with generic valuation v facing $(p^*, q^*)_t$ in Γ_R purchases at time τ if and only if $\tau \in \arg \max_t \delta^t (v - p_t^*)$ in any full-delivery PBE. Furthermore, if valuation v is generic with respect to $(p^*)_t$, then the payoff of a consumer with valuation v uniquely determines her purchase time. Therefore, $(Q)_t$ is the same under any full-delivery PBE with price-rationing path $(p^*, q^*)_t$ in Γ_R as under σ .

Next note that if the seller ever sets $p_\tau \neq p_\tau^*$, she receives zero continuation payoff. Since this is her minmax value in Γ_R , she cannot receive continuation payoff strictly less than this in the continuation game from $\tau + 1$ onward in Γ_R under a full-delivery PBE with price-rationing path $(p^*, q^*)_t$. Now we have seen that $(Q)_t$ is the same in any full-delivery PBE with price-rationing path $(p^*, q^*)_t$ in Γ_R as in σ , and by construction $(p)_t$ is the same as well, so the seller's on-path continuation payoff from $\tau + 1$ onward must be the same too, so in particular this continuation payoff must be nonnegative. This implies that setting $p_\tau \neq p_\tau^*$ on path is not a profitable deviation. Similarly, the fact that setting $q_t = q_t^*$ is optimal on path along $(p^*, q^*)_t$ implies that setting $q_t = q_t^*$ is optimal on path in σ , because the cost of delivery and on-path continuation values are identical, while the payoff of zero that results from deviating from the equilibrium path in σ is at least as bad as the payoff from deviating in any PBE of Γ_R . Also, the seller's off-path play is optimal because off-path price-setting does not affect her payoffs and off-path delivery imposes a positive cost at no benefit.

³²The case where there are multiple maximizers is irrelevant, as this occurs for a set of measure-zero consumers.

We next check that each consumer's play is optimal. It is again obvious that his off-path play is optimal, as paying is costly and yields no benefit when the seller sets $q_t = 0$. That his on-path play is optimal follows from the fact that the seller's strategy is full delivery. So σ is a full-delivery PBE of Γ .

The above argument shows that if a price–quantity path $(p, Q)_t$ is a best full-delivery PBE price–quantity path in Γ_R , then it is also a best full-delivery PBE price–quantity path in Γ . For the converse, suppose that $(p^*, Q^*)_t$ is a full-delivery PBE price–quantity path in Γ . Consider the following strategy profile in Γ_R :

1. The seller sets price path $(p^*)_t$, and $q_t = 1$ and $x_t = 1$, as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$, $q_t = 1$, and $x_t = 0$ for all future periods. In particular, the seller sets $x_t = 0$ in any period in which she has set $p_t \neq p_t^*$.
2. A consumer with valuation v who has not yet received the good at τ pays at τ if and only if the seller has never deviated from her prescribed strategy and $\tau \in \arg \max_t \delta^t (v - p_t^*)$.

It is easy to check that this is a PBE in Γ_R . Furthermore, since no other players condition play on an individual consumer's actions, a consumer with generic valuation v purchases at time τ under this strategy profile if and only if a consumer with this valuation purchases at τ in any full-delivery PBE in Γ with price–quantity path $(p^*, Q^*)_t$. This implies that the mass of consumers who purchase at each period under this profile is the same as the mass of consumers who purchase at each period in any full-delivery PBE in Γ with price–quantity path $(p^*, Q^*)_t$, which then implies that the seller's profit under this strategy profile is the same as under any full-delivery PBE in Γ with price–quantity path $(p^*, Q^*)_t$. This completes the proof. \square

APPENDIX B: NON-FULL-DELIVERY EQUILIBRIA

This appendix considers non-full-delivery PBE of the relational contracting model of [Section 5](#). We conjecture that optimal PBE of Γ are not fully-delivery PBE, although the difference in payoff between an optimal PBE and a best full-delivery PBE must converge to 0 as δ converges to 1, as argued in the text. This is because setting $x < 1$ allows the seller to sell to some lower-valuation consumers before higher-valuation consumers. This may be useful for the seller, as selling to low-valuation consumers before high-valuation consumers may be a way to increase continuation payoffs without increasing quantity sold today, allowing the seller to sell more quickly.

While a complete analysis of optimal (non-full-delivery) PBE is outside the scope of this paper, we show here that analogues of parts (iii) and (iv) of [Theorem 1](#) for non-full-delivery PBE can be established without reference to the model with rationing. We view these results as complementary to those in the text, because full-delivery PBE are of particular interest for reasons discussed in the text. The results in this appendix do not establish the existence of full-delivery equilibria that yield profits close to static monopoly

profits; we do not know how to establish this result without using the connection to the model with rationing developed in [Section 5](#).

Intuitively, we can prove analogues of parts (iii) and (iv) of [Theorem 1](#) directly for non-full-delivery PBE because we can use nondelivery to substitute for rationing. That is, instead of using rationing to ensure that only fraction γ of those consumers who demand the good at price p_t at time t are allowed to purchase at t , the seller can charge γp_t to each of these consumers in exchange for delivering the good to each of them with probability γ . With this idea in hand, the proof of parts (iii) and (iv) of [Theorem 1](#) follows easily from the proof of [Proposition 6](#) in [Section 6](#).

PROPOSITION 11. *There exists a strategy profile in Γ that is a non-full-delivery PBE for high enough δ under which the seller's payoff converges to her static monopoly payoff as $\delta \rightarrow 1$.*

PROOF. Consider the following strategy profile:

1. The seller sets $p_t = \gamma p^m$ and $x_t = \gamma$ for all t and for γ an arbitrary positive constant less than 1, as long as she has conformed to this strategy in the past. Otherwise, she sets $p_t = \bar{v}$ and $x_t = 0$ for all future periods, and, in particular, sets $x_t = 0$ in any period in which she has set $p_t \neq \gamma p^m$.
2. A consumer with valuation v who has not yet received the good pays if and only if $v \geq p^m$, and the seller has never deviated from her prescribed strategy.

At any period t along the equilibrium path, a consumer with valuation $v < p^m$ has continuation value 0, while a consumer with valuation $v \geq p^m$ who has not yet received the good has continuation value $(\delta\gamma/(1 - \delta(1 - \gamma)))(v - p^m) < (\gamma/(1 - \delta(1 - \gamma))) \times (v - p^m)$, so every consumer's play is optimal by the one-shot deviation principle. It is clear that the seller's off-path play and on-path price setting are optimal. It remains only to check that the seller has an incentive to deliver the good along the equilibrium path. This condition is

$$\sum_{\tau=1}^{\infty} \delta^{\tau} ((1 - \gamma)^{t+\tau} p_{t+\tau} - \gamma(1 - \gamma)^{t+\tau} c) \geq \gamma(1 - \gamma)^t c \quad \text{for all } t \geq 0.$$

For any t , this can be rewritten as inequality (1). Now if $\delta > c/p^m$, then there exists a positive γ that satisfies (1). The above strategy profile then yields profit $(\gamma/(1 - \delta(1 - \gamma)))D(p^m)(p^m - c)$ for the seller, which converges to $D(p^m)(p^m - c)$ as δ converges to 1. \square

For the analogue of part (iv) of [Theorem 1](#) for non-full-delivery PBE, we argue as in the discussion following [Proposition 6](#). Consider the strategy profile where the seller fixes the price of a γ chance of receiving the good at some given γp . Recall that

$$\gamma^*(p) \equiv \frac{\delta p - c}{\delta p}.$$

By the same argument that led to (1), $\gamma^*(p)$ is the greatest probability of receiving the good that the seller can credibly offer at price $\gamma^*(p)p$ in a PBE with fixed price and delivery probability. The best PBE profit for the seller with a constant price path at γp and a constant sales rate γ is, therefore,

$$\left(\frac{\gamma^*(p)}{1 - \delta(1 - \gamma^*(p))} \right) D(p)(p - c),$$

which can be rewritten as

$$\left(p - \frac{c}{\delta} \right) D(p).$$

Therefore, if the seller sets $p_t = \gamma p^m(c/\delta)$ and $x_t = \gamma^*(p^m(c/\delta)) = (\delta p^m(c/\delta) - c)/(\delta p^m(c/\delta))$ for all t on the equilibrium path, and off-path play is given as in the strategy profile in the proof of Proposition 11, the seller's profit is equal to the static monopoly profit when cost equals c/δ . Finally, the seller can achieve a strictly higher payoff than this by slightly raising price and delivery probability early on while keeping quantity delivered constant in every period, in analogy with the discussion preceding Proposition 7.

REFERENCES

- Ausubel, Lawrence M. and Raymond J. Deneckere (1987), "One is almost enough for monopoly." *RAND Journal of Economics*, 18, 255–274. [482]
- Ausubel, Lawrence M. and Raymond J. Deneckere (1989), "Reputation in bargaining and durable goods monopoly." *Econometrica*, 57, 511–531. [482]
- Bagnoli, Mark, Stephen W. Salant, and Joseph E. Swierzbinski (1989), "Durable-goods monopoly with discrete demand." *Journal of Political Economy*, 97, 1459–1478. [482]
- Bernheim, B. Douglas and Michael D. Whinston (1998), "Incomplete contracts and strategic ambiguity." *American Economic Review*, 88, 902–932. [480]
- Biehl, Andrew R. (2001), "Durable-goods monopoly with stochastic values." *RAND Journal of Economics*, 32, 565–577. [482]
- Board, Simon (2008), "Durable-goods monopoly with varying demand." *Review of Economic Studies*, 75, 391–413. [482]
- Bond, Eric W. and Larry Samuelson (1987), "The Coase conjecture need not hold for durable good monopolies with depreciation." *Economics Letters*, 24, 93–97. [482]
- Bull, Clive (1987), "The existence of self-enforcing implicit contracts." *Quarterly Journal of Economics*, 102, 147–159. [483]
- Bulow, Jeremy I. (1982), "Durable-goods monopolists." *Journal of Political Economy*, 90, 314–332. [482]
- Cabral, Luis M. B., David J. Salant, and Glenn A. Woroch (1999), "Monopoly pricing with network externalities." *International Journal of Industrial Organization*, 17, 199–214. [482]

- Coase, Ronald H. (1972), "Durability and monopoly." *Journal of Law and Economics*, 15, 143–149. [480, 482, 504]
- Deneckere, Raymond and Meng-Yu Liang (2006), "Bargaining with interdependent values." *Econometrica*, 74, 1309–1364. [482]
- Deneckere, Raymond and Meng-Yu Liang (2008), "Imperfect durability and the Coase conjecture." *RAND Journal of Economics*, 39, 1–19. [482]
- Denicolò, Vincenzo and Paolo G. Garella (1999), "Rationing in a durable goods monopoly." *RAND Journal of Economics*, 30, 44–55. [483]
- Dudine, Paolo, Igal Hendel, and Alessandro Lizzeri (2006), "Storable good monopoly: The role of commitment." *American Economic Review*, 96, 1706–1719. [482]
- Ellison, Glenn and Sara Fisher Ellison (2009), "Search, obfuscation, and price elasticities on the internet." *Econometrica*, 77, 427–452. [480]
- Fudenberg, Drew, David K. Levine, and Jean Tirole (1985), "Infinite-horizon models of bargaining with one-sided incomplete information." In *Game-Theoretic Models of Bargaining* (Alvin E. Roth, ed.), 73–98, Cambridge University Press, Cambridge. [482, 487]
- Fudenberg, Drew and Jean Tirole (1991), *Game Theory*. MIT Press, Cambridge, Massachusetts. [482]
- Gul, Faruk (1987), "Noncooperative collusion in durable goods oligopoly." *RAND Journal of Economics*, 18, 248–254. [482]
- Gul, Faruk, Hugo Sonnenschein, and Robert Wilson (1986), "Foundations of dynamic monopoly and the Coase conjecture." *Journal of Economic Theory*, 39, 155–190. [482, 487]
- Hahn, Jong-Hee (2006), "Damaged durable goods." *RAND Journal of Economics*, 37, 121–133. [486]
- Hart, Oliver D. and Jean Tirole (1988), "Contract renegotiation and Coasian dynamics." *Review of Economic Studies*, 55, 509–540. [480, 487, 504]
- Hörner, Johannes and Nicolas Vieille (2009), "Public vs. private offers in the market for lemons." *Econometrica*, 77, 29–69. [482]
- Inderst, Roman (2008), "Durable goods with quality differentiation." *Economics Letters*, 100, 173–177. [486]
- Kahn, Charles M. (1986), "The durable goods monopolist and consistency with increasing costs." *Econometrica*, 54, 275–294. [482]
- Karp, Larry (1996), "Depreciation erodes the Coase conjecture." *European Economic Review*, 40, 473–490. [482]
- Klein, Benjamin and Keith B. Leffler (1981), "The role of market forces in assuring contractual performance." *Journal of Political Economy*, 89, 615–641. [483]

- Laffont, Jean-Jacques and Jean Tirole (1988), "The dynamics of incentive contracts." *Econometrica*, 56, 1153–1175. [505]
- Levin, Jonathan (2003), "Relational incentive contracts." *American Economic Review*, 93, 835–857. [483]
- Malcomson, James M. (forthcoming), "Relational incentive contracts." In *Handbook of Organizational Economics* (Robert Gibbons and John Roberts, eds.), Princeton University Press, Princeton, New Jersey. [483]
- Mason, Robin (2000), "Network externalities and the Coase conjecture." *European Economic Review*, 44, 1981–1992. [482]
- Masten, Scott E. and Renata Kosová (2009), "Post-sale service and the limits of reputation." Unpublished paper, Ross School of Business, University of Michigan. [483]
- McAfee, R. Preston and Thomas Wiseman (2008), "Capacity choice counters the Coase conjecture." *Review of Economic Studies*, 75, 317–332. [482]
- Olsen, Trond E. (1992), "Durable goods monopoly, learning by doing and the Coase conjecture." *European Economic Review*, 36, 157–177. [482]
- Shapiro, Carl (1982), "Consumer information, product quality, and seller reputation." *Bell Journal of Economics*, 13, 20–35. [483]
- Shapiro, Carl (1983), "Premiums for high quality products as returns to reputations." *Quarterly Journal of Economics*, 98, 659–679. [483]
- Sobel, Joel (1991), "Durable goods monopoly with entry of new consumers." *Econometrica*, 59, 1455–1485. [482]
- Stokey, Nancy L. (1982), "Rational expectations and durable goods pricing." *Bell Journal of Economics*, 12, 112–128. [482]
- Van Cayseele, Patrick (1991), "Consumer rationing and the possibility of intertemporal price discrimination." *European Economic Review*, 35, 1473–1484. [483, 490]