Optimal Selling With Risk-Averse Agents

Tomer Siedner*

December 30, 2020

Abstract

We consider the basic setup of one seller, one buyer, and one good, when the agents exhibit *risk aversion*. We give necessary conditions that characterize the *optimal mechanism* that maximizes the seller's expected utility. In contrast to the risk-neutral case, where a single deterministic price is optimal, when the agents are risk averse the optimal mechanism is continuous and consists of a continuum of lotteries. However, the allocation function remains partly unchanged: the buyer types who get the good with probability one when the agents are risk averse are exactly those that buy the good in the risk-neutral case. We also introduce a regularity condition under which our characterization enables us to directly calculate the optimal mechanism. Finally, we prove that "one-price" mechanisms guarantee the seller a bounded fraction of the optimal utility and this fraction depends only on the distribution of the buyer's valuation for the good.

*Department of Economics and the Federmann Center for the Study of Rationality, the Hebrew University, Jerusalem. This paper was written as part of the author's dissertation, under the supervision of Sergiu Hart, whose advice, guidance, and patience are deeply appreciated. The author wishes also to thank Ilan Nehama for numerous discussions and suggestions that improved the paper tremendously, as well as Juan Carlos Carbajal, Murali Agastia, Ilan Kremer, Eyal Winter, Yannai Gonczarowski, Igal Milchtaich, and other seminar participants at the Hebrew University, for their valuable comments. The research was partially funded by ERC grant 249159.

Contents

1 Introduction

An agent – the seller – has an object she wants to sell. Another agent – the buyer – wants that object. What should the seller do? How should the seller manage the interaction between them? In this scenario, the seller cannot achieve a better expected revenue than via a "deterministic posted-price" mechanism.[1] The seller offers the good for a fixed price, and trade occurs only if the buyer is willing to pay at least this price. This famous result, first described in the early eighties by Myerson (1981); Riley and Samuelson (1981), and Riley and Zeckhauser (1983), hinges on several assumptions, one of them being that the agents are risk neutral. This assumption, however, is regularly violated in real life, where people often demonstrate risk aversion. On that account, our paper

[1] Sometimes called "one-price mechanism."

Payment	Utility loss
3	3
10	20
25	75
30	120

Table 1.1: Utility loss from specific payments

aims to analyze the same basic situation – one seller, one buyer, one good – with a small yet significant change: the agents in our setup are *risk averse*.

As mentioned above, when the agents are risk neutral, the deterministic posted-price mechanism is optimal. However, the seller can exploit any other attitude toward risk of the buyer (either risk aversion or risk loving) to screen between different buyer types, and thus extract more revenue. The following example demonstrates this for a risk-averse buyer.[2] Consider a risk-neutral seller facing a buyer who has two possible types: he is willing to pay for the good either 10 or 30 with equal probabilities. The revenue-maximizing mechanism, if the buyer is risk neutral, will set a price of 30 for the good, attaining an expected revenue of $(1/2) \cdot 30 = 15$. Assume now that the buyer is risk averse and, specifically, that his utility loss from payments is as in Table 1.1. It is easy to verify that with this risk-averse buyer a mechanism that offers the buyer to either pay 3 for a 33% chance of winning the good or buy the good for 29 will yield the seller an expected revenue of $(1/2) \cdot 3 + (1/2) \cdot 29 = 16$.

Furthermore, a risk-averse seller does not maximize her expected revenue, but rather her expected utility from the revenue. Such a risk-averse seller will forgo some of her expected revenue in order to be able to sell with higher probability. Thus, even if the buyer is risk neutral, the revenue-maximizing mechanism might not maximize the utility of the risk-averse seller. To illustrate this, assume that $u(x) = \min\{x, 100\}$; i.e., the seller has no need of more than

[2] See Evdokia et al. (2018) for the case of a risk-loving agent, where the optimal mechanism is a "randomized posted-price" mechanism.

100 units. Assume that the buyer values the object at either 100 or 300 with equal probability. The revenue-maximizing mechanism will set a price of 300, yielding an expected revenue of $(1/2) \cdot 300 = 150$ and an expected utility of $(1/2) \cdot 100 = 50$. The utility-maximizing mechanism, however, will set a price of 100, thus achieving a lower expected revenue of 100, but a higher expected utility of 100.

Although the case of risk-neutral agents has been studied intensely (see, e.g., Myerson, 1981; Riley and Samuelson, 1981; Riley and Zeckhauser, 1983; Klemperer, 1999), there are surprisingly few results for the case of risk-averse agents. Most of the literature on mechanism design under risk aversion either compares known auctions (e.g., Riley and Samuelson, 1981; Matthews, 1987; McAfee and McMillan, 1987a,b; Hu et al., 2010; Krishna, 2002), or analyzes what happens as the agents become more and more risk averse (e.g., Klemperer, 1999; Hu et al., 2010). A very few papers attempt to characterize the optimal mechanism, and the most notable among them is Maskin and Riley (1984). Other papers, in trying to characterize the optimal mechanism, make use of additional changes to the basic setup. Thus, for example, Eso and Futo (1999) allow for monetary transfers between the buyers and their mechanisms violate interim individual rationality. In another, more recent paper, Baisa (2017) characterizes the optimal mechanism for a setup that assumes risk aversion and a "good product."[3] In yet another recent paper, Kazumura et al. (2020) try to characterize the optimal mechanism among a subset of the possible mechanisms they deem "desirable".

However, to the best of our knowledge, our paper is the first to characterize the optimal mechanism in the setup described above, i.e., with a risk-averse seller facing a risk-averse buyer.

[3] As a result, Baisa's mechanism requires the buyers to submit their entire price-demand curve.

In light of the above, our contribution to the existing literature is threefold.

1. We prove that, given our setting, a unique optimal mechanism always exists, and establish three properties that are all necessary conditions for a mechanism to be optimal:

   (a) The optimal mechanism is continuous and offers a continuum of lotteries.[4]

   (b) The buyers who buy the good under the "revenue-maximizing" (RM) mechanism are exactly those who get the good for certain under the optimal mechanism. Thus, for the most eager buyers, the optimal allocation is the same as the revenue-maximizing allocation.

   (c) The optimal payment function must satisfy first-order conditions for maxima, which are essentially a continuous version of the Kuhn–Tucker conditions, tailored to our case.

2. In addition, when the buyer's valuation for the good is drawn from a probability distribution F (with a probability density function f) for which $x^2 f(x)$ is a strictly increasing function, finding the optimal mechanism is reduced to optimizing one parameter only. If, along with $x^2 f(x)$ being a strictly increasing function, the lowest buyer type has no value for the object, then our characterization completely identifies the optimal mechanism and provides the tools to calculate it directly.

3. We also consider the one-price utility-maximizing (OPUM) mechanism, which, among all one-price mechanisms, yields the seller the highest ex-

[4] This result, which is in contrast to the risk-neutral case, where the optimal mechanism has a single, deterministic price, seems surprising at first. However, we want to thank an anonymous referee for observing that following Manelli and Vincent (2007), this is to be expected, at least in the case of a risk-neutral buyer. As per their analysis, we maximize utility over a convex set: all incentive-compatible direct mechanisms. A linear functional, as in the risk-neutral case, is maximized at an extreme point: a one-price mechanism. However, a strictly concave function, as in the risk-averse case, is not maximized at an extreme point, and thus must include randomization.

pected utility. We provide a lower bound [5] on the fraction of the optimal utility that the OPUM mechanism can guarantee for the seller (see Proposition 26 for more details). Moreover, when the agents have CARA utility functions, the OPUM mechanism is shown to yield up to 99.5% of the optimal utility (see Section 5.2 for more details). Hence, the OPUM mechanism might be a good substitute when the optimal mechanism is hard to find.

Lastly, we represent our agents by additive separable concave utility functions. It is worth mentioning that our utility functions can also be interpreted in different ways, other than as risk aversion. Thus, for example, Che and Gale (1998) use the same utility functions to represent buyers who face financial constraints.

The paper is organized as follows. In Section 2 we define the setup and present a reduction to the case of a risk-neutral buyer. In Section 3 we define what are mechanisms and discuss revenue-maximizing mechanisms. In Section 4 we characterize the optimal mechanism when the buyer is risk neutral. In Section 5 we give an improved characterization of the optimal mechanism under a specific regularity condition. In Section 6 we attempt to determine when the OPUM mechanism is nearly as good as the optimal mechanism. In Section 7 we finally generalize our previous results to the case of a risk-averse buyer. In the last section we further generalize some of our results.

2 Preliminaries

As explained in the Introduction, we study the classic problem of characterizing the optimal mechanism when there is one seller, one buyer, and one good, with

[5] This lower bound depends only on the distribution of the buyer's valuation for the good, not on the utility functions of the agents.

a simple, yet significant, difference from the standard model: our agents are risk averse and the seller aims to maximize her expected utility.

2.1 Basic setup

In line with the literature, we adopt the framework of expected utility, and assume that the agents have vNM utility functions. Further, we assume that the seller has no costs and derives no utility from having the good (other than the revenue gained from the sale). We also assume that the buyer's utility from money and possession of the good is an additive separable utility function.[6] Thus, we represent the seller and the buyer via the utility functions $u(s)$ and $V(x,k,s) = k \cdot x + v(-s)$, respectively, where x denotes the utility the buyer derives from the good (his type[7]), k is a binary variable that has a value of 1 if the buyer receives the good and 0 otherwise, and s is the payment from the buyer to the seller (in dollars). The functions u and v are utility functions for money, which we assume, without loss of generality, are null at zero. We also assume that any utility function for money satisfy the following rather standard restrictions.

Assumption A.
A utility function for money:

1. is twice differentiable.
2. is strictly increasing.
3. has a finite right derivative at zero.

[6]The standard models assume quasi-linearity which is a special case of separable additive utility.

[7]Assuming additive separable utility function, the buyer's utility from the good does not depend on his current wealth. By contrast, when the buyer is risk averse, his willingness to pay does depend on his wealth.

That is why we choose to identify the buyer's type with his utility from the good rather than with his willingness to pay.

Still in line with the literature, we assume that the buyer's type is drawn from a probability distribution F (with a probability density function f) that is known to the seller, and satisfies the following assumptions.

Assumption B.

The probability distribution of the buyer types:

1. is atomless.
2. has full support on $[a, c]$, and is zero outside this interval.
3. is twice continuously differentiable.

We diverge from the standard models by taking the agents' utility functions for money, u and v, to be concave. As a result, our agents exhibit risk aversion rather than risk neutrality. Note that the analysis is far from trivial and the results are completely different compared to the standard literature, even though we introduced risk aversion through the minimal change possible and kept the separable additive utility assumption. Furthermore, as the next two sections show, it is enough that only one of the agents exhibits risk aversion, and already the analysis changes radically. Of course, it will be interesting to generalize our results to more general utility functions.[8]

2.2 Reduction to a risk-neutral buyer

Much is known about selling optimally when the agents are risk neutral. Therefore, we will now present a reduction that will enable us to use known results that are based on the assumption of a risk-neutral buyer, and do not take into account the seller's attitude toward risk. Such, for example, is the case with a number of observations Myerson (1981) made in regard to the optimal mechanism. Apparently, substituting the utility units (utils) of the buyer for money,

[8]For more general representations see, e.g., Maskin and Riley (1984); Matthews (1987); Baisa (2017).

the buyer becomes risk neutral while the seller remains risk averse. Hence, working with the buyer's utils instead of money, we can analyze a scenario with a risk-averse buyer as if the buyer were risk neutral. Thus, the payment function determines for each buyer type how many utils he loses, and not how much money he has to pay.

Assume a risk-averse seller who has a single good and is represented by a utility function for money, u. Assume, in addition, that this seller is facing a risk-averse buyer whose valuation for the good is drawn from a probability distribution F and whose utility function is $V(x,k,s) = k \cdot x + v(-s)$, as explained above. Furthermore, assume that u and v satisfy Assumption A, and F satisfies Assumption B. Finally, our agents are risk averse; therefore, we will assume that u and v are strictly concave functions.

In order to obtain the reduction to the case of a risk-neutral buyer, we define $s_v = -v(-s)$, and so $V(x,k,s) = k \cdot x + v(-s) = k \cdot x - s_v$ and[9] $u(s) = u\left(-v^{-1}(-s_v)\right)$, which we denote by $\tilde{V}(x,k,s_v)$ and $u_v(s_v)$, respectively. Evidently, $\tilde{V}$ is linear in s_v and, in addition, $u_v(s_v)$ satisfies Assumption A and is strictly concave in s_v (see Claim 42 in the Appendix for the proof). Consequently, we can denote our agents by $u_v(s_v)$ and $\tilde{V}(x,k,s_v)$, which is exactly the case of a risk-averse seller who faces a risk-neutral buyer. Naturally, since $u_v(s_v) = u(s)$, the expected utility the seller derives from a given mechanism is the same whether we denote the seller by u_v and the mechanism by s_v or whether we denote the seller by u and the mechanism by s.

Remark 1. In fact, we can weaken the concavity requirement considerably. As long as u and v satisfy

$$\frac{v''(v(-s))}{v'(v(-s))} < -\frac{u''(s)}{u'(s)},$$

[9] The function v is strictly increasing, and hence the inverse function, v^{-1}, always exists.

$u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$ is strictly concave in s_v (see Claim 43 in the Appendix for the proof). Evidently, we can allow for one of the agents to be risk neutral or even risk loving, i.e., to have a linear or even convex utility function for money. Consequently, we can account, to some extent, for a buyer who is "loss averse" and thus incorporate elements from prospect theory. (Kahneman and Tversky, 1979; Tversky and Kahneman, 1991, 1992).

In light of the above, we can now define our basic scenario as a "general selling problem":

Definition 2. A *general selling problem* is a triple of functions $SP = (u, v, F)$ such that:

1. The seller's utility function for money, u, satisfies Assumption A.
2. The buyer's utility function for money, v, satisfies Assumption A.
3. The probability distribution of the buyer types, F, satisfies Assumption B.
4. $u_v(s) = u\left(-v^{-1}(-s)\right)$ is a strictly concave function.

When the buyer is risk neutral, i.e., the buyer's utility for money is the identity and $V(x, k, s) = k \cdot x - s$, we can omit the buyer's utility function for money and define a "reduced selling problem":

Definition 3. A *reduced selling problem* is a pair of functions $SP = (u, F)$ such that:

1. The seller's utility function for money, u, satisfies Assumption A.
2. The probability distribution of the buyer types, F, satisfies Assumption B.
3. $u(s)$ is a strictly concave function.

Our paper aims to maximize the expected utility the seller can achieve from selling the good, given a general selling problem (u, v, F). Given the reduction shown above, we can analyze (u, v, F) as the reduced selling problem (u_v, F), where $u_v(s) = u\left(-v^{-1}(-s)\right)$.

2.2.1 Example of using the reduction

Assume that we are given a general selling problem (u, v, F) where $F = x\mathbf{1}_{x\in[0,1]}$ and $u(s) = v(s) = 1 - e^{-\alpha s}$; i.e., the buyer types are distributed uniformly over $[0, 1]$ and both the seller and the buyer have a constant absolute risk aversion (CARA) utility function with the same coefficient. It can be easily verified that $v^{-1}(y) = -\alpha^{-1} ln(1 - y)$ and thus $u_v(s_v) = u\left(-v^{-1}(-s_v)\right) = s_v/(1 + s_v)$, which is clearly a strictly increasing concave function. Therefore, we can analyze $\left(1 - e^{-\alpha s}, 1 - e^{-\alpha s}, x\mathbf{1}_{x\in[0,1]}\right)$ as the reduced selling problem $\left(s_v/(1 + s_v), x\mathbf{1}_{x\in[0,1]}\right)$.

As for mechanisms, under (u_v, F), a one-price mechanism with a price of σ will be $\tilde{\mu} = (\tilde{q}, s_v)$, where $\tilde{q}(x) = \mathbf{1}_{x\geq\sigma}$ and $s_v(x) = \sigma \cdot \mathbf{1}_{x\geq\sigma}$. In the original representation, i.e., under (u, v, F), the mechanism will be $\mu = (q, s)$, where $q(x) = \mathbf{1}_{x\geq\sigma}$ and $s(x) = -v^{-1}(-\sigma) \cdot \mathbf{1}_{x\geq\sigma} = \alpha^{-1} ln(1 + \sigma)$. Of course, by construction, both mechanisms are identical from the seller's point of view.

3 The Reduced Problem with a Risk-Neutral Buyer

As explained above, we will now analyze the special case of a risk-neutral buyer as a first step toward maximizing the expected utility of a seller who faces a risk-averse buyer. Thus, in Sections 3 through 6 we will focus on analyzing the selling problem (u, F).

Admittedly, this reduction is done mainly to simplify the analysis. Nonetheless, we think that this special case holds some interest on its own. Take for example a small high-tech company that is looking to be sold, a small biotech company that is developing a new drug and wants to sell it to a big pharmaceutical company that has the resources to further develop it, or maybe a small company that has a unique know how, can purify water sources, and wants to sell her services to local governments. All these small companies might effectively face only a single buyer (the last example constitutes a single-dimensional environment, see Daskalakis, 2015 for more details). Moreover, these are small companies, that tend to be risk averse, facing huge entities, which tend to be risk neutral. Lastly, these small companies, by virtue of having a special and unique product, have market power and can dictate the terms of negotiations, at least to some degree.

Remark 4. When we switch from money to the utils of the buyer, the buyer's utility from having the good coincides with his willingness to pay for it. Hence, in the special case of a risk-neutral buyer, the buyer's type can be taken to represent the buyer's willingness to pay, as in the standard literature. Even though that doesn't change the analysis itself, it seems more plausible that the seller may have some idea of the buyer's willingness to pay for the good rather than how much utility the buyer can derive from having the good.

3.1 Mechanisms

A mechanism is a selling procedure that the seller can choose in order to sell her good. Our goal in this paper is to characterize the mechanism that maximizes the seller's expected utility (henceforth "the optimal mechanism"). As we assume that the buyer is risk neutral, we can use the celebrated "Revelation Principle" (Gibbard, 1973; Green and Laffont, 1977; Dasgupta et al., 1979; Myerson, 1981; see also the book of Krishna, 2002) and consider only "direct" mechanisms that are incentive compatible (IC) and individually rational (IR). A direct mechanism μ is a pair of functions $\mu = (q, s) : [a, c] \to [0, 1] \times \mathbb{R}$. The *allocation function*, q, determines the buyer's[10] chances of winning the good, $q(x)$, and the *payment function, s,* determines how much he has to pay, $s(x)$. Thus, if the buyer's valuation of the good is x, he will pay $s(x)$ for a lottery in which he can win the good with probability $q(x)$. The mechanism μ satisfies IC if the buyer is (weakly) better off reporting his true value than any other value, i.e., $x \cdot q(x) - s(x) \geq x \cdot q(y) - s(y)$ for any $x, y \in [a, c]$. In our setup, s must be a non-decreasing function in order to satisfy IC. The mechanism μ satisfies IR if the buyer's expected utility from the mechanism is never negative, i.e., $x \cdot q(x) - s(x) \geq 0$. In our setup, when IC is satisfied, IR translates to the requirement that $a \cdot q(a) - s(a) \geq 0$.

Note that in the optimal mechanism it must be that the utility of the smallest buyer type is zero. Furthermore, since the buyer has a quasi-linear utility function, we can use another result by Myerson (1981) by which the optimal payment function s defines, and can be defined by, the optimal allocation function q,

$$\begin{aligned} q(x) &= \frac{s(a)}{a} + \int_a^x \frac{s'_+(t)}{t} dt \\ s(x) &= x \cdot q(x) - \int_a^x q(t)\, dt, \end{aligned} \tag{3.1}$$

where $s'_+(t)$ is the right derivative at t. Henceforth, whenever we refer to either q or s as the mechanism itself, we assume that the other function is defined using the above equations.

For every mechanism μ, we define *a buyer's utility payoff* function to be $b(x) = x \cdot q(x) - s(x)$. As is shown in Hart and Nisan (2013, 2017), b must be a convex function that satisfies $0 \leq b'_+ \leq 1$. Moreover, any function $b : [a, c] \to [0, c]$ that is convex and has derivatives between zero and one uniquely defines an IC-IR direct mechanism[11] through $q = b'_+$, $s = x \cdot b'_+ - b$, as illustrated in Figure 3.1.

[10]The buyer here buys a lottery, rather than the good itself. This is in contrast to the risk-neutral case where the discussion can be restricted to deterministic mechanisms.

[11]We restrict ourselves to seller-favorable mechanisms, as defined in Hart and Reny (2015).

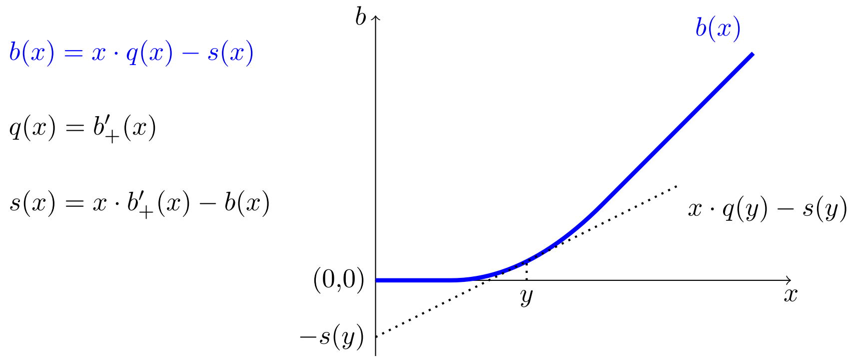

Figure 3.1: Calculating $s(x)$ and $q(x)$ given $b(x)$.

Finally, similarly to Hart and Nisan (2013, 2017), when we maximize the expected utility of the seller, we can assume, w.l.o.g., that $s(a) \geq 0$. Otherwise, it is easy to verify that $\hat{s}(x) \coloneqq s(x) - s(a)$ is an IC-IR mechanism that yields a higher expected utility than s, where $\hat{s}(x) \geq 0$.

3.1.1 Revenue-maximizing mechanisms

Myerson's and others' (Myerson, 1981; Riley and Zeckhauser, 1983) classic result proves that in our setup there is a revenue-maximizing "posted-price" mechanism; i.e., the good is offered for a given price in a take-it-or-leave-it offer. As shown in the introduction, under risk aversion, maximizing the revenue may not coincide with maximizing the utility of the seller. In particular, when the seller is risk averse, a one-price mechanism cannot be an optimal mechanism. To see this, consider the following simple example.

Example 5. Let F be the uniform distribution over $[0, 1]$, and let the seller's utility function, u, be a strictly concave utility function. For tractability reasons, we express mechanisms as their associated buyer's expected utility payoff. Thus, a one-price mechanism with a price of z (the solid green line in Figure 3.2) is

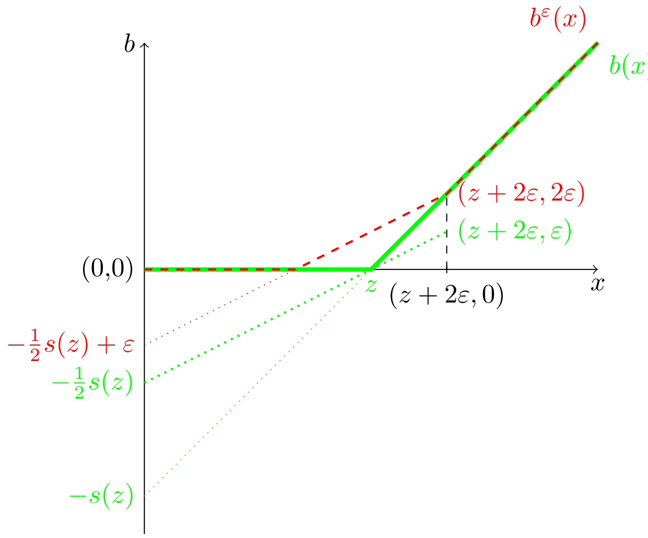

Figure 3.2: Under risk aversion, a one-price mechanism cannot, in general, be an optimal mechanism.

denoted by the buyer's utility payoff function

$$b(x) = \begin{cases} 0 & 0 \le x \le z \\ x - z & z \le x \le 1. \end{cases}$$

Rochet (1985) and Manelli and Vincent (2007) proved that $b(z) = z \cdot q(z) - s(z)$, $b'_{+}(z) = q(z) = 1$, and that at $(0, -s(z))$ the y-axis meets the line that goes through z with a slope of $b'_{+}(z)$ (the dotted green line in Figure 3.2).

Let us now analyze the mechanism b^{ε} (the dashed red line in Figure 3.2). Note that it is parallel to the line that goes through $(z, 0)$ and $(0, -s(z)/2)$ in the small interval $[z - 2\varepsilon, z + 2\varepsilon]$, and it is equal to b outside this interval. Under the original mechanism b, if the buyer type is in $(z - 2\varepsilon, z)$ then the seller is paid nothing, while if the buyer type is in $(z, z + 2z)$, the seller receives $s(z)$. In other words, the seller faces a lottery on this interval: receive no pay or $s(z)$ with equal probability. As in Figure 3.2, b^{ε} replaces this lottery with a guaranteed payment of $s(z)/2-\varepsilon$. Since u is concave, there is a small enough ε such that the seller prefers the constant payment over the lottery, even if this means losing some money in expectation.

To summarize, a "corner" in the function b means the seller faces a lottery. By "cropping the corner" we replace this lottery with a constant payment, albeit at the cost of diminishing the seller's expected revenue. Thus, for a risk-neutral seller, cropping would spell bad news. Our seller, however, is risk averse, and so cropping, if done correctly, improves her expected utility.

There is nothing unique about the corner in Example 5. Whenever $u'' < 0$, any corner in b can be cropped. Hence, when u'' is always negative, the buyer's utility payoff function from the optimal mechanism has no corners, making it continuously differentiable. Put differently, when the seller's utility function is strictly concave, as in our setup, we expect the optimal mechanism (i.e., the functions q and s) to be continuous (for a formal proof, see Section 4.3.1).

Note that when the smallest buyer type possible is itself a revenue-maximizing price, i.e., the buyer types are distributed on $[a, c]$ and the price a maximizes the expected revenue, there is no corner. In such a case, the optimal mechanism is a one-price mechanism.

3.1.2 Notations and conventions.

First, throughout the paper, derivatives of monotone functions (such as q, s, b) are right derivatives, which always exist for monotone functions.

Second, we restrict the class of mechanisms to seller-favorable mechanisms. This means that for every buyer type x, $s(x)$ will be as high as possible without violating IC. Formally, it means that q and s are right-continuous. Moreover, as shown in Hart and Reny (2015), this is done without loss of generality, since we are looking for the optimal mechanism.

Third, we denote by $\mathbb{M}$ the set of IC and IR seller-favorable mechanisms $\mu = (q, s)$ that satisfy $s(a) \geq 0$. We let $U(\mu, F) = \mathbb{E}_{F(x)}[u(s(x))]$ and we let $URev(F) = sup_{\mu \in \mathbb{M}} U(\mu, F)$ be the supremum on the expected utility the seller may achieve by using IC-IR mechanisms.

Fourth, RM price denotes revenue-maximizing price, i.e., a price set by a posted-price revenue-maximizing mechanism.

Last, we define $r^* = min\{z | z \in argmax\{z \cdot (1 - F(z))\}\}$. Thus, when the buyer is risk neutral, r^* is the minimal RM price.[12]

4 Optimal Mechanisms

In this section we characterize the unique mechanism that yields the seller the maximal expected utility among all IC and IR mechanisms, given a selling problem (u, F). Thus, an optimal mechanism is a solution to the maximization problem:

$$argmax_{(s,q) \in \mathbb{M}} \int_a^c u(s(t)) f(t) \, dt.$$

Remark 6. Note that if q is the optimal allocation function, it must be that $q(c) = 1$. Otherwise $\hat{q}(x) := q(x) / q(c)$ will be an IC-IR mechanism that yields higher expected utility to the seller.

A useful decomposition of an optimal mechanism

In our proofs we will make use of a decomposition of the optimal mechanism μ into one-price mechanisms. For this decomposition, let $\mu = (q, s)$ be an optimal mechanism. Then q may be viewed as a cumulative distribution function[13] and

[12] When the buyer is risk averse, r^* may no longer be a revenue-maximizing price. See Remark 30 for more details.

[13] Indeed, q is non-negative by definition, non-decreasing by IC, and $q(c) = 1$ by Remark 6 above.

we have

$$\begin{aligned}
q(x) &= \int_a^x dq(t) &&= \int_a^c \mathbf{1}_{x\geq t} dq(t)\\
s(x) &= \int_a^x t dq\,(t) &&= \int_a^c t\mathbf{1}_{x\geq t} dq(t)\\
b(x) &= \int_a^x (x-t)\, dq\,(t) &&= \int_a^c [x-t]_+ dq(t).
\end{aligned}$$

For more on this decomposition see Hart and Reny (2017).

We will now prove that a unique optimal mechanism always exists and characterize it.

4.1 Existence and uniqueness of the optimal mechanism

Proposition 7. *An optimal mechanism always exists for any selling problem* (u, F).

The proof is quite standard and is therefore relegated to Appendix 9.2.

Proposition 8. *The optimal IC-IR direct (seller-favorable) mechanism for a given selling problem* (u, F) *is unique.*

Due to the strict concavity of the seller's utility function, a convex combination of any two mechanisms that differ on a non-empty interval yields higher expected utility to the seller. Hence, any two optimal mechanisms must be equal almost everywhere. The seller-favorable requirement ensures that our mechanisms are right-continuous. Therefore, any two optimal mechanisms must be identical. It is a standard argument, and so the formal proof is relegated to Appendix 9.3.

4.2 Characterization of the optimal mechanism

Before we can characterize the optimal mechanism, we need the following concept of monotonicity:

Definition 9. A non-decreasing (non-increasing) function h is said to be *strictly increasing (decreasing) around* x if for any $x' < x$ and $x'' > x$ it holds that $h(x'') > h(x')$ $(h(x'') < h(x'))$.

Note that the notion of the function h being *strictly increasing around* x generalizes the notion of h being strictly increasing in a small neighborhood of x. Our notion allows h to be constant either right before or right after x, or even on both sides if h jumps at x.

We can now state our main theorem for the reduced selling problem.

Theorem 10. *Given a selling problem* (u, F)*, a unique optimal IC-IR mechanism* $\mu = (q, s)$ *always exists and it must satisfy the following conditions:*

1. *The functions* q *and* s *are continuous.*
2. *There is a constant* $\lambda \geq 0$ *s.t.* $x \int_x^c u'(s(t)) f(t)\, dt \leq \lambda$ *for every* x*, with equality when* q *is strictly increasing around* x*.*
3. $q(x) = 1$ *if and only if*[14] $x \geq r^*$*.*

Remark 11. Using the reduction from Section 2.2, this result immediately extends to general selling problems (u, v, F); see Theorem 33 in Section 7.1.

Remark 12. Property 1 uses the strict concavity of the utility function. However, the property can be somewhat generalized to weakly concave utility functions (see Section 8.1).

Remark 13. Property 2 is essentially a continuous version of the Kuhn–Tucker theorem, tailored to our maximization problem. Indeed, as discussed in Section 4.3.2 below, if we allowed our allocation function q to have only a finite number of values, our λ would become the Kuhn–Tucker multiplier.

[14]Recall that in a reduced selling problem, r^* is the minimal price set by a posted-price revenue-maximizing mechanism, which always exists.

Remark 14. Property 3 means that the buyer types that would have bought the good in an RM mechanism are exactly those that get the good with probability one under the optimal mechanism. It also means that $s(r^*)$ is the maximum of s, and hence that $URev(F) \leq u(s(r^*)) \leq u(r^*)$.

Note that, as expected, when u is linear, the optimal mechanism coincides with the revenue-maximizing mechanism, as they both become the posted-price mechanism with the price of r^*, for which Property 3 clearly holds. In addition, Property 2 holds as well, since when u is linear, then $x\int_x^c u'(s(t))\,f(t)\,dt$ becomes $x(1-F(x))$ and, of course, $x(1-F(x))$, being the revenue from setting a price of x for the good, is maximized at r^*. By contrast, Property 1, which relies heavily on the strict concavity of u, is violated.

4.3 Proof of Theorem 10

In this section we will prove Theorem 10. Throughout the section we will assume that we are given a selling problem (u, F).

4.3.1 Proof of Property 1: Continuity of q and s

In order to prove the continuity of q and s, we assume a discontinuity in the payment function for some buyer type y, and show that we can improve the mechanism. For the seller, any buyer type x where the payment function changes is essentially a lottery between the payments made by the different buyer types in a small neighborhood of x. However, our seller is risk averse and strictly prefers the expectation of the lottery – which is not offered to her when there is a discontinuity in the payment function. Thus, as in example 5, we can improve the mechanism by replacing the lottery with a constant payment that is close to the lottery's expectation. In other words, if q jumps at y, we replace

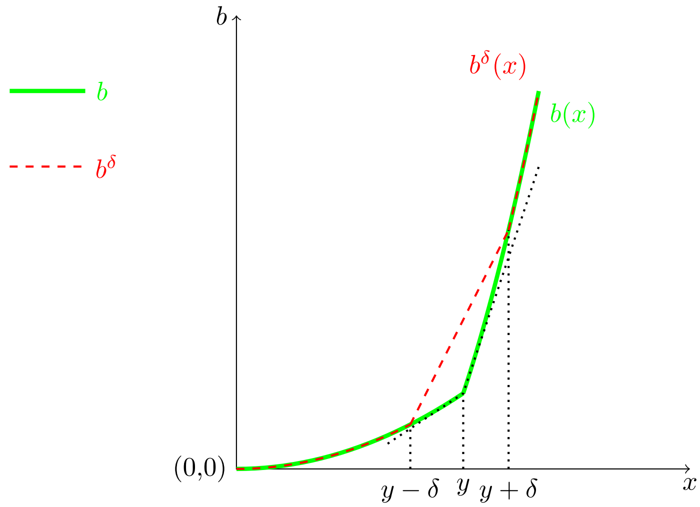

Figure 4.1: $q = b'$ jumps at y, and q^δ (which gives rise to b^δ) yields higher expected utility to the seller.

the mechanism, in a small interval around y, with a constant payment that approaches $\left(q\left(y^-\right)+q\left(y^+\right)\right)/2$ (see Figure 4.1).

Proof. Let $\mu = (q, s)$ be a mechanism, where b is the buyer's utility payoff function, and assume that q has a discontinuity in y, i.e., $q(y^-) < q(y^+)$. By eq. (3.1), $s\left(y^+\right) - s\left(y^-\right) = y\left(q\left(y^+\right) - q\left(y^-\right)\right) > 0$.

For every $\delta > 0$, let $\mu^\delta = (q^\delta, s^\delta)$ be the mechanism obtained from μ by linearly interpolating b in the interval $(y-\delta, y+\delta)$ (see Figure 4.1). Since b^δ, the buyer's utility payoff function, is convex and has derivatives between zero and one, it follows that μ^δ is an IC-IR mechanism. Since q^δ and s^δ are constant on $(y-\delta, y+\delta)$, we may denote those constants by $\overline{q}$ and $\overline{s}$, respectively. We have

$$\begin{aligned}
\bar{q} &= \frac{b\left(y+\delta\right)-b\left(y-\delta\right)}{2\delta} = \frac{b\left(y+\delta\right)-b\left(y\right)}{2\delta} + \frac{b\left(y\right)-b\left(y-\delta\right)}{2\delta} \\
&\underset{\delta\to 0^+}{\longrightarrow} \frac{q(y^+)+q(y^-)}{2}
\end{aligned}$$

and

$$\begin{aligned}\bar{s} &= \bar{q}\cdot(y+\delta)-b(y+\delta) \\ &\underset{\delta\to 0^+}{\longrightarrow} \frac{q(y^+)+q(y^-)}{2}\cdot y-b(y)=\frac{s(y^+)+s(y^-)}{2}. \end{aligned} \tag{4.1}$$

Recall that $U(\mu,F)=\mathbb{E}_{F(x)}\left[u\left(s\left(x\right)\right)\right]$. It follows that

$$\Delta^{\delta}:=U\left(\mu^{\delta},F\right)-U\left(\mu,F\right)=\int_{y-\delta}^{y+\delta}\left(u\left(\bar{s}\right)-u\left(s\left(t\right)\right)\right)f\left(t\right)dt.$$

Let us split the integral at y, and estimate the two parts separately:

$$\begin{aligned}\frac{1}{\delta}\int_{y-\delta}^{y}\left(u\left(\bar{s}\right)-u\left(s\left(t\right)\right)\right)f\left(t\right)dt &\geq \frac{1}{\delta}\int_{y-\delta}^{y}\left(u\left(\bar{s}\right)-u\left(s\left(y^-\right)\right)\right)f\left(t\right)dt \\ &\underset{\delta\to 0^+}{\longrightarrow} \left(u\left(\frac{s\left(y^+\right)+s\left(y^-\right)}{2}\right)-u\left(s\left(y^-\right)\right)\right)f\left(y\right),\end{aligned}$$

where the inequality holds because $s\left(t\right)\leq s\left(y^-\right)$ for $t<y$ and u is monotonic and for the limit we use eq. (4.1) and the continuity of $f\left(x\right)$. Similarly,

$$\begin{aligned}\frac{1}{\delta}\int_{y}^{y+\delta}\left(u\left(\bar{s}\right)-u\left(s\left(t\right)\right)\right)f\left(t\right)dt &\geq \frac{1}{\delta}\int_{y}^{y+\delta}\left(u(\bar{s})-u\left(s\left(y+\delta\right)\right)\right)f\left(t\right)dt \\ &\underset{\delta\to 0^+}{\longrightarrow} \left(u\left(\frac{s\left(y^+\right)+s\left(y^-\right)}{2}\right)-u\left(s\left(y^+\right)\right)\right)f\left(y\right),\end{aligned}$$

where the limit holds as above and for the inequality we use the monotonicity of s and u. Therefore,

$$\underset{\delta\to 0^+}{liminf}\frac{1}{\delta}\Delta^{\delta}\geq\left(2u\left(\frac{s\left(y^+\right)+s\left(y^-\right)}{2}\right)-u\left(s\left(y^+\right)\right)-u\left(s\left(y^-\right)\right)\right)f\left(y\right)>0,$$

which is strictly positive by the strict concavity of u, and so $\Delta^{\delta}>0$ for $\delta>0$ small enough. Hence, μ cannot be an optimal mechanism. Naturally, this means that the optimal allocation function, q, must be continuous, and in turn the optimal pricing function, s, must also be continuous. $\square$

Remark 15. Note that for our proof to hold, it is enough that $u''(\hat{s}) < 0$ for some $\hat{s} \in (s(y^-), s(y^+))$. This implies that $2u((s(y^+) + s(y^-)/2)) > u(s(y^+))+u(s(y^-))$, which in turn implies that $\liminf_{\delta \to 0^+} \Delta^\delta/\delta$ is strictly positive and that μ cannot be optimal. Therefore, even if we assume weak concavity, if $\hat{s} \in (s(a), s(c))$ and $u''(\hat{s}) < 0$, then there must be some $\hat{x} \in (a, c)$ such that $s(\hat{x}) = \hat{s}$.

4.3.2 Proof of Property 2: There is a constant λ such that $x\int_x^c u'(s(t)) f(t)\, dt \leq \lambda$ with equality whenever q is strictly increasing around x

Before we delve into the proof of this property, it would be helpful to have some insight regarding this formula $x\int_x^c u'(s(t)) f(t)\, dt$. To this end, we limit our mechanism to be a step function with a finite number of steps. Assume that $\mu = (q, s)$ is an optimal mechanism and that the allocation function q can only change at $a = x_1, x_2, ..., x_n = c$. Define[15] $s_i = s(x_i)$, $\Delta_i^s = s_i - s_{i-1}$, $\Delta_i^q = q(x_i) - q(x_{i-1})$ and $\Delta_i^F = F(x_{i+1}) - F(x_i)$ for $i = 1, 2, ..., n$. Eq. 3.1 tells us that $\Delta_i^s = x_i \Delta_i^q$ and so $s_i = \sum_{j=1}^i \Delta_j^s = \sum_{j=1}^i x_j \Delta_j^q$. Therefore, our problem is to find the $\{\Delta_i^q\}_{i=1}^n$ that maximize

$$U = \sum_{i=1}^n u(s_i)\, \Delta_i^F = \sum_{i=1}^n u\left(\sum_{j=1}^i x_j \Delta_j^q\right) \Delta_i^F,$$

s.t. $\Delta_i^q \geq 0$ and $\sum_{i=1}^n \Delta_i^q \leq 1$.

Since

$$\frac{\partial s_i}{\partial \Delta_k^q} = x_k \mathbf{1}_{i \geq k},$$

[15] We let $s_0 = q_0 = 0$.

the Kuhn–Tucker conditions here are $\sum_{i=1}^{n}\Delta_i^q=1$ and, for $k=1,2,...,n$,

$$\lambda\geq x_k\sum_{i=k}^{n}u^{'}\left(s_i\right)\Delta_i^F$$

with equality when $\Delta_k^q>0$, which is exactly our formula for this finite-step function.

We will now prove Property 2. Note that the Kuhn–Tucker theorem cannot be used in our proof, as our problem has an infinite number of constraints.

Proof. Let $\lambda=\int_a^c s\left(t\right)u'\left(s\left(t\right)\right)f\left(t\right)dt$. Then, by Claim 16 below, $x\int_x^c u'\left(s\left(t\right)\right)f\left(t\right)dt\leq\lambda$ and, by Claim 17, there is an equality whenever q is strictly increasing around x. □

Claim 16. $y\int_y^c u'\left(s\left(t\right)\right)f\left(t\right)dt\leq\int_a^c s\left(t\right)u'\left(s\left(t\right)\right)f\left(t\right)dt$ for any $y\in[a,c]$.

Proof. Let μ be an optimal mechanism with a buyer's utility payoff function b, and let $y\in[a,c]$. We define

$$b_y\left(x\right)=\left(x-y\right)_+=\begin{cases}0 & 0\leq x<y\\ x-y & y\leq x\leq c\end{cases},$$

i.e., b_y is the one-price mechanism with a price of y. We also define $b_{y,\varepsilon}=\left(1-\varepsilon\right)b+\varepsilon b_y$ (see Figure 4.2). Recall that we restrict ourselves to mechanisms in which $s\left(x\right)\geq0$ and thus it's easy to see that $b_{y,\varepsilon}$ defines an IC-IR mechanism $\mu_{y,\varepsilon}$ with the payment function

$$s_{y,\varepsilon}\left(x\right)=\begin{cases}\left(1-\varepsilon\right)s\left(x\right) & a\leq x<y\\ \left(1-\varepsilon\right)s\left(x\right)+\varepsilon y & y\leq x\leq c.\end{cases}$$

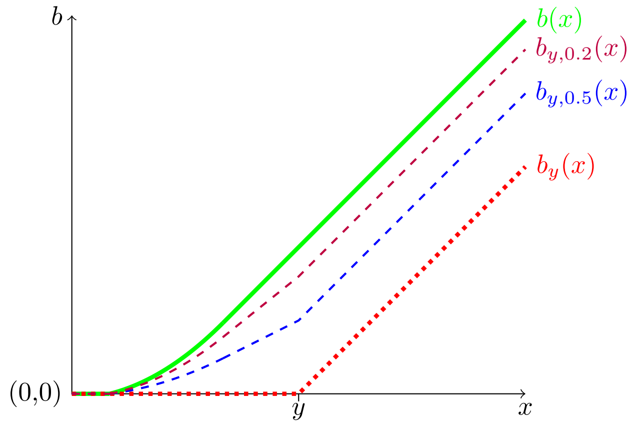

Figure 4.2: Defining $b_{y,\varepsilon}$

The optimality of μ means that its expected utility cannot be lower than that of $\mu_{y,\varepsilon}$, and hence

$$\int_a^c \left[u\left(s_{y,\varepsilon}\left(t\right)\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt \leq 0.$$

We next split the integral at y,

$$\int_a^y \left[u\left(s_{y,\varepsilon}\left(t\right)\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt + \int_y^c \left[u\left(s_{y,\varepsilon}\left(t\right)\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt \leq 0,$$

and using the definition of $s_{y,\varepsilon}$, we have

$$\int_a^y \left[u\left(\left(1-\varepsilon\right) s\left(t\right)\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt + \int_y^c \left[u\left(\left(1-\varepsilon\right) s\left(t\right) + \varepsilon y\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt \leq 0. \tag{4.2}$$

Since u is concave, we know that $u(s+h) - u(s) \geq h \cdot u'(s+h)$ for any h, and hence

$$u((1-\varepsilon)s(t)) - u(s(t)) \geq -\varepsilon s(t) \cdot u'((1-\varepsilon)s(t))$$
$$u((1-\varepsilon)s(t) + \varepsilon y) - u(s(t)) \geq \varepsilon(y - s(t)) \cdot u'((1-\varepsilon)s(t) + \varepsilon y).$$

Applying these two inequalities to eq. (4.2), we get

$$-\varepsilon \int_a^y s(t) \cdot u'((1-\varepsilon)s(t)) f(t)\,dt + \varepsilon \int_y^c (y - s(t)) \cdot u'((1-\varepsilon)s(t) + \varepsilon y) f(t)\,dt \leq 0.$$

Dividing by ε and rearranging, we have

$$y \int_y^c u'((1-\varepsilon)s(t) + \varepsilon y) f(t)\,dt \leq \int_a^y s(t) \cdot u'((1-\varepsilon)s(t)) f(t)\,dt + \int_y^c s(t) \cdot u'((1-\varepsilon)s(t) + \varepsilon y) f(t)\,dt.$$

We can then take ε to zero to get

$$y \int_y^c u'(s(t)) f(t)\,dt \leq \int_a^c s(t)\, u'(s(t)) f(t)\,dt.$$

Note that the right-hand side of the last equation is a constant. □

Claim 17. If $\mu = (q, s)$ is an optimal mechanism and q is strictly increasing around y, then $y \int_y^c u'(s(t)) f(t)\,dt = \int_a^c s(t)\, u'(s(t)) f(t)\,dt$.

To prove this claim, we assume that q is an optimal mechanism that is strictly increasing around y. Using the decomposition introduced in Section 4, we build a slightly perturbed version of the allocation function. We take a small ε, keep the perturbed allocation function constant on $(y - \varepsilon, y + \varepsilon)$, and uniformly distribute $q(y+\varepsilon) - q(y-\varepsilon)$, namely, the increase in q inside this interval, over $[a,c]$. Note that q is strictly increasing around y, and hence $q(y+\varepsilon) - q(y-\varepsilon)$. Comparing the perturbed mechanism to the optimal mechanism proves our claim.

Proof. Let $\mu = (q,s)$ be an optimal mechanism, and let y be such that q is strictly increasing around y. Then for every $\delta > 0$ we have $\rho := q(y+\delta) - q(y-\delta) > 0$. Define a new mechanism $\tilde{\mu} = (\tilde{q}, \tilde{s})$ by

$$\tilde{q}(x) := \begin{cases} \frac{1}{1-\rho} q(x) & a \leq x \leq y-\delta \\ \frac{1}{1-\rho} q(y-\delta) & y-\delta \leq x \leq y+\delta \\ \frac{1}{1-\rho}\left(q(x) - \rho\right) & y+\delta \leq x \leq c. \end{cases} \tag{4.3}$$

Using the decomposition introduced in Section 4, we get

$$\begin{aligned} \tilde{s}(x) &= \int_a^c t\mathbf{1}_{x\geq t} d\tilde{q}(t) = \frac{1}{1-\rho}\left(\int_a^c t\mathbf{1}_{x\geq t} dq(t) - \int_{y-\delta}^{y+\delta} t\mathbf{1}_{x\geq t} dq(t)\right) \\ &= \frac{1}{1-\rho} s(x) - \frac{1}{1-\rho}\int_{y-\delta}^{y+\delta} t\mathbf{1}_{x\geq t} dq(t). \end{aligned}$$

When $x \leq y-\delta$, the last integral equals 0, and when $x \geq y-\delta$ it is non-negative and at most $\int_{y-\delta}^{y+\delta} (y+\delta)\, dq(t) = (y+\delta)\rho$. Therefore,

$$\frac{1}{1-\rho} s(x) \geq \tilde{s}(x) \geq \frac{1}{1-\rho} s(x) - \mathbf{1}_{x\geq y-\delta}\frac{\rho}{1-\rho}(y+\delta),$$

or, after rearranging,

$$\frac{\rho}{1-\rho}\left(s(x) - \mathbf{1}_{x\geq y-\delta}(y+\delta)\right) \leq \tilde{s}(x) - s(x) \leq \frac{\rho}{1-\rho} s(x). \tag{4.4}$$

Remember that q, being optimal, is continuous[16] (see proof in Section 4.3.1 above). Therefore, as $\delta \to 0$ we have that $\rho \longrightarrow 0$, which implies that $\tilde{s} \to s$ by eq. 4.4.

By the concavity of u, we have

$$\begin{aligned} u\left(\tilde{s}\left(x\right)\right) - u\left(s\left(x\right)\right) &\geq u'\left(\tilde{s}\left(x\right)\right) \cdot \left(\tilde{s}\left(x\right) - s\left(x\right)\right) \\ &\geq \frac{\rho}{1-\rho} u'\left(\tilde{s}\left(x\right)\right) \cdot \left(s\left(x\right) - \mathbf{1}_{x \geq y-\delta}\left(y+\delta\right)\right). \end{aligned} \tag{4.5}$$

By the optimality of μ, it follows that $U\left(\tilde{\mu}, F\right) \leq U\left(\mu, F\right)$, i.e.,

$$0 \geq \int_a^c \left[u\left(\tilde{s}\left(t\right)\right) - u\left(s\left(t\right)\right)\right] f\left(t\right) dt.$$

By eq. (4.5), we have

$$\begin{aligned} 0 \quad \geq \quad & \frac{\rho}{1-\rho} \int_a^c u'\left(\tilde{s}\left(t\right)\right) s\left(t\right) f\left(t\right) dt \\ & -\frac{\rho}{1-\rho}\left(y+\delta\right) \int_{y-\delta}^c u'\left(\tilde{s}\left(t\right)\right) f\left(t\right) dt. \end{aligned}$$

Multiplying by $(1-\rho)/\rho$ and letting $\delta \to 0$ yields, by the bounded convergence theorem,[17]

$$0 \geq \int_a^c u'\left(s\left(t\right)\right) s\left(t\right) f\left(t\right) dt - y \int_y^c u'\left(s\left(t\right)\right) f\left(t\right) dt,$$

which, together with Claim 16, concludes the proof. □

[16] We need the continuity of q here, and hence this proof only holds when u is strictly concave.

[17] We use the requirement that u' be bounded. Hence, this result does not necessarily hold for utility functions with an unbounded derivative, such as CRRA utility functions.

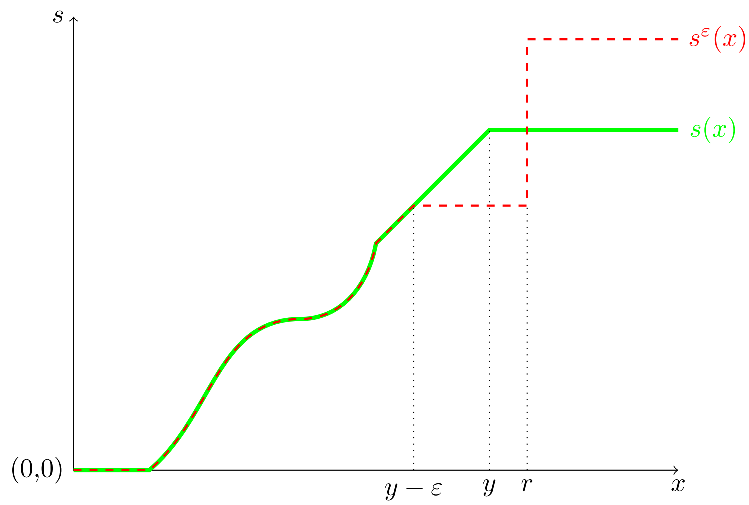

Figure 4.3: Building s^{ε} from s under the assumption that $y < r$

4.3.3 Proof of Property 3: $q(x) = 1$ if and only if $x \geq r^*$

Let $y = inf\{x|q(x) = 1\}$; then, Property 3 can be rewritten as $r^* = y$. Recall that the optimal mechanism is continuous (see Section 4.3.1), and hence $q(y) = 1$ and $y = min\{x|q(x) = 1\}$. The proof proceeds in two steps. We first prove by way of contradiction that $r^* \leq y$ and then use Property 2 of Theorem 10 to prove that $r^* \geq y$, thus completing the proof.

Claim 18. $r^* \leq y$.

In order to prove this claim, we show how the optimal mechanism can be improved in the case where $y < r^*$. Briefly, if $y < r^*$, we can improve the optimal mechanism by lowering the price paid by buyers whose type is in $(y - \varepsilon, r^*)$, while raising the price paid by buyers whose type is in (r^*, c) (see Figure 4.3). Since r^* is the minimal RM price, we have strictly increased the expected revenue. It turns out that when this change is small enough, the increase in expected revenue also guarantees higher expected utility to the seller, despite the concavity of the seller's utility function.

Proof. Assume that μ is optimal and that $y < r^*$. Using the decomposition introduced in Section 4, we define[18]

$$q^{\varepsilon}(x) = \begin{cases} q(x) & a \leq x \leq y-\varepsilon \\ q(y-\varepsilon) & y-\varepsilon \leq x < r^* \\ 1 & r^* \leq x \leq c, \end{cases}$$

$$s^{\varepsilon}(x) = \int_a^c t 1_{x \geq t} dq^{\varepsilon}(t) = \begin{cases} s(x) & a \leq x \leq y-\varepsilon \\ s(y-\varepsilon) & y-\varepsilon \leq x \leq r^* \\ s(y-\varepsilon) + r^*(1-q(y-\varepsilon)) & r^* \leq x \leq c, \end{cases}$$

and so

$$\begin{aligned} U(s^{\varepsilon}, F) - U(s, F) = & \int_{y-\varepsilon}^{r^*} [u(s(y-\varepsilon)) - u(s(t))] f(t) dt \\ & + \int_{r^*}^{c} [u(s^{\varepsilon}(r^*)) - u(s(y))] f(t) dt. \end{aligned}$$

By the concavity and monotonicity of u, we have

$$\begin{aligned} U(s^{\varepsilon}, F) - U(s, F) \geq & \quad u'(s(y-\varepsilon)) [s(y-\varepsilon) - s(y)] (F(r^*) - F(y-\varepsilon)) \\ & + u'(s^{\varepsilon}(r^*)) [s(y-\varepsilon) + r^*(1-q(y-\varepsilon)) - s(y)] (1-F(r^*)) \\ \geq & \quad u'(s(y-\varepsilon)) [s(y-\varepsilon) - s(y)] (1-F(y-\varepsilon)) \\ & + u'(s^{\varepsilon}(r^*)) r^* (1-q(y-\varepsilon)) (1-F(r^*)). \end{aligned}$$

Recall that $s(y) - s(y-\varepsilon) = \int_{y-\varepsilon}^{y} t dq(t) \leq y(1-q(y-\varepsilon))$. Thus, we can divide the last equation by $(1-q(y-\varepsilon))$, which is greater than zero by the

[18] By definition of y, $q(y-\varepsilon) < 1 = q(y)$.

definition of y, to get

$$\frac{U(s^{\varepsilon},F)-U(s,F)}{1-q(y-\varepsilon)}\geq -u'(s(y-\varepsilon))\,y\,(1-F(y-\varepsilon))+u'(s^{\varepsilon}(r^*))\,r^*\,(1-F(r^*)).$$

As ε goes to zero, both $u'(s(y-\varepsilon))$ and $u'(s^{\varepsilon}(r^*))$ go to[19] $u'(s(y))$, and so the right-hand side of the last inequality becomes

$$u'(s(y))\left[r^*(1-F(r^*))-y(1-F(y))\right].$$

Note that $x(1-F(x))$ is the expected revenue one gets when the good is assigned a price of x. Since r^* is the minimal price that maximizes the expected revenue, it must be that $r^*(1-F(r^*))-y(1-F(y))>0$. Thus, when ε is small enough, $U(s^{\varepsilon},F)-U(s,F)>0$, in contradiction to the optimality of s, which, in turn, refutes the assumption that $y<r^*$. $\square$

Claim 19. $y\leq r^*$

Proof. First, let $G(x)=x\int_x^c u'(s(t))f(t)\,dt=x(1-F(x))h(x)$, where

$$h(x):=\int_x^c\frac{f(t)}{1-F(x)}u'(s(t))\,dt=\mathbb{E}_{f(x)}\left[u'(s(t))\,|t\geq x\right].$$

By the definition of y, q is strictly increasing around y. Thus, by Property 2 of Theorem[20] (10), G is maximal at y and, in particular, $G(y)\geq G(r^*)$. Hence,

$$y(1-F(y))\,h(y)\geq r^*(1-F(r^*))\,h(r^*).$$

Recall that r^*, being an RM price, is a maximum point of $x(1-F(x))$. Hence, $y(1-F(y))\leq r^*(1-F(r^*))$, which implies that $h(y)\geq h(r^*)$.

[19]We rely here on the continuity of s. Hence our proof of Claim (18) does not hold when u is weakly concave.

[20]Property 2 relies on the continuity of s and the boundedness of $u^{'}$. Hence, our proof that $y\leq r$ holds only when u is strictly concave and $u^{'}$ is bounded.

Finally, note that $u^{'}(s)$ is a non-increasing function and that, by the definition of y, if $x < y$ then $q(x) < q(y)$, and so $s(x) < s(y)$ and hence $u^{'}(s(x)) > u^{'}(s(y))$. Therefore, h, being the expectation of $u^{'}(s)$ taken on $[x, c]$, is strictly decreasing on $[a, y]$. Consequently, $h(y) \geq h(r^*)$ implies that $y \leq r^*$. $\square$

5 The Regularity Condition

Theorem 10 gives us a characterization of the optimal mechanism. This characterization, however, only gives necessary conditions for a mechanism to be optimal, and there may be mechanisms that satisfy this property that are not optimal. In this chapter we show that when $x^2 f(x)$ is strictly increasing, our characterization enables us to find the optimal mechanism by optimizing one parameter only. Moreover, if, in addition to $x^2 f(x)$ being strictly increasing, the smallest buyer type has no value for the good, i.e., $a = 0$, then there is only one mechanism that satisfies Theorem 10. Furthermore, Theorem 10 also gives us tools to calculate this mechanism directly. Since an optimal mechanism always exists, this mechanism must be the optimal mechanism.

5.1 Refining the characterization by assuming that $x^2 f(x)$ is strictly increasing

Let us look at Property 2 of Theorem 10:

$$x \int_x^c u'(s(t)) f(t)\, dt \leq \lambda.$$

Note that since u' is strictly decreasing, the higher $s(t)$ is, the lower $u'(s(t))$ is. This suggests that in the optimal mechanism we want λ to be as low as possible. Property 2, however, only holds on the integral level. The following corollary will enable us to make pointwise arguments by switching from an integral equation to a differential equation.

Corollary 20. *If s is strictly increasing around y, then $y^2 u'(s(y)) f(y) = \lambda$, where λ is the constant from Property 2 of Theorem 10.*

Proof. First, let us recall that $G(x) := x \int_x^c u'(s(t)) f(t)\, dt$ and is differentiable with

$$G^{'}(x) = \frac{G(x)}{x} - xu'(s(x)) f(x). \tag{5.1}$$

If s is strictly increasing around y, then IC requires that q be strictly increasing around y. By Theorem 10, $G(y) = \lambda$ and λ is the maximum of G, which together imply that $G^{'}(y) = 0$. Plugging both $G(y) = \lambda$ and $G^{'}(y) = 0$ into eq. (5.1), we get

$$\lambda = u'(s(y)) y^2 f(y). \tag{5.2}$$

$\square$

Remark 21. Note that eq. (5.2) is the Euler–Lagrange equation. Indeed, our goal is to maximize the functional $U(s(t)) = \int_a^c u(s(t)) f(t)\, dt$. Thus, using the buyer's utility payoff function b, we can denote $L(t, b(t), b'(t)) = u(t \cdot b'(t) - b(t)) f(t)$. The Euler–Lagrange equation states that[21] $L_2' = {}^{d}\!/\!{}_{dt} L_3'$ or, in our case, $2w(t) + tw'(t) = 0$, where $w(t) = u'(s(t)) f(t)$. The solution to this differential equation is $w(t) = \lambda / t^2$, which gives us $\lambda = u'(s(t)) t^2 f(t)$. Note, however, that we maximize over convex functions only, which is probably why in our case the Euler–Lagrange equation holds only when q is strictly increasing around y.

Note that when $y^2 u'(s(y)) f(y) = \lambda$, or $u^{'}(s(y)) = \lambda / y^2 f(y)$, it follows that the lower λ is, the higher $s(y)$ is.[22] Unfortunately, this equation holds only when s is strictly increasing around y. Thus, in general, higher λ might

[21] We let L_n' denote the derivative of L with respect to its n-th argument.
[22] Since $u'(s)$ is a non-increasing function.

yield higher expected utility to the seller, by shifting the places where s changes. This is why we need to introduce our regularity condition, and for the rest of this section we will assume that $x^2 f(x)$ is an increasing function.

Next, recall that in our setup, u is a strictly concave function, and so we can define $\phi = u^{'-1}$ to be the inverse function of the derivative of u. Define $z = min\{r^*, inf\{x|s(x) > s(a)\}\}$, and by continuity $s(z) = s(a)$. We are now ready to prove that in order to find the optimal mechanism, we only need to optimize two parameters: $s(a)$ and z.

Corollary 22. *Given a selling problem (u, F), if $g(x) := x^2 f(x)$ is an increasing function, then the optimal mechanism has the following structure:*

$$s(x) = \begin{cases} s(a) & a \leq x \leq z \\ \phi\left(\frac{u'(s(a))g(z)}{g(x)}\right) & z \leq x \leq r^* \\ \phi\left(\frac{u'(s(a))g(z)}{g(r^*)}\right) & r^* \leq x \leq c, \end{cases}$$

for some $z \in [a, r^*]$ and $s(a) \in [0, a]$.

Remark. Clearly, if $x \in [a, z]$ then $s(x) = s(a)$ (by the definition of z) and $x \in [r^*, c]$ implies that $s(x) = s(r^*)$ (Property 3 of Theorem 10). Thus, we only need to prove what happens when $x \in [z, r^*]$.

Proof. We start by proving that if g is a strictly increasing function, then on $[z, r^*]$ s is strictly increasing and, by Corollary 20, $u'(s(x))g(x) = \lambda$.

First, assume that $s(\hat{x}_1) = s(\hat{x}_2)$, $\hat{x}_1 < \hat{x}_2$ and $(\hat{x}_1, \hat{x}_2) \subset [z, r^*]$ and let $x_1 = inf\{x|s(x) = s(\hat{x}_1)\}$ and $x_2 = sup\{x|s(x) = s(\hat{x}_1)\}$. From the definitions of x_1 and x_2, it follows that s is strictly increasing around x_1 and around x_2. By Corollary 20, we have that $u'(s(x_1))g(x_1) = \lambda = u'(s(x_2))g(x_2)$. Hence, as $s(x_1) = s(x_2)$ due to the continuity of s, we have that $g(x_1) = g(x_2)$. This, of course, is in contradiction to g being strictly increasing.

Thus, on $[z, r^*]$, s is a strictly increasing function, and, by Corollary 20, $u'(s(x))\,g(x) = \lambda$ or, alternatively, $s(x) = \phi(\lambda/g(x))$. Evaluating λ at z implies that $\lambda = u'(s(a))\,g(z)$, which implies that $s(x) = \phi(u'(s(a))\,g(z)/g(x))$. Thus, the optimal mechanism must have the following form:

$$s(x) = \begin{cases} s(a) & x \le z \\ \phi\left(\frac{u'(s(a))g(z)}{g(x)}\right) & z \le x \le r^* \\ \phi\left(\frac{u'(s(a))g(z)}{g(r^*)}\right) & r^* \le x. \end{cases}$$

□

Thus, if $x^2 f(x)$ is an increasing function, finding the optimal mechanism reduces to optimizing over two variables: $s(a)$ and z. We can, however, do better than that. Theorem 10 gives us another equation: $q(r^*) = 1$. Combining this equation with $s'(x) = xq'(x)$ (which holds for the optimal mechanism[23]), we have

$$1 = q(r^*) = q(a) + \int_z^{r^*} q'(t)\,dt = \frac{s(a)}{a}\mathbf{1}_{a>0} + \int_z^{r^*} \frac{s'(t)}{t}dt.$$

Next, $s(x) = \phi(u'(s(a))\,g(z)/g(x))$ on $[z, r^*]$ and hence[24]

$$s'(x) = -\phi'\left(\frac{u'(s(a))\,g(z)}{g(x)}\right)\frac{u'(s(a))\,g(z)\,g'(x)}{g(x)^2}.$$

Combining the last two equations, we get

$$1 = \frac{s(a)}{a}\mathbf{1}_{a>0} - u'(s(a))\,g(z)\int_z^{r^*} \phi'\left(\frac{u'(s(a))\,g(z)}{g(t)}\right)\frac{g'(t)}{tg(t)^2}dt,$$

[23] It can easily be derived from eq. (3.1).
[24] ϕ is a decreasing function and hence $s'(x) > 0$, as expected.

which gives us a connection between $s(a)$ and z, thus decreasing the level of freedom of our optimization problem.

We can do still better, however. It turns out that fixing one parameter determines the optimal mechanism completely, as our next theorem demonstrates.

Theorem 23. *Given a selling problem* (u, F), *where* $g(x) = x^2 f(x)$ *is an increasing function, denote* $m(\sigma) = min\left\{z|\sigma/a\mathbf{1}_{a>0} + \int_z^{r^*} \frac{s'_\sigma(t)}{t} dt \leq 1\right\}$, *where* $s_\sigma(t) = \phi(u'(\sigma) g(z)/g(t))$. *In this case, the optimal mechanism has the following structure:*

$$s(x) = \begin{cases} \sigma & a \leq x \leq m(\sigma) \\ \phi\left(\frac{u'(\sigma)g(m(\sigma))}{g(x)}\right) & m(\sigma) \leq x \leq r^* \\ \phi\left(\frac{u'(\sigma)g(m(\sigma))}{g(r^*)}\right) & r^* \leq x \leq c, \end{cases}$$

where $\sigma \in [0, a]$.

Proof. We already know that if $z = inf\{x|s(x) > s(a)\}$, and $x \in [z, r^*]$, then $s(x) = \phi(u'(\sigma) g(z)/g(x))$. We also know that the lower z is, the lower $g(z)$ is, and so the higher $s(x)$ is (since ϕ is a decreasing function). Consequently, the optimal mechanism will have the smallest possible[25] z. □

Remark 24. Similarly, we can show that if g is an increasing function, and s is the optimal mechanism, then given z, we have $s(a) = max\left\{\sigma|\sigma/a\mathbf{1}_{a>0} + \int_z^{r^*} (s'(t)/t)\, dt \leq 1\right\}$, where $s(x) = \phi(u'(\sigma) g(z)/g(x))$.

Thus, given $s(a)$, the optimal mechanism is completely determined. It can be seen that we reduced the optimization problem to optimizing one parameter only. We next show that in the special case where $a = 0$, we don't need to optimize at all.

[25] z is bounded from below by the requirement that $q(c) \leq 1$.

Corollary 25. *Given a selling problem* (u,F)*, such that* $g(x) \coloneqq x^2 f(x)$ *is an increasing function and* $a=0$*, the optimal mechanism is*

$$s(x)=\begin{cases}0 & 0\leq x\leq m\\ \phi\left(\frac{u'(0)g(m)}{g(x)}\right) & m\leq x\leq r^*\\ \phi\left(\frac{u'(0)g(m)}{g(r^*)}\right) & r^*\leq x\leq c,\end{cases}$$

where $m = min\left\{z|\int_z^{r^*}(s'(t)/t)\,dt\leq 1\right\}$ *and* $s(t)=\phi(u'(0)\,g(z)/g(t))$ *on* $[z,r^*]$.

Proof. Plug $a=0$ into Theorem 23. □

5.2 Finding the optimal mechanism: An example using a CARA seller

In this example we demonstrate the use of Corollary (25) and calculate the optimal mechanism when the seller has a constant absolute risk aversion (CARA) utility function, i.e., $u(s)=1-e^{-\alpha s}$, and the buyer types are distributed uniformly over $[0,1]$. We conclude by comparing the optimal mechanism to the RM mechanism.

We start by noting that since F is uniform over $[0,1]$, it follows that $f(x)=1$, $g(x)=x^2$ (which is an increasing function), and that the minimal RM price is $r=1/2$. We also note that $u'(s)=\alpha e^{-\alpha s}$, and hence $\phi(y)=u'^{-1}(y)=\alpha^{-1}\cdot ln(\alpha/y)$ and $\phi'(y)=-1/\alpha y$. Lastly, $u'(0)=\alpha$, and hence

$$\phi\left(\frac{u'(0)\,g(m)}{g(x)}\right)=\phi\left(\frac{\alpha m^2}{x^2}\right)=\frac{2}{\alpha}ln\left(\frac{x}{m}\right).$$

Plugging $r^* = 1/2$ and the last equation into Corollary 25, we get that the optimal mechanism is

$$s(x) = \begin{cases} 0 & x \leq m \\ \frac{2}{\alpha} ln\left(\frac{x}{m}\right) & m \leq x \leq 0.5 \\ \frac{2}{\alpha} ln\left(\frac{1}{2m}\right) & 0.5 \leq x, \end{cases}$$

where[26] $m = min\left\{z | \int_z^{0.5} \left[2/\alpha t^2\right] dt \leq 1\right\}$.

Solving $\int_m^{0.5} \left[2/\alpha t^2\right] dt = 1$ gives us

$$1 = \int_m^{0.5} \frac{2}{\alpha t^2} dt = -\frac{2}{\alpha} \frac{1}{t}\bigg|_m^{0.5} = \frac{2}{\alpha}\left(\frac{1}{m} - 2\right),$$

which implies that $m = 2/(\alpha + 4)$. Hence, the optimal mechanism is

$$s(x) = \begin{cases} 0 & 0 \leq x \leq \frac{2}{\alpha+4} \\ \frac{2}{\alpha} ln\left(\frac{\alpha+4}{2} x\right) & \frac{2}{\alpha+4} \leq x \leq \frac{1}{2} \\ \frac{2}{\alpha} ln\left(\frac{\alpha+4}{4}\right) & \frac{1}{2} \leq x \leq 1. \end{cases}$$

5.2.1 Comparative statics

Let us now compare the seller's expected utility from the optimal mechanism vs. the RM mechanism in the above example; i.e., the seller has a CARA utility function and the buyer type x is distributed uniformly over $[0, 1]$. Let us also compare the two mechanisms when the risk aversion coefficient α changes.

We first compute the seller's expected utility from the optimal mechanism: $U(s) = \int_a^c u(s(t))\, dt$.

[26]On the interval $[z, r^*]$, $s(t) = 2\alpha^{-1} ln(t/z)$ and so $s'(t) = 2/\alpha t$.

Expected utility	α=1	α=5	α=10	α=50	α=100
Optimal mechanism	0.2	0.56	0.71	0.93	0.96
RM mechanism	0.197	0.46	0.5	0.5	0.5
RM/optimal ratio	0.985	0.82	0.7	0.54	0.52

Table 5.1: Comparison of the optimal mechanism to the RM mechanism

$$u(s(x)) = 1 - e^{-\alpha s(x)} = \begin{cases} 0 & 0 \le x \le 2/(\alpha+4) \\ 1 - \frac{4}{(\alpha+4)^2 x^2} & 2/(\alpha+4) \le x \le 0.5 \\ 1 - \frac{16}{(\alpha+4)^2} & 0.5 \le x \le 1 \end{cases}$$

and

$$U(s) = \int_{2/(\alpha+4)}^{0.5} \left[1 - \frac{4}{(\alpha+4)^2 x^2}\right] dx + \frac{1}{2}\left(1 - \frac{16}{(\alpha+4)^2}\right) =$$

$$= \frac{1}{2} + \frac{8}{(\alpha+4)^2} - \frac{2}{\alpha+4} - \frac{2}{\alpha+4} + \frac{1}{2} - \frac{8}{(\alpha+4)^2} = 1 - \frac{4}{\alpha+4},$$

or $U(s) = \alpha/(\alpha+4)$. As for the RM mechanism, if the seller sells the good for a price of 0.5, the seller's expected utility is $0.5 \cdot u(1/2) = \left(1 - e^{-\alpha/2}\right)/2$.

Table (5.1) evaluates the optimal mechanism for different values of α. Note that the larger α is, the more risk averse the seller is. As expected, the optimal mechanism can yield the seller significantly more expected utility than the RM mechanism can. It also appears that the gap between the two mechanisms becomes more significant as the risk aversion of the seller increases.

6 Posted-Price Mechanisms

In the last section we calculated the optimal mechanism in a specific example. As expected, the optimal mechanism guaranteed the seller a much higher expected utility than the RM mechanism. On the other hand, the RM mechanism is much easier to calculate, especially if the regularity condition doesn't hold. In light of

Expected utility	α=1	α=5	α=10	α=50	α=100
Optimal mechanism	0.2	0.56	0.71	0.93	0.96
OPUM mechanism	0.199	0.54	0.69	0.9	0.94
OPUM/optimal ratio	0.995	0.964	0.972	0.968	0.98

Table 6.1: Comparison of the optimal mechanism to the one-price utility maximizing (OPUM) mechanism

the above, we are interested in finding a mechanism that is easy to calculate, but one that is also not too bad when compared to the optimal mechanism. Let us than look at "posted-price" mechanisms, i.e., mechanisms that have only a single price. For example, in the classic setup, when the agents are risk neutral, there is always an RM mechanism that is a posted-price mechanism. The one-price utility maximizing (OPUM) mechanism, which guarantees the seller the highest expected utility achievable by posted-price mechanisms, is another example of a posted-price mechanism.

Table 6.1 compares the utility of the OPUM mechanism to the optimal utility in the example presented in Section 5.2, i.e., when the seller has a CARA utility function, and the buyer types are distributed uniformly over $[0, 1]$. As seen in this table, in terms of the expected utility for the seller, the OPUM mechanism may sometimes be almost as good as the optimal mechanism. Furthermore, unlike the RM to optimal ratio, the OPUM to optimal ratio doesn't necessarily decrease as α grows large.

In this section we show that we can bound the OPUM to optimal ratio from below. Moreover, we show that this bound does not depend on the utility function of the seller, but solely on the probability distribution of the buyers. We cannot, however, provide a universal lower bound on the OPUM to optimal ratio. This is contrary to the multi-good scenario, where simple mechanisms can guarantee a constant fraction of the maximal revenue (see for example Hart and Reny, 2017; Hart and Nisan, 2017; Babaioff et al., 2014).

Before we proceed, let us note that in this section s is an optimal mechanism, and ρ is the price of the OPUM mechanism, i.e., $u(\rho)(1-F(\rho)) \geq u(p)(1-F(p))$ for any other price p. Recall that $URev(F)$ is the expected utility of the seller from the optimal mechanism, and define $SURev(F) = u(\rho)(1-F(\rho))$ to be the expected utility from the OPUM mechanism.

Let us also note that in this section we assume that the OPUM mechanism exists, which is not necessarily true outside our specified domain. Take for example the case where $u(s) = ln(s)$ and $x \sim U[0,1]$, which lies outside our domain. In this case, the expected utility from setting a price of p is $p \cdot ln(0) + (1-p) \cdot ln(p)$, which, of course, has no maximal point.

We can now prove the following proposition.

Proposition 26. *Given a selling problem (u, F), the OPUM mechanism attains at least $(1-F(r^*))$ of the optimal expected utility, i.e., $SURev(F)/URev(F) \geq 1-F(r^*)$.*

Proof. Let s be an optimal mechanism. As stated in Remark 14, $URev(F) \leq u(r^*)$.

The posted-price mechanism with a price of r^* yields the seller an expected utility of $u(r^*)(1-F(r^*))$, and hence

$$SURev(F) \geq u(r^*)(1-F(r^*)).$$

Combining the last two inequalities gives us $SURev(F) \geq u(r^*)(1-F(r^*)) \geq URev(F)(1-F(r^*))$ or

$$\frac{SURev(F)}{URev(F)} \geq 1-F(r^*).$$

□

Remark 27. This proof holds in a much broader domain than our setup, and so the OPUM mechanism may guarantee at least $1-F(r^*)$ of the optimal expected utility even when F and u don't meet our requirements.

When our Assumption Are violated, there may not be an optimal mechanism, and even if one exists, it may be hard to calculate. Thus, depending on our needs and on $F(r^*)$, and in light of the last remark, the OPUM mechanism may sometimes provide a good enough approximation to the optimal mechanism. Moreover, the OPUM mechanism can yield much higher expected utility than is suggested by the bound we attained on the OPUM to optimal ratio, namely, $1 - F(r^*)$. See for example Table 6.2, where this bound is equal to 1/2 while the OPUM mechanism yields around 99.5% of the optimal expected utility. Indeed, when $a = 0$, i.e., when $x \sim F[0, c]$, we can improve our bound, even if it still seems far from tight. The following two propositions give us two more bounds.

Proposition 28. *Given a selling problem* (u, F)*, when* $x \sim [0, c]$*, the* $SURev(F)/URev(F)$ *ratio is at least*

$$\frac{(1 - F(r^*))}{\left(1 - F\left(r^* \sqrt{u'(r^*)/u'(0)}\right)\right)}.$$

Proof. Let $x \sim F[0, c]$ and s be the optimal mechanism. Let

$$M(x) = \begin{cases} 0 & x < s(r^*) \\ s(r^*) & x \geq s(r^*). \end{cases}$$

Since $M(x)$ is a posted-price mechanism, it follows that

$$SURev(F) \geq \mathbb{E}[u(M(x))] = (1 - F(s(r^*)))\, u(s(r^*)).$$

Now, letting $z = inf\{x | s(x) > 0\}$, and noting that $s(x) = 0$ on $[a, z]$, we get

$$URev(F) = \int_z^c u(s(t))\, f(t)\, dt \leq (1 - F(z))\, u(s(r^*)).$$

Thus, we have

$$\frac{SURev\left(F\right)}{URev\left(F\right)} \geq \frac{1-F\left(s\left(r^{*}\right)\right)}{1-F\left(z\right)} \geq \frac{1-F\left(r^{*}\right)}{1-F\left(z\right)}. \tag{6.1}$$

Lastly, by Corollary 20,

$$\lambda = z^{2}u'\left(0\right) = r^{*2}u'\left(s\left(r^{*}\right)\right) \Rightarrow$$

$$z^{2} = r^{*2}u'\left(s\left(r^{*}\right)\right)/u'\left(0\right) \geq r^{*2}u'\left(r^{*}\right)/u'\left(0\right)$$

and $z \geq r^{*}\sqrt{u'\left(r^{*}\right)/u'\left(0\right)}$. Hence,

$$\frac{SURev\left(F\right)}{URev\left(F\right)} \geq \frac{1-F\left(r^{*}\right)}{1-F\left(r^{*}\sqrt{\frac{u'\left(r^{*}\right)}{u'\left(0\right)}}\right)}.$$

□

Our next bound is even more specific, and only holds when $x \sim U\left[0,c\right]$, i.e., $F\left(x\right) = x/c$. Also, even when it holds, it might be worse than our second bound, as is seen in Table 6.2.

Proposition 29. *Given a selling problem* $\left(u,F\right)$ *where* $x \sim U\left[0,c\right]$, *the OPUM mechanism attains at least* $\left(c-r^{*}\right)/\left(c-\left(c-r^{*}\right)u'\left(r^{*}\right)/u'\left(0\right)\right)$ *of the optimal expected utility.*

Proof. Substituting $F\left(x\right) = x/c$ into eq. (6.1), we have

$$\frac{SURev\left(F\right)}{URev\left(F\right)} \geq \frac{1-F\left(r^{*}\right)}{1-F\left(z\right)} = \frac{c-r^{*}}{c-z}.$$

Now, Theorem 10 implies that

$$z\int_{z}^{c} u'\left(s\left(t\right)\right)f\left(t\right)dt = \lambda \geq r^{*}\int_{r^{*}}^{c} u'\left(s\left(t\right)\right)f\left(t\right)dt$$

and, since $s(x) = s(r^*)$ on $[r^*, c]$, we have

$$z\int_z^{r^*} u'(s(t)) f(t)\,dt + zu'(s(r^*))(1 - F(r^*))$$

$$\geq r^* \cdot u'(s(r^*))(1 - F(r^*))$$

and, using the concavity of $u(s)$ and rearranging, it becomes

$$zu'(s(z))(F(r^*) - F(z)) \geq (r^* - z)\,u'(s(r^*))(1 - F(r^*)).$$

Next, using $F(x) = x/c$, we get

$$zu'(s(z))\frac{r^* - z}{c} \geq (r^* - z)\,u'(s(r^*))\frac{c - r^*}{c},$$

which can be rewritten as

$$z \geq \frac{u'(s(r^*))}{u'(s(z))}(c - r^*).$$

Lastly, $s(z) \geq 0$ and $s(r^*) \leq r^*$, and hence $z \geq (c - r^*)\,u'(r^*)/u'(0)$, and we have our bound, namely,

$$\frac{SURev(F)}{URev(F)} \geq \frac{c - r^*}{c - (c - r^*)\frac{u'(r^*)}{u'(0)}}.$$

□

Note that when $u'(0) \gg u'(r^*)$, the latter two bounds converge to our first bound, namely, $1 - F(r^*)$.

Table 6.2 shows the three bounds we derived for the OPUM/optimal ratio. The bounds are calculated for the example given in Section 5.2, where $x \sim U[0,1]$ and $u(s) = 1 - e^{-\alpha s}$, and hence $r^* = 0.5$ and $u'(s) = \alpha e^{-\alpha s}$.

Expected utility	α=1	α=5	α=10	α=50	α=100
OPUM/optimal ratio	0.995	0.964	0.972	0.968	0.98
First bound	0.5	0.5	0.5	0.5	0.5
Second bound	0.819	0.584	0.521	0.5	0.5
Third bound	0.718	0.521	0.502	0.5	0.5

Table 6.2: Comparing the three bounds to the actual OPT/OPUM ratio when $x \sim U[0,1]$ and $u(s) = 1 - e^{-\alpha s}$

Clearly, in that example the real ratio is much higher than the bounds suggest, and the question remains open: how well does the OPUM mechanism fare compared to the optimal mechanism, in terms of expected utility to the seller?

7 The General Problem with a Risk-Averse Buyer

In this section we use the reduction shown in Section 2.2 to generalize our previous results to the case of a risk-averse buyer. Recall that this reduction, i.e., from a selling problem with a risk-averse buyer to a selling problem with a risk-neutral buyer, is done merely by changing the payment units from money to utils of the buyer. Consequently, the optimal allocation function for the reduced problem is also the optimal allocation function for the original selling problem and our procedure is as follows.

1. Given a selling problem (u, v, F), use the reduction described in Section 2.2 to present it as (u_v, F).

2. Analyze (u_v, F) and characterize, or even find, the optimal mechanism $\tilde{\mu} = (\tilde{q}, s_v)$.

3. The optimal mechanism for (u, v, F) is $\mu = \left(\tilde{q}, -v^{-1}(-s_v)\right)$.

Remark 30. Before we continue, note that r^* appears in the optimal allocation function for (u_v, F), which, as mentioned above, remains the same for (u, v, F). Recall, however, that in our analysis the buyer's type represents his valuation for the good, rather than his willingness to pay for the good. These are equal when the buyer is risk neutral, but not when he is risk averse. Consequently, r^*, which is defined as $min\left\{z|z \in argmax\left\{z \cdot (1 - F(z))\right\}\right\}$, need not maximize the revenue under (u, v, F).

We can now prove our main result, which is a generalized version of Theorem 10 to the case of a risk-averse buyer.

7.1 The optimal mechanism when the buyer is risk averse

We start by proving that an optimal mechanism always exists.

Proposition 31. *Given a selling problem (u, v, F), define $u_v(s) = u\left(-v^{-1}(-s)\right)$ and let $\tilde{\mu} = (\tilde{q}, s_v)$ be the optimal mechanism for (u_v, F). Then $\mu = \left(\tilde{q}, -v^{-1}(-s_v)\right)$ is the unique optimal IC-IR direct mechanism for (u, v, F).*

Proof. It is easy to see that μ is an IC-IR direct mechanism.[27] In order to see that μ is optimal for (u, v, F), let s be any IC-IR direct mechanism for (u, v, F). Then we have

$$\begin{aligned}\int_a^c u(s(t)) f(t)\, dt &= \int_a^c u\left(-v^{-1}(v(-s(t)))\right) f(t)\, dt = \int_a^c u_v(-v(-s(t))) f(t)\, dt \\ &\leq \int_a^c u_v(s_v(t)) f(t)\, dt = \int_a^c u\left(-v^{-1}(-s_v(t))\right) f(t)\, dt,\end{aligned}$$

where the inequality holds due to the optimality of s_v under (u_v, F), and the last two equalities follow from the definition of u_v. Hence, μ is an optimal mechanism for (u, v, F). $\square$

Remark 32. Note that when the buyer is risk averse, eq. 3.1 no longer holds, and the relation between s and q changes.

[27] See Claim 35 below for a formal proof

To see that the optimal mechanism for (u, v, F) is unique, recall the uniqueness proof in Section 9.3. That proof did not make Assumption About the buyer preferences toward risk, and for this reason holds when the buyer is risk-averse. To summarize, μ is a unique, optimal, IC-IR direct mechanism that always exists for (u, v, F).

Theorem 33. *Given a selling problem (u, v, F), there always exists a unique optimal IC-IR direct mechanism $\mu = (q, s)$. This optimal mechanism satisfies the following conditions:*

1. *The functions q and s are continuous.*
2. *There is a constant $\lambda \geq 0$ s.t.*

$$x \int_x^c \frac{u'(s(t))}{v'(v(-s(t)))} f(t)\, dt \leq \lambda$$

   for every x, with equality when q is strictly increasing around x.

3. *$q(x) = 1$ if and only if $x \geq r^*$*

Proof. As per Proposition 31, let $\mu = (q, s) = \left(q, -v^{-1}(-s_v)\right)$ be the unique optimal mechanism for (u, v, F), where $(\tilde{q}, s_v)$ is the optimal mechanism for (u_v, F) and $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$. It is easy to see that the functions q and s satisfy Properties 1 and 3 of Theorem 33. As for Property 2, since s_v satisfies Property 2 of Theorem 10, there is a λ such that $x \int_x^c u_v'(s_v(t)) f(t)\, dt \leq \lambda$ with equality when $\tilde{q}$ is strictly increasing around x. Given the definition of u_v, it follows that $u_v'(s_v) = u'\left(-v^{-1}(-s_v)\right)/v'(-s_v) = u'(s)/v'(v(-s))$. We can now substitute for $u_v'(s_v)$ and get

$$x \int_x^c \frac{u'(s(t))}{v'(v(-s(t)))} f(t)\, dt \leq \lambda$$

with equality when q is strictly increasing around x. □

Remark 34. Note that if we assume that the buyer is risk neutral, i.e., that v is linear, then Theorem 33 becomes Theorem 10.

Claim 35. The mechanism $\mu = \left(\tilde{q}, -v^{-1}(-s_v)\right)$ is an IC-IR direct mechanism under (u, v, F) if and only if $\tilde{\mu} = (\tilde{q}, s_v)$ is an IC-IR mechanism under (u_v, F), where $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$.

Proof. First, μ and $\tilde{\mu}$ are both direct mechanisms. Second, under (u, v, F), a buyer with type x who faces the mechanism μ and reports y has a utility payoff of[28] $\tilde{q}(y) \cdot x + v(-s(y)) = \tilde{q}(y) \cdot x - s_v(y)$, which is exactly the utility payoff the same buyer will have after reporting y when facing the mechanism $\tilde{\mu}$ under (u_v, F). Consequently, μ is IR if and only if $\tilde{\mu}$ is IR, and μ is IC if and only if $\tilde{\mu}$ is IC. □

7.2 The regular case when the buyer is risk averse

As we saw in Section 5, when the buyer is risk neutral and $g(x) = x^2 f(x)$ is a strictly increasing function, then finding the optimal mechanism is easier, since it involves optimizing one parameter only. Moreover, when $a = 0$ and $g(x) = x^2 f(x)$ is a strictly increasing function, we don't even need to optimize: Theorem 10 provides us with the means to calculate it directly. When the buyer is risk averse and the regularity condition is met, we can combine Proposition 31 and Theorem 23 to prove the following.

Proposition 36. *Let (u, v, F) be a selling problem such that $g(x) = x^2 f(x)$ is an increasing function. We denote $u_v(s) = u\left(-v^{-1}(-s)\right)$, $\phi_{u_v} = u_v'^{-1}$, and $m(\sigma) = min\left\{z | \sigma / a \mathbf{1}_{a>0} + \int_z^{r^*} \frac{s'_\sigma(t)}{t} dt \leq 1\right\}$, where $s_\sigma(t) = \phi_{u_v}(u'(\sigma) g(z) / g(t))$. Then the optimal mechanism has the following structure:*

[28] Recall that $s_v(y) = -v(-s(y))$.

$$s(x)=\begin{cases}-v^{-1}(-\sigma) & a\le x\le m(\sigma)\\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(\sigma)g(m(\sigma))}{g(x)}\right)\right) & m(\sigma)\le x\le r^*\\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(\sigma)g(m(\sigma))}{g(r^*)}\right)\right) & r^*\le x\le c,\end{cases}$$

$$q(x)=\begin{cases}\sigma/a\mathbf{1}_{a>0} & a\le x\le m(\sigma)\\ \sigma/a\mathbf{1}_{a>0}+\int_{m(\sigma)}^{x}\frac{s_\sigma'(t)}{t}dt & m(\sigma)\le x\le r^*\\ 1 & r^*\le x\le c,\end{cases}$$

for some $\sigma\in[0,a]$.

Proof. Let $u_v(s_v)=u\left(-v^{-1}(-s_v)\right)$ and let s_v be the optimal mechanism for (u_v,F). By Proposition 31, $\mu=(q,s)=\left(\tilde{q},-v^{-1}(-s_v)\right)$ is the optimal mechanism for (u,v,F). Plugging in s_v as defined by Theorem 23, we get s as defined above. □

Corollary 37. *Let (u,v,F) be a selling problem such that $g(x)=x^2f(x)$ is an increasing function and $a=0$. Then the optimal mechanism is*

$$s(x)=\begin{cases}0 & 0\le x\le m\\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(0)g(m)}{g(x)}\right)\right) & m\le x\le r^*\\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(0)g(m)}{g(r^*)}\right)\right) & r^*\le x\le c,\end{cases}$$

$$q(x)=\begin{cases}0 & a\le x\le m\\ \int_m^x\frac{s_0'(t)}{t}dt & m\le x\le r^*\\ 1 & r^*\le x\le c,\end{cases}$$

where $u_v(s_v)=u\left(-v^{-1}(-s_v)\right)$, $\phi_{u_v}=u_v'^{-1}$, $m=min\left\{z|\int_z^{r^*}(s_0'(t)/t)\,dt\le 1\right\}$ and $s_0(t)=\phi_{u_v}(u_v'(0)\,g(z)/g(t))$ on $[z,r^*]$.

Proof. Plug $a = 0$ into Proposition 36. $\square$

7.3 Posted-price mechanisms when the buyer is risk averse

In Section 6 we proved that when the buyer is risk neutral, posted-price mechanisms can guarantee the seller at least a $(1 - F(r^*))$ fraction of the optimal utility. As mentioned in Remark 27, the proof of Proposition 26 is quite general, and it holds even when the buyer is risk averse.

7.4 Finding the optimal mechanism: An example using CARA agents

Let (u, v, F) be such that $u(s) = v(s) = 1 - e^{-\alpha s}$ and $F(x) = x\mathbf{1}_{x\in[0,1]}$; i.e., our agents have the same CARA utility functions and the buyer types are distributed uniformly over $[0, 1]$.

Since $g(x) = x^2 f(x) = x^2$ and $a = 0$, we can use Corollary 37. Hence, the optimal mechanism is

$$s(x) = \begin{cases} 0 & 0 \le x \le m \\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(0)g(m)}{g(x)}\right)\right) & m \le x \le r^* \\ -v^{-1}\left(-\phi_{u_v}\left(\frac{u_v'(0)g(m)}{g(r^*)}\right)\right) & r^* \le x \le c, \end{cases}$$

where $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$, $\phi_{u_v} = u_v^{'-1}$, $m = min\left\{z|\int_z^{r^*}(s_0'(t)/t)\,dt \le 1\right\}$, and $s_0(t) = \phi_{u_v}(u_v'(0)\,g(z)/g(t))$ on $[z, r^*]$.

It is straightforward to calculate that $-v^{-1}(-y) = \alpha^{-1}ln(1+y)$, $u_v(s_v) = s_v/(1+s_v)$ and $\phi_{u_v}(y) = u_v^{'-1}(y) = \left(1-\sqrt{y}\right)/\sqrt{y}$. Hence, $-v^{-1}(-\phi_{u_v}(y)) = -ln(y)/2\alpha$. We can also compute and substitute for $u_v'(0) = 1$, $r^* = 0.5$, and $g(x) = x^2$, and hence the optimal mechanism becomes

$$s(x)=\begin{cases}0 & 0\le x\le m\\ \frac{1}{\alpha}ln\left(\frac{x}{m}\right) & m\le x\le 0.5\\ \frac{1}{\alpha}ln\left(\frac{1}{2m}\right) & 0.5\le x\le 1,\end{cases}$$

where $m=min\left\{z|\int_z^{0.5}\left(s_0'(t)/t\right)dt\le 1\right\}$ and $s_0(t)=\frac{t-z}{z}$ on $[z,0.5]$.

When we solve $\int_m^{0.5}\left(s_0'(t)/t\right)dt=1$ for m, using $s_0'(t)=1/m$, we get that $e^{-m}=2m$ or $m\approx 0.352$. Hence, the optimal mechanism for (u,v,F) is

$$s(x)=\begin{cases}0 & 0\le x\le 0.352\\ \frac{1}{\alpha}ln\left(x\left(\alpha+2\right)\right) & 0.352\le x\le 0.5\\ \frac{1}{\alpha}ln\left(\frac{\alpha+2}{2}\right) & 0.5\le x\le 1,\end{cases}$$

$$q(x)=\begin{cases}0 & 0\le x\le 0.352\\ \frac{1}{m}ln\left(\frac{x}{m}\right) & 0.352\le x\le 0.5\\ 1 & 0.5\le x\le 1.\end{cases}$$

Remark 38. Note that m here is independent of α, which is to be expected, since $u_v(s)=s/(1+s)$ is independent of α. If, however, the buyer and the seller had different CARA coefficients, say $u(s)=1-e^{-\beta s}$ and $v(s)=1-e^{-\alpha s}$, then we would have had $u_v(s)=1-1/(1+s)^{-\beta/\alpha}$ and m would have been dependent on α and β.

8 Extending Our Setup

8.1 Weakly concave utility functions

If we are facing a selling problem (u, F) where u is only weakly concave, then the optimal mechanism may not be unique, and Theorem 10 is transformed as follows.

Proposition 39. *Given a selling problem (u, F), such that u is a weakly concave utility function, an optimal IC-IR mechanism s must satisfy the following conditions:*

1. Let $\hat{s} \in (s(a), s(c))$. If u' is strictly decreasing around $\hat{s}$, then there exists an $\hat{x} \in (a, c)$ s.t. $s(\hat{x}) = \hat{s}$.
2. There is a constant $\lambda \geq 0$ s.t. $x \int_x^c u'(s(t)) f(t)\, dt \leq \lambda$ for every[29] x.
3. If $x < r^*$ then $q(x) < 1$.

The proofs are basically the same as in the strictly concave case, and so they are omitted here. For the difference between the three properties here and in Theorem 10, see Remark 15, footnote 16, and footnote 19, respectively.

Next, we want to generalize Proposition 39 to the case of a risk-averse buyer. First, as explained in Remark 1, as long as $v''(v(-s))/v'(v(-s)) < -u''(s)/u'(s)$, our results hold since $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$ is strictly concave. If, however, $v''(v(-s))/v'(v(-s)) \geq -u''(s)/u'(s)$, Theorem 33 is no longer true, and instead we have the following proposition.

Proposition 40. *Given a selling problem (u, v, F) where $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$ is weakly concave, an optimal IC-IR mechanism s must satisfy the following conditions:*

[29] This property has no bite, as it is always true, regardless of whether s is optimal or not. Take for example $\lambda = c \cdot u'(0)$. We only state it here for the purpose of comparison with Theorem 10.

1. Let $\hat{s} \in (s(a), s(c))$. If u_v' is strictly decreasing around $-v(-\hat{s})$, then there exists an $\hat{x} \in (a,c)$ s.t. $s(\hat{x}) = \hat{s}$.

2. There is a constant $\lambda \geq 0$ s.t. $x \int_x^c u'(s(t)) f(t)\, dt \leq \lambda$ for every x.

3. If $x < r^*$ then $q(x) < 1$.

As above, Property 2 is always true and was stated for the purpose of comparison with Theorem 33. Furthermore, since the reduction from Section 2.2 doesn't affect the allocation function, Property 3 is clearly the same as in Proposition 39. Let us now prove Property 1.

Proof. Assume that $\hat{s} \in (s(a), s(b))$. Using the reduction from Section 2.2, let $s_v(x) = -v(-s(x))$ and it follows that $-v(-\hat{s}) \in (s_v(a), s_v(b))$. Hence, if u_v' is strictly decreasing around $-v(-\hat{s})$, there is an $\hat{x} \in (a,c)$ s.t. $s_v(\hat{x}) = -v(-\hat{s})$ or $-v(-s(\hat{x})) = -v(-\hat{s})$, and hence $s(\hat{x}) = \hat{s}$. □

8.2 Unbounded u' (or $u'(x) \underset{x \to 0}{\longrightarrow} \infty$)

If u' is unbounded, Theorem 10 is no longer true, and instead we have the following proposition.

Proposition 41. *Given a selling problem (u, v, F) where u' is unbounded, the optimal IC-IR mechanism $\mu = (q, s)$ satisfies the following conditions:*

1. The functions q, s are continuous.

2. There is a constant $\lambda \geq 0$ s.t. $x \int_x^c u'(s(t)) f(t)\, dt \leq \lambda$ for every[30] x.

3. If $x < r^*$ then $q(x) < 1$.

The proof of Proposition 41 follows the proof of Theorem 33, and is therefore omitted here. Note, however, that similar to the weakly concave case, the proof

[30] As mentioned, this property is always true, regardless of whether s is optimal or not, and we only state it here for the purpose of comparison with Theorem 33.

of Claim 17 is no longer valid (see footnote 17). This is why Properties 2 and 3 of Proposition 41 are the same as in Proposition 39 and differ from Properties 2 and 3 of Theorem 33.

References

Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *CoRR*, abs/1405.6146, 2014.

Brian Baisa. Auction design without quasilinear preferences. *Theoretical Economics*, 12(1):53 – 78, 2017.

Yeon-Koo Che and Ian Gale. Standard auctions with financially constrained bidders. *Review of Economic Studies*, 65(1):1–21, 1998.

Partha Dasgupta, Peter Hammond, and Eric Maskin. The implementation of social choice rules: Some general results on incentive compatibility. *Review of Economic Studies*, 46(2):185–216, 1979.

Constantinos Daskalakis. Multi-item auctions defying intuition? *SIGecom Exch.*, 14(1):41 – 75, November 2015.

Peter Eso and Gabor Futo. Auction design with a risk averse seller. *Economics Letters*, 65(1):71 – 74, 1999. ISSN 0165-1765.

Nikolova Evdokia, Yang Ger, and Pountourakis Emmanouil. Optimal mechanism design with risk-loving agents. In *14th Conference on Web and Internet Economics (WINE 2018)*, 2018.

Allan Gibbard. Manipulation of voting schemes: A general result. *Econometrica*, 41(4):587–601, 1973. ISSN 00129682, 14680262.

Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2): 427–438, 1977. ISSN 00129682, 14680262.

Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. *CoRR*, abs/1304.6116, 2013.

Sergiu Hart and Noam Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017. ISSN 0022-0531.

Sergiu Hart and Philip J. Reny. Maximal revenue with multiple goods: Non-monotonicity and other observations. *Theoretical Economics*, 10(3):893–922, 2015.

Sergiu Hart and Philip J. Reny. The better half of selling separately. *CoRR*, abs/1712.08973, 2017.

Audrey Hu, Steven A. Matthews, and Liang Zou. Risk aversion and optimal reserve prices in first- and second-price auctions. *Journal of Economic Theory*, 145(3):1188 – 1202, 2010. ISSN 0022-0531.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. ISSN 00129682, 14680262.

Tomoya Kazumura, Debasis Mishra, and Shigehiro Serizawa. Strategy-proof multi-object mechanism design: Ex-post revenue maximization with non-quasilinear preferences. *Journal of Economic Theory*, 188:105036, 2020. ISSN 0022-0531.

Paul Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999. ISSN 1467-6419.

Vijay Krishna. *Auction Theory*. Elsevier, 1 edition, 2002.

Alejandro Manelli and Daniel Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

Eric Maskin and John Riley. Optimal auctions with risk averse buyers. *Econometrica*, 52(6):1473–1518, 1984. ISSN 00129682.

Steven Matthews. Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica*, 55(3):633–646, 1987. ISSN 00129682.

R. Preston McAfee and John McMillan. Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43(1):1–19, 1987a. ISSN 0022-0531.

R. Preston McAfee and John McMillan. Auctions with entry. *Economics Letters*, 23(4):343–347, 1987b. ISSN 0165-1765.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

John Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–92, 1981.

John Riley and Richard Zeckhauser. Optimal selling strategies: When to haggle, when to hold firm. *The Quarterly Journal of Economics*, 98(2):267–289, 1983.

J. C. Rochet. The taxation principle and multi-time Hamilton–Jacobi equations. *Journal of Mathematical Economics*, 14(2):113–128, 1985. ISSN 0304-4068.

Amos Tversky and Daniel Kahneman. Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Economics*, 106(4): 1039–61, November 1991.

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, Oct 1992. ISSN 1573-0476.

9 Appendix

9.1 Proof of claims from Section 2

Claim 42. Assume that the buyer's and the seller's utility functions, $V(x,k,s) = xk + v(-s)$ and u, respectively, satisfy Assumption A, are null at zero, and are concave.[31] Then $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$ is strictly concave in s_v, satisfy Assumption A, and $u_v(0) = 0$.

Proof. Denote the buyer by $V(x,k,s) = xk + v(-s)$, as above. Let $s_v = -v(-s)$; i.e., the payment the buyer has to make "costs" him s_v utility units. Define now $u_v(s_v) = u\left(-v^{-1}(-s_v)\right) = u(s)$. Since v is strictly increasing, v^{-1} exists. Moreover, since v is a strictly concave function, v^{-1} is strictly convex, $v^{-1}(-s_v)$ is strictly convex, and $-v^{-1}(-s_v)$ is strictly concave. Finally, since u is concave, $u\left(-v^{-1}(-s_v)\right)$ is strictly concave. Similarly, we can show that u_v is also strictly increasing, twice differentiable, and has a finite derivative. We also have that $u_v(0) = u\left(-v^{-1}(0)\right) = u(0) = 0$. $\square$

Claim 43. Assume that u and v satisfy Assumption A. Then $u_v(s_v) = u\left(-v^{-1}(-s_v)\right)$ is strictly concave if and only if

$$\frac{v''(v(-s))}{v'(v(-s))} < -\frac{u''(s)}{u'(s)}.$$

[31] We assume that v is strictly concave, but we allow u to be weakly concave or even linear.

Proof. First, we take the derivative of u_v with regard to s_v,

$$u_v'(s_v) = \frac{u'\left(-v^{-1}(-s_v)\right)}{v'(-s_v)} = \frac{u'(s)}{v'(v(-s))}.$$

Since u and v are both strictly increasing functions, then $u'(s), v'(s) > 0$, and hence $u_v' > 0$. Consequently, u_v is a strictly increasing function.

Second, let's take the second derivative of u_v with regard to s_v,

$$u_v''(s_v) = \frac{u''(s)\, v'(v(-s)) + u'(s)\, v''(v(-s))}{v'(v(-s))^2}.$$

Since u_v is strictly increasing, the last equation shows us that u_v is strictly concave if and only if

$$u''(s)\, v'(v(-s)) + u'(s)\, v''(v(-s)) < 0.$$

Again, $u'(s), v'(s) > 0$, and so we can rearrange the last inequality in order to have that u_v is strictly concave if and only if

$$\frac{v''(v(-s))}{v'(v(-s))} < -\frac{u''(s)}{u'(s)}.$$

□

9.2 Proof of existence of the optimal mechanism

Proposition 44. *An optimal mechanism always exists for any selling problem (u, F).*

Proof. In order to prove that there exists an optimal mechanism, we will show that there exists a mechanism $\hat{\mu} = (\hat{q}, \hat{s}) \in \mathbb{M}$ s.t.[32] $URev(F) = \int_a^c u(\hat{s}(t))\, f(t)\, dt$.

[32]Recall that $URev(F)$ is the highest possible expected utility, given F.

We begin by noting that if $\mu = (q, s) \in \mathbb{M}$, then $s(x) \leq c$, hence $URev(F) \leq (c-a) \cdot u(c)$, and hence $URev(F)$ must be finite. Thus, if we take $\mathbb{R} = \{r | \exists \mu \in \mathbb{M} \, \text{s.t.} \, r = U(\mu, F)\}$, there is a sequence $r_i \in \mathbb{R}$ that converges to $URev(F)$. Let s_i be the corresponding payment functions, i.e., $r_i = \int_a^c u(s_i(t)) f(t) \, dt$, and let $b_i = x \cdot q_i - s_i$ be the corresponding buyer's utility payoff functions. These b_i are continuous, uniformly bounded (as $0 \leq b_i(x) \leq c$), and uniformly equicontinuous (as they are Lipschitz functions with Lipschitz constant 1), and hence there is a subsequence that converges uniformly (by the Arzelà–Ascoli theorem). We denote this limit by $\hat{b}$ and define $\hat{\mu} = (\hat{q}, \hat{s})$, where $\hat{q} = \hat{b}'$ and $\hat{s} = x\hat{b}' - \hat{b}$. It's easy to show that $\hat{b}$ is convex, non-decreasing, and has derivatives between zero and one, and that $\hat{s}(a) \geq 0$, and hence $\hat{\mu} \in \mathbb{M}$.

Next, we note that $u(s_i(x)) \leq u'(0)(s_i(x)) \leq u'(0) \cdot x$, where the first inequality is due to the concavity of u together with the monotonicity of s_i, and the second inequality is due to IR. By Lebesgue's dominated convergence theorem,[33] we have that

$$URev(F) = lim_{i \to \infty} \int_a^c u(s_i(t)) f(t) \, dt = \int_a^c u(\hat{s}(t)) f(t) \, dt.$$

Therefore, given a selling problem (u, F), an optimal mechanism always exists. □

9.3 Proof of the uniqueness of the optimal mechanism

Proposition 45. *The optimal IC-IR direct mechanism is unique.*

Proof. Assume that $\mu_1, \mu_2 \in \mathbb{M}$ are both optimal mechanisms and let $A = \{x | s_1(x) \neq s_2(x)\}$. If we define $\mu = (\mu_1 + \mu_2)/2$, then it is clearly an IC-IR mechanism, with the payment function $s = (s_1 + s_2)/2$. Consequently, if we let $I = [a, c]$ then

[33] The proof also holds when u' is unbounded, since in that case we can use $u(s_i(x)) \leq u(\varepsilon) + u'(\varepsilon) \cdot x - \varepsilon \cdot u'(\varepsilon)$ to the same effect.

$U(\mu, F) = \int_A u(s(t)) f(t)\,dt + \int_{I\setminus A} u(s(t)) f(t)\,dt$. Now, of course $\mu = \mu_1 = \mu_2$ on $I\setminus A$, and by the strict concavity of u it must be that $u(s(x)) > (u(s_1(x)) + u(s_2(x)))/2$ on A. We know, however, that μ_1 and μ_2 are optimal mechanisms, and hence A must be of zero measure, or else $U(\mu, F) > U(\mu_1, F)$. Furthermore, since $s_1 = s_2$ a.e. and they are both right-continuous (remember that w.l.o.g. we only consider seller-favorable mechanisms), it must be that $s_1 = s_2$. Thus, the optimal IC-IR direct mechanism must be unique. $\square$