The hold-up problem with flexible unobservable investments

Daniel Krähmer

June 19, 2024

Abstract

The paper studies the canonical hold-up problem with one-sided investment by the buyer and full bargaining power by the seller. The buyer can covertly choose any distribution of valuations at a cost and privately observes her valuation. I show that unlike in the well-understood case with linear costs, if investment costs are strictly convex, the buyer's equilibrium utility is strictly positive and total welfare is strictly higher than when valuations are public information, thus alleviating the severity of the hold-up problem. In fact, when costs are mean-based or display decreasing risk, the equilibrium outcome might be efficient.

Keywords: Information Design, Hold-Up Problem, Flexible Investment, Unobservable Information

JEL: C61, D42, D82

Universität Bonn (kraehmer@hcm.uni-bonn.de); I thank Francesc Dilme, Volker Nocke, Doron Ravid, Roland Strausz, Tymon Tatur, Jonas von Wangenheim for very valuable feedback. Special thanks go to Mathijs Janssen. I gratefully acknowledge financial support from the German Research Foundation (DFG) through Germany's Excellence Strategy EXC 2126/1–390838866 and CRC TR 224.

1 Introduction

The hold-up problem arises in a situation where one party (the buyer) can sink a relation-specific, non-contractible investment that increases the surplus she can realize jointly with another party (the seller) before bargaining over the division of the surplus takes place. To fix ideas, consider the classical application of a vertical relationship where a manufacturer who operates in a product market relies on a critical production input supplied by a monopolistic seller. Investments by the manufacturer that increase her profit on the product market, such as investments in product or process innovations to improve product quality, therefore increase her valuation for the supplier's input. The hold-up problem leads to inefficiently low investment if the seller can extract some of the gains from trade ex post, because a rational buyer then anticipates that she cannot fully appropriate the additional surplus generated by her investment.

In this article, I shed new light on the hold-up problem by allowing the buyer's investment to be flexible, that is, the buyer can, at a cost, choose *any* probability distribution over her valuation for the seller's good. My main results relate the severity of the hold-up problem to the shape of the cost of investment. I argue that the hold-up problem becomes less severe with the degree of cost convexity. Interestingly, I identify natural cost environments where the hold-up problem is completely eliminated in the sense that it causes no inefficiency. In particular, this is the case when costs are sufficiently convex and depend only on the mean of the distribution or decrease in its riskiness.

I consider the hold-up problem when the seller has full bargaining power ex post, and the buyer's choice of an investment distribution is covert and unobservable for the seller.[1] While these assumptions stack the deck in favour of the hold-up problem to be severe, I focus on the case that the buyer's realized valuation remains her private information. As is well known, if, in contrast, the buyer's valuation is observable by the seller, the seller will extract all gains from trade, and in anticipation of this, the buyer will not invest at all. Thus, what potentially creates investment incentives in my setting with privately observed valuations is the buyer's prospect to obtain information rents ex post.

Throughout, I impose two assumptions on the cost function. First, costs are monotone in the sense that if one distribution first order stochastically dominates (FOSD) another, then it is more costly, a natural assumption in the investment context here. Second, costs are convex. This intuitively captures increasing marginal costs of investment. Consider a distribution of valuations which is a weighted average of a "high-valuation" and a "low-valuation" distribution, that is, the former dominates the latter in terms of FOSD. Convexity then implies that marginally in-

[1]As explained in more detail below, this is the key conceptual difference between my paper and Condorelli and Szentes (2020) where the buyer can commit to a distribution. The buyer's lack of commitment may be viewed as a limitation of the parties' contractual possibilities.

creasing the weight on the high-valuation distribution gets more costly the higher the weight. In the context of investments in product quality through R&D, this, for example, reflects that it is increasingly difficult to push the research frontier, due, for example, to the scarcity of new ideas.[2]

A special case is that costs are linear which arises when the buyer's choice of distribution corresponds to choosing a mixed strategy over deterministic valuations. For linear costs, Gul (2001) shows that even if the buyer's valuation is her private information, the hold-up problem remains maximally severe in the sense that the same welfare outcomes obtain as in the case when the buyer's valuation is observable, that is, the buyer's equilibrium utility is zero, and total welfare is equal to the zero investment gains from trade.[3] My first main result shows that, in contrast, when costs are strictly convex, then the buyer's equilibrium utility is strictly positive and total welfare exceeds the zero investment gains from trade. Thus, the hold-up problem becomes less severe with convex costs. Moreover, I show that both the buyer's utility and total welfare might increase as investment costs increase.[4]

My second set of results connects the severity of the hold-up problem to the "risk properties" of the cost function. Costs are said to be decreasing (resp. increasing) in risk when costs are decreasing (resp. increasing) as the distribution gets more risky in the sense of a mean preserving spread. I show that when costs are decreasing in risk (or depend only on the mean of the distribution), then the first-best outcome occurs in equilibrium if investment costs are sufficiently convex. Thus, even though the seller has full bargaining power ex post, the hold-up problem causes no inefficiency. In contrast, when costs display increasing risk, then there is always socially inefficient under-investment (unless costs are so high that zero investment is efficient).

The risk properties of the cost function capture, in a stylized way, salient features of the application in question. For example, in the context of a manufacturer's investments in quality improvements, the risk properties of the cost function intuitively capture whether it is more or less costly to generate a large quality improvement with a relatively small probability, or an, on average equal, intermediate improvement more reliably. In this sense, costs that are increasing in risk are likely to apply to manufacturers operating in mature product markets where process innovations that yield predictable quality improvements are less costly than product innovations characterized by rare, radical quality jumps. Reversely, costs that are decreasing in risk might more adequately apply to young product markets with high innovation rates. In this sense, my

[2]While it is intuitive to connect convexity to a notion of decreasing returns to scale in production, in this paper, I simply take a reduced form approach and impose assumptions on the cost function directly. Hébert (2018) provides a microfoundation for a number of convex cost functions within the class of divergences that emerge as the result of a dynamic effort provision process.

[3]I replicate this result in a slightly more general form.

[4]My analysis does not explicitly cover concave investment costs, in part for the technical reason that this leads to optimizing convex functions over convex domains. Based on the example in Section 2, I conjecture though that the case with concave costs resembles the linear case qualitatively.

results suggest that the hold-up problem might be more severe in mature as opposed to young markets.

The basic reason behind my first result lies in the nature of the buyer's commitment problem when choosing a distribution. As highlighted by Condorelli and Szentes (2020), the buyer would like to commit to a distribution so as to induce the seller to choose a low price. However, since the buyer's investment is covert in my setting, she cannot commit. Thus, if the seller were to choose the low price from the commitment outcome, incentives would arise for the buyer to secretely deviate to a different distribution. The strength of the deviation incentives is the difference between marginal benefits and marginal costs. The marginal benefits from a deviation are constant, because the benefit of choosing a distribution is simply the buyer's share of the *expected* ex post trading surplus which is linear in the distribution.

Now, if investment costs are linear, marginal costs are constant, too. Thus, the buyer's net deviation incentives are constant, that is, the same for any distribution. In fact, if the seller were to choose the low price from the commitment outcome, the buyer would want to deviate to the "highest possible" distribution, and in this sense, the buyer's commitment problem is most pronounced when costs are linear.[5] Because the buyer's deviation incentives must be zero in equilibrium, it follows that the buyer is indifferent between all distributions in equilibrium (much like in a standard mixed strategy equilibrium). In particular, her equilibrium utility is then equal to her utility from not investing which is (normalized to) zero.

If, on the other hand, costs are strictly convex, marginal costs are increasing, and thus larger deviations become more costly. In other words, convexity of costs attenuates the buyer's commitment problem, and this force moves the equilibrium outcome closer to the commitment outcome, thus increasing the buyer's utility (and total welfare[6]). In particular, if the cost function is the sum of a linear and a scaled strictly convex part, then the buyer's utility (and thus total welfare) increases if the convex part is scaled up from zero to positive. In this sense, the buyer's utility (and total welfare) increases with investment costs.

Section 2 illustrates this reasoning in a simple example where the buyer's distribution can take on two possible valuations. In the general setting, where the buyer's possible valuations form an interval, I assume that the cost function is smooth in that it admits a Gateaux derivative. This offers a remarkably tractable framework which allows me to make the above marginal utility reasoning precise in my setting with flexible investments. As a methodological point, I provide an equilibrium characterization in terms of the Gateaux derivative of the cost function.[7] As is

[5]This is the reason why, with linear costs, the commitment outcome of Condorelli and Szentes (2020) cannot be sustained when the seller does not observe the buyer's distribution, as in Gul (2001).

[6]As I show, in my setting the seller's profit is always equal to the zero investment gains from trade. Thus, welfare is simply the buyer's utility up to a constant.

[7]Technically, I use a result in Georgiadis et al. (2024) which characterizes the optima of a concave functional that

well known, the Gateaux derivative is a directional derivative that extends the notion of a partial derivative, or gradient, to functions of functions. Accordingly, the Gateaux derivative of the cost function captures the marginal cost change if the buyer moves marginally from one distribution in the direction of another distribution.

For my second set of results which connect the risk properties of the cost function to the severity of the hold-up problem, I draw on insights of Cerreia-Vioglio et al. (2017) who characterize the risk properties of the cost function in terms of the shape of its Gateaux derivative. I also scale the cost function with a parameter that captures the magnitude of both costs and marginal costs. I show that in all cases, whether costs depend only on the mean, or are decreasing or increasing in risk, total welfare and the buyer's equilibrium utility increase as the cost parameter increases from zero to a critical value, consistent with the intuition that a more convex cost function (higher marginal costs) alleviates the buyer's commitment problem.

Moreover, if costs depend only on the mean or are decreasing in risk, and the cost parameter is sufficiently large, the equilibrium outcome is efficient, that is, the hold-up problem disappears completely.[8] The intuition behind this result is easiest to see when costs are decreasing in risk. Given a pricing strategy of the seller, costs being decreasing in risk implies that mean preserving investment distributions that are more spread out are cheaper. As a result, the buyer puts all probability mass only on the smallest and largest possible valuations. In turn, the seller's best response is to choose either a high price equal to the largest possible valuation, or a low price equal to the lowest possible valuation. In fact, the seller best responds with the low price if the buyer puts relatively little probability weight on the large valuation. This is indeed optimal for the buyer if it is relatively costly to invest in large valuations, which in my parametrization occurs if the cost parameter is sufficiently large. But note when the price is equal to the lowest possible valuation, trade is always efficient ex post, and the buyer is the residual claimant of the efficient surplus, leading her to choose the efficient investment distribution. In other words, if the cost parameter is sufficiently large, equilibrium investment is efficient.

This reasoning does not apply when costs are increasing in risk. In this case, because given that the buyer prefers a minimally risky distribution, her best response, if the seller were to set the low price, would be to concentrate all mass on a single valuation. But this is inconsistent with equilibrium because, anticipating that the buyer puts mass on a single valuation, the seller would then fully hold up the buyer, and the buyer would rather not invest. Instead, equilibrium involves a mixed pricing strategy by the seller. Therefore, equilibrium trade is inefficient ex post, and the equilibrium outcome then differs from the first-best.

admits a Gateaux derivative.

[8]When costs depend only on the mean, there might be multiple equilibria. I focus on Pareto-optimal equilibria, which are also buyer-optimal because in any equilibrium, the seller's profit is equal to the zero investment gains from trade.

While my focus in this paper is on value generating investments, a hold-up problem also arises when an initially uninformed buyer can invest in information acquisition. In an extension, I show how my model can be adapted to include information acquisition. Interestingly, in my extension, the comparative statics with respect to investment and information acquisition costs are opposed to one another. While higher (marginal) investment costs attenuate the hold-up problem, higher (marginal) information acquisition costs aggravate the hold-up problem. The intuition behind this difference is that higher information acquisition costs diminish the buyer's ability to extract surplus in the form of information rents.

Related Literature

My paper is related to the abovementioned papers by Gul (2001) and Condorelli and Szentes (2020). Like Gul (2001), I consider the hold-up problem with unobservable investments, but allow for convex investment costs. Absent linear costs, the standard logic behind mixed strategy equilibria that is employed in Gul (2001) does not apply, and I use methods based on Gateaux differentiability of the cost function to derive equilibria.[9] Non-linear costs are also considered in Condorelli and Szentes (2020), who, in contrast to my paper, consider the case when the buyer can commit to an investment distribution (or, equivalently, the seller can observe the distribution but not the realized valuation). One insight of my analysis is that even if the buyer cannot commit, the commitment outcome may obtain in equilibrium, namely in those cases where my equilibrium outcome coincides with the first-best.[10]

From a technical point of view, I use a result in Georgiadis et al. (2024) which provides a first-order condition that characterizes the maximum of a concave, Gateaux-differentiable functional. Georgiadis et al. (2024) study the optimal design of a wage contract in a moral hazard problem with flexible effort choice. By contrast, I apply their first-order condition not to a design problem but to derive explicit equilibria of the hold-up game.[11]

My paper shares with Ravid et al. (2022) the feature that the seller does not observe the distribution of the buyer's valuation. The key difference is that in Ravid et al. (2022), the buyer's ex ante choice is to acquire information about, rather than invest in, her valuation. Ravid et al.

[9]Gul (2001) does not only consider the problem with one-shot ultimatum bargaining as I do here, but shows that efficiency can be restored by an appropriate dynamic bargaining protocol.

[10]There is a literature which, in settings with linear or parametric investment costs, studies the effects of partial observability of the buyer's investment. Dilme (2019), in a setting where investment increases both parties' valuation and the informed party makes a take-it or leave-it offer ex post, shows that the non-investing party gets less than when investment is observable. Hermalin and Katz (2009), allowing for stochastic returns to investment, show that the buyer's investment incentives can go up or down when the seller obtains a signal about the buyer's valuation. Lau (2008), in a framework with deterministic investment returns, shows that better information about the buyer's valuation benefits the seller but not the buyer.

[11]For other work that studies optimal contract design with flexible effort choice but specific cost functions, see Diamond (1998), Hébert (2018), or Barron et al. (2020). Caroll (2015) allows for flexible effort choices to study robust contract design.

(2022) show that in the limit when information acquisition costs are small, the buyer is worse off than when information is for free. In contrast, I obtain welfare results away from the limit because the information acquisition constraint that the distribution be Bayesian consistent with a prior is missing from my framework and allows for more explicit equilibrium characterizations. In an extension, I show how my framework can be used to speak to information acquisition in the case that the buyer's true valuations can take on only two values.[12]

The paper is organized as follows. The next section presents an example with two possible buyer valuations. Section 3 presents the general model. Sections 4 and 5 contain the key equilibrium and welfare results. Section 6 derives explicit results when costs are mean-based, decreasing or increasing in risk, respectively. Section 7 discusses a connection to information acquisition. Section 8 concludes. All formal proofs are in the appendix.

2 Example

This section presents a simple example to illustrate some of the key intuitions of the paper. There is a buyer and a seller. The buyer can have two possible valuations $v \in \{\underline{v}, \bar{v}\}$ for the seller's good where $0 < \underline{v} < \bar{v}$. Ex ante, the buyer chooses a probability $f \in [0,1]$ with which the high valuation $\bar{v}$ occurs. Ex post, after the valuation is realized, the seller makes a take-it or leave-it offer by choosing a price p. If the buyer rejects, both parties obtain zero payoff. If the buyer accepts, her payoff is valuation minus price, and the seller's payoff is the price.

Without loss, the seller chooses a price equal either to $\underline{v}$ or $\bar{v}$. Allowing for mixed strategies, let $h \in [0,1]$ be the probability with which he chooses the high price. The buyer's cost of investing f is $C(f) = \ell f + 1/2 \cdot \kappa f^2$, where $\kappa \geq 0$ and $\ell \in (0, \bar{v} - \underline{v})$. Hence, the buyer's and the seller's utility from the combination (f,h) are[13]

$$U(f,h) = f(1-h)(\bar{v}-\underline{v}) - \ell f - \frac{1}{2}\kappa f^2, \quad \Pi(f,h) = hf\bar{v} + (1-h)\underline{v}. \tag{1}$$

If $\kappa = 0$, then costs are linear. This can be interpreted in the sense that the buyer has two pure strategies: choose valuation $\underline{v}$ at cost 0, and choose valuation $\bar{v}$ at cost ℓ. The linear costs $C(f) = \ell f$ therefore correspond to the cost of the mixed strategy where valuation $\underline{v}$ (resp. $\bar{v}$) is chosen with probability $1-f$ (resp. f).

The first-best investment level f^{FB} maximizes the total surplus $(1-f)\underline{v} + f\bar{v} - C(f)$ and is

[12]For the case with buyer commitment and information acquisition, see Roesler and Szentes (2017).
[13]I assume that when indifferent, the buyer accepts.

thus given by

$$f^{FB} = \min\left\{\frac{\bar{v} - \underline{v} - \ell}{\kappa}, 1\right\}. \tag{2}$$

Next, I discuss equilibrium. I begin with the benchmark case that the buyer's valuation v becomes publicly observable before the seller sets the price. In this case, the seller will always choose $p = v$. As a consequence, the buyer's ex post surplus is zero, and she will therefore choose zero investment $f = 0$. The seller's profit is $\underline{v}$, the buyer's utility is 0, and total welfare is $\underline{v}$.

Next, I consider the case with unobservable investment where the seller observes neither the distribution f nor the valuation v.[14] Consider first the seller's best response to f: if $f < \underline{v}/\bar{v}$, the seller optimally chooses the low price ($h = 0$). If $f = \underline{v}/\bar{v}$, the seller is indifferent between the high and low price ($h \in [0, 1]$). And if $f > \underline{v}/\bar{v}$, the seller optimally chooses the high price ($h = 1$).

Consider now the case with linear costs ($\kappa = 0$). In this case, the same welfare outcomes obtain as in the case with observable valuation. To see this, note that there is no equilibrium where the seller sets a deterministic price. The reason is that if the seller were to charge the low price ($h = 0$), the buyer's utility from investment f is $f(\bar{v}-\underline{v})-\ell f$. Thus, her marginal investment benefit is $\bar{v} - \underline{v}$, and her marginal investment cost is ℓ. Since $\ell < \bar{v} - \underline{v}$, the buyer's best response would be "full" investment $f = 1$. At that point, however, charging the low price would not be a best response for the seller.

Likewise, if the seller were to charge the high price ($h = 1$), the buyer's utility from investment f is $0 - \ell f$ and her best response would be "zero" investment $f = 0$. At that point, however, charging the high price would not be a best response for the seller (recall that $\underline{v} > 0$).

Therefore, there is only a mixed strategy equilibrium where the buyer chooses $f = \underline{v}/\bar{v}$ so as to keep the seller indifferent between both prices, and the seller chooses h so as to keep the buyer indifferent between zero investment ($f = 0$) and full investment ($f = 1$). The left panel of Figure 1 illustrates the best response functions and the equilibrium outcome.[15]

Because the buyer is indifferent, her equilibrium utility is zero (her utility from zero investment). As the seller is indifferent, her equilibrium profit is $\underline{v}$ (her profit from charging the price $\underline{v}$), and total equilibrium welfare is $\underline{v}$. Therefore, payoffs are identical to the case with observable investment: While there is now positive investment in equilibrium and the buyer obtains an ex post information rent, her rent is fully dissipated by her ex ante investment expenditures.

[14]The case with observable investment and unobservable valuation corresponds to the setting in Condorelli and Szentes (2020) where the buyer can commit to a distribution of valuations.

[15]More formally, given h, the buyer maximizes $f(1-h)(\bar{v}-\underline{v}) - \ell f$. Thus, her best response is $f = 0$ if $h > (\bar{v} - \underline{v} - \ell)/(\bar{v} - \underline{v})$, and $f = 1$ if $h < (\bar{v} - \underline{v} - \ell)/(\bar{v} - \underline{v})$, and she is indifferent otherwise. Since $\ell < \bar{v} - \underline{v}$, the only intersection of the buyer's and seller's best responses is at $f = \underline{v}/\bar{v}$ and $h = (\bar{v} - \underline{v} - \ell)/(\bar{v} - \underline{v})$.

Moreover, the positive investment does not improve welfare because trade is not efficient ex post.

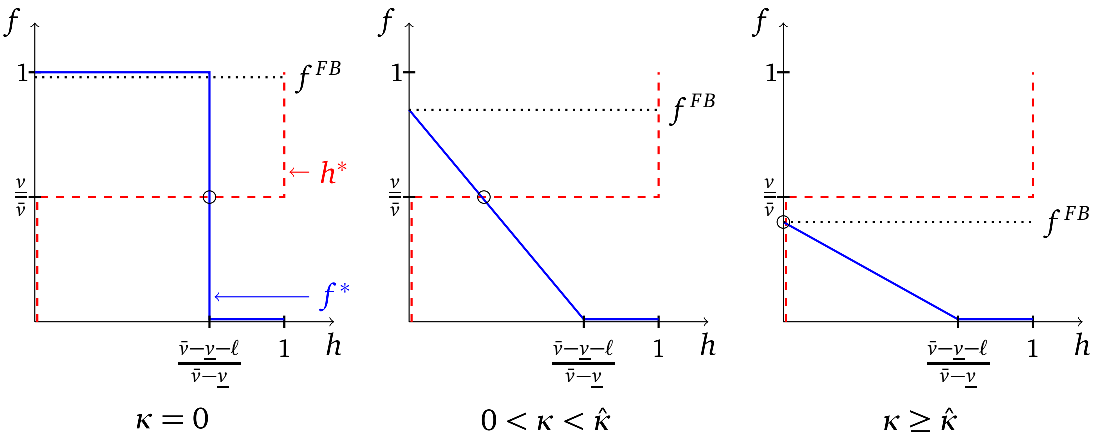

Figure 1: The figure shows the best response functions of the buyer (f^*, blue, solid) and the seller (h^*, red, dashed) and the first-best investment (f^{FB}, dotted). The buyer's best response to $h = 0$ is equal to the first-best investment because if $h = 0$, the buyer is the residual claimant of the full surplus. The left panel shows the case with linear cost ($\kappa = 0$). The center and right panels show the cases with strictly convex cost ($\kappa > 0$) when the equilibrium outcome does (right) or does not (center) coincide with the first-best.

Consider now what changes if κ is larger than zero so that investment costs are convex. If the seller charges the low price ($h = 0$), the buyer's marginal investment benefit is still $\bar{v} - \underline{v}$ but her marginal costs $\ell + \kappa f$ are now increasing in investment f. Therefore, if κ is sufficiently large, the buyer's best response to the low price is below $\underline{v}/\bar{v}$ so that charging the low price remains indeed optimal for the seller. More specifically, observe that if the seller charges the low price, trade is efficient ex post, and the buyer is the residual claimant of the efficient surplus. Thus, the buyer's best response to the low price is always equal to the first-best investment. Therefore, whenever f^{FB} is below $\underline{v}/\bar{v}$, the first-best outcome f^{FB} obtains in equilibrium, and the hold-up problem causes no inefficiency. Formally, it follows from (2) that $f^{FB} \leq \underline{v}/\bar{v}$ if and only if $\kappa \geq \hat{\kappa} = \bar{v}/\underline{v} \cdot (\bar{v} - \underline{v} - \ell)$. Thus, the first-best outcome prevails if $\kappa \geq \hat{\kappa}$. This case is illustrated in the right panel of Figure 1.[16]

What happens if $\kappa \in (0, \hat{\kappa})$? In this case, the first-best outcome does not obtain in equilibrium, because marginal costs are too low for the buyer not to deviate to an investment larger than $\underline{v}/\bar{v}$

[16]Formally, the buyer now maximizes $f(1-h)(\bar{v}-\underline{v}) - \ell f - 1/2 \cdot \kappa f^2$. By a straightforward calculation, letting $h_0 = (\bar{v}-\underline{v}-\ell-\kappa)/(\bar{v}-\underline{v})$ and $h_1 = (\bar{v}-\underline{v}-\ell)/(\bar{v}-\underline{v})$, her best response is

$$f^*(h) = \begin{cases} 1 & if \quad h < h_0 \\ \frac{\bar{v}-\underline{v}-\ell}{\kappa} - \frac{\bar{v}-\underline{v}}{\kappa} h & if \quad h \in [h_0, h_1] \\ 0 & if \quad h > h_1 \end{cases}$$

In particular, f^* is decreasing resulting in a unique equilibrium.

if the seller were to choose the low price ($h = 0$). This is illustrated in the central panel of Figure 1. Similar to the case with linear cost, in equilibrium the buyer chooses $f = \underline{v}/\bar{v}$ so as to keep the seller indifferent, and the seller chooses h so as to render the buyer's choice of $\underline{v}/\bar{v}$ optimal, that is, $(1-h)(\bar{v}-\underline{v}) = C'(\underline{v}/\bar{v})$. In contrast to the case with linear cost, however, the buyer now obtains strictly positive utility. The reason is that up to the equilibrium investment level f, her marginal costs are strictly lower than her marginal benefits due to convexity of the costs and linearity of the benefits. Formally, the buyer's equilibrium utility in the range $\kappa \in (0, \hat{\kappa})$ calculates to

$$U_B = f(1-h)(\bar{v}-\underline{v}) - C(f) = \frac{1}{2}\kappa\left(\frac{\underline{v}}{\bar{v}}\right)^2. \quad (3)$$

Notice that this is increasing in κ. The reason is that in the range $\kappa \in (0, \hat{\kappa})$, the direct effect of facing higher investment costs is outweighed by the indirect strategic effect that the seller reduces the price (reduces h) as κ increases. Moreover, as the seller is indifferent between the low and the high price, her profit is $\underline{v}$, irrespective of κ. Thus, total welfare is $U_B + \underline{v}$.

Overall, three insights emerge from these observations:

(i) The buyer's equilibrium utility and total welfare are non-monotone in costs. They increase up to $\hat{\kappa}$ and then decrease, as illustrated by Figure 2. The blue (solid) curve plots the buyer's utility as a function of κ.

(ii) Because the buyer's utility is strictly positive for $\kappa > 0$, total welfare is strictly larger than in the benchmark case when the buyer's valuation is public information.

(iii) For values $\kappa \geq \hat{\kappa}$, the equilibrium outcome is efficient, and the hold-up problem disappears, as illustrated by Figure 2. The red (dashed) curve plots first-best welfare. (Note that for $\kappa \leq \bar{v} - \underline{v} - \ell$, we have that $f^{FB} = 1$. Thus, welfare is linear in κ in this range.)

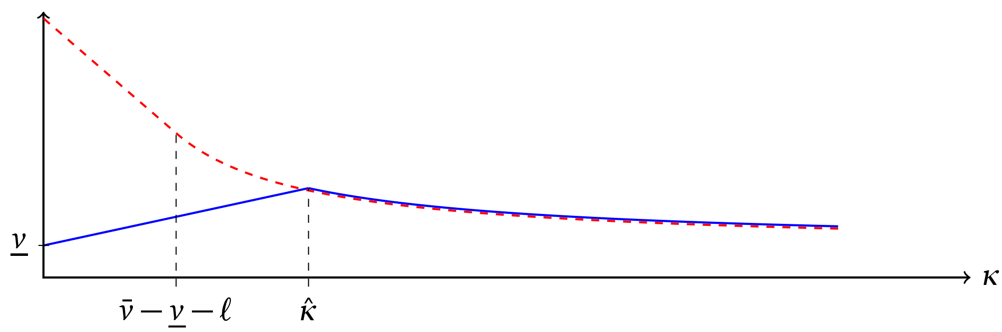

Figure 2: The figure shows total equilibrium welfare $U_B + \underline{v}$ (blue) and first-best welfare (red, dashed) as a function of κ for the values $\underline{v} = 1$, $\bar{v} = 2$, $\ell = 1/2$, $\hat{\kappa} = 1$.

I now turn to the general model where the buyer's valuation can take any value in an interval.

I will show that (i) and (ii) carry over to the general setting if investment costs are strictly convex. Even though part (iii) might appear special to the setting with two valuations (as the seller effectively only chooses among two prices), I show that it does carry over to the general setting if the cost function is "mean-based" or displays "decreasing risk".

3 Model

There is a seller who has a good, and there is a buyer who can invest in her valuation for the good by choosing a cumulative distribution function (cdf) F over the set of possible valuations $V = [\underline{v}, \bar{v}]$, $0 \leq \underline{v} < \bar{v}$, at cost $C(F)$. Let $\mathscr{F}$ denote the set of all cdf's over V. The timing is as follows: The seller chooses a price p, and the buyer simultaneously chooses a cdf $F \in \mathscr{F}$. Then the buyer privately observes her realized valuation v and decides to accept or reject to trade at the price p. If she rejects, both parties get zero. If she accepts, the buyer's payoff is $v - p$, and the seller's payoff is p.

The buyer's strategy specifies a cdf F and a decision to accept or reject, contingent on p. A (mixed) strategy for the seller is a cdf over prices. In a perfect Bayesian equilibrium (henceforth: equilibrium), the buyer accepts any price $p < v$ and rejects any price $p > v$, and her choice of cdf is optimal given the seller's pricing strategy, and the seller's strategy is optimal given the buyer's choice of cdf and acceptance/rejection decision.

It is a standard argument that in any equilibrium, the buyer accepts with probability 1 when indifferent ($p = v$) and that the seller never chooses a price strictly below the buyer's smallest possible valuation $\underline{v}$. Moreover, it is weakly dominated for the seller to choose a price strictly above the buyer's largest possible valuation $\bar{v}$. To analyze the initial stage of the game, I therefore focus on seller strategies H that are cdf's from the set $\mathscr{F}$.

The buyer's expected share of the trading surplus when valuation v has realized and before having observed the price is

$$\overline{H}(v) = \int_{\underline{v}}^{v} (v - p)\, dH(p) = \int_{\underline{v}}^{v} H(p)\, dp, \tag{4}$$

where the second equality follows from integration by parts. The ex ante expected utilities for the buyer and seller are given by[17]

$$U(F, H) = \int_V \overline{H}(v)\, dF - C(F), \qquad \Pi(H, F) = \int_V (1 - F(p^-))p\, dH(p). \tag{5}$$

[17] $F(p^-)$ denotes the left limit $\lim_{q \uparrow p} F(q)$.

With abuse of language, I refer to a combination $(F,H) \in \mathscr{F}^2$ as an equilibrium when F and H are mutual best responses given U and Π.

I next state the assumptions on the cost function that I impose throughout the paper.

A1 *$C : \mathscr{F} \to \mathbb{R}$ is convex and continuous*[18].

A2 *C is Gateaux differentiable, that is, for all $F, \tilde{F} \in \mathscr{F}$, the "Gateaux differential" of F in the direction of $\tilde{F}$, given as the limit*

$$\delta C(F;\tilde{F}-F) = \lim_{\epsilon\to 0}\frac{1}{\epsilon}[C(F+\epsilon(\tilde{F}-F))-C(F)] \tag{6}$$

exists. Moreover, there is a continuous "Gateaux derivative" $c_F : V \to \mathbb{R}$ so that

$$\delta C(F;\tilde{F}-F) = \int_V c_F(v)\, d(\tilde{F}-F). \tag{7}$$

A3 *$c_F(v)$ is strictly increasing in v for all $F \in \mathscr{F}$.*

Convexity means literally that the cost of the weighted average of two distributions is smaller than the weighted average of the cost of each distribution. I explain the economic significance of convexity below.

A2 captures that costs are smooth. Intuitively, the Gateaux differential $\delta C(F;\tilde{F}-F)$ approximates the cost change $C(\tilde{F})-C(F)$ if one moves from F in the direction of $\tilde{F}$. Expression (7) means that the Gateaux differential is linear (in $\tilde{F}-F$) and thus amounts to a linear approximation of cost changes—analogous to a first-order Taylor approximation with c_F corresponding to the gradient of C at the point F.[19,20]

A3 implies that C is monotone in the sense that if one distribution dominates another in terms of FOSD, then it is more costly. In fact, monotonicity of C is equivalent to c_F being increasing

[18]More precisely, C is assumed to be continuous in the topology of weak convergence, that is, if F_n converges weakly to F, then $C(F_n)$ converges to $C(F)$. Because the set of cdfs $\mathscr{F}$ is compact, continuity of C ensures existence of various optimizers below.

[19]To see the analogy, suppose $C : \mathbb{R} \to \mathbb{R}$ is uni-dimensional. Then by multiplying and dividing the Gateaux-differential with $\tilde{F}-F$, it can be written as

$$\delta C(F;\tilde{F}-F) = \lim_{\epsilon\to 0}\frac{1}{\epsilon(\tilde{F}-F)}[C(F+\epsilon(\tilde{F}-F))-C(F)]\cdot(\tilde{F}-F) = C'(F)\cdot(\tilde{F}-F),$$

and hence $\delta C(F;\tilde{F}-F) \approx C(\tilde{F})-C(F)$. It is well-known though that, in general, the Gateaux differential need not be linear, but only homogeneous.

[20]It is well-known that the Gateaux derivative is unique only up to a constant. The reason is that F and $\tilde{F}$ in (7) are both cdf's. Thus adding a constant to c_F does not affect the right hand side.

for all $F \in \mathscr{F}$ (see Cerreia-Vioglio et al., 2017). While A3 is therefore slightly stronger than monotonicity of C, it allows me to avoid some uninteresting case distinctions.[21]

Convexity together with monotonicity implies increasing marginal costs of investment in the following sense. Consider the cost $\Phi(\tau) = C(\tau F + (1-\tau)G)$ of the weighted average of the distributions F and G where F FOSD G. Since C is monotone and convex, so is Φ.[22] Thus, the cost of a marginal increase of the weight τ increases with the weight. For example, in the context of R&D investment, where τ may reflect the probability of a "breakthrough", this captures that it becomes increasingly difficult to increase the breakthrough probability, for example, because new ideas become increasingly scarce.

A3 also implies that C is uniquely minimized by the distribution which places full mass on the lowest valuation $\underline{v}$. I shall refer to this distribution as F_{min} and normalize its cost to zero, ensuring that investment costs are non-negative:[23]

$$F_{min} = \mathbb{1}_{[\underline{v},\bar{v}]}, \quad C(F_{min}) = 0. \tag{12}$$

F_{min} can thus be interpreted as the default distribution that arises if the buyer does "not invest".[24]

Finally, note that continuity of C implies that there is a well-defined first-best distribution that maximizes the total surplus $\int_V v\, dF - C(F)$. If unique, I denote the first-best distribution as F^{FB}.

An important special case is the class of linear cost functions. C is linear if and only if $C(F) = \int_V c(v)\, dF(v)$ for a continuous function $c : V \to \mathbb{R}$. In this case, the Gateaux derivative is $c_F = c$ for all F. Linearity arises in a setting where the buyer can choose a valuation $v \in V$ at cost $c(v)$, and V is the set of the buyer's pure strategies (see, e.g., Gul, 2001). $C(F)$ is then the cost of the mixed strategy that randomizes over V according to F.

A useful benchmark for the analysis is the case in which the buyer's valuation becomes public information before the seller chooses the price. In equilibrium, the seller then chooses the price equal to the valuation and extracts the entire trading surplus. Anticipating this, the buyer chooses

[21]A3 implies that C is strictly monotone in the sense that if F FOSD G and does not coincide with G almost everywhere, then $C(F) > C(G)$.

[22]To see convexity of Φ, let $\zeta \in (0,1)$, then

$$\begin{aligned}
\Phi(\zeta\tau + (1-\zeta)\tau') &= C(G + (\zeta\tau + (1-\zeta)\tau')(F-G)) && (8)\\
&= C(\zeta(G + \tau(F-G)) + (1-\zeta)(G + \tau'(F-G))) && (9)\\
&\leq \zeta C(G + \tau(F-G)) + (1-\zeta)C(G + \tau'(F-G)) && (10)\\
&= \zeta\Phi(\tau) + (1-\zeta)\Phi(\tau'). && (11)
\end{aligned}$$

[23]$\mathbb{1}$ denotes the indicator function.

[24]That F_{min} minimizes C follows from monotonicity and since any distribution FOSD F_{min}. Uniqueness follows from the fact that c_F is strictly increasing. I omit the details.

the default distribution F_{min}. The resulting utilities and welfare are

$$U^{PUB} = 0, \quad \Pi^{PUB} = \underline{v}, \quad W^{PUB} = \int_V v \, dF_{min} = \underline{v}. \tag{13}$$

4 Equilibrium Analysis

My first proposition is the main equilibrium characterization of the paper.

Proposition 1 *(i) There is an equilibrium.*

(ii) (F, H) is an equilibrium if and only if there are λ and $\pi \geq \underline{v}$ such that

$$\overline{H}(v) - c_F(v) - \lambda \leq 0 \quad \forall v \in V, \tag{14}$$

$$\overline{H}(v) - c_F(v) - \lambda = 0 \quad \forall v \in supp(F), \tag{15}$$

$$(1 - F(p^-))p - \pi \leq 0 \quad \forall p \in V, \tag{16}$$

$$(1 - F(p^-))p - \pi = 0 \quad \forall p \in supp(H). \tag{17}$$

Part (i) follows from a standard fixed point argument along the same lines as in the existence proof in Ravid et al. (2022, footnote 22). Part (ii), more precisely, the conditions (14) and (15) which characterize the buyer's best response in terms of the Gateaux derivative, are somewhat non-standard. To shed light on part (ii), it is easiest to first consider the conditions (16) and (17). These conditions represent the familiar conditions for a (mixed) pricing strategy by the seller to be a best response to F: Any price in the support of the strategy must yield the same profit $\pi = (1 - F(p^-))p$, and any price outside the support must not yield a higher profit.

The conditions (14) and (15) are analogous conditions for the buyer. In fact, consider the special case of linear C, and recall the interpretation of C as the cost of a mixed strategy when the buyer can choose a valuation v at cost $c(v)$. In this case, the buyer's utility from the pure strategy v is $\overline{H}(v) - c(v)$, and the conditions (14) and (15) therefore represent the conditions for a (mixed) strategy by the buyer to be a best response to H. The significance of part (ii) is that the same formal conditions characterize the buyer's best response even when the Gateaux derivative c_F is not constant in F.

Notice, however, that when c_F is not constant in F, (14) and (15) describe F only implicitly, because F appears on both sides. To see this more clearly, note that (14) and (15) imply that a point in $supp(F)$ is a maximizer of $\overline{H}(v) - c_F(v)$. Therefore, F is a solution to (14) and (15) if and only if its support satisfies

$$supp(F) \subseteq \arg\max_{v \in V} \overline{H}(v) - c_F(v). \tag{18}$$

To establish (14) and (15), I can use Proposition 1 in Georgiadis et al. (2024) which shows that F maximizes $U(G,H) = \int_V \overline{H}(v)dG - C(G)$ over G, and is thus a best response to H, if and only if F is the solution to the first-order condition[25]

$$\int_V \overline{H}(v) - c_F(v)\, dF \geq \int_V \overline{H}(v) - c_F(v)\, dG \quad \forall G \in \mathscr{F}. \tag{19}$$

With $\lambda = \int_V \overline{H}(v) - c_F(v)\, dF$, this writes

$$\int_V \overline{H}(v) - c_F(v) - \lambda\, dF = 0 \quad and \quad \int_V \overline{H}(v) - c_F(v) - \lambda\, dG \leq 0 \quad \forall G \in \mathscr{F}. \tag{20}$$

Because F and G are cdf's, this is equivalent to (14) and (15).

Finally, it is noteworthy that Proposition 1 does not use the monotonicity assumption A3.

5 Welfare Analysis

This section contains the key welfare results for general cost functions. The main result shows how the buyer's equilibrium utility and total welfare depend on the convexity of the cost function. To set the stage, I first show that in any equilibrium, the seller's profit is equal to $\underline{v}$ and thus coincides with his profit in the case when valuations are public.[26]

Proposition 2 *The seller's equilibrium profit is $\underline{v}$.*

The argument is by contradiction. If the seller's profit was strictly larger than $\underline{v}$, then since the seller is indifferent between all prices in the support of the pricing distribution, the price $\underline{v}$ is not in the support. Thus, the smallest price, say p_ℓ, in the pricing distribution is strictly larger than $\underline{v}$. This implies that p_ℓ cannot be in the support of the buyer's valuation distribution, because the buyer would benefit from redistributing probability mass from $v = p_\ell$ to $v = \underline{v}$. This follows from the buyer's best response condition and the fact that c_F is strictly increasing. But if $v = p_\ell$ is not

[25]A similar characterization of the optimum of a concave functional in terms of its Gateaux derivative appears in Luenberger (1997, Lemma 1, p. 227).

[26]Proposition 2 does rely on A3. Without further assumptions on c_F, equilibrium profits are not uniquely pinned down because there can be multiple equilibria with different profits. This can be illustrated already in the linear case. Suppose that $\underline{v} = 0$, and that c is strictly positive on the interval $(\underline{v}, \bar{v})$ and $c(\underline{v}) = c(\bar{v}) = 0$. In this case, all distributions that place mass only on $\underline{v}$ or $\bar{v}$ are costless for the buyer. There are therefore multiple equilibria: the buyer places probability f on $\bar{v}$ and $1-f$ on $\underline{v}$ for some $f \in [0,1]$, and the seller chooses $p = \bar{v}$. The seller's profit is $f\bar{v}$. This observation generalizes, and it can be shown that, in general, the seller's profit is bounded by

$$\min\left(\arg\min_{v \in V} c_F(v)\right) \leq \Pi \leq \max\left(\arg\min_{v \in V} c_F(v)\right). \tag{21}$$

in the support of the buyer's distribution, it cannot be an optimal price for the seller, because a slight price increase would make him better off.

Since the seller's profit is $\underline{v}$, total equilibrium welfare is $W = U_B + \underline{v}$ and thus pinned down by the buyer's equilibrium utility. The next proposition, characterizes the buyer's equilibrium utility.[27]

Proposition 3 *Let (F, H) be an equilibrium.*

(i) The buyer's equilibrium utility is

$$U_B = \int_V c_F(v)\, dF(v) - C(F) - c_F(\underline{v}). \tag{22}$$

(ii) If C is linear, then $U_B = 0$.

(iii) If C is strictly convex and $C(F) \neq 0$, then $U_B > 0$.

The proof of part (i) shows that λ in part (ii) of Proposition 1 is equal to $-c_F(\underline{v})$. Once this is established, the expression for U_B is immediate from plugging (15) into (5). Part (ii) is immediate from (22), the definition of linear costs, as well as the normalization $C(F_{min}) = 0$ and the fact that $C(F_{min}) = \int c(v) d\mathbb{1}_{[\underline{v},\bar{v}]} = c(\underline{v})$ in the linear case. Note that part (ii) replicates the result in Gul (2001) that when costs are linear, the welfare outcomes when the buyer's valuation is her private information coincide with those in the public information benchmark. The result here is slightly more general because Gul (2001) only allows for convex c.

The argument behind part (iii), perhaps the key result of the paper, follows from familiar marginal benefit versus marginal cost reasoning which carries over to my setting with flexible investments, and it is instructive to elaborate on it. Imagine that, instead of choosing among all cdf's, the buyer chooses a uni-dimensional investment level $\tau \in [0, 1]$ by selecting a convex combination

$$F^{\tau} = \tau F + (1 - \tau) F_{min} \tag{23}$$

of the equilibrium distribution F and the default distribution F_{min}. Given the seller's equilibrium distribution H, the buyer's benefit from investing τ is

$$\Psi(\tau) = \int_V \overline{H}(v)\, dF^{\tau} = \tau \int_V \overline{H}(v)\, d(F - F_{min}) + \int_V \overline{H}(v)\, dF_{min}, \tag{24}$$

[27] Proposition 3 also holds without A3. The only difference in (22), the term $c_F(\underline{v})$ has to be replaced by $\min_{v \in V} c_F(v)$ (which is well-defined since c_F is continuous by assumption).

and her cost is $\Phi(\tau) = C(F^\tau)$.

Because the buyer's equilibrium choice F (in the equilibrium where she chooses among all cdf's) is F^1, her equilibrium utility is

$$U_B = \Psi(1) - \Phi(1). \tag{25}$$

I now use a marginal benefit and cost argument to show that this expression is strictly positive. Indeed, because the investment benefit $\Psi(\tau)$ is linear in τ, marginal benefits $\Psi'(\tau) = \int_V \overline{H}(v)\, d(F - F_{min})$ are constant. On the other hand, note that

$$F^{\tau+\epsilon} = F^\tau + \epsilon(F - F_{\min}). \tag{26}$$

Therefore, marginal costs at τ correspond exactly to the Gateaux differential at F^τ in the direction $F - F_{min}$. Formally:

$$\Phi'(\tau) = \lim_{\epsilon \to 0} \frac{1}{\epsilon}[C(F^\tau + \epsilon(F - F_{\min})) - C(F^\tau)] = \delta C(F^\tau, F - F_{min}) = \int_V c_{F^\tau}(v)\, d(F - F_{\min}). \tag{27}$$

Now, the fact that F is a best response implies that at $\tau = 1$, marginal benefits (weakly) exceed marginal costs, because otherwise, the buyer could profitably deviate by marginally lowering τ. Formally, the equilibrium conditions (14) and (15) imply that at $\tau = 1$:

$$\begin{aligned}
\Psi'(1) - \Phi'(1) &= \int_V \overline{H}(v) - c_F(v)\, d(F - F_{min}) && (28)\\
&= \int_V \lambda\, dF - \int_V \overline{H}(v) - c_F(v)\, dF_{min} && (29)\\
&= -\int_V \overline{H}(v) - c_F(v) - \lambda\, dF_{min} && (30)\\
&\geq 0. && (31)
\end{aligned}$$

Crucially, because C is strictly convex, so is Φ,[28] and thus Φ' is strictly increasing. Therefore, because marginal benefits are constant, marginal benefits strictly exceed marginal costs for all $\tau < 1$: $\psi'(\tau) > \phi'(\tau)$.

Finally, note that the buyer can guarantee herself a weakly positive utility by investing $\tau = 0$ which corresponds to choosing the default distribution $F_{min} = F^0$.[29]

Therefore, because the buyer's marginal utility is strictly positive up to $\tau = 1$, her overall

[28] See footnote 22.

[29] In fact, this would result in utility $\Psi(0) - \Phi(0) = \int_V \overline{H}(v)\, dF_{min} - C(F_{min})$ which is weakly positive because $\overline{H}$ is positive and $C(F_{min}) = 0$ by (12).

utility is strictly positive. Formally:[30]

$$U_B = \Psi(1) - \Phi(1) = \int_0^1 \Psi'(\tau) - \Phi'(\tau)\, d\tau + \Psi(0) - \Phi(0) > 0. \tag{32}$$

The next corollary uses the results obtained so far to derive welfare changes when a strictly convex part is added to a linear cost function. Part (i) says that the buyer's equilibrium utility and total welfare may increase as costs increase. Part (ii) says that with strictly convex costs, the welfare when valuations are private is strictly larger than in the benchmark case with public valuations. In this sense, the hold-up problem is alleviated with unobservable investments.

Corollary 1 *Consider a cost function that is a combination of a linear and a strictly convex cost function, that is,* $C(F) = \int_V \ell(v)\, dF + \kappa\Gamma(F)$ *where* $\kappa \geq 0$ *and* Γ *is strictly convex and satisfies A2-A3. Suppose there is* $\tilde{\kappa} > 0$ *and an equilibrium* $(F_{\tilde{\kappa}}, H_{\tilde{\kappa}})$ *with* $C(F_{\tilde{\kappa}}) \neq 0$*. Then we have:*

(i) The buyer's equilibrium utility and total welfare is strictly larger at $\tilde{\kappa}$ *than at* $\kappa = 0$*.*

(ii) Total welfare when valuations are private information is strictly larger than when they are public information at $\tilde{\kappa}$ *but the same at* $\kappa = 0$*:* $W\mid_{\kappa=\tilde{\kappa}} > W^{PUB}\mid_{\kappa=\tilde{\kappa}}$ *and* $W\mid_{\kappa=0} = W^{PUB}\mid_{\kappa=0}$*.*

Part (i) is a direct implication of Propositions 2 and 3. Part (ii) follows from (i) and (13). In the next section, I impose more structure on the cost function that allows me to perform comparative statics over an entire interval of cost parameters κ.

6 Cost specifications

From now on, I assume that C can be written as

$$C(F) = \kappa\Gamma(F), \tag{33}$$

where $\kappa > 0$ and Γ is strictly convex and satisfies A1-A3 and has Gateaux derivative γ_F. The parameter κ scales both costs and marginal costs and will be used as a comparative statics parameter.[31]

Moreover, I shall distinguish cost functions according to their "risk properties" which characterize how costs change when the distribution gets more risky in the mean preserving spread

[30] For the special case that the cost function is of the form $C(F) = \tilde{C}(\int c(v)\, dF)$, $\tilde{C}$, c both convex, Georgiadis et al. (2024) analogously show that in a flexible moral hazard problem, the agent obtains a strictly positive rent if and only if $\tilde{C}$ is not linear.

[31] The analysis carries over without much substantial change, but requires more case distinctions, if $C(F)$ is of the form $\int_V \ell dF + \kappa\Gamma(F)$ as in Corollary 1 as long as $c_F = \ell + \kappa\gamma_F$ satisfies A3 and, respectively, A4, A5, or A6.

sense. By Cerreia-Vioglio et al. (2017), the risk properties of a cost function are closely connected to the shape of the Gateaux derivative. I distinguish between the following cases.

A4 $\Gamma(F) = \Gamma_0(M_F)$ *depends only on the mean* M_F *of* F *where* $\Gamma_0 : [\underline{v}, \bar{v}] \to \mathbb{R}$ *is strictly increasing and strictly convex and differentiable.*

A5 γ_F *is strictly concave for all* $F \in \mathscr{F}$.

A6 γ_F *is strictly convex and differentiable for all* $F \in \mathscr{F}$.

Under A4, all that matters is the mean of a distribution but not its risk. Notice that the Gateaux derivative $\gamma_F(v) = \Gamma_0'(M_F)v$ is linear, thus excluding the other two cases. Moreover, strict monotonicity and convexity of Γ_0 ensure that Γ satisfies A3 and is strictly convex. (Differentiability is only imposed to simplify the exposition.)

By Cerreia-Vioglio et al. (2017), A5 implies that Γ (and thus C) is decreasing in risk, that is, $\Gamma(F) \leq \Gamma(G)$ if F is a mean preserving spread of G. Likewise, A6 implies that Γ is increasing in risk.[32] (Differentiability of γ_F in A6 is only imposed to simplify the exposition.) As indicated in the introduction, the risk properties of the cost function capture salient features of the application in question. Note also that the risk properties and convexity are distinct properties.[33]

Throughout this section, I assume that $\underline{v} > 0$,[34] and I denote by T_f the distribution that places mass $(1-f)$ on $\underline{v}$ and mass f on $\bar{v}$, $f \in [0,1]$.

The main results of this section characterize equilibria and welfare as a function of κ. I show that in all three cases, there is a critical $\hat{\kappa}$ so that the buyer's equilibrium utility and total welfare increase in κ for all $\kappa \in (0, \hat{\kappa})$, and welfare under private information is strictly larger than when the buyer's valuation becomes public information. Moreover, under A4 and A5, equilibrium welfare coincides with first-best welfare for κ larger than $\hat{\kappa}$. In other words, the hold-up problem causes no inefficiency in this case. In contrast, under A6 the equilibrium outcome differs from the first-best (unless costs are so high that the default distribution is first-best).

I begin the analysis with noting a general property of the buyer's equilibrium distribution that follows directly from the fact that the seller's equilibrium profit is $\underline{v}$ by Proposition 2. Define for

[32] In fact, Cerreia-Vioglio et al. (2017) show that C being decreasing (increasing) in risk is *equivalent* to c_F being concave (convex) for all F. I impose "strictness" in A5 and A6 to simplify the exposition, as this rules out multiplicity of equilibria at various places.

[33] To see this, consider the class of "moment-based" cost functions $\Gamma(F) = \tilde{\Gamma}(\int c(v)\, dF)$ with Gateaux derivative $\gamma_F(v) = \tilde{\Gamma}'(\int c(v)\, dF)c(v)$. Γ is convex if $\tilde{\Gamma}$ is convex, while the risk properties are determined by the shape of c.

[34] The case $\underline{v} = 0$ is special in that the seller cannot guarantee himself a positive profit by charging a price equal to the lowest possible valuation. In this case, there are "mis-coordination" equilibria where the buyer does not invest and the seller charges a high price. The assumption $\underline{v} > 0$ rules out these uninteresting equilibria. Gul (2001) imposes a similar assumption.

$\pi, \beta \in [\underline{v}, \bar{v}]$, the distribution[35]

$$G_{\pi}^{\beta}(v) = \begin{cases} 0 & if \quad v < \pi \\ 1 - \pi/v & if \quad v \in [\pi, \beta) \\ 1 & if \quad v \geq \beta \end{cases} . \tag{34}$$

G_{π}^{β} is known as an "equal revenue distribution" because if the seller faces "demand" G_{π}^{β}, then any price $p \in [\pi, \beta]$ gives him the same revenue π. Expressed differently, a distribution F allows the seller to get profit larger than π if it is located below $G_{\pi}^{\bar{v}}$ at some point v, because the price $p = v$ yields the seller profit $v(1 - F(v^{-})) > v(1 - G_{\pi}^{\bar{v}}(v)) = \pi$. This observation readily implies that because the seller's equilibrium profit is $\underline{v}$, the buyer's equilibrium distribution must be located (weakly) above $G_{\underline{v}}^{\bar{v}}$:

Lemma 1 *Let (F, H) be an equilibrium. Then F is first order stochastically dominated by $G_{\underline{v}}^{\bar{v}}$, that is, $F(v) \geq G_{\underline{v}}^{\bar{v}}(v)$ for all $v \in V$.*

6.1 Mean-based costs

In this section, I assume that A4 holds. I begin by characterizing the first-best distribution. (Recall that $\Gamma_0' > 0$ and $\Gamma_0'' > 0$ by assumption.)

Lemma 2 *Let A4 hold, and define*

$$\kappa_0 = \frac{1}{\Gamma_0'(\bar{v})}, \quad \kappa_1 = \frac{1}{\Gamma_0'(\underline{v})}, \quad M^{FB}(\kappa) = \begin{cases} \bar{v} & if \quad \kappa \leq \kappa_0 \\ \Gamma_0'^{-1}(1/\kappa) & if \quad \kappa \in (\kappa_0, \kappa_1) \\ \underline{v} & if \quad \kappa_1 \leq \kappa \end{cases} . \tag{35}$$

Then any distribution F with mean $M_F = M^{FB}(\kappa)$ is first-best.

To understand the lemma, notice that a first-best distribution maximizes the total surplus

$$\int_V v \, dF - C(F) = M_F - \kappa \Gamma_0(M_F), \tag{36}$$

which depends only on the mean, because costs are mean-based. Because $M_F \in [\underline{v}, \bar{v}]$, the claim follows from the first-order condition for the maximizer of (36). Note also that because Γ_0 is strictly convex, the maximizer is unique whereas the first-distribution is unique only for κ outside

[35] This distribution plays a key role also in Gul (2001), Roesler and Szentes (2017), Condorelli and Szentes (2020), Ravid et al. (2022).

the interval (κ_0, κ_1). For $\kappa \in (\kappa_0, \kappa_1)$, any distribution F with $M_F = M^{FB}(\kappa)$ is a first-best distribution. Finally, observe that M^{FB} is strictly decreasing within (κ_0, κ_1) due to strict convexity of Γ_0.

Next, I turn to equilibrium. The next proposition shows that there are two types of (candidate) equilibria. Subsequently, I characterize the conditions under which each of them occurs.

Proposition 4 *Under A4, (F, H) is an equilibrium only if $F(v) \geq G_{\underline{v}}^{\bar{v}}(v)$ for all $v \in V$, and $H = T_h$ with $h < 1$, that is, the seller choose the price $p = \underline{v}$ with probability $1-h$ and the price $p = \bar{v}$ with probability h. Moreover:*

(i) (F, T_h) is an equilibrium with $h > 0$ if and only if $F(\bar{v}^-) = 1 - \underline{v}/\bar{v}$, $\kappa\Gamma_0'(M_F) < 1$, and $h = 1 - \kappa\Gamma_0'(M_F)$.

(ii) (F, T_h) is an equilibrium with $h = 0$ if and only if F is a first-best distribution.

The key insight of the proposition is that in any equilibrium the seller randomizes between (at most) the two prices $\underline{v}$ and $\bar{v}$. To see the reason, recall that given a pricing distribution H of the seller, the buyer's objective is

$$\int_V \overline{H}(v)\, dF - \kappa\Gamma_0(M_F). \tag{37}$$

Now, $\overline{H}(v) = \int_{\underline{v}}^{v} v - p\ dH(p)$ is the buyer's expected trading surplus when her valuation is v. Observe that $\overline{H}$ is convex because as the buyer's valuation increases from v to $v + dv$, her trading surplus at prices $p \leq v$ increases linearly by dv, and, in addition, trade possibly occurs more often, that is, also for prices $p \in (v, v + dv]$. Therefore, because costs depend only on the mean, the buyer always (weakly) benefits from increasing the spread of her valuation distribution in a mean-preserving way.

Now suppose that the seller were to use an interior price $p \in (\underline{v}, \bar{v})$ in her pricing strategy. Then, $\overline{H}$ is strictly convex around p, and because costs are mean-based, the buyer's objective is therefore uniquely maximized by a distribution that puts probability only on the most extreme valuations $\underline{v}$ and $\bar{v}$. But this is inconsistent with equilibrium, because clearly for a price to be optimal for the seller, it must be in the support of the buyer's distribution.

While, therefore, the seller only randomizes between $\underline{v}$ and $\bar{v}$, he cannot charge $p = \bar{v}$ with probability $h = 1$. The reason is that if $p = \bar{v}$ the buyer's utility is $0 - C(F)$ and her best response would be the default distribution. But then $p = \bar{v}$ would yield zero profit and is not a best response for the seller. Hence, $h < 1$ in equilibrium.

Part (i) of Proposition 4 describes a (candidate) equilibrium where the seller is indifferent between $\underline{v}$ and $\bar{v}$. For this to be the case, the buyer needs to put probability mass $\underline{v}/\bar{v}$ on the highest valuation $\bar{v}$, hence $1-F(\bar{v}^-) = \underline{v}/\bar{v}$. Moreover, the buyer obtains utility $(1-h)(M_F - \underline{v}) - \kappa\Gamma_0(M_F)$. The first-order condition for the (mean of the) buyer's best response thus implies $1-h = \kappa\Gamma_0'(M_F)$.

Part (ii) describes an efficient (candidate) equilibrium where the seller charges the low price $\underline{v}$ with probability 1 ($h = 0$). Ex post trade is then efficient, and the buyer is the residual claimant of the efficient surplus. Thus, choosing a first-best distribution is a best response.

When costs are mean-based, there can be multiple equilibria where the buyer's distribution has different means. To see this, consider condition (i) of Proposition 4. Given the seller chooses price $\bar{v}$ with probability h, the mean of the buyer's best response is pinned down by the first order condition $1-h = \kappa\Gamma_0'(M_F)$. On the other hand, for the seller to be indifferent between the prices $\underline{v}$ and $\bar{v}$, all that is needed is that the buyer's distribution puts mass $\underline{v}/\bar{v}$ on $\bar{v}$. As long as this is the case, *any* $h \in [0,1]$ is a best response for the seller. Because the requirement to put mass $\underline{v}/\bar{v}$ on $\bar{v}$ is satisfied by various distributions with different means, there are multiple combinations of mutual best responses.[36]

Even though there might be multiple equilibria, in any equilibrium, the seller's profit is $\underline{v}$ by Proposition 2. Therefore, an equilibrium which maximizes the buyer's utility also maximizes total welfare, and is also Pareto-optimal. It is therefore natural to focus on Pareto-optimal equilibria which I characterize next.

To do so, recall from Lemma 2 that any distribution with mean $M^{FB}(\kappa)$ is a first-best distribution. As κ increases, $M^{FB}(\kappa)$ increases from $\underline{v}$ to $\bar{v}$, and so there is a (unique) $\hat{\kappa}$ where the first-best mean $M^{FB}(\hat{\kappa})$ is equal to the mean of the equal-revenue distribution $G_{\underline{v}}^{\bar{v}}$. In other words, at the value $\hat{\kappa}$, the equal-revenue distribution $G_{\underline{v}}^{\bar{v}}$ is a first-best distribution.[37] It is straightforward to verify that $\hat{\kappa} \in (\kappa_0, \kappa_1)$.

Proposition 5 *Under A4, we have:*

(i) If $\kappa < \hat{\kappa}$, then:

(a) There is a unique Pareto-optimal equilibrium (F,H) with $F = G_{\underline{v}}^{\bar{v}}$ and $H = T_h$ with $h = 1-\kappa\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$.

(b) In this equilibrium, the buyer's utility U_B and total welfare are strictly increasing in κ.

(ii) If $\kappa \geq \hat{\kappa}$, then:

(a) There is a first-best distribution F so that (F,H) with $H = T_h$ and $h = 0$ is a Pareto-optimal equilibrium.

[36]Among all distributions F that put mass $\underline{v}/\bar{v}$ on $\bar{v}$ and satisfy $F \geq G_{\underline{v}}^{\bar{v}}$, the distribution $T_{\underline{v}/\bar{v}}$ has the smallest mean, and the distribution $G_{\underline{v}}^{\bar{v}}$ has the largest mean. Therefore, any distribution F with $G_{\underline{v}}^{\bar{v}} \leq F \leq T_{\underline{v}/\bar{v}}$ and $\kappa\Gamma_0'(M_F) < 1$ is consistent with equilibrium.

[37]Formally, $\hat{\kappa} = 1/\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$.

(b) In any Pareto-optimal equilibrium, total welfare is equal to the first-best.

(iii) If $\kappa < \kappa_1$, then in any Pareto-optimal equilibrium, total welfare is strictly larger than in the benchmark case when valuations are public information.

The proposition says that as the (marginal) cost parameter κ increases up to $\hat{\kappa}$, the buyer's utility and total welfare increase. If κ is above $\hat{\kappa}$, equilibrium is first-best, that is, the hold-up problem causes no inefficiency. (Notice that since $\hat{\kappa} < \kappa_1$, the hold-up problem does not disappear for the uninteresting reason that costs are so high that the default distribution is efficient.) Finally, if the default distribution is not first-best ($\kappa < \kappa_1$), equilibrium welfare is larger than in the benchmark with public information.

To understand the results, consider first $\kappa < \hat{\kappa}$. In this case, there is no efficient equilibrium as in part (ii) of Proposition 4. The reason is that (marginal) costs are too low to sustain an efficient equilibrium: if the seller were to choose the low price $\underline{v}$ ($h = 0$), the buyer would best respond with a first-best distribution. But for $\kappa \leq \hat{\kappa}$, any first-best distribution is "too large" for the seller to best respond with the low price. More formally, since $M_{G_{\underline{v}}^{\bar{v}}} < M^{FB}$ if $\kappa < \hat{\kappa}$, no first-best distribution can be first order stochastically dominated by $G_{\underline{v}}^{\bar{v}}$ and thus be part of an equilibrium (recall Lemma 1). However, it is straightforward to verify that the combination $F = G_{\underline{v}}^{\bar{v}}$ and $H = T_h$ with $h = 1 - \kappa\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$ satisfies part (i) of Proposition 4 and is thus an equilibrium. It is uniquely Pareto-optimal because among all buyer distributions that are equilibrium candidates, $G_{\underline{v}}^{\bar{v}}$ has the maximal mean. This explain part (a).

The intuition behind part (b) is that as κ increases within the range $(0, \hat{\kappa})$, there are two effects on the buyer's utility. On the one hand, there is a price effect, since the seller increases the probability of charging the low price $p = \underline{v}$ as κ increases. One way to think about the price effect is that as κ increases, the buyer's marginal costs increase. Thus the buyer gains commitment not to put too much mass on high valuations, inducing the seller to lower the price. On the other hand, the buyer faces higher investment costs. In the range $(0, \hat{\kappa})$, the price effect outweighs the cost effect, and therefore U_B increases with κ. In fact, since costs are mean-based, the buyer effectively chooses a (uni-dimensional) mean $M \in [\underline{v}, \bar{v}]$. Since the buyer's utility is $(1-h)(M-\underline{v})-\kappa\Gamma_0(M)$, the marginal benefit of doing so is the probability of the low price $1-h$. Thus, given $1-h$, the buyer increases the mean until marginal costs are equal to marginal benefits: $1-h = \kappa\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$. On the other hand, for "inframarginal" units $M < M_{G_{\underline{v}}^{\bar{v}}}$, marginal benefits are strictly larger than marginal costs because costs are strictly convex, thus generating strictly positive utility for the buyer. Now, as κ increases, the difference between marginal benefits and marginal costs for the inframarginal units becomes more pronounced. As a result, the buyer's utility as well as total welfare increase in κ.

Next, consider the case $\kappa \geq \hat{\kappa}$. (Marginal) costs are now sufficiently high to sustain an effcient equilibrium as in part (ii) of Proposition 4. Indeed, because $M_{G_{\underline{v}}^{\bar{v}}} \geq M^{FB}$ if $\kappa \geq \hat{\kappa}$, it follows from an intermediate value argument that there is an equal revenue distribution $G_{\underline{v}}^{\beta}$ with largest support point $\beta < \bar{v}$ such that $M_{G_{\underline{v}}^{\beta}} = M^{FB}$. Since costs are mean-based, $G_{\underline{v}}^{\beta}$ is thus a first-best distribution, and thus a best response by the buyer to the low price. Moreover, since $G_{\underline{v}}^{\beta}$ is an equal revenue distribution, charging the low price for sure is a best response for the seller. Clearly, the equilibrium is Pareto-optimal, and this explains part (ii) of the proposition.

Finally, part (iii) is a direct consequence of the fact that the buyer's utility is strictly positive for all $\kappa < \kappa_1$.

6.2 Strictly concave Gateaux derivative

In this section, I assume A5. The analysis is almost identical to the case with mean-based costs for two reasons. First, I show that because costs are decreasing in risk, the first-best distribution puts mass only on the extreme valuations $\underline{v}$ and $\bar{v}$. Second, for the same reason as with mean-based costs, in any equilibrium, the seller will randomize between $\underline{v}$ and $\bar{v}$. Under A5, this will imply that also the buyer's best response puts mass only on $\underline{v}$ and $\bar{v}$. However, since distributions that put mass only on $\underline{v}$ and $\bar{v}$ can be described by their mean alone, the setting becomes effectively mean-based. As a consequence, Proposition 4 and 5 carry over almost literally. I begin with the characterization of the first-best distribution.

Lemma 3 *Under A5, there is a unique first-best distribution given by $F^{FB} = T_{f^{FB}}$ where f^{FB} minimizes $f(\bar{v}-\underline{v})-C(T_f)$ over $f \in [0,1]$.*

To understand the intuition, consider the problem of maximizing total surplus for a given mean:

$$\max_F \int_V v\, dF - C(F) = M_F - C(F) \quad s.t. \quad M_F = M. \tag{38}$$

Because C is decreasing in risk, F minimizes costs by maximizing risk, that is, by putting mass only on the extreme valuations $\underline{v}$ and $\bar{v}$. Thus, the solution to (38) is a distribution T_f and induces surplus $f(\bar{v}-\underline{v})-C(T_f)$. (Uniqueness of the solution follows formally from strict concavity of γ_F which implies that C is "strictly" decreasing in risk.)

I next derive necessary conditions for equilibrium. First, observe that for the same reason as with mean-based costs, in any equilibrium, the seller randomizes only between the prices $\underline{v}$ and $\bar{v}$. Moreover, the buyer's best response is now in the class of distributions T_f which put mass only on the extreme valuations $\underline{v}$ and $\bar{v}$. The reason is the same as in the previous paragraph, because if the seller randomizes between $\underline{v}$ (with $1-h$) and $\bar{v}$ (with h), the buyer's expected utility is

$(1-h)(M_F-\underline{v})-C(F)$, and thus she benefits by increasing the spread of the distribution (strictly so since γ_F is strictly concave).

Finally, by Lemma 1, in equilibrium F is first order stochastically dominated by $G_{\underline{v}}^{\bar{v}}$. For a distribution $F=T_f$, this is the case if and only if $f \leq \underline{v}/\bar{v}$ or, equivalently, $T_f \geq T_{\underline{v}/\bar{v}}$. Finally, the same argument as with mean-based costs implies $h<1$ in equilibrium. The following proposition summarizes.

Proposition 6 *Under A5, (F,H) is an equilibrium only if $H=T_h$ with $h<1$, and $F=T_f$ with $F \geq T_{\underline{v}/\bar{v}}$.*

By Lemma 3 and Proposition 6, the search for first-best and the buyer's equilibrium distributions can be restricted to distributions T_f. Moreover, the mean $M=(1-f)\underline{v}+f\bar{v}$ of such a distribution uniquely pins down $f=\frac{M-\underline{v}}{\bar{v}-\underline{v}}$. Thus, instead of searching for first-best or equilibrium values of f, one can as well search for first-best or equilibrium values of $M \in [\underline{v},\bar{v}]$. With this change of variables, the setting becomes effectively mean-based, and the results from the previous section apply. Specifically, define the cost of T_f with mean M as

$$\Gamma_0(M)=\Gamma(T_{\frac{M-\underline{v}}{\bar{v}-\underline{v}}}). \tag{39}$$

Γ_0 is then increasing and strictly convex with derivative[38]

$$\Gamma_0'(M)=\frac{\gamma_{T_{\frac{M-\underline{v}}{\bar{v}-\underline{v}}}}(\bar{v})-\gamma_{T_{\frac{M-\underline{v}}{\bar{v}-\underline{v}}}}(\underline{v})}{\bar{v}-\underline{v}}. \tag{41}$$

The characterization of the mean of the first-best distribution is then identical to Lemma 2. Likewise, given that the seller chooses T_h in equilibrium, the buyer's utility from T_f in terms of its mean is $(1-h)(M-\underline{v})-\kappa\Gamma_0(M)$ which is formally identical to the case with mean-based costs. Therefore, Proposition 4 and Proposition 5 carry over verbatim with the only difference that the distribution $G_{\underline{v}}^{\bar{v}}$ is replaced by the distribution $T_{\underline{v}/\bar{v}}$. The reason is that now the buyer's equilibrium distribution is in the class of distributions $F=T_f \geq T_{\underline{v}/\bar{v}}$.

Remark Because under A4 and A5, the equilibrium outcome is first-best for $\kappa \geq \hat{\kappa}$, the equilibrium outcome also coincides with the buyer-optimal commitment outcome as analyzed in Condorelli and Szentes (2020). The reason is that seller's profit when the buyer has commitment is never

[38] To see this, note that $T_{f+\epsilon}=T_f+\epsilon(T_1-T_0)$, and thus

$$\frac{d}{df}\Gamma(T_f)=\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left[\Gamma(T_f+\epsilon(T_1-T_0))-\Gamma(T_f)\right]=\delta\Gamma(T_f;T_1-T_0)=\int_V \gamma_{T_f}\, d(T_1-T_0)=\gamma_{T_f}(\bar{v})-\gamma_{T_f}(\underline{v}). \tag{40}$$

The expression for $\Gamma_0'(M)$ now follows with the chain rule.

smaller than $\underline{v}$, because charging the price $\underline{v}$ guarantees profit $\underline{v}$. Therefore, if $\kappa \geq \hat{\kappa}$, then even with commitment, the buyer cannot attain higher utility than her equilibrium utility $W^{FB} - \underline{v}$. In other words, commitment has no value for the buyer.

6.3 Strictly convex Gateaux derivative

In this section, I assume A6. A somewhat different picture emerges in this case compared to the previous sections. The first-best distribution is now deterministic, and in equilibrium, the seller now randomizes over an interval of prices, and equilibria will not be efficient (unless costs are so high that the default distribution is first-best). I begin by characterizing the first-best.

Lemma 4 *Under A6, a first-best distribution is a deterministic distribution $F^{FB} = \mathbb{1}_{[v^{FB},\bar{v}]}$ that puts all mass on v^{FB} where v^{FB} maximizes $v - \kappa\Gamma(\mathbb{1}_{[v,\bar{v}]})$ over $v \in [\underline{v}, \bar{v}]$. In particular, the default distribution is first-best if and only if*

$$\kappa \geq \kappa_1 = \frac{1}{\gamma'_{\mathbb{1}_{[\underline{v},\bar{v}]}}(\underline{v})}. \tag{42}$$

Intuitively, consider again problem (38). Because now C is increasing in risk, F minimizes costs by minimizing risk, that is, by concentrating all mass on a single point. (That this is strictly optimal follows formally from the strict convexity of γ_F.) Thus, the solution to (38) is deterministic. Moreover, the total surplus induced by a distribution that puts all mass on v is $v - \kappa\Gamma(\mathbb{1}_{[v,\bar{v}]})$.

The characterization for when the default distribution is first-best will be used below when I compare equilibrium and first-best. The formal derivation follows from a general characterization of the first-best distribution that I give in the appendix.

I next characterize equilibrium. Recall that when costs are mean-based or decreasing in risk, the seller randomizes between $\underline{v}$ and $\bar{v}$ in equilibrium. This is now no longer the case. The reason is that given this strategy by the seller, the buyer's expected utility is $(1-h)(M_F - \underline{v}) - C(F)$. Since costs are now increasing in risk, the same argument as in the previous paragraph implies that the buyer would optimally want to concentrate her valuation on a single point. But then it is not a best response for the seller to randomize between $\underline{v}$ and $\bar{v}$.

To state equilibrium formally, note that strict convexity of γ_F implies that the function $v - \kappa\gamma_F(v)$ has a unique maximizer $v^*(F)$ on $[\underline{v}, \bar{v}]$.

Proposition 7 *Under A6, (F, H) is an equilibrium if and only if*

$$F = G_{\underline{v}}^{v^*(F)}, \quad H(p) = \kappa\gamma'_F(p)\mathbb{1}_{[\underline{v},v^*(F))}(p) + \mathbb{1}_{[v^*(F),\bar{v}]}(p). \tag{43}$$

Proposition 7 says that the buyer's distribution is an equal revenue distribution and that the seller's pricing distribution is essentially equal to the (ordinary) derivative of the Gateaux derivative. In particular, the supports $[\underline{v}, v^*(F)]$ of both distributions are convex and identical.

The proof of the proposition shows that (only) strategies of the form (43) satisfy the equilibrium characterization in Proposition 1. This can be verified using the same arguments as in the proof of Proposition 1 in Gul (2001). The difference is that Gul (2001) considers the case with linear cost so that γ_F does not depend on F.[39] While this does not matter for the formal argument, note, however, that whenever costs are not linear, the equilibrium is characterized only implicitly because F depends on $v^*(F)$.

A more explicit characterization can be obtained by noting that $v^*(F) = \beta$ is the upper support bound of an equal revenue distributions $F = G_{\underline{v}}^{\beta}$. Thus, equilibrium is characterized by the solutions to the equation $v^*(G_{\underline{v}}^{\beta}) = \beta$. Since $v^*(F)$ maximizes $v - \kappa\gamma_F(v)$, $v^*(F)$ is the solution to the respective first order condition. This pins down the equilibrium value of β as stated next.

Proposition 8 *Under A6, (F, H) is an equilibrium if and only if*

$$F = G_{\underline{v}}^{\beta}, \quad H(p) = \kappa\gamma'_{G_{\underline{v}}^{\beta}}(p)\mathbb{1}_{[\underline{v},\beta)}(p) + \mathbb{1}_{[\beta,\bar{v}]}(p), \tag{44}$$

and β is any value that satisfies:

$$\kappa\gamma'_{G_{\underline{v}}^{\beta}}(\beta) = 1, \quad \textit{or} \quad \beta = \underline{v} \textit{ and } \kappa\gamma'_{G_{\underline{v}}^{\underline{v}}}(\underline{v}) \geq 1, \quad \textit{or} \quad \beta = \bar{v} \textit{ and } \kappa\gamma'_{G_{\underline{v}}^{\bar{v}}}(\bar{v}) \leq 1. \tag{45}$$

In general, there might be multiple solutions β to (45) and thus multiple equilibria. A sufficient condition for there to be a unique solution β is that $\gamma'_{G_{\underline{v}}^{\beta}}(v)$ is strictly increasing in β for all $v \in V$. To see this, note that the convexity of γ then implies that the function $\gamma'_{G_{\underline{v}}^{x}}(x)$ is strictly increasing in x, and an intermediate value argument implies that (45) has a unique solution.[40] One class of cost functions which satisfies this property is the class of "moment-based" cost functions $\Gamma(F) = \tilde{\Gamma}(\int_V c(v)\, dF)$ where $\tilde{\Gamma}$ is differentiable, increasing and convex, and c is differentiable, increasing and strictly convex.[41]

For the purpose of comparative statics with respect to κ, I shall now assume that equilibrium is unique. Recall that κ_1 is the critical value from which on the default distribution becomes first-best.

[39] As in Gul (2001), differentiability of γ_F is not needed for the result. In this case, the derivative of γ_F in (43) has to be replaced by the right derivative.

[40] Recall that a differentiable convex function is also continuously differentiable.

[41] In this case, $\gamma_F(v) = \tilde{\Gamma}'(\int_V c(v) dF) c(v)$, and $\gamma'_F(v) = \tilde{\Gamma}'(\int_V c(v) dF) c'(v)$. As F increases in the FOSD sense, so does $\tilde{\Gamma}'(\int_V c(v)\, dF)$ as well as $\gamma'_F(v)$. Because $G_{\underline{v}}^{\beta}$ increases in the FOSD sense as β increases, the claim follows.

Proposition 9 *Suppose that A6 holds, and let*

$$\hat{\kappa} = \frac{1}{\gamma'_{G^{\bar{v}}_{\underline{v}}}(\bar{v})}. \tag{46}$$

Moreover, suppose that there is a unique equilibrium. Then, we have:

(i) The buyer's utility U_B and total welfare is strictly increasing in κ for $\kappa \in (0, \hat{\kappa})$.

(ii) Total welfare when valuations are private information is strictly larger than when they are public for all $\kappa < \kappa_1$.

(iii) First-best welfare is strictly larger than equilibrium welfare for all $\kappa < \kappa_1$, that is, unless the default distribution is first-best.

The intuition behind part (i) is similar to that behind part (i) of Proposition 5. As κ increases in the range $(0, \hat{\kappa})$, the direct effect of facing higher investment costs is outweighed by the indirect strategic effect that the seller reduces the price (in the first order sense). More precisely, observe that for values $\kappa \leq \hat{\kappa}$, the buyer chooses $F = G^{\bar{v}}_{\underline{v}}$ in equilibrium. Therefore, the Gateaux derivative $\kappa\gamma_{G^{\bar{v}}_{\underline{v}}}$ becomes steeper as κ increases. By Proposition 7, this means that the seller's pricing distribution decreases in the first order sense. This price effect outweighs the direct cost effect because costs Γ are strictly convex.

Part (ii) follows from the fact that for $\kappa < \kappa_1$ the equilibrium distribution is different from the default distribution. Therefore, by Proposition 3, the buyer's utility is strictly positive, and thus total welfare is strictly larger than when valuations are private information.

While parts (i) and (ii) are analogous to the case with linear or strictly concave Gateaux derivative from the previous sections, part (iii) is different. The reason is that unless the buyer's equilibrium distribution is the default distribution, the buyer's and the seller's distribution have the same (non-degenerate) interval support $[\underline{v}, \beta]$, leading to trade inefficiencies ex post. The equilibrium outcome is therefore not efficient.

I conclude this section with an example that sheds light on the difference between equilibrium and first-best welfare.

6.3.1 Example

Consider the second moment $Q_F = \int_V v^2 \, dF$, and define

$$\Gamma(F) = \frac{1}{4}Q_F^2 - \frac{1}{4}\underline{v}^4 \tag{47}$$

with Gateaux derivative

$$\gamma_F(v) = \frac{1}{2} Q_F v^2. \tag{48}$$

Γ is evidently convex. Since F_{min} places all mass on $\underline{v}$, the second term in (47) ensures that $\Gamma(F_{min}) = 0$, in line with the normalization in (12).

With an eye on applying Proposition 8, note that since $G_{\underline{v}}^{\beta}$ has density $\underline{v}/v$ on $[\underline{v}, \beta)$ and a mass point of size $\underline{v}/\beta$ at β, we have

$$Q_{G_{\underline{v}}^{\beta}} = \int_{\underline{v}}^{\beta} v^2 \frac{\underline{v}}{v^2}\, dv + \beta^2 \cdot \frac{\underline{v}}{\beta} = 2\underline{v}\beta - \underline{v}^2 \tag{49}$$

so that $\gamma'_{G_{\underline{v}}^{\beta}}(v) = [2\underline{v}\beta - \underline{v}^2]v$. Condition (45) for the equilibrium value β therefore writes

$$\kappa[2\underline{v}\beta^2 - \underline{v}^2\beta] = 1, \quad \text{or} \quad \beta = \underline{v} \text{ and } \kappa\underline{v}^3 \geq 1, \quad \text{or} \quad \beta = \bar{v} \text{ and } \kappa[2\underline{v}\bar{v}^2 - \underline{v}^2\bar{v}] \leq 1. \tag{50}$$

Thus, I obtain the following equilibrium characterization where I calculate the buyer's utility using (22).

Lemma 5 *Let Γ be given by (47). Let $\hat{\beta}$ be the positive solution to the quadratic equation*

$$2\underline{v}\kappa\hat{\beta}^2 - \underline{v}^2\kappa\hat{\beta} - 1 \;=\; 0. \tag{51}$$

Define $\hat{\kappa} = \frac{1}{\underline{v}\bar{v}(2\bar{v}-\underline{v})}$ and $\kappa_1 = \frac{1}{\underline{v}^3}$. Then the equilibrium value β in Proposition 8 is

$$\beta = \begin{cases} \bar{v} & if \quad \kappa \leq \hat{\kappa} \\ \hat{\beta} & if \quad \kappa \in (\hat{\kappa}, \kappa_1) \\ \underline{v} & if \quad \kappa \geq \kappa_1 \end{cases}. \tag{52}$$

Moreover, the buyer's equilibrium utility is $U_B = \kappa\underline{v}^2\left(\beta - \underline{v}\right)^2$.

The blue solid line in Figure 3 illustrates the typical shape of equilibrium welfare $U_B + \underline{v}$ as a function of κ. It increases linearly in the range $\kappa \in (0, \hat{\kappa})$ and then decreases. At κ_1, investment costs become prohibitive, and $F = F_{min}$ in equilibrium.

Next, I characterize the first-best which follows straightforwardly from Lemma 4.

Lemma 6 *Let Γ be given by (47). Let $\kappa_0 = 1/\bar{v}^3$. Then the first-best distribution is $F^{FB} = \mathbb{1}_{[v^{FB}, \bar{v}]}$*

where

$$v^{FB} = \begin{cases} \bar{v} & if \quad \kappa \leq \kappa_0 \\ \kappa^{-\frac{1}{3}} & if \quad \kappa \in (\kappa_0, \kappa_1) \\ \underline{v} & if \quad \kappa \geq \kappa_1 \end{cases} . \tag{53}$$

Figure 3 plots the first-best welfare (red, dashed line). The difference between first-best and equilibrium welfare is largest at $\kappa = 0$ (corresponding to linear costs) and then decreases with κ until it becomes zero at the level κ_1 where costs become prohibitive so that zero investment is efficient.

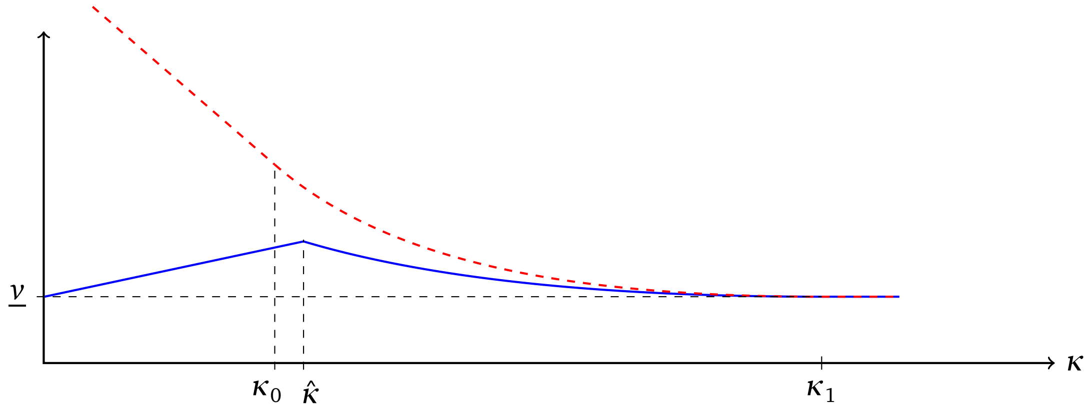

Figure 3: Total equilibrium welfare (blue) and first-best welfare (red, dashed) for the cost specification (47) as a function of κ for the values $\underline{v} = 1$, $\bar{v} = 3/2$, $\kappa_0 = 1/12$, $\kappa_1 = 1/2$.

7 Investment and information acquisition

My model can be extended to allow an interpretation where it is costly for the buyer not only to invest in, but also to learn about, her valuation. Ravid et al. (2022) consider a framework where the buyer's true valuation is a value θ in a compact interval Θ. Initially, the buyer only has a prior belief $F_0 \in \mathscr{F}$ over Θ, but she can acquire a signal about her valuation at a cost before trading. Since the buyer's preferences for the good are linear in the valuation, a signal corresponds to a distribution $F \in \mathscr{F}$ of posterior means that is a mean preserving contraction (MPC) of the prior F_0. Optimizing over a functional subject to the MPC constraint is, by now, a well-studied problem when the functional is linear (Dworczak and Martini, 2019, Kleiner et al., 2021), but is difficult when the functional is non-linear such as when information acquisition costs are non-linear.

The approach presented in this paper can however be applied to include information acquisition when one restricts the space Θ of the buyer's true valuations to consist of only two possible values, as I now illustrate. Suppose that $\Theta = \{\underline{v}, \bar{v}\}$. A prior then corresponds simply to a mean

$M_0 \in V = [\underline{v}, \bar{v}]$, and any signal corresponds to a cdf $F \in \mathscr{F}$ over posterior means $v \in V$ with the simplified MPC constraint that F has mean M_0, that is, $\int_V v dF = M_0$.

Specifically, suppose that without investing, the buyer's valuation is equal to the lowest possible valuation $\underline{v}$ with probability 1. The buyer can invest to increase the mean M_F of the valuation distribution F at a cost $\rho\varphi(M_F)$, $\rho \geq 0$, where φ is strictly increasing and convex. The mean M_F then corresponds to the prior, and the buyer can learn about the true valuation given the prior at a cost. Specifically, consider a strictly convex function $r : V \to \mathbb{R}$ and let $\sigma \int_V r(v) - r(M_F)\, dF$, $\sigma \geq 0$, be the ("posterior-separable") cost of information acquisition. Since r is convex, information acquisition costs increase in the mean preserving spread order, or equivalently, in Blackwell informativeness. Moreover, acquiring no information, which corresponds to choosing the degenerate distribution $\mathbb{1}_{[M,\bar{v}]}$ that places probability 1 on M, is costless.

The cost function

$$C(F) = \rho\varphi(M_F) + \sigma \int_V r(v) - r(M_F)\, dF \tag{54}$$

therefore combines the cost of investing and the cost of information acquisition. Suppose that $\rho\varphi - \sigma r$ is strictly convex. Then C is convex, and if r is strictly increasing, then the Gateaux derivative

$$c_F(v) = [\rho\varphi'(M_F) - \sigma r'(M_F)]v + \sigma r(v) \tag{55}$$

is strictly increasing and strictly convex. Therefore, Proposition 8 applies unchanged with $\kappa\gamma_F$ replaced by c_F. Moreover, as explained after the statement of Proposition 8, equilibrium is unique if $c'_{G^\beta_{\underline{v}}}(v)$ is strictly increasing in β for all $v \in V$. This is satisfied here because $\rho\varphi - \sigma r$ is strictly convex.

I now use these observations to obtain the following comparative statics result.

Lemma 7 *Let C be given by (54). Let φ, r be strictly increasing, differentiable, and strictly convex, and let $\rho\varphi - \sigma r$ be strictly convex. Then there are $\hat{\rho}, \hat{\sigma}(\rho)$ so that for all (ρ, σ) with $\rho < \hat{\rho}$ and $\sigma < \hat{\sigma}(\rho)$, the buyer's equilibrium utility and total welfare is strictly increasing in ρ and strictly decreasing in σ.*

The formal reason behind this result is similar to the reason behind Propositions 5 and 9. As ρ increases or σ decreases, the convex part of the cost function C goes up. Thus, an increase in ρ or a decrease in σ corresponds to an increase in κ in Propositions 5 and 9.

Economically, Lemma 7 suggests that the comparative statics with respect to investment and information acquisition costs are opposed to one another. While higher investment costs (higher

ρ) attenuate the hold-up problem, higher information acquisition costs (higher σ) aggravate the hold-up problem. The intuition behind this difference is that higher information acquisition costs diminish the buyer's ability to extract surplus in the form of information rents.

8 Conclusion

In this paper, I reconsider the hold-up problem with unobservable investments when the buyer's investment is flexible. I argue that the severity of the hold-up problem critically depends on the shape of investment costs. I show that compared to the case with linear costs, the buyer's equilibrium utility and equilibrium welfare is higher with convex costs. When, in addition, costs are mean-based or decreasing in risk and costs are sufficiently convex, the hold-up problem disappears in the sense that the equilibrium outcome is efficient.

I have argued that my results are relevant for applications to the extent that the shape of the cost function captures salient features of such applications. For example, in the context of a manufacturer's investments in quality improvement, while convexity captures increasing marginal costs of quality improvements, the risk properties of the cost function capture, in a stylized way, whether making relatively small but certain quality improvements is more or less costly than making large quality improvements with small probability. While in this paper, I adopt a reduced form approach and impose assumptions on the cost function directly, it is an interesting avenue for future work to micro-found the shape of the cost function building on a specific application. On a related note, my results suggest that, if feasible, the investing party has incentives to design the shape of the cost function in an attempt to create commitment.

Finally, the framework presented in this paper is sufficiently tractable to extend the model to allow for competition between multiple buyers or multiple sellers or both. More generally, the framework is also portable to other applications with flexible pre-investments.

A Appendix

Proof of Proposition 1 I only show part (ii). (Part (i) follows with the same arguments as the proof in footnote 22 in Ravid et al., 2022.) Note first that it is a standard argument that the (mixed) strategy H is a best response to F for the seller if and only (16) and (17) hold, where π is the seller's best response profit. Because the seller can guarantee himself the profit $\underline{v}$ by choosing the price $\underline{v}$ with probability 1, we have $\pi \geq \underline{v}$.

That the buyer's best response to H is characterized by (14) and (15) is shown in the main text. QED

Proof of Proposition 2 Note first that by setting the price $\underline{v}$ with probability 1, the seller can ensure profit $\underline{v}$. Thus, $\underline{v} \leq \Pi$. To see that $\underline{v} \geq \Pi$, assume to the contrary that $\underline{v} < \Pi$. Let $p_\ell = \min supp(H)$ be the lower support bound of the seller's pricing distribution.

I first show that $p_\ell \notin supp(F)$. Indeed, since $\underline{v} < \Pi = (1 - F(p_\ell^-))p_\ell$, we have $\underline{v} < p_\ell$. Because $\overline{H}(\underline{v}) = \overline{H}(p_\ell) = 0$ by definition of $\overline{H}$, and since $c_F(\underline{v}) < c_F(p_\ell)$ by A3, it follows

$$\overline{H}(\underline{v}) - c_F(\underline{v}) - \lambda = 0 - c_F(\underline{v}) - \lambda > 0 - c_F(p_\ell) - \lambda = \overline{H}(p_\ell) - c_F(p_\ell) - \lambda. \tag{56}$$

Therefore, (14) and (15) imply that $p_\ell \notin supp(F)$.

Now distinguish two cases:

(a) $F(p_\ell^-) < 1$. Since $F(\bar{v}) = 1$ and $p_\ell \notin supp(F)$, this implies that $p_\ell < \bar{v}$. Therefore, since $p_\ell \notin supp(F)$, there is $q > p_\ell$ with $F(p_\ell^-) = F(q)$, and hence the seller could increase profits by increasing the price from p_ℓ to q, contradicting that $p_\ell \in supp(H)$.

(b) $F(p_\ell^-) = 1$. Then the seller's profit from price p_ℓ is zero, and hence, since $p_\ell \in supp(H)$, we have that $\Pi = 0$. This contradicts that $\underline{v} < \Pi$. QED

Proof of Proposition 3 To show the generality of the proposition, I prove it without invoking A3. In this case, $c_F(\underline{v})$ in (22) is replaced by $\min_{v\in V} c_F(v)$. Moreover, without A3, F_{min} is defined as the distribution that minimizes $C(F)$.

As to (i). I first show that $\lambda = -\min_{v\in V} c_F(v)$. Indeed, let $v_\ell = \min supp(F)$, $p_\ell = \min supp(H)$ be the lower support bounds.

Observe first that $v_\ell \leq p_\ell$. Otherwise, if $p_\ell < v_\ell$, then $F(p_\ell^-) = 0$ so that the seller's profit at p_ℓ is $(1 - F(p_\ell^-))p_\ell = p_\ell$. But since $F(v_\ell^-) = 0$ by definition of v_ℓ, the seller could strictly increase his profit by deviating to the price $p = v_\ell$. By (16) and (17), this contradicts that p_ℓ is in $supp(H)$.

Next, I show that $v_\ell \in \arg\min_{v\in V} c_F(v)$. To the contrary, suppose $c_F(v_\ell) > c_F(\hat{v})$ where $\hat{v} \in \arg\min_{v\in V} c_F(v)$. Because $v_\ell \leq p_\ell$, we have that $\overline{H}(v_\ell) = 0$. Therefore, because (trivially) $\overline{H}(\hat{v}) \geq 0$, we have

$$\overline{H}(v_\ell) - c_F(v_\ell) - \lambda < \overline{H}(\hat{v}) - c_F(\hat{v}) - \lambda. \tag{57}$$

By (14) and (15), this contradicts that v_ℓ is in $supp(F)$.

Therefore, because $v_\ell \leq p_\ell$ implies $\overline{H}(v_\ell) = 0$, and because $c_F(v_\ell) = \min_{v\in V} c_F(v)$, we infer from (15) that

$$0 = \overline{H}(v_\ell) - c_F(v_\ell) - \lambda = -\min_{v\in V} c_F(v) - \lambda, \tag{58}$$

as desired.

To see the expression for U_B, recall from (5) that $U_B(H,F) = \int_V \overline{H}(v)\, dF(v) - C(F)$. Thus, plugging in $\overline{H}$ from (15) yields (22) with $c_F(\underline{v})$ replaced by $\min_{v\in V} c_F(v)$.

As to (ii). On the one hand, when C is linear, we have $\int_V c_F(v)\, dF = \int_V c(v)\, dF = C(F)$. Thus, $U_B = -\min_{v\in V} c(v)$ by (22). On the other hand, the default distribution F_{min} minimizes $C(F) = \int c(v)\, dF$ and thus places full mass on points v where c is minimal. Thus, $C(F_{min}) = \min_{v\in V} c(v)$. Finally, recall that $C(F_{min}) = 0$ by assumption. Taken together, this implies $U_B = 0$.

As to (iii). The proof is given in the main text. QED

Proof of Corollary 1 The proof is given in the main text. QED

Proof of Lemma 1 The proof is given in the main text. QED

Proof of Lemma 2 The proof is given in the main text. QED

Proof of Proposition 4 Let (F,H) be an equilibrium. To see the first part of the statement, note that F is first order stochastically dominated by $G_{\underline{v}}^{\bar{v}}$ by Lemma 1. Moreover, to see that the seller's pricing distribution is of the form T_h, recall from (18) that in equilibrium every valuation in the buyer's support maximizes the function $\overline{H}(v) - c_F(v)$. Because $\overline{H}$ is convex by definition, and c_F is linear for mean-based costs, either (1) $\overline{H}(v) - c_F(v)$ is maximized at a corner point $\underline{v}$ or $\bar{v}$, or (2) any $v \in [\underline{v}, \bar{v}]$ is a maximizer, and $\overline{H}$ is, in fact, linear itself. In case (1), because the buyer has no valuation in $(\underline{v}, \bar{v})$, setting a price in $(\underline{v}, \bar{v})$ is strictly suboptimal. Case (2) implies, by definition, that $\overline{H}' = H$ is constant on $(\underline{v}, \bar{v})$, and thus H has no support point in $(\underline{v}, \bar{v})$. Finally, the argument that $h < 1$ in equilibrium is in the main text.

It remains to show (i) and (ii).

As to (i). We have that $h \in (0,1)$ in equilibrium if and only if the seller is indifferent between the prices $\underline{v}$ and $\bar{v}$ which is the case if and only if $F(\bar{v}^-) = 1 - \underline{v}/\bar{v}$. Further, F is a best response by the buyer if and only if (14) and (15) hold. Note that $\overline{H}(v) = (1-h)v$ (since $H = T_h$), and that $c_F(v) = \kappa\Gamma_0'(M_F)v$. Hence, (14) and (15) write

$$(1-h)v - \kappa\Gamma_0'(M_F)v - \lambda \;=\; 0 \quad \forall v \in supp(F), \tag{59}$$

$$(1-h)v - \kappa\Gamma_0'(M_F)v - \lambda \;\leq\; 0 \quad \forall v \in V. \tag{60}$$

Since $F(\bar{v}^-) = 1 - \underline{v}/\bar{v}$, F has a mass point of mass $\underline{v}/\bar{v}$ at $\bar{v}$. Because $\underline{v}/\bar{v} < 1$, F has at least one other point in its support, and thus (59) is true for at least two points in V. Since the function on the left hand side is linear in v, it follows that (59) is actually true for all points in V. Thus, (59) and (60) are equivalent to

$$(1-h)v - \kappa\Gamma_0'(M_F)v - \lambda \;=\; 0 \quad \forall v \in V. \tag{61}$$

This is equivalent to $h = 1 - \kappa\Gamma_0'(M_F)$ and $\lambda = 0$. Thus, $h \in (0, 1)$ is equivalent to $\kappa\Gamma_0'(M_F) < 1$. This completes the proof of (i).

As to (ii). To see the "only if"-part, let (F, T_h) be an equilibrium with $h = 0$, that is, $p = \underline{v}$ with probability 1. Then the buyer's utility is $M_F - \underline{v} - \kappa\Gamma_0(M_F)$ which is equal to the full surplus minus the constant $\underline{v}$. Therefore, the buyer's best response F is a first-best distribution.

To see the "if"-part, let F be a first-best distribution with $F(v) \geq G_{\underline{v}}^{\bar{v}}(v)$ for all $v \in V$. Then, in particular $F(\bar{v}^-) \geq G_{\underline{v}}^{\bar{v}}(\bar{v}^-) = G_{\underline{v}}^{\bar{v}}(\bar{v}) = 1 - \underline{v}/\bar{v}$. In other words, the mass on $\bar{v}$ is weakly less than $\underline{v}/\bar{v}$. Thus, the seller weakly prefers $p = \underline{v}$ over $p = \bar{v}$. Hence, $h = 0$ is a best response by the seller. Moreover, given $h = 0$, F is a best response by the buyer as argued in the previous paragraph, and this completes the proof. QED

Proof of Proposition 5 As to (i),(a). Let $\kappa \leq \hat{\kappa}$. This implies that $M_{G_{\underline{v}}^{\bar{v}}} < M^{FB}$. Then there is no first-best distribution F^{FB} that is first order stochastically dominated by $G_{\underline{v}}^{\bar{v}}$. Hence, $F^{FB} \not\geq G_{\underline{v}}^{\bar{v}}$. Hence, by part (ii) of Proposition 4, there is no equilibrium with $h = 0$.

Next, I show that $(G_{\underline{v}}^{\bar{v}}, T_h)$ with $h = 1 - \Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$ is an equilibrium. By part (i) of Proposition 4, it is sufficient to show that $\kappa\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}}) < 1$. Indeed, because $\underline{v} \leq M_{G_{\underline{v}}^{\bar{v}}} < M^{FB}$, Lemma 2 implies that either $M^{FB} = \bar{v}$ and thus $\kappa\Gamma_0'(\bar{v}) = \kappa\Gamma_0'(M^{FB}) < 1$, or $M^{FB} = \Gamma_0'^{-1}(1/\kappa)$ and thus $\kappa\Gamma_0'(M^{FB}) = 1$. Since Γ_0 is strictly convex, Γ_0' is strictly increasing, and so the fact that $M_{G_{\underline{v}}^{\bar{v}}} < M^{FB}$ implies that $\kappa\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}}) < 1$, as desired.

To complete the proof of part (a), I now argue that $(G_{\underline{v}}^{\bar{v}}, T_h)$ with $h = 1 - \Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})$ is uniquely Pareto-optimal. Indeed, let $(\tilde{F}, T_{\tilde{h}})$ be another equilibrium. As remarked above, there is no equilibrium with $\tilde{h} = 0$. Hence, $\tilde{h} = 1 - \Gamma_0'(M_{\tilde{F}})$ by part (i) of Proposition 4. The buyer's utility in $(\tilde{F}, T_{\tilde{h}})$ is

$$\tilde{U}_B = (1 - h)(M_{\tilde{F}} - \underline{v}) - \kappa\Gamma_0(M_{\tilde{F}}) = \kappa[\Gamma_0'(M_{\tilde{F}})(M_{\tilde{F}} - \underline{v}) - \Gamma_0(M_{\tilde{F}})]. \tag{62}$$

The derivative with respect to $M_{\tilde{F}}$ is $\kappa\Gamma_0''(M_{\tilde{F}})(M_{\tilde{F}} - \underline{v})$ which is positive since Γ_0 is strictly convex. Therefore, the Pareto-optimal equilibrium maximizes $M_{\tilde{F}}$. Recall from Proposition 4 that $\tilde{F} \geq G_{\underline{v}}^{\bar{v}}$ in any equilibrium. Therefore, $M_{\tilde{F}} < M_{G_{\underline{v}}^{\bar{v}}}$ for any $\tilde{F} \neq G_{\underline{v}}^{\bar{v}}$, and hence $(G_{\underline{v}}^{\bar{v}}, T_h)$ is uniquely Pareto-optimal.

To complete the proof of part (i), it remains to show (b). Note that since the default distribution puts all mass on $\underline{v}$, we have $\Gamma_0(\underline{v}) = 0$ by (12). Therefore, (62) implies that the buyer's equilibrium utility is

$$U_B = \kappa[\Gamma_0'(M_{G_{\underline{v}}^{\bar{v}}})(M_{G_{\underline{v}}^{\bar{v}}} - \underline{v}) - (\Gamma_0(M_{G_{\underline{v}}^{\bar{v}}}) - \Gamma_0(\underline{v}))]. \tag{63}$$

Note that the term in the square brackets is strictly positive because strict convexity of Γ_0 implies

$\Gamma_0'(x)(x-\underline{v}) > \Gamma_0(x)-\Gamma_0(\underline{v})$ for all $x > \underline{v}$. Therefore, U_B is strictly increasing in κ.

As to (ii),(a). I first show that there is an equilibrium (F, T_h) with $h = 0$ and a first-best distribution F. Indeed, $\kappa \geq \hat{\kappa}$ implies $M_{G_{\underline{v}}^{\bar{v}}} \geq M^{FB} \geq \underline{v} = M_{G_{\underline{v}}^{\underline{v}}}$. Thus, an intermediate value argument delivers that there is $\beta \leq \bar{v}$ so that $M_{G_{\underline{v}}^{\beta}} = M^{FB}$, and hence $G_{\underline{v}}^{\beta}$ is a first-best distribution by Lemma 2. Moreover, $G_{\underline{v}}^{\beta} \geq G_{\underline{v}}^{\bar{v}}$ by definition, and thus $(G_{\underline{v}}^{\beta v}, T_h)$ with $h = 0$ is an equilibrium by part (ii) of Proposition 4.

Since the seller chooses $p = \underline{v}$ with probability 1 in this equilibrium, and $G_{\underline{v}}^{\beta}$ is a first-best distribution, the buyer extracts the residual first-best surplus $W^{FB} - \underline{v}$. Since the seller gets $\underline{v}$ in any equilibrium by Proposition 2, there is no equilibrium in which the buyer gets a higher utility. It follows that $(G_{\underline{v}}^{\hat{v}}, T_0)$ is a Pareto-optimal equilibrium.

Part (b) is obvious, and this completes the proof of part (ii).

Finally, part (iii) follows from part (iii) Proposition 3 and (13) together with the observation that for $\kappa < \kappa_1$, we have $C(F) \neq 0$ for the equilibrium distributions F from part (i) and (ii). QED

Proof of Lemma 3 To prove the lemma, I first characterize the first-best distribution for general c_F.

Lemma A.1 *F^{FB} maximizes $\int_V v\, dG - C(G)$ if and only if there is λ^{FB} such that*

$$v - c_{F^{FB}}(v) - \lambda^{FB} \leq 0 \quad \forall v \in V, \tag{64}$$

$$v - c_{F^{FB}}(v) - \lambda^{FB} = 0 \quad \forall v \in supp(F^{FB}). \tag{65}$$

The proof of Lemma A.1 is identical to the proof that establishes the best response conditions (14) and (15) for the buyer. The only difference is that in the objective function, the buyer's expected gross benefit $\int_V \overline{H}(v)\, dF$ is replaced by $\int_V v\, dF$. QED

I can now prove Lemma 3. Because c_F is strictly concave, $v - c_F(v)$ is strictly convex and thus maximized at the points $\underline{v}$ or $\bar{v}$. Thus, a distribution that satisfies (64) and (65) must be a two-point distribution T_f for some f, and the optimal f^{FB} maximizes $\int_V v\; dT_f - C(T_f) = f(\bar{v} - \underline{v}) + \underline{v} - C(T_f)$. Since C is strictly convex by assumption, $C(T_f)$ is strictly convex in f, and thus f^{FB} is unique. QED

Proof of Proposition 6 The proof is given in the main text except for the claim that $F = T_f$ in any equilibrium. To see this, note that by (18), any point in the support of the buyer's equilibrium distribution maximizes $\overline{H}(v) - \kappa\gamma_F(v)$. Because of convexity of $\overline{H}$ and strict concavity of γ_F, the only possible maximizers are $\underline{v}$ or $\bar{v}$. Thus, $F = T_f$. QED

Proof of Lemma 4 By Lemma A.1, a first-best distribution is characterized by the conditions (64) and (65). Because γ_F is strictly convex for all F, the function $v - \kappa\gamma_F(v)$ has a unique maximizer for all F. Therefore, the conditions (64) and (65) can be satisfied only for a distribution F^{FB} which is a degenerate distribution $\mathbb{1}_{[v^{FB},\bar{v}]}$ for some point v^{FB}. Since a degenerate distribution that puts all mass on v generates total welfare $v - \kappa\Gamma(\mathbb{1}_{[v,\bar{v}]})$, v^{FB} maximizes this expression.

To see (42), note that by (64) and (65), the default distribution $\mathbb{1}_{[\underline{v},\bar{v}]}$ is first-best if and only if the only point in its support, $\underline{v}$ maximizes $v - \kappa\gamma_{\mathbb{1}_{[\underline{v},\bar{v}]}}(v)$. This the case if and only if $1 - \kappa\gamma'_{\mathbb{1}_{[\underline{v},\bar{v}]}}(\underline{v}) \leq 0$, that is, (42). QED

Proof of Proposition 7 That (F, H) as stated in the proposition satisfies the equilibrium conditions in Proposition 1 can be verified using the arguments as in the proof of Proposition 1 in Gul (2001). I omit the details. QED

Proof of Proposition 8 Because $v^*(G^\beta_{\underline{v}})$ maximizes $v - \kappa\gamma_{G^\beta_{\underline{v}}}(v)$, we have that $v^*(G^\beta_{\underline{v}})$ is given as the solution v^* to the first order condition

$$\kappa\gamma'_{G^\beta_{\underline{v}}}(v^*) = 1, \quad \text{or} \quad v^* = \underline{v} \text{ and } \kappa\gamma'_{G^\beta_{\underline{v}}}(\underline{v}) \geq 1, \quad \text{or} \quad v^* = \bar{v} \text{ and } \kappa\gamma'_{G^\beta_{\underline{v}}}(\bar{v}) \leq 1. \tag{66}$$

By Proposition 7, the equilibrium value of β is given by $v^*(G^\beta_{\underline{v}}) = \beta$. Inserting this in (66) yields the claim. QED

Proof of Proposition 9 As to (i). Since $\kappa < \hat{\kappa}$, Proposition 8 implies that $F = G^{\bar{v}}_{\underline{v}}$ in equilibrium. By part (i) of Proposition 3, the buyer's equilibrium utility is thus

$$U_B \;=\; \kappa\left(\int_V \gamma_{G^{\bar{v}}_{\underline{v}}}(v)\, dK^{\bar{v}}_{\underline{v}}(v) - \Gamma(G^{\bar{v}}_{\underline{v}}) - \min_{v\in V}\gamma_{G^{\bar{v}}_{\underline{v}}}(v)\right). \tag{67}$$

It follows with the same arguments as in part (iii) of Proposition 3 that the term in brackets is strictly positive. Thus, since $G^{\bar{v}}_{\underline{v}}$ is independent of κ, U_B is strictly increasing in κ for all $\kappa < \hat{\kappa}$.

As to (ii). Recall that F_{min} is uniquely given by the distribution $\mathbb{1}_{[\underline{v},\bar{v}]}$ that places mass 1 on $\underline{v}$. For $\kappa < \kappa_1$, Proposition 8 implies that $F = G^\beta_{\underline{v}}$ with $\beta > \underline{v}$ in equilibrium. By part (iii) of Proposition 3, it follows that $U_B > 0$. Hence, since seller profit is $\underline{v}$ by Proposition 2, total welfare $U_B + \underline{v}$ is strictly larger than $\underline{v} = W^{PUB}$.

As to (iii). By Lemma 4 and Proposition 8, the equilibrium distribution F coincides with the first-best if and only if $F = G^{\underline{v}}_{\underline{v}} = \mathbb{1}_{[\underline{v},\bar{v}]}$ and $F^{FB} = \mathbb{1}_{[\underline{v},\bar{v}]}$. The former is the case if $\kappa \geq \kappa_1$, and the latter is the case if $\underline{v}$ maximizes $v - \kappa\gamma_{\mathbb{1}_{[\underline{v},\bar{v}]}}(v)$ which is also equivalent to $\kappa \geq \kappa_1$. Therefore, first-best welfare is strictly larger than equilibrium welfare if and only if $\kappa < \kappa_1$. QED

Proof of Lemma 5 The characterization of the equilibrium value β follows by straightforward algebra from Proposition 8 and (50).

It remains to calculate U_B. By (22) and (48):

$$\int_V c_F(v)\, dF(v) \;=\; \frac{1}{2}\kappa Q_F \int_V v^2\, dF(v) = \frac{1}{2}\kappa Q_F^2, \tag{68}$$

and $\min_{v\in V} c_F(v) = c_F(\underline{v}) = \frac{1}{2}\kappa Q_F \underline{v}^2$. Hence, by (22),

$$\begin{aligned} U_B &= \int_V c_F(v)\, dF(v) - C(F) - \min_{v\in V} c_F(v) &(69)\\ &= \frac{1}{2}\kappa Q_F^2 - \frac{1}{4}\kappa Q_F^2 + \frac{1}{4}\kappa \underline{v}^4 - \frac{1}{2}\kappa Q_F \underline{v}^2 &(70)\\ &= \frac{1}{4}\kappa \left(Q_F^2 + \underline{v}^4 - 2Q_F\underline{v}^2\right) &(71)\\ &= \frac{1}{4}\kappa \left(Q_F - \underline{v}^2\right)^2. &(72)\end{aligned}$$

Plugging in $Q_F = 2\underline{v}\beta - \underline{v}^2$ for $F = G_{\underline{v}}^{\beta}$ from (49) yields the claim. QED

Proof of Lemma 6 By Lemma 4, v^{FB} maximizes

$$v - \kappa\Gamma(\mathbb{1}_{[v,\bar{v}]}) = v - \frac{1}{4}\kappa v^4 + \frac{1}{4}\kappa \underline{v}^4. \tag{73}$$

Expression (53) now follows from the first order condition for this maximization problem. QED

Proof of Lemma 7 Notice first that since c_F is strictly increasing so that A3 holds, the default distribution puts all mass on $\underline{v}$. Therefore, by (12),

$$C(F_{min}) = \varphi(\underline{v}) = 0. \tag{74}$$

Moreover, by Proposition 2, A3 also implies that the seller's profit is $\underline{v}$ so that total welfare is $U_B + \underline{v}$. Hence, it is enough to show the claim for U_B only. To do so, define

$$\hat{\rho} = \frac{1}{\varphi'(M_{G_{\underline{v}}^{\bar{v}}})}, \qquad \hat{\sigma}(\rho) = \frac{1 - \rho\varphi'(M_{G_{\underline{v}}^{\bar{v}}})}{r'(\bar{v}) - r'(M_{G_{\underline{v}}^{\bar{v}}})}. \tag{75}$$

By definition, we then have for all (ρ, σ) with $\rho < \hat{\rho}$ and $\sigma < \hat{\sigma}(\rho)$ that

$$c'_{G_{\underline{v}}^{\bar{v}}}(\bar{v}) \leq 1, \tag{76}$$

and hence, Proposition 8 implies that $G_{\underline{v}}^{\bar{v}}$ is the buyer's equilibrium distribution. By part (i) of

Proposition 3, because $c_{G_{\underline{v}}^{\bar{v}}}(v)$ is minimized at $\underline{v}$, the buyer's equilibrium utility is thus

$$\begin{aligned}
U_B &= \int_V c_{G_{\underline{v}}^{\bar{v}}}(v)\, dK_{\underline{v}}^{\bar{v}}(v) - C(G_{\underline{v}}^{\bar{v}}) - \min_{v \in V} c_{G_{\underline{v}}^{\bar{v}}}(v) && (77)\\
&= \rho\left[\int_V \varphi'(M_{G_{\underline{v}}^{\bar{v}}}) v\, dK_{\underline{v}}^{\bar{v}}(v) - \varphi(M_{G_{\underline{v}}^{\bar{v}}}) - \varphi(\underline{v})\right] && (78)\\
&\quad -\sigma\left[\int_V r'(M_{G_{\underline{v}}^{\bar{v}}}) v\, dK_{\underline{v}}^{\bar{v}}(v) - r(M_{G_{\underline{v}}^{\bar{v}}}) - r(\underline{v})\right] && (79)\\
&\quad -\sigma r(\underline{v}) && (80)\\
&= \rho\left[\varphi'(M_{G_{\underline{v}}^{\bar{v}}}) M_{G_{\underline{v}}^{\bar{v}}} - \varphi(M_{G_{\underline{v}}^{\bar{v}}}) - \varphi(\underline{v})\right] - \sigma\left[r'(M_{G_{\underline{v}}^{\bar{v}}}) M_{G_{\underline{v}}^{\bar{v}}} - r(M_{G_{\underline{v}}^{\bar{v}}})\right]. && (81)
\end{aligned}$$

Because $\varphi(\underline{v}) = 0$ by (74), strict convexity of φ and r implies that the square brackets are strictly positive. Thus, the buyer's utility is strictly increasing in ρ and strictly decreasing in σ. QED

References

Barron, Daniel, George Georgiadis, and Jeroen Swinkels. Optimal contracts with a risk-taking agent. Theoretical Economics 15.2 (2020): 715-761.

Carroll, Gabriel. Robustness and linear contracts. American Economic Review 105.2 (2015): 536-563.

Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci. Stochastic dominance analysis without the independence axiom. Management Science 63.4 (2017): 1097-1109.

Diamond, Peter. Managerial incentives: On the near linearity of optimal compensation. Journal of Political Economy 106.5 (1998): 931-57.

Dilme, Francesc. Pre-trade private investments. Games and Economic Behavior 117 (2019): 98-119.

Dworczak, Piotr, and Giorgio Martini. The simple economics of optimal persuasion. Journal of Political Economy 127.5 (2019): 1993-2048.

Condorelli, Daniele, and Balazs Szentes. Information design in the hold-up problem. Journal of Political Economy 128.2 (2020): 681-709.

Georgiadis, George, Doron Ravid, and Balazs Szentes. Flexible moral hazard problems. Econometrica 92.2 (2024): 387-409.

Gul, Faruk. Unobservable investment and the hold-up problem. Econometrica 69.2 (2001): 343-376.

Hébert, Benjamin. Moral hazard and the optimality of debt. The Review of Economic Studies 85.4 (2018): 2214-2252.

Hermalin, Benjamin E., and Michael L. Katz. Information and the hold-up problem. The RAND Journal of Economics 40.3 (2009): 405-423.

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack. Extreme points and majorization: Economic applications. Econometrica 89.4 (2021): 1557-1593.

Lau, Stephanie. Information and bargaining in the hold-up problem. The RAND Journal of Economics 39.1 (2008): 266-282.

Luenberger, David G. Optimization by vector space methods. John Wiley & Sons, 1997.

Ravid, Doron, Anne-Katrin Roesler, and Balazs Szentes. Learning before trading: on the inefficiency of ignoring free information. Journal of Political Economy 130.2 (2022): 346-387.

Roesler, Anne-Katrin, and Balazs Szentes. Buyer-optimal learning and monopoly pricing. American Economic Review 107.7 (2017): 2072-2080.