Possibly-Final Offers

Nolan H. Miller
John F. Kennedy School of Government
Harvard University
Cambridge, MA 02138
nolan_miller@harvard.edu

Nikita E. Piankov
Analysis Group, Inc.
Boston, MA 02199
npiankov@analysisgroup.com

Richard J. Zeckhauser
John F. Kennedy School of Government
Harvard University
Cambridge, MA 02138
richard_zeckhauser@harvard.edu

A price-setting seller faces a buyer with unknown reservation value. We show that if the buyer is sufficiently risk averse, the seller can benefit from employing a Possibly-Final Offer (PFO) strategy. In a PFO, if the buyer rejects the seller's initial offer the seller sometimes terminates the interaction. If the seller does not terminate, he follows up with a subsequent, more attractive offer. As the buyer's risk aversion increases, the seller's expected profit under the optimal PFO approaches the full-information profit. These results extend to contexts with endogenous commitment, multiple types of buyers, multidimensional objects, and nonseparable utility functions.

1. Introduction

When two parties bargain over the sale of a good and one party rejects the other's offer, a critical question arises: Was that offer final, or will an improved offer be forthcoming? This uncertainty may lead the offer recipient to accept less-advantageous current terms to avoid the possibility that the offerer will terminate the negotiation following a

Previous versions of this paper circulated under the title, "When to Haggle." We thank Chris Avery, Drew Fudenberg, Oliver Hart, David Laibson, Guhan Subramanian, participants at the Harvard Seminar on Research on Behavior and Markets, a referee, a co-editor, and the editor for helpful comments. Philipp Schnabl provided excellent research assistance.

Journal of Economics & Management Strategy, Volume 15, Number 3, Fall 2006, 789–819

rejection. A risk-averse offer recipient may make significant concessions to avoid even a modest threat of termination.[1] This paper shows how the offering party may be able to take advantage of this effect to improve his expected outcome by using a Possibly-Final Offer (PFO), a strategy that couples an initial offer with a positive probability that the offerer will walk away following a rejection rather than make a second, more attractive offer.

We illustrate the power of PFOs in the presence of risk-averse buyers by studying a buyer–seller interaction in which the seller makes offers to a buyer whose value for the object for sale (either high or low) is unknown to the seller. Absent randomization, the seller's optimal strategy takes one of two forms. He can either sell to all buyers at the low-value buyers' reservation value, or he can sell only to the high-value buyers at their reservation value. The fundamental law of demand dictates the outcome: The seller can either sell to some buyers at a high price or to all buyers at a low price, but he cannot do both simultaneously. Implementing deterministic mechanisms over time cannot help the seller.[2] If the seller offers a high price followed by a lower one, all buyers will simply wait for the lower price.

A PFO strategy requires probabilistic play by the seller. In the two-type case, the seller makes an initial price offer. If that offer is rejected, he terminates the negotiation with a positive probability less than one; with the complementary probability he makes an improved (truly final) offer.[3] When compared to the deterministic strategy of selling to all buyers at a price equal to the low-type buyer's reservation value, making the seller's initial offer "possibly-final" offers the benefit of allowing him to charge a higher initial price to high-valued buyers. The cost of this strategy is that, in the event of a rejection, the seller terminates the transaction with some probability, walking away from the possibility of completing the sale via a second, lower-priced offer. The main insight of this paper is that as buyer risk aversion increases, the balance tips in favor of the former effect: the seller is able to significantly raise the initial price charged to high-valued buyers using only a small threat of

1. For clarity, throughout the paper, we treat the seller as male and the buyer as female.

2. Implicitly, we are talking about time periods short enough so that time preference does not play a significant role.

3. In our model, the seller initially does not know the buyer's willingness-to-pay, but it remains fixed over time, leading to the optimality of a declining price schedule. Thus, our environment differs from that of Das Varma and Vettas (2001), which studies a model in which the buyer's willingness-to-pay is independently and identically distributed over time and show that the seller's optimal prices increase, and Rustichini and Wolinsky (1995), which considers the case in which willingness-to-pay follows a Markov process and show that over time the seller will move between high and low prices in order to learn the current state of demand.

termination.[4] Thus, for sufficiently risk averse buyers, PFOs outperform deterministic strategies and may even be Pareto improving.[5,6] In the limit, as buyers become infinitely risk averse, the seller is effectively able to charge each buyer type her reservation value while terminating the transaction without making a sale with only arbitrarily small probability. Almost all of buyers' surplus is thereby extracted.

Two recent mergers illustrate PFOs at play. When Marathon Oil acquired Pennaco in a cash deal in December 2000, it made an initial offer at $17, which was rejected. Such an outcome was inevitable; first offers in such circumstances are rarely presented as or thought to be final offers. Its second offer, a week later, was $19. Pennaco's Board instructed its Chair/CEO to push for $20 or at least $19.50. Marathon stated, however, that $19 was its "absolute, final, best, top offer." Pennaco accepted. In effect, Marathon made one serious offer, and framed it as a final offer. Had Pennaco rejected, we do not know whether an improved offer would have been forthcoming. Clearly Pennaco feared it would not.[7]

The takeover of German telecommunications giant Mannesmann (M) by Vodafone (V) presents a quite different picture. On November 15, 1999, V offered 43.7 of its shares for each Mannesmann share. After a quick rejection, V came back four days later with a 53.7 share hostile offer (value of 240 euro). Ten days later, M rejected this offer as well. Three weeks later, V noted publicly that its 53.7 share offer was final, and would expire on February 7, 2000. On February 4, M agreed to a friendly merger providing 58.9 V shares for each M share. Because V's shares rose over this period, the value per M share was 353 euros on this date. Despite a firmly announced final offer, V still increased its payment in shares by a further 10%, with monetary value per M share soaring by nearly 40% from the date of the "final" offer. Thus concluded

4. Our results do not depend on exploiting differences in buyer risk aversion. Rather, they rely on (in the two-type case) risk aversion for the higher-valued buyer, and the difference in willingness-to-pay between the high- and low-valued buyers. Hence, these results apply even when all types have the same attitude toward risk.

5. When buyers are risk neutral, we show that PFO strategies cannot benefit the seller. This is an illustration of Riley and Zeckhauser's (1983) "no haggling" result, which shows that it cannot pay to "haggle" when buyers are risk neutral. The seller should commit to a firm price and refuse to revise his offer, or indeed entertain any offer from a rejecting buyer.

6. PFOs may also be superior to deterministic strategies for intermediate levels of risk aversion. We show that, for any level of buyer risk aversion, there is a range of buyer populations (mixtures of high- and low-valued buyers) for which PFOs outperform deterministic strategies.

7. Some Pennaco shareholders launched an unsuccessful suit alleging their Board did not meet its fiduciary duties. The potential for such suits makes it almost inevitable that in corporate acquisitions first offers are refused, implying that PFOs can only start with second offers. In re Pennaco Energy, 787 A.2d 691 (Del. ch. 2001).

the largest merger of the twentieth century.[8] Facing a PFO, Pennaco blinked; Mannesmann rejected, and Vodafone offered more.[9]

The problem of selling to risk-averse buyers is considered by Matthews (1983), and by Maskin and Riley (1984). Matthews considers a monopolist who sells to heterogeneous, risk-averse buyers. The optimal mechanism, he shows, has the buyer who pays more up front getting a greater probability of receiving the object. Maskin and Riley (1984) examine optimal auctions when buyers are risk averse and show that they impose risk on all but the most eager buyers.

Our results add to the Matthews (1983) and Maskin and Riley (1984) studies of optimal mechanisms with risk-averse buyers in three important ways. First, because the optimal mechanisms they consider involve buyers making payments even when they do not receive the object, neither captures a key feature of practical buyer–seller interactions: should the buyer and seller not reach an agreement, no transfers are made. Second, both Matthews (1983) and Maskin and Riley (1984) make specific assumptions about the form of the buyer's utility function and the type of buyer heterogeneity. The present paper places no restrictions on utility functions beyond requiring risk aversion. Third, the prior papers dealt with unidimensional objects, whereas the present paper extends to the case of multidimensional objects.

Our PFO mechanism has a parallel to an auction with a random reserve price. The idea that such reserve prices may play a role in increasing seller revenue has gained some attention recently in the auctions literature. In a related paper, Li and Tan (2000) show that when buyers are sufficiently risk averse, the seller can increase his expected profit in a first-price auction by using a hidden, random reserve price instead of an announced, certain one. The intuition behind the desirability of a hidden reserve is the same as that which drives our results. The risk of losing the object because of bidding below the reserve price drives the risk-averse bidder to increase his bid. Such secret reserve prices are employed in real-world auctions. Bajari and Hortacsu (2003) empirically investigate the role of secret reserve prices on eBay, and Li and Perrigne (2003) and Perrigne (2003) analyze their role in timber auctions.[10]

8. See *The Financial Times*, February 3, 2000, February 4, 2000, and February 5, 2000 for details.

9. Knowledgeable observers believe that there was significant emotion and some hostility within Mannesmann while the "final offer" was open, and that Vodafone may have raised its offer primarily to dampen emotion and earn goodwill, that is, to secure a benevolent victory. Part of the uncertainty surrounding PFOs involving mergers is whether the acquirer will respond to such concerns.

10. We discuss the connection between our pricing mechanism and an auction in Section 2.4.

Mechanisms that capitalize on risk have also been shown to be useful in other contexts, at least in theory, because randomization can relax incentive-compatibility constraints. For example, there is a substantial literature on optimal taxation that shows that the government can sometimes benefit from implementing a random taxation schedule (see Weiss, 1976, and Brito et al., 1995 and references therein), and Arnott and Stiglitz (1988) show that randomization may be useful in insurance problems. However, the use of random mechanisms has often been criticized as being impractical.[11] This criticism argues that, in order to be useful in contracting, random mechanisms must be credible and verifiable. However, from a single realization it is impossible to detect whether a party is following a particular random strategy. Such inferences are only possible in the long run, once sufficient statistical evidence has been amassed. However, because most interactions take place over much shorter time frames, it may be difficult for principals to credibly commit to using stochastic mechanisms. If so, that may prevent them from reaping their benefits.

Our mechanism resists this criticism. Because both buyer and seller return to the status quo should the seller terminate the transaction following an initial rejection, enforceability (and therefore verifiability) is not an issue. True, in order for the seller to benefit from the PFO, buyers must believe there is some likelihood that the seller will not offer a lower price following rejection, but this is essentially the same issue that arises when the seller follows the deterministic strategy of charging a high price and selling only to high-valued buyers, a practice that some sellers certainly succeed in employing. Further, our analysis shows that not only does the best PFO outperform deterministic strategies, but also that, provided buyers are sufficiently risk averse, small deviations from deterministic strategies are also beneficial for the seller. Thus it is not necessary for the seller to commit to a particular termination probability to benefit from using a PFO. Often it is enough merely to sow the suspicion that termination is possible. The question then becomes whether the seller can find ways to convince buyers that chance plays a role in the process. A number of commonly observed features of buyer–seller interactions help to achieve this goal, such as delegation, automatic discounting, fostering competition among buyers, keeping stock low, and others. We explore a number of these examples in greater detail in Section 4.

An additional criticism of random mechanisms is that they are too complicated to be useful in practical situations. The mechanisms proposed in the optimal auctions and taxation literatures are highly complex and rely on bidders/taxpayers having a nuanced

11. See, for example, Laffont and Martimort (2002).

understanding of incentive constraints in order to generate their benefits. Our model illustrates that, in the presence of risk aversion, even simple random mechanisms can be beneficial. And, by implementing our random mechanism over time, we believe that it becomes easier for buyers to understand and for sellers to employ. Finally, as our examples illustrate, buyer uncertainty about whether the seller will come back with a new offer plays an important role in many real-world sales situations.

This paper proceeds as follows. Section 2 describes the game between seller and buyer contracting over the price of a standardized good. It shows that a PFO strategy is optimal whenever the buyer is risk averse, at least for some distributions of buyers. We then extend the analysis in several ways, considering endogenous commitment by the seller, multiple types, and multidimensional objects, and relaxing the assumption that buyers' utility functions are additively separable in money. Section 4 discusses ways in which sellers may operationalize PFO strategies. Section 5 concludes. The proofs are presented in the Appendix.[12]

2. Possibly-Final Offers: The Theory

2.1 Timing

A risk-neutral seller offers a single, indivisible object for sale. The seller makes the buyer an offer, and, should that offer be rejected, may either make a more attractive second offer or terminate the interaction. With two types of buyers, no more than two rounds of offers will be necessary, and we therefore model the interaction as a two-stage game, represented schematically in Figure 1. In the first stage, the seller makes an offer to the buyer, which she may either accept or reject. If the buyer accepts, the transaction takes place according to the offer and the game ends. If the buyer rejects, the game ends with probability y. That is, with probability y, the initial offer is final. We will often refer to y as the seller's "walk-away" probability. With probability $1 - y$, the game continues to stage 2, where the seller makes another (truly final) offer and the buyer has another opportunity to accept or reject it.

A strategy for the seller consists of first and second prices p_1 and p_2, and the probability y that the seller's initial offer is final. Formally, the seller follows a PFO strategy if he chooses $0 < y < 1$. We refer to strategies with $y = 0$ or $y = 1$ as deterministic strategies. For each buyer type, the buyer's strategy specifies, for each possible seller's strategy, whether the buyer accepts or rejects each of the seller's offers. Our solution concept

12. Several of the longer proofs and some supporting calculations are in an on-line appendix available from the authors.

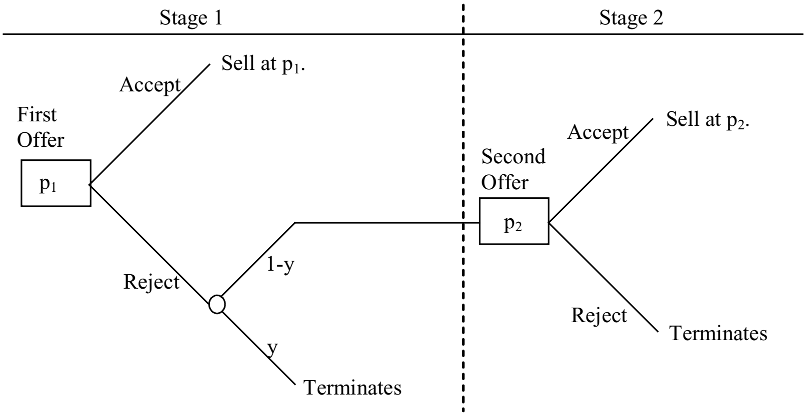

FIGURE 1. THE SEQUENCE OF THE GAME

is subgame-perfect Nash equilibrium. Hence we require each type of buyer to play a best response to the seller's strategy.

Although we describe the buyer–seller interaction as taking place over time, our extensive form is equivalent to a game with no time element at all. In effect, the seller offers the buyer the choice between two possible contracts. The first contract offers to sell the item for sure for a price, p_1. The second offers a lottery where with probability $1 - y$ the buyer may purchase the item for price p_2, and with probability y no offer is made (or the price is infinite). We sacrifice no generality by modeling the situation as being formally equivalent to a static problem, because our results do not depend on the time preferences of the buyers. In addition, we argue that the crucial element that motivates the use of PFO strategies in the real world is risk aversion, not time preference.[13] In many buyer–seller interactions, the time lapse between the first and second offers is very short, usually less than a day, and frequently just a few minutes or seconds. The analysis of this paper could be extended to include time preference. However, doing so would complicate the model without providing additional insights.

We assume that sellers are able to commit to their strategies. Such commitments have two components. First, it must be that if the buyer rejects the seller's first offer, she cannot come back later and accept it. That is, the seller commits to the fact that once rejected, p_1 is "off the table." Second, the seller commits to terminating the game with a positive probability, which implies that sometimes a second offer is not made even though buyer and seller would both benefit were there one. Although some real-world sellers may be unable to make

13. As we discuss in the conclusion, loss aversion may also play a role, and our results extend to the case of loss-averse buyers.

such commitments, later in the paper we discuss several buyer–seller situations where sellers do make such commitments, and we show how some commonly observed sales practices may help them to do so. In the extensions, we present a simple model with endogenous commitment.

2.2 Model

In our basic model, a seller offers a single unit of a good for sale to a buyer. The seller is assumed to be risk neutral and to choose his strategy to maximize expected profit. Each buyer knows her own reservation value, but the seller knows only the distribution of buyer types.

There are two types of buyers. With probability $0 < \mu < 1$ the buyer is HIGH. Otherwise she is LOW. Throughout the paper, we will use the subscript h to denote HIGH, l to denote LOW, and $t \in \{h, l\}$ to stand for a generic type.

HIGH receives utility $u_h(q, p) = \theta_h q + v_h(w_h - pq)$, where $q = 1$ if she purchases the object and 0 if she does not, and p is the price paid for the object. HIGH's utility for after-expenditure money, $v_h()$, is assumed to be strictly increasing, strictly concave, and twice differentiable, with $v_h(w_h) = 0$. With probability $1 - \mu$, the buyer is LOW and has utility function $u_l(q, p) = \theta_l q + v_l(w_l - pq)$, where $q \in \{0, 1\}$ equals 1 if she purchases the object and 0 otherwise, and $v_l()$ is strictly increasing with $v_l(w_l) = 0$.

Let r_h and r_l be the reservation prices of HIGH and LOW, respectively, defined as

$$v_t(w_t - r_t) = -\theta_t \quad \text{for } t \in \{h, l\}. \tag{1}$$

We are interested in examining the effect of increasing the buyer's risk aversion while holding willingness-to-pay constant. Thus we take the buyers' reservation prices r_t as our primitives, and let θ_t be defined implicitly according to (1). We assume that HIGH is willing to pay more for the object than LOW, that is, $r_h > r_l$.

In the absence of PFOs, the seller's optimal contract takes one of two forms. The largest price at which both HIGH and LOW will buy is r_l, and thus the maximum profit the seller can earn while selling to both types is r_l. On the other hand, if the seller is willing to sell only to HIGH, he can charge a price as high as r_h. Because the probability of the buyer being HIGH is μ, the seller expects profit μr_h in this case. Finally, selling to both types offers higher profit than selling only to HIGH whenever

$$r_l \geq \mu r_h.$$

This implies that the seller will choose to sell only to both types of buyers whenever $\mu \leq \frac{r_l}{r_h}$, and will sell only to HIGH otherwise. Throughout the paper we will refer to this as the "deterministic" case.

2.3 Results

The seller's objective is to choose p_1, p_2, and y to maximize his expected profit. Without loss of generality, we assume that the seller designs the offers so that the first offer is acceptable to HIGH but not to LOW, while the second offer, if made, is acceptable to both types.[14]

The seller's problem is, in essence, a monopolistic screening problem, with the added feature that the mechanisms available to the seller have been expanded to include those where the seller can commit to making some offers only probabilistically. We adopt the standard approach to such problems, and write the seller's problem as a constrained-maximization problem:

$$\max_{p_1, p_2, y} \mu p_1 + (1-y)(1-\mu) p_2$$

$$\text{s.t.} \quad \theta_h + v_h(w_h - p_1) \geq 0, \tag{2}$$

$$\theta_l + v_l(w_l - p_2) \geq 0, \tag{3}$$

$$\theta_h + v_h(w_h - p_1) \geq (1-y)(\theta_h + v_h(w_h - p_2)), \text{ and} \tag{4}$$

$$(1-y)(\theta_l + v_l(w_l - p_2)) \geq \theta_l + v_l(w_l - p_1). \tag{5}$$

The constraints are the standard constraints in a screening problem. Conditions (2) and (3) are the participation constraints. HIGH must prefer the first offer to the status quo, (2), and LOW must prefer the second offer, if one is made, to the status quo, (3). Conditions (4) and (5) are the incentive-compatibility constraints. HIGH must prefer the first offer to the uncertain prospect of a second offer, (4), and LOW must prefer to wait for a second offer rather than accept the first offer, (5). Again, we assume that if the buyer is indifferent between buying and not, she buys.

Before continuing, note that if the seller could observe the buyer's type, he would charge r_h to HIGH and r_l to LOW, earning *ex ante* expected profit $\mu r_h + (1-\mu) r_l$. However, such a scheme violates (4).

14. We ignore the case where the first offer is rejected, because if that were to occur it would be optimal to make a second offer that is either accepted by all types or only by HIGH. The case where the second offer is accepted by both types is equivalent to one where $p_1 = p_2$, p_2 is acceptable to LOW, and the second offer is made with probability 1, that is, $y = 0$. The case where the second offer is accepted only by HIGH is equivalent to one where that offer is made first, and no second offer is made, $y = 1$. Similarly, because the case where LOW rejects the second offer is equivalent to one where the second offer is made with zero probability, we can ignore that case as well.

As is standard in screening problems, we begin by arguing that constraints (3) and (4) bind at the seller's optimal solution, and that (2) and (5) can be safely ignored. Optimal values of p_1, p_2, and y are denoted with asterisks. LOW's participation constraint, (3), clearly binds. Hence $p_2^* = r_l$. Constraint (4) can be rewritten as

$$\begin{aligned} v_h(w_h - p_1) &\geq -\theta_h + (1-y)(\theta_h + v_h(w_h - p_2)) \\ &\geq y(-\theta_h) + (1-y)v_h(w_h - p_2) \\ &\geq yv_h(w_h - r_h) + (1-y)v_h(w_h - r_l). \end{aligned} \tag{6}$$

Because the objective function increases in p_1, this constraint must bind at the optimum. Because $v_h()$ is strictly increasing, this assures that $r_l \leq p_1^* \leq r_h$, and $p_2^* = r_l$ and $r_l \leq p_1^* \leq r_h$ imply (2) and (5) are satisfied at the optimum.

By (6), when p_1 is set at its profit-maximizing level, $w_h - p_1^*$ is the certainty equivalent of a lottery offering $w_h - r_h$ with probability y and $w_h - r_l$ with probability $(1-y)$. Formally, implicitly define $p_1(y)$ according to

$$v_h(w_h - p_1(y)) \equiv yv_h(w_h - r_h) + (1-y)v_h(w_h - r_l). \tag{7}$$

It is easily shown that $p_1(y)$ is strictly increasing and strictly concave in y, and that $p_1(0) = r_l$ and $p_1(1) = r_h$.

Based on the preceding arguments, the seller's problem can be written as an unconstrained-optimization problem:

$$\max_{0 \leq y \leq 1} \mu p_1(y) + (1-y)(1-\mu)r_l. \tag{8}$$

Let y^* be the walk-away probability that maximizes the seller's expected profit. The first-order condition is given by

$$\mu p_1'(y^*) - (1-\mu)r_l \begin{cases} \leq 0 & \text{if } y^* = 0 \\ = 0 & \text{if } 0 < y^* < 1 \\ \geq 0 & \text{if } y^* = 1 \end{cases}. \tag{9}$$

A PFO strategy is optimal if there exists a y^* with $0 < y^* < 1$ such that (9) holds with equality. If $y^* = 0$, then the seller adopts the deterministic strategy of charging $p_1^* = p_2^* = r_l$ and selling to all buyers immediately, while if $y^* = 1$ the seller adopts the deterministic strategy of selling only to HIGH at price r_h.

Whether a PFO strategy is desirable in a particular situation depends on the relationship between buyer's risk aversion and the mix of high- and low-valued buyers in the population. The next several propositions explore this relationship.

Proposition 1: *If HIGH is risk averse, a PFO offers higher profit than deterministic contracting when* $\mu = \frac{r_l}{r_h}$.

Proof. All proofs are presented in the Appendix.

Proposition 2: *For any μ such that $0 < \mu < 1$, if HIGH is sufficiently risk averse, using a PFO increases the seller's expected profit*: $0 < y^* < 1$. *As HIGH becomes infinitely risk averse, the seller's expected profit converges in the limit to the full-information profit of* $\mu r_h + (1 - \mu)r_l$.

The same basic intuition underlies Propositions 1 and 2. To understand the trade-offs the seller faces, consider the seller making a second offer of r_l with probability 1, that is, $y = 0$. HIGH's incentive-compatibility constraint implies that the highest first price that HIGH will accept will be r_l, because HIGH can always wait for the second offer. Suppose, now, that the seller introduces a slight probability $\varepsilon > 0$ of walking away after an initial rejection. By walking away with positive probability, he sacrifices some expected profit due to sales not made to LOW. However, because there is a positive probability that the seller's initial offer is the final offer, the risk of termination makes rejecting the first offer in hopes of a better second offer less attractive to HIGH (i.e., it relaxes her incentive-compatibility constraint), and thus increases the maximum first offer that HIGH is willing to accept, $p_1(\varepsilon) > r_l$. Thus, in moving from a deterministic offer to a PFO, the two countervailing quantities that must be weighed against each other are the decrease in profit due to selling to LOW only part of the time, $-\varepsilon(1 - \mu)\, r_l$, and the increase in profit from selling to HIGH at the higher price, $\mu(p_1(\varepsilon) - r_l)$. If $p_1(\varepsilon)$ is small, then the loss on sales forgone outweighs the gain due to increasing the first-round price, while when $p_1(\varepsilon)$ is large the gain from the higher price outweighs the loss due to missed sales.

When $\mu = \frac{r_l}{r_h}$ and HIGH is risk neutral, the seller is exactly indifferent between selling only to HIGH at price r_h, selling to both HIGH and LOW at r_l, or using a PFO strategy with $0 < y < 1$, $p_1 = p_1(y)$, and $p_2 = r_l$. Because introducing risk aversion increases HIGH's willingness to pay in the first round (i.e., $p_1(y)$), any amount of risk aversion tips the balance in favor of a PFO strategy when $\mu = \frac{r_l}{r_h}$. For other values of μ, deterministic strategies do strictly better than PFO strategies against a risk-neutral buyer. However, as risk aversion increases, $p_1(y)$ increases, eventually approaching r_h as risk aversion becomes infinite. Consequently, for any μ strictly between 0 and 1 there is a level of risk aversion high enough that PFOs outperform deterministic strategies, and as HIGH becomes infinitely risk averse, the seller is able to induce HIGH to pay a price near r_h even when the chance of the initial offer being final is small (i.e., y is near 0), the result being that the seller's expected profit approaches the full-information maximum, $\mu r_h + (1 - \mu)r_l$.

Proposition 2 brings to mind Corollary 1 in Matthews (1983), which shows that when the Arrow–Pratt coefficient of the buyers' constant absolute risk aversion (CARA) utility functions goes to infinity, the seller's expected profit converges to the full-information maximum. However, Proposition 2 applies whenever the buyers are sufficiently risk averse, regardless of the form of the utility function, and thus extends the Matthews result beyond the CARA case. Although Proposition 2 is stated with only two types, we show in Proposition 7 that this assumption is not crucial.

The intuition underlying Proposition 2 is quite robust, and would apply across a wide variety of environments. For example, suppose that the buyer's value is not fixed. Rather, it is either 8 or 10, with transitions between the two being governed by a Markov process. If the buyer's value were known, the seller would wait for the buyer to have a value of 10 and then sell her the object for a price of (nearly) 10. If the seller is patient and the buyer is sufficiently risk averse, the seller can earn nearly this much profit by quoting a price near 10 and threatening to walk away with a small probability. In this case, the buyer purchases the object at this high price the first time her value is 10 rather than risk losing it while waiting for a lower price.

The arguments in support of Propositions 1 and 2 also establish that, for any μ, given sufficient risk aversion, small deviations from the optimal deterministic contract increase profit. Thus, these results imply that if the optimal deterministic strategy is to sell only to high-valued buyers, then the seller can increase profit by introducing a small probability of continuation (because the buyer can still charge very risk-averse high-valued buyers almost their reservation price while also making some sales to low-valued buyers). Similarly, if the optimal deterministic strategy is to sell to all buyers, then if the seller can convince the buyers that there is a small probability of termination following a rejection of his initial offer, he can increase the initial price high-valued buyers will accept sufficiently to offset any losses due to lost opportunities to sell to low-valued buyers. Thus, it is not necessary for the seller to be able to commit to a particular PFO in order to benefit, and even PFOs that are very similar to the optimal deterministic strategies can benefit the seller.

In addition to potentially increasing the seller's profit, PFO strategies may also be Pareto improving. In the absence of randomization, the seller's optimal strategy is either to sell to all buyers at the low-type buyer's reservation value or to sell to only high-valued buyers at their reservation value. In the latter case, all buyers earn zero surplus. If in such a case a PFO is optimal for the seller, then by definition it increases the seller's profit. Because the initial price offered by the seller under the optimal PFO is lower than the high-valued buyer's reservation value, high-valued buyers earn a positive surplus under the optimal PFO. Finally, because a low-valued buyer either is not offered the object or purchases it at her reservation value, the low-valued buyer continues to earn zero surplus. Although our mechanism is quite different, this phenomenon is similar in spirit to the one investigated by Deneckere and McAfee (1996) who illustrate how manufacturers offer high- and low-quality versions of their product as a means of price discrimination and show that such strategies may be Pareto improving.[15]

The converse to the first part of Proposition 2 is also true. If HIGH is sufficiently tolerant of risk, the seller maximizes profit either by selling to all buyers immediately, that is, setting $p_1 = p_2 = r_l$, and $y = 0$, or by selling only to HIGH buyers, that is, setting $p_1 = r_h$, and $y = 1$.[16] This leads to Proposition 3, a restatement of the Riley–Zeckhauser (1983) "no haggling" result in the current environment.[17]

Proposition 3: *If HIGH is risk neutral, the seller cannot increase his expected profit by using a PFO strategy rather than the optimal deterministic strategy.*

Propositions 4 and 5 establish the natural comparative statics of our problem.

Proposition 4: *For any μ such that $0 < \mu < 1$, the seller's expected profit is nondecreasing in the level of HIGH's risk aversion. If a PFO strategy is optimal, then an increase in HIGH's risk aversion strictly increases the seller's expected profit.*

Proposition 4 holds the population mix fixed and asks what happens when HIGH's utility function changes. A related question is how, holding the buyers' utility functions fixed, the optimal contract changes as the proportion of HIGHs in the population changes. In other words, for which values of μ does using a PFO have the greatest potential, and how does the set of values of μ for which PFOs are optimal expand as the buyers become more risk averse?

As we already knew, if HIGH is risk neutral, using a PFO cannot improve profit. In this case, the seller sells to both HIGH and LOW

15. They use the term "damaged goods" to refer to such a cost-quality situation. One of their prime examples is an Intel computer chip whose capabilities were constrained to produce a lower-quality chip. A probabilistic second offer is in some sense a good that is "damaged" to make it less appealing to high-value buyers, and therefore to promote price discrimination.

16. This is true for any μ strictly between 0 and 1, although for $\mu = \frac{r_l}{r_h}$ haggling will always be optimal unless HIGH is risk neutral.

17. Proposition 3 also applies if the buyers are risk loving.

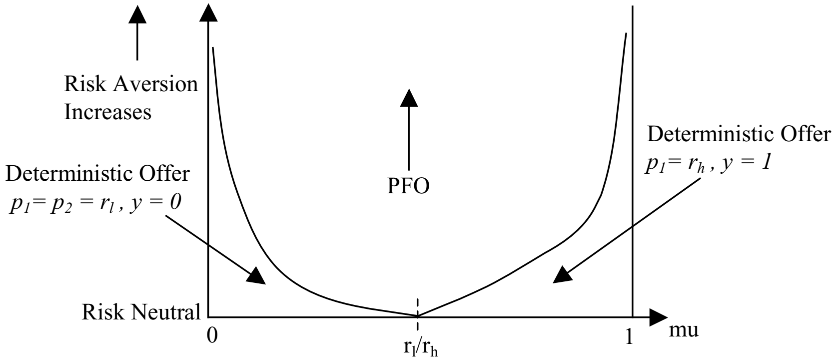

FIGURE 2. OPTIMAL STRATEGY, RISK AVERSION, AND BUYER MIX

for low μ and sells only to HIGH for high μ, switching as μ crosses the critical level, $\mu = \frac{r_l}{r_h}$. Once risk aversion enters, this balance point remains critical; it is the first place where the optimality of a PFO emerges. Once HIGH becomes strictly risk averse, it becomes optimal to use a PFO for $\mu = \frac{r_l}{r_h}$ and for a closed interval around it. Further, increases in risk aversion strictly spread the range over which using a PFO is optimal. In the limit, it is optimal to use a PFO for all strictly mixed populations μ such that $0 < \mu < 1$. Figure 2 summarizes the answers to these questions. Proposition 5 derives them.

Proposition 5: *If HIGH is risk averse, then*

(5a) the set of μ for which it is optimal for the seller to use a PFO is a closed interval, $[\mu_0, \mu_1]$, where $0 < \mu_0 < \frac{r_l}{r_h} < \mu_1 < 1$.
(5b) increasing HIGH's risk aversion decreases μ_0 and increases μ_1.

2.4 Relationship to Auctions

In our basic model, the seller names the price. However, the same ideas are relevant where the buyer names the price, as in an auction. Consider, for example, a first-price sealed-bid auction. The threat to walk away from negotiations in our basic game translates into a strategy where the seller uses a secret, random reserve price.[18]

To make the connection as transparent as possible, consider an auction with a single risk-averse buyer whose value for the object is r_h with probability μ or r_l with probability $(1 - \mu)$. Suppose the seller

18. See Li and Tan (2000) for an analysis of the role of secret reserve prices in auctions.

announces before the buyer bids that the reserve price is $p_1(y)$ with probability y and r_l with probability $(1 - y)$. LOW will never bid more than her value, so her best strategy is to bid r_l. Because there is no competition, HIGH faces the choice between bidding $p_1(y)$ and winning the object with probability 1, or bidding r_l and winning the object with probability $1 - y$. Bidding $p_1(y)$ is preferred whenever

$$v_h(w - p_1(y)) + \theta_h \geq y(0) + (1 - y)(v_h(w - r_l) + \theta_h),$$

or, by the same reasoning used at (6),

$$v_h(w - p_1(y)) \geq y v_h(w - r_h) + (1 - y)v_h(w - r_l). \tag{10}$$

By the definition of $p_1(y)$, (10) holds with equality, and thus it is a best response for HIGH to bid $p_1(y)$.[19]

Using the connection between our PFO problem and a first-price auction with a random reserve price, we can show that the random reserve price scheme outlined above is optimal (from the seller's perspective) when there is one bidder and the buyer cannot be charged unless she wins the object. The argument begins by noting that when the buyer only pays the seller if she wins the object, a general simultaneous mechanism consists of a bid-probability-of-winning pair for HIGH and LOW. Recalling that we earlier argued that any mechanism that gives HIGH the object with probability less than 1 is dominated by one that gives it for sure, such general mechanisms correspond to the mechanisms we consider in our PFO problem. Thus, the seller-optimal PFO mechanism corresponds to the seller-optimal general simultaneous mechanism, the only difference being that the optimal PFO is implemented sequentially. However, as we noted above, the PFO problem we consider is formally equivalent to a simultaneous mechanism. Finally, because the random-reserve first-price auction described above is equivalent to the optimal PFO, it follows that a random-reserve first-price auction is optimal (in the special case of a single buyer who only pays if she is awarded the object).

Although our selling mechanism is most closely linked to an auction with a single bidder, the same sorts of effects arise in auctions with multiple bidders. The randomness in the seller's reserve price imposes risk over and above the risk imposed by other buyers' bidding behavior that, in the case of risk-averse buyers, will induce them to bid more aggressively and may increase the seller's expected revenue.[20]

19. If desired, HIGH's indifference can be broken by letting the reserve price be $p_1(y) - \varepsilon$ with probability y and r_l with probability $(1 - y)$.

20. In a similar vein, McAfee and McMillan (1987) show that when bidders are risk averse, uncertainty can also increase seller revenues.

Similarly, the same effects would seem to arise in a Dutch auction, due to its correspondence to the first-price sealed-bid auction. Thus, the seller in such an auction may benefit from employing a stochastic reserve price, that is, a probabilistic rule for stopping the auction if the price drops too far. Faced with such a risk of "losing" the object, bidders will be induced to behave more aggressively. If the gain from the increase in the auction price outweighs the loss from the increased likelihood of the object going unsold, the seller will benefit from such a practice.

3. Extensions to the Basic Model

We now extend the analysis in three ways. First, we consider a model in which the seller makes walking away from a sale credible by potentially selling the object to another buyer, where the process of searching for the outside buyer is costly and stochastic. Second, we show in the context of the model from Section 2 that the results are robust to the introduction of additional buyer types. In fact, provided that all buyers are sufficiently risk averse, the optimal selling strategy is a sequence of Possibly-Final Offers, a "PFO cascade." Third, we show that the results also hold for more general specifications, allowing for the object to be multidimensional, including the purchase of more than one unit, and relaxing the assumption that the buyer's utility function is additively separable in money. Each of the three extensions stands alone.

3.1 Outside Opportunities and Costly Search

Our basic model assumes that the seller can commit to walking away from the buyer at no cost. In this section, we incorporate the cost of commitment into the model by assuming that, in addition to the buyer with unknown valuation, whom we will refer to as the "current" buyer, there are also "outside" buyers with known valuation p_0. The outside buyer's value for the object is less than HIGH's, $p_0 < r_h$.

The game we consider in this section is essentially the same as the one considered in our basic model, with the search aspect added. At the start of the game (stage 0), the seller undertakes costly investments that determine the probability y that an outside buyer will be available if the current buyer rejects the seller's initial offer. These investments could take the form of advertising, direct marketing, etc. At stage 1, the seller offers the current buyer price p_1. If the current buyer rejects, then, if an outside buyer is available, the seller offers the object to her at price p_0.[21] If there is no outside buyer available, the seller returns to the current buyer with a lower price, p_2.

To locate an outside buyer, the seller must make an *ex ante* investment. The greater the investment, the more likely it is that an outside buyer will be available to sell in case the current buyer rejects the seller's initial offer. Let $c(y)$ be the cost associated with locating an outside buyer with probability y. We assume that $c(y)$ is strictly increasing and strictly convex in y. In order to focus on the interesting case, we assume that $c'(0) > p_0 - r_l$. This assumption is trivially satisfied if $p_0 \leq r_l$. When $p_0 > r_l$, the seller would rather sell to the outside buyer than to LOW, which gives him an incentive to increase y. Due to this incentive, the seller might choose to make a probabilistic offer (i.e., choose $0 < y < 1$) even if he knew he faced a LOW current buyer. Assuming $c'(0) > p_0 - r_l$ ensures that this incentive alone is not sufficient to induce the seller to adopt a PFO strategy. Thus, as in our basic model, if the seller adopts a PFO strategy he will do so because threatening to walk away from the current buyer increases the HIGH current buyer's willingness-to-pay for the object immediately.

We also assume that the seller is able to commit to sell to the outside buyer at price p_0 whenever such a buyer is found. This assumption is without loss of generality when $p_0 > r_l$, because in this case the seller would rather sell to the outside buyer than to LOW. While the seller will prefer to sell to LOW rather than the outside buyer when $p_0 < r_l$, assuming that the seller is able to commit to sell to the outside buyer facilitates comparison with our basic model (because that model corresponds to the case where $p_0 = 0$). Hence, we do not rule out this case *ex ante*. Indeed, in cases where $0 < p_0 < r_l$, the seller's cost of walking away is lower than in our basic model. Thus, allowing for outside buyers with low reservation values actually lessens the severity of the seller's commitment problem.

We begin by deriving the full information benchmark. If the seller can observe the buyer's type and the current buyer is HIGH, the seller charges her price r_h. If the current buyer is LOW, our assumption that $c'(0) > p_0 - r_l$ implies that the seller's best strategy is to charge price r_l and sell to the current buyer with probability 1 rather than attempt to attract an outside buyer and then sell to the current buyer at price r_l if one does not materialize. Thus, as in our basic model, the seller's full information profit is $\mu r_h + (1 - \mu) r_l$.

21. Because the outside buyer's value is known, the seller should charge p_0 whenever he sells to the outside buyer.

If the seller cannot identify the current buyer's type, then he must offer a screening contract as in our earlier analysis. The seller's optimization problem is written as

$$\max_{p_1, p_2, y} \mu p_1 + (1-\mu)[y p_0 + (1-y) p_2] - c(y)$$

$$\begin{aligned} \text{s.t.} \quad & \theta_h + v_h(w_h - p_1) \geq 0, \\ & \theta_l + v_l(w_l - p_2) \geq 0, \\ & \theta_h + v_h(w_h - p_1) \geq (1-y)(\theta_h + v_h(w_h - p_2)), \text{ and} \\ & (1-y)(\theta_l + v_l(w_l - p_2)) \geq \theta_l + v_l(w_l - p_1). \end{aligned}$$

Importantly, notice that the changes in this model affect only the seller's objective function. The current buyer's participation and incentive compatibility constraints are unchanged.

Next we characterize the seller's optimal deterministic selling strategy. Due to the potential presence of the outside buyer, there are three possible deterministic offers that may be optimal. If the seller chooses not to search for an outside buyer, then it can either set its price at r_h and sell only to HIGH or at r_l and sell to all current buyers, regardless of type. These strategies offer expected profit μr_h and r_l, respectively. If the seller searches for an outside buyer, he must choose $y = 1$, in which case its optimal selling scheme is to charge price r_h initially and, if that offer is rejected, sell to the outside buyer at price p_0. The seller's expected profit under this alternative is $\mu r_h + (1-\mu)p_0 - c(1)$. Under our assumption that $c'(0) > p_0 - r_l$, it is straightforward to show that this strategy offers less than the full information profit. As before, PFO strategies correspond to the case where there is a positive probability strictly less than 1 that the current buyer will receive a second offer if the first offer is refused.

Because the current buyer's participation constraints and incentive-compatibility constraints are the same as in our basic model, the maximum first-stage price that HIGH will be willing to pay is still given by $p_1(y)$, as defined in (7). It is straightforward to show that $p_1(y) > yp_0 + (1-y)r_l$, in which case the seller's preferred selling mechanism involves quoting price $p_1(z)$ initially (which is accepted if the current buyer is HIGH), and, if this offer is rejected, selling to the outside buyer at price p_0 if available, and selling to the current buyer at price r_l otherwise.

The seller's problem can therefore be written as

$$\max_{0 \leq y \leq 1} \mu p_1(y) + (1-\mu)[(yp_0 + (1-y)r_l)] - c(y).$$

Taking the derivative, the optimal level of search satisfies

$$\mu p_1'(y^*) + (1-\mu)(p_0 - r_l) - c'(y^*) \begin{cases} \leq 0 & \text{if } y^* = 0 \\ = 0 & \text{if } 0 < y^* < 1 \\ \geq 0 & \text{if } y^* = 1. \end{cases}.$$

That is, for an interior solution, the seller invests resources to generate an outside buyer (i.e., increases y) until the expected marginal increase in profit from selling to HIGH equals the marginal cost of search less the marginal impact of increasing the likelihood of selling to the outside buyer at price p_0 instead of selling at stage 2 at price r_l.

Proposition 6: *In the costly search case, for any μ such that $0 < \mu < 1$, if HIGH is sufficiently risk averse using a PFO increases the seller's expected profit*: $0 < y^* < 1$. *As HIGH becomes infinitely risk averse, the seller's expected profit converges in the limit to the full-information profit of $\mu r_h + (1-\mu)r_l$.*

Adding the possibility of selling to an outside buyer introduces two new wrinkles to the problem. On the one hand, it is costly for the seller to generate an outside buyer, and this cost tends to drive him toward a small y. On the other hand, in the event that the seller walks away from the current buyer, the seller receives p_0, whereas in our basic model he received no profit. This factor gives the seller an incentive to raise y. In the limiting case we consider in the proposition, the seller is able to induce HIGH to pay almost her reservation price using an arbitrarily small threat of selling to the outside buyer, and thus the search cost does not significantly impact the results. For more moderate levels of risk aversion, the seller would have to balance the beneficial impact of HIGH's willingness-to-pay in the first period against the cost of finding an outside buyer and the mitigating fact that walking away from the current buyer now results in non-negative profit p_0.

Before going on, one case in particular bears mentioning. When $p_0 > r_l$, the seller would rather sell the object to the outside buyer than to LOW. Because of this, the seller's commitment to probabilistically walk away from the current buyer following an initial rejection is credible. If the buyer rejects the initial offer, then with probability y the seller has a better opportunity and credibly takes it. Otherwise, the seller returns to the current buyer and sells a lower price. Thus, the addition of search costs and an outside buyer endogenizes the commitment on the part of the seller that we assumed in our basic model. Interestingly, the seller's threat to walk away will be credible as long as $p_0 > r_l$. However, the existence of a buyer with value p_0 induces a very risk-averse HIGH buyer to be willing to pay (nearly) $r_h > p_0$ for the object because, from the point of view of HIGH, the item being sold to the outside buyer (in which case HIGH earns zero surplus) is equivalent to being offered the object at price r_h. The type of situation we discuss here, where an outside buyer is used to induce a current buyer to pay a higher price, plays a critical role in a number of the examples we discuss in Section 4.[22]

3.2 Selling to Multiple Types: PFO Cascades

Consider the model of Section 2, where the seller offers a single unit of a good to a buyer whose reservation price is unknown. Suppose there are $n > 2$ types of buyers, let r_k be the type-k buyer's reservation price, with $r_1 > r_2 > \cdots > r_n$.[23] Let μ_k be the prior probability of a buyer being type k, with $\mu_k > 0$ and $\sum_{k=1}^{n} \mu_k = 1$. A buyer of type k has initial wealth w_k and utility function

$$u_k(q, p) = \theta_k q + v_k(w_k - p),$$

where $q = 1$ if the buyer purchases the object and 0 otherwise, and p is the price she pays should she buy the object. For simplicity, we will let $v_k() = v()$ and $w_k = w$ for all k, although the results do not depend on this assumption.

The seller once again offers a sequence of prices. Let p_k be the kth price offered. Let y_k be the probability that, if the buyer rejects the seller's offer of p_{k-1}, the seller will walk away rather than make a kth offer. Once again we consider screening contracts where p_1 is accepted only by type 1, p_2, if a second offer is made, is acceptable to types 1 and 2, but type 1 prefers p_1, and thus p_2 is accepted only by type 2, and so on. Clearly, it is optimal to offer p_1 with probability 1, that is, $y_1 = 0$. Hence the seller's objective is to

$$\max_{y_2,\ldots,y_n,p_1,\ldots,p_n} \sum_{k=1}^{n} \left(\prod_{j=1}^{k} (1 - y_j) \right) \mu_k p_k. \tag{11}$$

Each type of buyer has a participation constraint and a set of incentive-compatibility constraints. The participation constraints are written as

$$\theta_k + v(w - p_k) \geq 0, \quad k = 1, \ldots, n. \tag{12}$$

22. A version of this model where the outside buyer's willingness-to-pay is endogenous is available from the authors.

23. This would also be the appropriate model in the two-type case if HIGH's and LOW's valuations were determined by a fundamental plus noise, and hence unknown. For example, if HIGH and LOW are equally likely, and HIGH's value is equally likely to be 4, 5, or 6, and LOW's value is equally likely to be 1, 2, or 3, this corresponds to a six-type case where the buyer's value is equally likely to be 1, 2, 3, 4, 5, or 6.

The incentive-compatibility constraints can be divided into those that ensure type k does not want to accept the first $k-1$ offers

$$\left(\prod_{j=t+1}^{k}(1-y_j)\right)(\theta_k + v(w-p_k)) \geq \theta_k + v(w-p_t),$$
$$k = 2, \ldots, n, \quad \text{and } t = 1, \ldots, k-1, \tag{13}$$

and those that ensure that type k prefers the kth offer to any subsequent probabilistic offer:

$$\theta_k + v(w-p_k) \geq \left(\prod_{j=k+1}^{t}(1-y_j)\right)(\theta_k + v(w-p_t)),$$
$$k = 1, \ldots, \; n-1 \quad \text{and } t = k+1, \ldots, n. \tag{14}$$

The optimal contract involves a PFO if the optimal solution involves $0 < y_k^* < 1$ for some k. The optimal contract involves a PFO cascade if $0 < y_k^* < 1$ for all $1 < k \leq n$.

Proposition 7: *When buyers are sufficiently risk averse, the solution to the multiple-type problem involves a PFO cascade. Further, as the buyers become infinitely risk averse, the seller's expected profit converges in the limit to the full-information maximum.*

The key observation underlying Proposition 7 is that each buyer is indifferent between the game ending and buying the object at her reservation price. Thus, just as in the two-type case, as long as there is a chance of the game terminating following a rejection, as buyers become very risk averse, their willingness-to-pay for the object rather than risk termination approaches their reservation value. Thus, the seller is able to extract this value even if the threatened chance of termination is small.

3.3 Multidimensional Objects and General Utility Functions

Consider a problem where an offer can relate to both the price of the object and its characteristics, x, a vector of attributes that may include quantity, various measures of quality, and other aspects of the object relevant for the buyers' valuations. Let $u_h(x, w_h - p)$ be HIGH's utility function and $u_l(x, w_l - p)$ be LOW's, and assume that $u_t(0, w_t) = 0$, that $u_t(x, w_t)$ is strictly quasi-concave for $t \in \{h, l\}$, and that $\frac{\partial^2 u_h(x,w)}{\partial w^2} < 0$ for all (x, w) (i.e., that HIGH is strictly risk averse over wealth).

Let reservation price, $r_t(x)$, satisfy $u_t(x, w_t - r_t(x)) = 0$ for $t = h, l$. Assume that $r_h(x) > r_l(x)$ for all x. Thus HIGH is willing to pay more for object x than is LOW, and any object-price pair that is just acceptable to LOW will offer strictly positive utility to HIGH.

The seller produces object x according to cost function $c(x)$. The seller's problem is

$$\max_{x_1, p_1, x_2, p_2} \mu(p_1 - c(x_1)) + (1 - y)(1 - \mu)(p_2 - c(x_2))$$

$$\text{s.t.} \quad u_h(x_1, w_h - p_1) \geq (1 - y)u_h(x_2, w_h - p_2), \tag{15}$$

$$u_h(x_1, w_h - p_1) \geq 0, \tag{16}$$

$$(1 - y)u_l(x_2, w_l - p_2) \geq u_l(x_1, w_l - p_1), \text{ and} \tag{17}$$

$$u_l(x_2, w_l - p_2) \geq 0. \tag{18}$$

Denote the solution to this problem by $(x_1^*, p_1^*, x_2^*, p_2^*, y^*)$ and indicate the seller's optimized profit by π^*.

Let $(x_1^F, p_1^F, x_2^F, p_2^F)$ be the pair of offers that would maximize profit with full information, assumed for simplicity to be unique. That is, (x_1^F, p_1^F) maximizes $p_1 - c(x_1)$ subject to $u_h(x_1, w_h - p_1) \geq 0$, and (x_2^F, p_2^F) maximizes $p_2 - c(x_2)$ subject to $u_l(x_2, w - p_2) \geq 0$. Under our assumptions, $(x_1^F, p_1^F, x_2^F, p_2^F)$ is unique and represents the theoretical maximum (i.e., full-information) profit those buyers can yield. Let π^F stand for the seller's profit in this case.

Let $(x_1^D, p_1^D, x_2^D, p_2^D, y^D)$ solve the seller's problem in the deterministic case where there is incomplete information but PFOs are not permitted (i.e., $y \in \{0, 1\}$). Again, assume $(x_1^D, p_1^D, x_2^D, p_2^D)$ is unique. Let π^D stand for the seller's profit in this case, and note that under our assumptions, $\pi^D < \pi^F$, because $(x_1^D, p_1^D, x_2^D, p_2^D)$ must satisfy (15), but $(x_1^F, p_1^F, x_2^F, p_2^F)$ violates it.

Proposition 8: *When HIGH is sufficiently risk averse, the seller's profit-maximizing contract involves a PFO, that is, $0 < y^* < 1$. Further, the seller's expected profit is nondecreasing with HIGH's risk aversion, and strictly increasing with risk aversion when the seller uses a PFO. Finally, as HIGH becomes infinitely risk averse the seller's expected profit ultimately converges to the full-information profit, $\pi^* \to \pi^F$.*

The key step in the proof is to note that the lottery HIGH faces when she rejects the first offer is equivalent to a lottery where, with probability y, HIGH is offered the object/price combination that would maximize the seller's profit given full information about HIGH's utility function (giving HIGH 0 utility), and with probability $1 - y$ the buyer is offered the same object at a lower price. The proof then proceeds as in the proof of Proposition 2. As HIGH's risk aversion increases to infinity (i.e., as $u_h()$ becomes more concave along the monetary dimension), the expected utility of this lottery decreases to zero for any y strictly between 0 and 1, and the seller is able to extract more and more of HIGH's surplus at ever decreasing cost in terms of lost sales to LOW.

3.4 PFOs Using Price and Quality or Quantity[24]

An important special case of multidimensional PFOs arises where the seller determines both the price and quality of the object to be sold. All results in this section apply if quantity, that is, the number of units sold to the buyer, is substituted for quality. Because Proposition 8 applies, we know that a PFO strategy is optimal for any mix of buyers when HIGH is sufficiently risk averse. So, rather than prove the general result again, we instead make plausible assumptions about the form of the buyers' utility functions that allow us to explicitly derive the optimal contract and analyze its comparative statics.

Let q be the quality of the object, measured in terms of the dollars required to produce that level of quality, and let p be the purchase price of the object. The analysis is the same if q represents quantity instead of quality. As before, there are two types of buyers, HIGH and LOW, and the probability that the buyer is HIGH is μ, where $0 < \mu < 1$. HIGH has utility function $u_H(q, p) = (f(q) - p)^{\frac{1}{b}}$, where f is a strictly increasing, strictly concave function. LOW has utility function $u_L(q, p) = (af(q) - p)^{\frac{1}{b}}$, where $0 < a < 1$. Thus, as specified, any offer (q, p) that is acceptable to LOW is also acceptable to HIGH, which is the application of the assumption that $r_h(x) > r_l(x)$ from Section 3.3. The risk aversion of a buyer over surplus, $f(q) - p$, is captured by $b > 1$, where higher values of b correspond to more risk-averse buyers. To simplify the analysis, we assume that $f(q) = \sqrt{q}$ and all propositions and corollaries are proved for this case.[25]

If the buyer accepts offer (q, p), the risk-neutral seller earns profit $p - q$. It is straightforward to show that if the seller were able to identify the buyer's type, that is, there were full information, the seller would offer HIGH the price $p_1^F = \frac{1}{2}$ and quality $q_1^F = \frac{1}{4}$, and LOW the price $p_2^F = \frac{a}{2}$ and quality $q_2^F = \frac{a^2}{4}$.

The structure of the game when the seller doesn't know the buyer's type is as in the earlier sections, although offers now comprise quality-price pairs. The seller makes a first offer (q_1, p_1), which the buyer may

24. Supporting computations and proofs of the propositions in this subsection are available from the authors.

25. The same qualitative results hold as long as $f(q)$ is strictly increasing and strictly concave, $f(0) = 0$, and $f'(0) > \frac{1}{a}$.

either accept or reject. If the buyer rejects this initial offer, the seller walks away with probability y or makes a second offer (q_2, p_2) with probability $1 - y$. As previously, the seller's first offer is tailored to be accepted by HIGH but not by LOW. The second offer is designed so that if it is made, LOW accepts.

The active constraints in this game are LOW's participation constraint and HIGH's incentive-compatibility constraint. Thus the seller's problem is written as

$$\begin{aligned} \max_{p_1, s_1, p_2, s_2, z} \quad & \mu(p_1 - q_1) + (1 - y)(1 - \mu)(p_2 - q_2), \\ \text{s.t.} \quad & (f(q_1) - p_1)^{\frac{1}{b}} \geq (1 - y)(f(q_2) - p_2)^{\frac{1}{b}}, \text{ and} \\ & (af(q_2) - p_2)^{\frac{1}{b}} \geq 0. \end{aligned} \tag{19}$$

As before, we begin by considering deterministic contracts. As is usual in screening problems, LOW's contract is designed to offer her zero surplus, while HIGH's contract is designed so that HIGH is just indifferent between the two offers. Denote the optimal deterministic contract with the superscript D, and let π^D be the optimized level of the seller's profit. Proposition 9 summarizes the optimal deterministic contract. Supporting computations and proofs for all propositions and corollaries are presented at the end of the Appendix.

Proposition 9: *If $\mu \leq a$, then the optimal deterministic contract is*

$$\left(q_1^D, p_1^D\right) = \left(\frac{1}{4}, \frac{1 - a - \mu a + a^2}{2(1 - \mu)}\right), \quad \textit{and} \tag{20}$$

$$\left(q_2^D, p_2^D\right) = \left(\left(\frac{a - \mu}{2(1 - \mu)}\right)^2, \frac{a(a - \mu)}{2(1 - \mu)}\right), \tag{21}$$

and $\pi^D = \frac{1}{4}\frac{\mu - 2\mu a + a^2}{1 - \mu}$. If $\mu > a$, the optimal deterministic contract is

$$\left(q_1^D, p_1^D\right) = \left(\frac{1}{4}, \frac{1}{2}\right), \quad \textit{and} \quad \left(q_2^D, p_2^D\right) = (0, 0),$$

and $\pi^D = \frac{\mu}{4}$.

As in our basic model, when μ is sufficiently large, it is optimal for the seller to contract only with HIGH. When the proportion of HIGHs in the population is relatively small, though, it becomes optimal for the seller to contract with both HIGH and LOW. However, because the seller now has two instruments available, he can make different offers to HIGH and LOW even without randomization. First, he offers a high quality at a high price, which is acceptable to HIGH but not to LOW. The second offer is a lower quality at a lower price, which is (just) acceptable to LOW, but because HIGH values quality more, it is inferior to the first offer from HIGH's point of view. Varying quality provides an instrument to discriminate between the types.

Randomization via PFO strategies provides an additional discrimination instrument and thereby complements quality variation. Recall the seller's problem, (19). Proposition 10 establishes that a PFO strategy is optimal if HIGH is sufficiently risk averse.

Proposition 10: *When PFOs are permitted, the seller prefers a PFO strategy to the optimal deterministic contract, provided that HIGH is sufficiently risk averse. When* $b \geq b^* = \frac{1}{2}(\frac{a(1-\mu)}{\mu(1-a)} + 1)$, *the optimal PFO contract is given by*

$$\left(q_1^*, p_1^*\right) = \left(\frac{1}{4}, \frac{1}{2} - \left(\frac{a(1-\mu)}{(2b-1)\mu(1-a)}\right)^{\frac{b}{b-1}} (1-a)a\frac{b-1}{2b-1}\right), \tag{22}$$

$$\left(q_2^*, p_2^*\right) = \left(\left(a\frac{b-1}{2b-1}\right)^2, a^2\frac{b-1}{2b-1}\right), \quad \textit{and} \tag{23}$$

$$y^* = 1 - \left(\frac{a(1-\mu)}{(2b-1)\mu(1-a)}\right)^{\frac{1}{b-1}}. \tag{24}$$

Otherwise, the seller maximizes expected profit by offering the optimal deterministic contract, (20) and (21).[26]

When PFOs are permitted, the seller gains an additional instrument, randomization, that he can use to separate the types. Making a second offer only probabilistically relaxes HIGH's incentive-compatibility constraint, but does so at the cost of reducing the expected revenue from selling to LOW. However, the magnitude of this cost decreases with HIGH's risk aversion. When HIGH is sufficiently risk averse, it becomes optimal to use randomization, at least to some extent, and in response the seller is able to increase the price from HIGH, and increase both the price and quality to LOW, along with the associated profit margin.

The minimum level of risk aversion for which a PFO is optimal, which is captured by b^*, depends on the parameters of the problem in an intuitive way. It is decreasing in μ because, as μ increases, the loss due to haggling decreases and the benefit increases, because there are fewer LOWs and more HIGHs in the population. And, it is increasing in a because, as a decreases, the potential profit earned by selling to LOW decreases. Thus, for any fixed b, a PFO becomes less costly and so is more likely to be profitable.

26. Note that if $\mu > a$, $b^* < 1$, in which case a PFO is always optimal.

For a given level of risk aversion, the optimal PFO contract offers a lower first price and a lower second quality and price than the full-information optimum. This distortion arises as the cost of separating the two types. However, as HIGH becomes infinitely risk averse this distortion vanishes, and the cost approaches zero.

Corollary 1: *As b approaches infinity, the optimal PFO contract converges to the full-information optimum,* $(q_1^F, p_1^F, q_2^F, p_2^F, y) = (\frac{1}{4}, \frac{1}{2}, \frac{a^2}{4}, \frac{a^2}{2}, 0)$.

4. Uncertainty-Inducing Sales Practices

Sellers employ a number of practices that capture the spirit of PFO strategies. When studying the examples, it is important to keep two features of our results in mind. The first is that, when buyers are sufficiently risk averse, the seller benefits from even small deviations from the optimal deterministic strategy. Thus, in order to benefit from a PFO, the seller needs to impose some uncertainty on the buyer, but it is not necessary that he commit to the optimal PFO. For example, when the optimal deterministic strategy is to sell to all buyers at the LOW's reservation value, convincing sufficiently risk-averse buyers that there is even a small probability of termination will raise the seller's expected profit. Second, in order to induce the buyer to pay a higher initial price, the seller needs to convince the buyer that there is a chance that, if she rejects the offer, she will not get another opportunity to purchase the item at a lower price. However, this does not imply that the seller must destroy (or permanently keep) the item if the buyer rejects his initial offer.

Sellers frequently threaten that if the buyer does not purchase the item immediately the item will be sold to another buyer in order to sow the seeds of uncertainty in buyers' minds. For example, car salespeople will often tell a buyer interested in a particular vehicle that another buyer is coming in later that day to look at that same car. In this case, the possibility that the car will be unavailable later is meant to induce the buyer to purchase immediately rather than wait and try to win better terms. If the seller is unable to make an alternative sale within a reasonable amount of time, he will later sweeten the deal.[27] Similarly, real estate agents often begin showing a house by having an

27. This strategy works best for one-of-a-kind items. An antique silver dealer, who has many hard-to-find or unique items, may say in effect: "This is my price for this hard-to-find piece. I expect that one of my regular customers may purchase it in the next few weeks. If you come back then, and it has not sold, I will offer you a better price." Robert Stern, Silstar, London Silver Vaults, personal communication, March 2005.

open house, in hopes that buyers will perceive that there are many others interested in the same house. Again, the uncertainty about whether the item will be available at a later date often leads risk-averse buyers to pay substantially more earlier, despite the fact that more favorable terms may become available later.

A similar strategy employed by some retailers is to intentionally keep supplies low. A buyer faced with the chance to purchase a "one of a kind" item knows that her opportunity will be lost if another buyer purchases the item before she can. Thus she may be unwilling to wait to see if the price declines in the future. Filene's Basement uses a PFO-type strategy at its flagship Boston store. Each item offered for sale is marked with the date on which it was originally made available. Items are then automatically marked down to 75% of the original price after 14 days of going unsold, 50% of the initial price after 21 days, and 25% of the initial price after 28 days.[28] Again, the possibility of lower prices in the future is coupled with the very real chance that the item may be purchased by someone else before the price drop occurs. Many other retailers follow versions of this strategy, offering items as sale or clearance prices, but only if they do not sell out at their regular prices.

A similar practice is prevalent on eBay, where many sellers of goods on eBay offer "Buy It Now" prices on items. An item is generally available at its Buy It Now price until the first bid is made on the item. A buyer who values the item above the Buy It Now price can purchase it, or wait in the hope that the eventual auction price will be lower. However, by not purchasing immediately the buyer takes the risk that some other buyer might Buy It Now or that the winning auction price will be higher than the Buy It Now price. These risks increase the attractiveness of the Buy It Now price, and may improve the seller's expected outcome.

To use a nonsales example, when a new Ph.D. is given an assistant professorship offer, she does not know whether the department or dean has other potential hires waiting in queue. Thus, if the department is limited to only make one offer at a time, then the offer recipient does not know whether a rejection will lead the department to improve the offer or move on to the next candidate on the list.

Sellers may also try to foster uncertainty about whether improved offers will be made by employing agents to sell for them. For example, the owners of car dealerships sell via salespeople, and it is often unclear just how much discretion the salesperson has to make a deal. Often, following a rejection of his initial price, the salesman will claim to "check" with his manager about whether it is possible to improve the deal or not. Uncertainty about the manager's reaction to a rejected offer

28. Items remaining unsold after 35 days are given to charity.

can serve to make a PFO strategy effective. (Presumably, managers have a stronger reason to maintain a reputation for not continuing to cut prices than would an individual salesman.) Agency arrangements also play an important role in corporate takeover negotiations, where CEOs work out the terms of the deal, but the final agreement must ultimately be approved by the board of directors. Thus, a CEO making a board-approved initial offer may benefit from giving the impression that the board is unlikely to approve better terms.

5. Conclusion

A seller who sets the terms in a buyer–seller interaction, which the buyer then accepts or rejects, will prefer a PFO strategy to any deterministic selling strategy, provided that the buyer is sufficiently risk averse. Our results hinge on sellers exploiting buyers' risk aversion. Elsewhere, it has been argued that, at least when stakes are small to moderate, behavior that appears to be caused by risk aversion may actually be caused by loss aversion.[29] However, buyer loss aversion only serves to reinforce our results. That is, loss aversion makes PFOs more likely to be beneficial. Beginning from a particular reference point, loss averse agents weigh losses much more than equivalent-sized gains. For example, Tversky and Kahneman (1991) find people to be indifferent between an equal probability of a loss of $\$x$ and a gain of $\$2x$, and \$0 for sure. In the context of our PFO game, the loss incurred when a buyer rejects an initial offer and the seller walks away is likely to loom large, all the more so because the high-valued buyer could have secured the item by accepting the seller's initial offer. Thus, loss aversion also drives buyers to accept higher initial offers, and consequently reinforces risk aversion in making PFOs effective tools for sellers.[30]

Although our analysis addresses the case where the buyer is risk averse over monetary outcomes, the results generalize to any case in which the buyer's utility function exhibits curvature (or a break in slope due to loss aversion). For example, PFO strategies would also be optimal in a situation where buyers have linear utility for money, but concave utility for quality. PFO strategies reap benefits because imposing risk on the buyer enhances the seller's ability to discriminate between buyers who differ on willingness-to-pay. A buyer with "curvature" is willing to sacrifice expected value to avoid losing a deal, and the threat of breaking

29. See Kahneman and Tversky (1979) for a description of loss aversion, originally presented in the context of Prospect Theory. See Rabin (2000) for an elegant discussion of risk aversion and loss aversion over moderate-stakes gambles.

30. Indeed, our results regarding the desirability of PFOs also hold when buyers are sufficiently loss averse but risk neutral or when buyers are both risk and loss averse.

off negotiations deters the high-valued buyer more than the low-valued one. In more elaborate formulations, curvature could complement other instruments for separation, such as differences in time preference or distaste for negotiation. Although this paper posits that the seller sets the terms of the deal, its results transfer seamlessly to the case where the buyer facing a risk-averse seller of unknown type sets the terms. Thrusting risk on risk-averse responders, as PFO strategies demonstrate, can be an effective tool for extracting surplus.

Appendix

Proof of Proposition 1. Let $\pi(y) = \mu p_1(y) + (1-y)(1-\mu)r_l$. When $\mu = \frac{r_l}{r_h}$, then $\pi(0) = \pi(1) = r_l$. Because $p_1(y)$ is strictly concave, so is $\pi(y)$. Therefore, $\pi(y)$ has a unique, interior maximizer. □

Proof of Proposition 2. A PFO is superior to deterministic contracting when there exists a y such that

$$\mu p_1(y) + (1-y)(1-\mu)r_l > \max\{\mu r_h, r_l\}.$$

An increase in a decision maker's risk aversion is equivalent to a decrease in her certainty equivalent for any lottery.[31] Hence, increasing risk aversion corresponds to a pointwise increase in $p_1(y)$. Let $p_1^n(y)$ be a sequence of strictly decreasing, strictly concave functions such that $p_1^n(0) = r_l$, $p_1^n(1) = r_h$, $p_1^{n+1}(y) > p_1^n(y)$ pointwise for all n, and $\lim_n p_1^n(y) = r_h$ for $y \in (0, 1)$. Hence $p_1^n(y)$ corresponds to increasingly risk-averse versions of HIGH. Let $y_n > 0$ be a sequence of walk-away probabilities with $\lim y_n = 0$.

$$\lim_{n\to\infty} \mu p_1^n(y_n) + (1-y_n)(1-\mu)r_l = \mu r_h + (1-\mu)r_l.$$

Because the seller's profit converges to the full information profit for any such sequence y_n, it also converges for the optimal sequence. □

Proof of Proposition 3. If HIGH is risk neutral, $p_1(y) = yr_h + (1-y)r_l$, and so $p_1'(y) = r_h - r_l$, and (9) becomes

$$\mu(r_h - r_l) - (1-\mu)r_l = \mu r_h - r_l,$$

which does not depend on y. Hence at the optimum, either $y^* = 0$, and the seller sells to all buyers at price r_l, or $y^* = 1$, and the seller sells to only HIGH at price r_h. □

31. See Mas-Colell et al. (1995), proposition 6.C.2.

Proof of Proposition 4. Let $v_{h1}()$ and $v_{h2}()$ be two monetary utility functions for HIGH such that $v_{h2}()$ is more risk averse than $v_{h1}()$. If the seller uses a PFO strategy against neither $v_{h1}()$ nor $v_{h2}()$, the expected profit is the same in both cases, and the result follows. Let $p_{1k}()$ be the reservation-price function when HIGH's utility function is $v_{hk}()$. Let $y_k \in [0, 1]$ be the optimal probability of a second offer when HIGH has utility function $v_{hk}()$. Expected profit when HIGH has utility function $v_{h1}()$ is given by

$$\begin{aligned}\mu p_{11}(y_1) + (1 - y_1)(1 - \mu)r_l &\le \mu p_{12}(y_1) + (1 - y_1)(1 - \mu)r_l \\ &\le \mu p_{12}(y_2) + (1 - y_2)(1 - \mu)r_l.\end{aligned}$$

The first inequality follows from $p_{12}(y) > p_{11}(y)$ and $y_1 \in [0, 1]$. The second inequality follows from the fact that y_1 is feasible but not optimal when HIGH has utility function v_{h2}. When $y_1 \in (0, 1)$, the first inequality is strict. □

Proof of Proposition 5. Available from the authors.

Proof of Proposition 6. A PFO is superior to deterministic contracting when there exists a y such that

$$\begin{aligned}&\mu p_1(y) + (1 - \mu)(yp_0 + (1 - y)r_l) - c(y) \\ &\quad > \max\{\mu r_h, \mu r_h + (1 - \mu)p_0 - c(1), r_l\}.\end{aligned}$$

Using the same argument as in the proof of Proposition 2, for any fixed $y > 0$, the left-hand side converges to $\mu r_h + (1 - \mu)(yp_0 + (1 - y)r_l) - c(y)$ as HIGH becomes infinitely risk averse. Letting y be arbitrarily small, this converges to $\mu r_h + (1 - \mu)r_l$, the full information profit. □

Proof of Proposition 7. Available from the authors.

Proof of Proposition 8. Available from the authors.

References

Arnott, R. and J. Stiglitz, 1988, "Randomization with Asymmetric Information," *RAND Journal of Economics*, 19, 344–362.

Bajari, P. and A. Hortacsu, 2003, "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 32, 329–355.

Brito, D., J. Hamilton, S. Slutsky, and J. Stiglitz, 1995, "Randomization in Optimal Income Tax Schedules," *Journal of Public Economics*, 56, 189–223.

Das Varma, G. and N. Vettas, 2001, "Optimal Dynamic Pricing with Inventories," *Economics Letters*, 72, 335–340.

Deneckere, R. and R.P. McAfee, 1996, "Damaged Goods," *Journal of Economics and Management Strategy*, 5, 149–174.

Kahneman, D. and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.

Laffont, J.J. and D. Martimort, 2002, *The Theory of Incentives: The Principal-Agent Model*, Princeton: Princeton University Press.

Li, H. and G. Tan, 2000, "Hidden Reserve Prices with Risk Averse Buyers," Mimeo, USC Department of Economics.

Li, T. and I. Perrigne, 2003, "Timber Sale Auctions with Random Reserve Prices," *Review of Economics and Statistics*, 85, 189–200.

Mas-Colell, A., M. Whinston, and J. Green, 1995, *Microeconomic Theory*, New York: Oxford University Press.

Maskin, E. and J. Riley, 1984, "Optimal Auctions with Risk Averse Buyers," *Econometrica*, 52, 1473–1518.

Matthews, S., 1983, "Selling to Risk Averse Buyers with Unobservable Tastes," *Journal of Economic Theory*, 30, 370–400.

McAfee, R. P. and J. McMillan, 1987, "Auctions with a Stochastic Number of Bidders," *Journal of Economic Theory*, 43, 1–19.

Perrigne, I., 2003, "Random Reserve Prices and Risk Aversion in Timber Sale Auctions," Mimeo, Pennsylvania State University.

Rabin, M., 2000, "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 68, 1281–1292.

Riley, J. and R. Zeckhauser, 1983, "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, 98, 267–289.

Rustichini, A. and A. Wolinsky, 1995, "Learning About Variable Demand in the Long Run," *Journal of Economic Dynamics and Control*, 19, 1283–1292.

Tversky, A. and D. Kahneman, 1991, "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *The Quarterly Journal of Economics*, 106, 1039–1061.

Weiss, L., 1976, "The Desirability of Cheating Incentives and Randomness in the Optimal Income Tax," *Journal of Political Economy*, 84, 1343–1352.