Mechanism Design with Aftermarkets: Cutoff Mechanisms

JOB MARKET PAPER

*Piotr Dworczak**

January 28, 2017

Abstract

I study a mechanism design problem of allocating a single good to one of several agents. The mechanism is followed by an aftermarket, that is, a post-mechanism game played between the agent who acquired the good and third-party market participants. The designer has preferences over final outcomes, but she cannot redesign the aftermarket. However, she can influence its information structure by disclosing information elicited by the mechanism, subject to providing incentives for agents to report truthfully.

I identify a class of allocation and disclosure rules, called cutoff rules, that are implementable regardless of the form of the aftermarket and the underlying distribution of types. A mechanism can be guaranteed to be truthful in all cases *only* if it implements a cutoff rule. Cutoff mechanisms are tractable, and admit an indirect implementation that often makes them easy to use in practice. Sufficient conditions are given for particularly simple designs, such as a second-price auction with disclosure of the price, to be optimal within the class of cutoff mechanisms.

The theory is illustrated with applications to the design of auctions followed by bargaining or resale markets, and to the optimal level of post-transaction transparency in financial over-the-counter markets.

Keywords: Mechanism Design, Information Design, Disclosure, Auctions, OTC Markets
JEL codes: D44, D47, D82, D83.

*Stanford University, Graduate School of Business, E-mail: dworczak@stanford.edu, Website: http://web.stanford.edu/~dworczak. I am indebted to my advisors, Andy Skrzypacz and Paul Milgrom, for their continuous support and guidance. I am also grateful to Darrell Duffie and Michael Ostrovsky for comments and serving on my dissertation committee. I would like to thank Mohammad Akbarpour, Eric Budish, Jeremy Bulow, Gabriel Carroll, Yossi Feinberg, Matthew Gentzkow, Ben Golub, Matthew Jackson, Maciej Kotowski, David Kreps, Nicolas Lambert, Shengwu Li, Giorgio Martini, Paul Oyer, Alessandro Pavan, Andrés Perlroth, Bartosz Redlicki, Peter Reiss, Tim Roughgarden, Ilya Segal, Vasiliki Skreta, Yang Song, Takuo Sugaya, Zhe Wang, Robert Wilson, Milena Wittwer, Anthony Zhang, seminar participants at Paul Milgrom's group meeting, Stanford theory lunch, Stanford market design seminar, Ce^2 Workshop 2016, 15th Annual Columbia/Duke/MIT/Northwestern IO Theory Conference, and 2016 European Winter Meeting of the Econometric Society for helpful comments and discussions.

Contents

1 Introduction

"The game is always bigger than you think." This phrase succinctly captures a prevalent feature of practical mechanism design problems – they can rarely be fully understood without the wider market context. When a seller designs an auction, she should not ignore future resale or bargaining opportunities which might influence bidders' endogenous valuations for the object. A dealer in a financial over-the-counter market understands that a counterparty in a transaction may not be the final holder of the asset. Yet, most theoretical models analyze the design problem in a vacuum.[1]

In this paper, I revisit the canonical mechanism design problem of allocating an object to one of several agents. Unlike in the standard model, the mechanism is followed by an *aftermarket*, defined as a post-mechanism game played between the agent who acquired the object and other market participants (*third parties*). The aftermarket is beyond the control of the mechanism designer but she may have preferences over equilibrium outcomes of the post-mechanism game, either directly (e.g. when the designer wants to maximize efficiency) or indirectly through the impact on agents' endogenous valuations (e.g. when the designer wants to maximize revenue).

Although the mechanism designer is unable to redesign the aftermarket, she can influence its information structure by releasing information elicited by the mechanism. As a result, the design problem is augmented with an additional choice variable – a *disclosure rule*. For example, if a bidder who wins an object engages in bargaining over acquisition of complementary goods after the auction, the designer must decide how much information about bids to reveal after the auction. The choice of a disclosure rule impacts the bargaining position of the bidder in the aftermarket.

The resulting structure of the problem can be described as a composition of mechanism and information design. In the first step, the mechanism elicits information from the agents in an incentive compatible way to determine the allocation and transfers. In the second step, the mechanism discloses information to other market participants in order to induce the optimal information structure in the aftermarket. The two parts of the problem interact non-trivially: The amount of information elicited by the mechanism determines the amount of information available for disclosure. Conversely, disclosure influences the incentives of agents to reveal their private information to the mechanism.

The paper addresses two questions: *(1)* How much information *can* be elicited and revealed (implementability), and *(2)* how much information *should* be elicited and revealed

[1] Some notable exceptions are discussed in the literature review.

to maximize the designer's objective function (optimal mechanisms).

The presence of an aftermarket exacerbates the well-recognized caveat of the mechanism design theory, which is that implementation of the optimal mechanism often hinges on the ability of the designer to fully exploit the details of the environment. Motivated by the Wilson doctrine, I introduce a class of mechanisms, called *cutoff mechanisms*, that are to some extent immune to this problem. My contribution is threefold. First, I prove several characterization theorems for the class of cutoff mechanisms, which provide reasons to restrict attention to this class. Second, I build a theory of optimal cutoff mechanisms. Third, I apply the theory to applications ranging from the optimal transparency of financial over-the-counter (OTC) markets to the design of auctions that are followed by bargaining or resale markets.

Suppose that the designer considers some allocation and disclosure rule. Implementing that rule is possible only if there exist transfers such that the resulting direct mechanism provides incentives for agents both to participate and to report truthfully. These incentives in the mechanism depend on the agents' values from acquiring the object, which are influenced by payoffs from the continuation game. Those, in turn, depend on the aftermarket protocol and the beliefs of aftermarket participants. Hence, the set of implementable allocation and disclosure rules in the first-stage mechanism varies with the aftermarket and the prior distribution of agents' types.

Cutoff mechanisms emerge from the analysis of allocation and disclosure rules, which I call *cutoff rules*, that are implementable regardless of the distribution of types and the form of the aftermarket. To understand how a cutoff mechanism works, consider allocating an object to one of $N \geq 1$ agents. In order to receive the object, an agent must outbid some threshold, which I refer to as the *cutoff* – be it a bid of another agent or a (possibly random and personalized) reserve price set by the seller. A cutoff mechanism only reveals information about the cutoff. That is, conditional on the type profile and the cutoff, the signal from a cutoff mechanism does not depend on the buyer's type. For example, if the object is allocated to the highest bidder in an auction, the relevant cutoff is the second highest bid, and the message sent by the auctioneer after the auction should only depend on the realization of the second highest bid.

The first main result of the paper establishes a surprising property that *only* mechanisms that implement a cutoff rule are incentive-compatible regardless of the distribution of types and the form of the aftermarket. This result implies that a cutoff mechanism is necessary if the designer lacks sufficient knowledge of these details, or for other reasons wants the mechanism to be implemented robustly. While transfers may need to depend on the distribution and the aftermarket, they can often be found endogenously in an indirect mechanism through equilibrium strategies of the agents.

Cutoff mechanisms have several other attractive properties. First, they are characterized by a monotonicity property, related to Myerson monotonicity, which makes them tractable for theoretical analysis. Second, they are the only mechanisms that provide good ex post incentives for agents in the sense that truthful reporting is a dominant strategy even *after* the signal is released to the aftermarket. Third, cutoff mechanisms correspond exactly to monotone equilibria of a class of simple dynamic auctions (resembling clock auctions) in which some garbling of the history of bidding is revealed.

In a companion paper, for a class of problems in which the third party takes a binary action, I derive conditions under which cutoff mechanisms are optimal. The conditions require that the preferences of the agent and the third party are sufficiently misaligned. This is satisfied, for example, in certain resale games, in which the initial stage is similar to a broker or wholesaler buying an object for resale.

The second major component of the paper is the analysis of *optimal* cutoff mechanisms. The amount of information that can be revealed by a cutoff mechanism depends on the allocation rule. For example, if the probability of winning the object is constant in the type of an agent, the cutoff is deterministic, and no information can be revealed by any cutoff mechanism. In contrast, when the probability of winning is strictly increasing in the type, disclosure of the cutoff may provide information about the type of the winner that affects aftermarket outcomes. This interplay between the allocation rule and disclosure results in optimal mechanisms that depend on the number of participating agents.

If the designer contracts with just one agent, a cutoff mechanism reveals information only about the realization of stochastic elements of the mechanism, such as a random reserve price. If the allocation rule is fixed, it may benefit the designer to disclose such (optimally garbled) information. However, if the allocation and disclosure rule are chosen jointly, a strong conclusion holds: For virtually any objective function that the designer may have and regardless of the aftermarket protocol, there always exists an optimal mechanism that reveals no information.

With more than one agent, it may be strictly optimal to disclose information also when the designer chooses both the allocation and the disclosure rule. Intuitively, this is because competition between agents can be used to elicit truthful reports even when the cutoff is informative about the winner's type. I provide sufficient conditions for optimality of simple mechanisms, such as an auction that allocates to the highest type subject to a reserve price, followed by either no disclosure, or revelation of the price paid by the winner.

The final contribution of the paper is to apply the methods and results described above to a number of real-life design problems. First, I consider optimal post-transaction transparency for bilateral transactions in financial OTC markets. A dealer (mechanism designer) chooses a mechanism to sell an asset to an intermediary (agent) who resells to a customer (third party) in the aftermarket. Neither the welfare-maximizing nor the profit-maximizing mechanisms release any signals. Thus, the model helps to explain why regulation that aims to increase post-transaction transparency may sometimes be undesirable.[2] I also consider an extension of the baseline model in which the initial seller has private information, such as information about the cost of providing the asset. In this case, the optimal mechanism only reveals information about that random cost. This result may be viewed as theoretical support for the use of financial benchmarks, such as LIBOR,[3] that are designed to disclose dealers' costs, without disclosing details of individual transactions.

The theory has policy implications for the design of trading platforms in the OTC market. Platforms allow the seller to elicit offers from multiple dealers. An important question concerns the extent to which these offers should be observed by other market participants.[4] If the dealer who buys the asset tries to resell it in the aftermarket, then in order to guarantee that it is always the most efficient dealer who trades (regardless of the distribution of dealers' costs of intermediation, or the bargaining protocol in the aftermarket), no information about that dealer's offer should be disclosed. However, the offers of dealers who do not trade can typically be disclosed without negatively influencing incentives, even if private information is correlated among dealers (as is likely, for example, when information pertains to the quality of the asset).

Finally, I consider an application to the design of auctions followed by bargaining or resale. The winning bidder engages in post-auction negotiations. For example, an agent who wins a contract in a procurement auction may negotiate with a subcontractor, or a winner in a spectrum auction may bargain with a third-party company over roaming agreements. The gains from negotiating depend on the type of the winner, which is initially unknown to the subcontractor or the third party. I show that optimal disclosure takes a simple form that depends on the structure of the aftermarket. Under certain conditions, the designer can implement the optimal mechanism without any knowledge of the distribution of agents' types, and with only limited knowledge about the aftermarket.

[2] Trading in many segments of the financial OTC market, including the market for some corporate bonds in the US, is subject to TRACE rules which require certain transaction data to be publicly disclosed.

[3] The London Interbank Offered Rate is the average of estimated interest rates at which large banks can borrow funds in the interbank money market.

[4] The SIFMA (2016) report shows a large variety in trading and disclosure rules adopted by US bond market platforms, especially regarding order visibility and the protection of the privacy of the winner.

The structure of the paper is summarized in the Table of Contents. The main ideas are first introduced in the context of a simple model, and then extended to a general setting. Some additional results are provided in an Online Appendix.

2 Literature review

This paper combines mechanism design with information design. In a seminal paper, Myerson (1981) solves the problem of allocating a single asset in a mechanism design framework. The designer is allowed to choose an arbitrary mechanism, subject to incentive-compatibility and participation constraints. In contrast, as surveyed by Bergemann and Morris (2016b), information design takes the mechanism (or game) as given and considers optimization over information structures. In the model that I study, the principal designs the mechanism and the information structure jointly.[5]

My analysis makes use of the concavification argument introduced by Aumann and Maschler (1995), and adopted to the Bayesian persuasion model by Kamenica and Gentzkow (2011). The argument establishes that in a large class of information design problems, optimization can be performed in the space of distributions over posterior beliefs, and thus the optimal payoff is a concave closure of the function that maps each posterior belief into the corresponding expected payoff for the designer. One of the main methodological contributions of my paper is to find a connection between the mechanism design problem and the concavification result via the introduction of cutoffs.[6] Because the model allows for an arbitrary aftermarket, the information design problem embedded in my problem is more general than Bayesian persuasion. Thus, the analysis is connected to multi-player generalizations of Bayesian persuasion, e.g. Bayes Correlated Equilibrium of Bergemann and Morris (2016a).[7]

With regard to the structure of the problem, the literature most closely related to my model is a series of papers by Calzolari and Pavan on sequential agency. In a sequential agency problem, the agent contracts with multiple principals, and upstream principals decide how much information to reveal to downstream principals. Thus, the allocation rule and the disclosure rule are designed jointly.

[5] To the best of my knowledge, the first paper to formally study a model with that structure is Myerson (1982) who provides a version of the Revelation Principle for an extended setting in which agents interact after the mechanism.

[6] Kolotilin, Li, Mylovanov and Zapechelnyuk (2015) combine mechanism design with Bayesian persuasion in a different context by studying a model in which the agent reports private information to the designer who then communicates her private information to the agent.

[7] Because signals are public in my model, the mechanism designer can only access a subset of BCE of the post-mechanism game.

A large part of the sequential agency literature deals with the problem of finding the correct revelation principle for that setting. If downstream principals do not observe the decisions made by upstream principals, the agent's message space may have to be enriched to include reports other than the type. This problem is addressed in Calzolari and Pavan (2008) and Calzolari and Pavan (2009), and in papers referenced therein. These issues are orthogonal to my theoretical analysis in Section 4 which models the aftermarket as a "black-box" game (a mapping from types and beliefs into final expected payoffs). However, because an agency problem is a special case of a game, the above considerations become relevant in solving examples of my model.

Calzolari and Pavan (2006b) show in a two-stage sequential agency model with one agent that, under certain conditions, it is optimal to reveal no information in the upstream mechanism. This conclusion is similar to my result about optimality of no-revelation in one-agent problems. However, the results are not related otherwise. None of the three economic assumptions of the main theorem of Calzolari and Pavan (2006b) are assumed in my analysis. For example, the upstream principal in Calzolari and Pavan (2006b) has no direct preferences over the outcome of the second stage. Although this is allowed by my model, I focus on exactly opposite cases when the principal cares about the final allocation (e.g. because she maximizes total surplus). Moreover, in my model, the preferences of the agent and the third party are typically not separable in the outcomes of the two stages.

Finally, Calzolari and Pavan (2006a) consider a model of a revenue-maximizing monopolist selling an object to an agent who can later resell to a third party. They study a simple setting with binary types which allows them to derive a closed-form solution. They show that it is sometimes optimal to distort the allocation and send explicit signals to influence the outcome of the second-stage game. Introducing cutoffs as a way to represent allocation and revelation in a mechanism provides a structural insight into the trade-offs in Calzolari and Pavan (2006a).[8] My model is more general in that it allows *(i)* an arbitrary objective function, *(ii)* multiple agents, *(iii)* arbitrary second-stage game, and *(iv)* an arbitrary discrete or continuous type space.

A large literature analyzes the consequences of resale after auctions (e.g. Gupta and Lebrun, 1999, Zheng, 2002, Haile, 2003, Hafalir and Krishna, 2008, Hafalir and Krishna, 2009, Zhang and Wang (2013)). The structure of the problem is similar to my model,

[8] Out of four mechanisms that can be optimal in the baseline model of Calzolari and Pavan (2006a) (depending on parameters), three are cutoff mechanisms, and one is not. Their paper and mine are complementary in that Calzolari and Pavan (2006a) provide an example of a problem where using a cutoff mechanism is suboptimal in the Bayesian sense, while this paper points out that implementability of their optimal mechanism relies on detailed knowledge of the setting. I comment on this further in the companion paper which establishes conditions under which cutoff mechanisms are optimal.

except that the second stage is a game between the bidders, rather than between the winning bidder and a third party. This makes the analysis of the problem qualitatively different. In the literature on auctions with resale, the revelation rule is either *(i)* made redundant by assuming an information structure in the resale stage (e.g. types are revealed, as in Gupta and Lebrun, 1999)[9], *(ii)* fixed for the purpose of the analysis (as in Haile, 2003 who assumes that all bids are revealed), or *(iii)* only relevant to the extent that it permits implementing the optimal allocation in an equilibrium of the auction (as in Zheng, 2002, where the optimal allocation and payoff are known ex-ante, and no revelation rule can increase the payoff of the mechanism designer). In contrast, the disclosure rule plays an active role in my model, and in particular interacts non-trivially with the optimal allocation rule. In a recent paper, Carroll and Segal (2016) consider a model where the auctioneer does not know the resale protocol and maximizes revenue in the worst case (the designer in my model maximizes a Bayesian objective function).[10]

A number of papers analyze the consequences of post-auction interaction between bidders. Zhong (2002), Goeree (2003), Katzman and Rhodes-Kropf (2008), and Zhang (2014) examine the effect of different bid announcement policies on revenue in standard auctions followed by Bertrand, Cournot, or other forms of competition. Lauermann and Virág (2012) consider a model where bidders exercise a common outside option after the auction. Giovannoni and Makris (2014) model the aftermarket through reduced-form reputational concerns. In the above papers, information disclosure has only local effects in the sense that the post-auction interaction does not rule out existence of a monotone equilibrium even when information is revealed. Dworczak (2015) analyzes the consequences of different bid disclosure rules in a setting where bidders trade other units of the object after the mechanism, and explicitly constructs non-monotone equilibria using a discrete type space. All of these papers (with the exception of Zhang, 2014, and one section in Dworczak, 2015) compare a small number of fixed auction formats (e.g. first-price, second-price) and announcement rules (e.g. full revelation of bids, revelation of the winning bid). Instead, this paper proposes a mechanism design approach in which the designer can choose an arbitrary allocation and disclosure rule.

The literature on information revelation in auctions is rich but has focused on disclosing information before or during the auction (seminal examples include Milgrom and Weber, 1982a, Milgrom and Weber, 1982b, and Eső and Szentes, 2007). My analysis focuses on ex-post revelation of information elicited by the mechanism.

[9] In Zhang and Wang (2013), one of the two potential buyers in the mechanism has a known value.

[10] The designer in Carroll and Segal (2016) does not control the information leakage to the aftermarket – all information is revealed before the resale stage.

The presence of aftermarkets has been cited as an important motivation for studying mechanisms (auctions) with allocative and informational externalities, for example in Jehiel, Moldovanu and Stacchetti (1996), Jehiel and Moldovanu (2001), and Jehiel and Moldovanu (2006). These papers focus on the case when all market participants are in the mechanism and hence do not explicitly consider the effects of ex-post information disclosure. I model the aftermarket as an interaction with third parties who are not in the mechanism which leads to a different structure and economics of the design problem.

3 Simple model of resale

In this section, I consider a simplified version of the model. There is one agent in the mechanism, and the aftermarket is a resale game with a single third party. The third party has a higher value than the agent, and the bargaining protocol is take-it-or-leave-it offers. This setting allows me to highlight the main ideas without obscuring the picture with additional details associated with the general case.[11] In Section 4, I extend the analysis to multi-agent mechanisms, continuous type spaces, and a general aftermarket.

In addition to highlighting the main ideas of the paper, the single-agent resale model is also of independent interest. In Section 5.1, I present an application to the issues of transparency and information disclosure in financial over-the-counter markets based on the model presented in this section (modified to incorporate some realistic features).

3.1 Model

A seller (mechanism designer) owns an indivisible object that she can allocate to an agent. The agent has value $\theta \in \Theta$ for holding the object, where Θ is a finite subset of non-negative real numbers. The agent's type is distributed according to a prior probability mass function f. If the agent acquires the object, she can resell it to a third party with value $v(\theta)$, where $v : \Theta \to \mathbb{R}$ is some non-decreasing function. I assume that $v(\theta) > \theta$, for all $\theta \in \Theta$.

The market game consists of two stages: *(1)* implementation of the mechanism, and *(2)* the aftermarket. In the first stage, the seller chooses and publicly announces a direct mechanism (x, π, t), where $x : \Theta \to [0, 1]$ is an allocation function, $\pi : \Theta \to \Delta(\mathcal{S})$ is a

[11] This simple model is similar to Calzolari and Pavan (2006a). I allow for more than two types of the agent, and interdependence in the value of the third party. Calzolari and Pavan (2006a) assume that the (private) value of the third party is binary, and may sometimes be lower than the value of the agent.

signal function with some finite signal set $\mathcal{S}$, and $t:\ \Theta \to \mathbb{R}$ is a transfer function.[12] If the agent reports $\hat{\theta}$, she receives the good with probability $x(\hat{\theta})$, and pays $t(\hat{\theta})$. Conditional on selling the good, the designer draws and publicly announces a signal $s \in \mathcal{S}$ according to distribution $\pi(\cdot|\,\hat{\theta})$.

In the second stage, if the good was allocated, the third party observes the signal realization s, and Bayes-updates her beliefs. The signal is the only source of information about the outcome of the mechanism for the third party.[13] I let f^s denote the updated belief over the agent's type,

$$f^s(\theta) = \frac{\pi(s|\,\theta)x(\theta)f(\theta)}{\sum_{\tau\in\Theta}\pi(s|\,\tau)x(\tau)f(\tau)},\ \theta \in \Theta.$$

With probability $\eta \in [0,\,1]$, the third party makes a take-it-or-leave-it offer to the agent, and with probability $1-\eta$, the agent makes a take-it-or-leave-it offer to the third party. If the good was not allocated in the mechanism, there is no aftermarket (this assumption is relaxed in Section 7.2).

Both the agent and the third party are expected-utility maximizers with quasi-linear utility. An equilibrium of the aftermarket is defined as a mixed-strategy Bayesian Nash equilibrium of the post-mechanism offer game described above. There could potentially exist multiple equilibria. In this section, I restrict attention to equilibria that maximize the probability of trade in the aftermarket.[14] In Section 4, I discuss how to handle equilibrium multiplicity more generally. Let $u(\theta;\,\bar{f})$ denote the continuation payoff of an agent with type θ conditional on acquiring the good in the first-stage mechanism, given posterior belief $\bar{f}$ held by the third party. For example, when $\eta = 1$ (the third party is the proposer) and $p^\star(\bar{f})$ is a profit-maximizing price quoted by the third party given belief $\bar{f}$, then $u(\theta;\,\bar{f}) = \max\{\theta,\,p^\star(\bar{f})\}$.

The payoff of the mechanism designer is normalized to zero if the good is not allocated, and is given by $V(\theta;\,\bar{f})$ otherwise, where $V:\ \Theta\times\Delta(\Theta)\to\mathbb{R}$ is assumed upper semi-continuous in the second argument. $V(\theta;\,\bar{f})$ is the expected payoff to the mechanism

[12] Focusing on direct mechanisms is without loss of generality, by the Revelation Principle. However, restricting attention to direct mechanisms will be consequential for some practical properties of mechanisms discussed in Section 3.6.

[13] It is irrelevant whether the third party observes that the agent acquired the good in the mechanism because she is going to condition on this event when making or accepting an offer in the aftermarket.

[14] If there are multiple equilibria with this property, I assume that the designer can make the selection, although this is not essential for the results as long as the selection is upper-hemi continuous. Without restricting off-equilibrium beliefs, the lemons market may have a continuum of equilibria with inefficiently low trade. For my analysis, it would also be enough to assume that a pure-strategy equilibrium is selected whenever it exists.

designer conditional on allocating the good, type of the agent being θ, and posterior belief $\bar{f}$ held by the third party in the aftermarket. For example, if $\eta = 1$, and the designer wants to maximize total surplus, we have $V(\theta;\, \bar{f}) = v(\theta)\mathbf{1}_{\{\theta \leq p^\star(\bar{f})\}} + \theta\mathbf{1}_{\{\theta > p^\star(\bar{f})\}}$. Although the payoff of the designer does not explicitly depend on transfers in the mechanism, the formulation includes expected revenue maximization as a special case. This is because $(\theta,\, \bar{f})$ pins down (via the underlying equilibrium of the aftermarket) the final allocation and hence transfers paid by the agent in the mechanism.[15]

3.2 Implementability

In order to find optimal mechanisms, I first characterize mechanisms that are feasible. I call a pair $(x,\, \pi)$, consisting of an allocation and disclosure rule, a *mechanism frame*.

Definition 1 (Implementability)**.** A mechanism frame $(x,\, \pi)$ is *implementable* if there exist transfers t such that the agent participates and reports truthfully in the first-stage mechanism, taking into account the continuation payoff from the aftermarket:

$$\sum_{s \in \mathcal{S}} u(\theta;\, f^s)\pi(s|\,\theta)x(\theta) - t(\theta) \geq 0, \tag{IR}$$

$$\theta \in \underset{\hat{\theta}}{\operatorname{argmax}} \sum_{s \in \mathcal{S}} u(\theta;\, f^s)\pi(s|\,\hat{\theta})x(\hat{\theta}) - t(\hat{\theta}), \tag{IC}$$

for all $\theta \in \Theta$.

Equations (IR) and (IC) are the standard participation and incentive-compatibility constraints. In a one-stage allocation problem (without the aftermarket), Myerson (1981) proves that x is implementable if and only if x is non-decreasing. Thus, the set of feasible mechanisms is independent of the distribution of types f which makes optimization tractable in many cases. With the aftermarket, whether $(x,\, \pi)$ is implementable or not depends on the prior distribution f. The distribution f impacts the continuation payoff $u(\theta;\, f^s)$ of the agent by affecting the equilibrium of the aftermarket through the posterior belief f^s held by the third party for any signal s.

[15] With a continuous type space, which is allowed by the general model of Section 4, this follows from the payoff equivalence theorem (see e.g. Milgrom, 2001). For example, with $\eta = 1$, revenue maximization corresponds to choosing

$$V(\theta;\, \bar{f}) = p^\star(\bar{f})\mathbf{1}_{\{\theta \leq p^\star(\bar{f})\}} + J(\theta)\mathbf{1}_{\{\theta > p^\star(\bar{f})\}},$$

where $J(\theta)$ is the virtual surplus function. I work with a finite type space in Section 3 so exact payoff equivalence does not hold. However, it is well known (see for example Dworczak and Zhang, 2015) that for a fixed $(x,\, \pi)$, the set of implementing transfers is a complete lattice with a unique highest element.

3.3 Cutoff mechanisms

In this section, I identify a class of allocation and disclosure rules that generalize the Myerson monotonicity property to the two-stage setting, and show that they are implementable regardless of the underlying distribution of types f. I first introduce the notion of a random-cutoff representation of an allocation rule. To simplify exposition, I assume throughout that $x(\max(\Theta)) = 1$, i.e. the highest type always receives the good (this assumption is relaxed in the general model).

3.3.1 Random-cutoff representation of an allocation rule

Fix a non-decreasing allocation rule $x(\theta)$ on Θ. Using the fact that $x(\max(\Theta)) = 1$, let c_x denote a random variable on Θ with cumulative distribution function $x(\theta)$.[16] By definition, $x(\theta) = \mathbb{P}(\theta \geq c_x)$. Thus, the allocation rule $x(\theta)$ can be implemented by drawing a cutoff from the distribution of c_x, and giving the good to the agent if and only if the reported type θ is greater than the realized cutoff. I will call c_x a *random-cutoff representation* of x. Let dx denote the probability mass function of c_x.[17]

Conversely, fix a random variable c distributed according to a cdf G with support on Θ. Then, $G(\theta)$ is a non-decreasing allocation rule on Θ, and $c = c_G$. That is, c is a random-cutoff representation of allocation G.

Summing up, there is a one-to-one correspondence between a subset of allocation rules and random cutoffs: Non-decreasing allocation rules are cdfs of random cutoffs.

For a fixed allocation rule x, the support of the random cutoff c_x is the set of points in Θ at which x increases strictly. In particular, degenerate (deterministic) cutoffs correspond to allocation rules that give the object with probability one to all types above some threshold. I let C denote the smallest space containing all possible realizations of cutoffs.[18]

3.3.2 Definition of a cutoff mechanism

In a cutoff mechanism, the signal distribution depends on the realization of the random cutoff representing the allocation rule, rather than on the reported type directly.

[16] I abuse notation slightly because cdfs are defined on the entire real line while x is defined on Θ. However, x can be trivially extended to a step function on the real line.
[17] I abuse notation by using dx (which is typically used for continuous measures) in a discrete setting. This allows me to keep consistent notation throughout all sections.
[18] Formally, C is the union of supports of c_x taken over all non-decreasing allocation rules x. C coincides with Θ under the special assumption $x(\max(\Theta)) = 1$.

Definition 2 (Cutoff rule and cutoff mechanism)**.** A mechanism frame (x, π) is a *cutoff rule* if x is non-decreasing, and the signal π can be represented as

$$\pi(s|\,\theta)x(\theta) = \sum_{c\leq\theta} \gamma(s|\,c)dx(c), \tag{3.1}$$

for each $\theta \in \Theta$ and $s \in \mathcal{S}$, for some signal function $\gamma : C \to \Delta(\mathcal{S})$.

A mechanism (x, π, t) is a *cutoff mechanism* if (x, π) is a cutoff rule.

In a cutoff mechanism (x, π, t), the agent reports θ, and the seller draws a cutoff c from distribution dx. If $\theta \geq c$, the agent acquires the good in exchange for a transfer, and the designer draws a signal to be announced from the distribution $\gamma(\cdot|\,c)$. If $\theta < c$, the agent does not acquire the good (it is irrelevant which signal is sent in this case). The signal is informative about the type of the agent because the third party conditions on the event that the agent acquired the good, i.e. $\theta \geq c$.

Although γ is defined on the entire set C, the properties of π depend only on how γ is defined on the support of dx. Signals sent conditional on realizations c with $dx(c) = 0$ are irrelevant because they occur with probability zero.

3.3.3 Characterization of cutoff mechanisms

Proposition 1 establishes a key property of cutoff rules.

Proposition 1. *A cutoff rule is implementable for any prior distribution of types f.*

A formal proof of Proposition 1 will be provided in the next section (see Lemmas 1 and 2). The key intuition is that under a cutoff rule the report of the agent does not directly influence the signal. The agent can change the outcome only by manipulating the probability with which she acquires the good. The continuation payoff $u(\theta;\,\bar{f})$ is non-decreasing in θ, for any underlying equilibrium of the aftermarket (this follows from a standard strategy-stealing argument, or direct inspection). This implies a single-crossing property in the first-stage mechanism and hence the existence of transfers that rule out profitable misreporting aimed at changing the probability of acquiring the good.

Implementability of cutoff rules regardless of the distribution f is reminiscent of why pivot mechanisms are dominant-strategy truthful. In a pivot mechanism, the report of an agent doesn't influence the transfer the agent pays, except when it changes the allocation. In a cutoff mechanism, the report doesn't influence the signal, except when it changes the allocation.

Under mild restrictions on the structure of the aftermarket, cutoff rules are the *only* mechanism frames that can be always implemented, regardless of the distribution f.

Proposition 2*a*. *Suppose $\eta > 0$ (the third party is sometimes the proposer). If (x, π) is implementable for every distribution of types f, then (x, π) is a cutoff rule.*

The assumption $\eta > 0$ implies that beliefs about the type of the agent influence the outcome of the aftermarket. If the agent always makes the offer, it is possible, for example when $v(\theta)$ is constant, that beliefs held by the third party are irrelevant. Then, each π is trivially implementable.

When $\eta = 0$, Proposition 2*a* remains true if the lemons problem is severe enough that beliefs about the type of the agent are always relevant, even if the prior distribution f is concentrated on a small subset of the type space. I say that the lemons condition *is locally severe* if $\max_{\hat{\theta}<\theta} v(\hat{\theta}) < \theta$, for all $\theta \in \Theta$.

Proposition 2*b*. *Suppose that $\eta = 0$, and the lemons condition is locally severe. If (x, π) is implementable for every distribution of types f, then (x, π) is a cutoff rule.*

To gain intuition for Propositions 2*a* and 2*b*, note that incentive-compatibility puts a constraint on the informativeness of implementable signal structures. Low types of the agent have a relatively stronger preference (relative to high types) for signals that lead to higher resale prices. However, a high resale price can only occur in the aftermarket under a signal that is more likely when the type of the agent is relatively high. If signals are too informative, this difference across types in willingness to pay for signals cannot be undone with transfers.

Cutoff rules are always implementable because they use the allocation function as a leverage. In particular, the amount of information that can be revealed in a cutoff mechanism depends on the allocation function through the distribution of cutoffs that it induces. If all types of the agent receive the good with the same probability (the allocation function is flat), the cutoff is degenerate and uninformative about the type. Propositions 2*a* and 2*b* imply that the conflict of incentives described in the preceding paragraph leads to impossibility of robust information disclosure in this case.[19] The "steeper" the allocation function, i.e. the higher the differences in probabilities of acquiring the good between high and low types, the more informative the realization of the cutoff is about the type of the agent. Consequently, more information can be revealed to the aftermarket without compromising incentives in the first-stage mechanism.

[19] In the companion paper, I show that this impossibility result is even stronger in that no information disclosure is possible for any fixed distribution of types f.

3.3.4 Proofs and discussion

To prove Propositions 1, 2*a*, and 2*b*, I state two important lemmas.

Lemma 1. *A mechanism frame (x, π) is a cutoff rule if and only if*

$$\pi(s|\,\theta)x(\theta) \text{ is non-decreasing in } \theta,\ \forall s \in \mathcal{S}. \tag{M}$$

Lemma 2. *If (x, π) satisfies condition (M), then (x, π) is implementable for every distribution of types f. Conversely, under the assumptions of Proposition 2a (or 2b), if a mechanism frame (x, π) is implementable for every distribution of types f, then (x, π) satisfies condition (M).*

Property (M) is a direct analog of the Myerson monotonicity condition from one-stage allocation problems. In the two-stage problem, implementability for every distribution requires that the probability of allocating the good and sending signal s is non-decreasing in the type, for any signal $s \in \mathcal{S}$. This is a stronger condition than Myerson monotonicity in that it implies that $x(\theta)$ is non-decreasing.[20]

Cutoff mechanisms possess the monotonicity property (M) directly by definition. If (x, π) is a cutoff rule, then it satisfies condition (M), and by Lemma 2, it is implementable for every distribution of types f, establishing Proposition 1.

Conversely, suppose that (x, π) is implementable for every distribution f, under the assumptions of Proposition 2*a* (or 2*b*). By Lemma 2, (x, π) satisfies condition (M). Lemma 1 implies that (x, π) is a cutoff rule, yielding Proposition 2*a* (and 2*b*).

Lemma 2 contains the key economic intuition for why *only* cutoff rules are implementable for every distribution of types, and thus I sketch its proof below. Lemma 1 is technical, and its proof is relegated to Appendix A.

Proof of Lemma 2. I first prove that under the assumptions of Proposition 2*a* (the analogous case 2*b* is considered in Appendix A.2), condition (M) is necessary for a mechanism frame (x, π) to be implementable for every distribution of types f. It is enough to show that $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$ for any two adjacent types $\theta > \hat{\theta}$.[21] Since (x, π) is assumed implementable for every distribution of types f, condition (IC) has to hold for θ and $\hat{\theta}$. In particular, type θ cannot find it profitable to report $\hat{\theta}$, and vice versa. When the two

[20] To see this, it is enough to add up $\pi(s|\,\theta)x(\theta)$ across all $s \in \mathcal{S}$. I show in the companion paper that if we only require implementability for a fixed distribution f, it is possible to implement some decreasing allocation rules $x(\theta)$ by appropriately disclosing information to the aftermarket.

[21] Formally, $\hat{\theta}$ is the largest type strictly smaller than θ (which exists due to finitiness of Θ).

resulting inequalities are added, transfers cancel out, and we obtain

$$\sum_{s\in\mathcal{S}} \left[u(\theta;\, f^s) - u(\hat{\theta};\, f^s)\right] \left[\pi(s|\,\theta)x(\theta) - \pi(s|\,\hat{\theta})x(\hat{\theta})\right] \geq 0. \tag{3.2}$$

In Appendix A.2, I prove existence of a distribution f with the following properties: *(i)* when the third party makes an offer, she offers price θ after seeing signal s if and only if $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$; otherwise, she offers price $\hat{\theta}$, *(ii)* when the agent makes an offer, and signal s satisfies $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$, trade takes place with probability one at a price above θ in equilibrium. The distribution f that achieves these two properties puts all mass on $\{\hat{\theta},\, \theta\}$, and is such that in the absence of additional information, the third party is indifferent between offering price θ and $\hat{\theta}$.

Given a distribution f with properties *(i)* and *(ii)*, we can observe that $u(\theta;\, f^s) = u(\hat{\theta};\, f^s)$ exactly when $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$ because both types resell with probability one. In the opposite case $\pi(s|\,\theta)x(\theta) < \pi(s|\,\hat{\theta})x(\hat{\theta})$, we have $u(\theta;\, f^s) > u(\hat{\theta};\, f^s)$ because with some strictly positive probability (equal to at least $\eta > 0$), only the low type $\hat{\theta}$ resells, at a price below the value of the high type θ. Inequality (3.2) becomes

$$\sum_{\{s\in\mathcal{S}:\, \pi(s|\,\theta)x(\theta)<\pi(s|\,\hat{\theta})x(\hat{\theta})\}} \alpha_s \left[\pi(s|\,\theta)x(\theta) - \pi(s|\,\hat{\theta})x(\hat{\theta})\right] \geq 0, \tag{3.3}$$

where $\alpha_s \equiv u(\theta;\, f^s) - u(\hat{\theta};\, f^s)$ is strictly positive for each s in the summation. We have obtained that a sum of strictly negative terms is non-negative. This is only possible when the set of indices in the sum is empty: $\{s \in \mathcal{S} : \pi(s|\,\theta)x(\theta) < \pi(s|\,\hat{\theta})x(\hat{\theta})\} = \emptyset$. Thus, condition (M) holds for every signal s.

To prove that condition (M) is sufficient for implementability for every distribution f, I use a condition for checking implementability in arbitrary type and allocation spaces from Dworczak and Zhang (2015).[22] Given a set of types and their final allocations, the assignment is implementable if and only if the matching between types and final allocations is efficient (see Dworczak and Zhang, 2015, for details and formal definitions). Because $u(\theta;\, f^s)$ is non-decreasing in θ for any f^s and any equilibrium in the aftermarket, matching efficiency is implied by pairwise stability – total surplus cannot be increased by swapping the allocations of some pair of types. Condition (3.2) is sufficient (and necessary) to ensure that θ and $\hat{\theta}$ cannot profitably swap their final allocations. Thus, it is enough to prove that inequality (3.2) holds for all $\theta > \hat{\theta}$. The fact that $u(\theta;\, f^s)$ is non-decreasing in θ

[22] See Rochet (1987) for the classical formulation of the implementability condition in arbitrary type spaces.

implies that the first square bracket is non-negative in each term of the sum in (3.2), and condition (M) implies that the second square bracket is non-negative. Thus, inequality (3.2) always holds, regardless of the underlying prior distribution f. □

Lemma 2 can be used to extend Proposition 1. The proof of sufficiency of condition (M) only used the fact that $u(\theta;\, f^s)$ is non-decreasing in θ. This implies that cutoff rules are implementable not only for every distribution of types but also for a large class of bargaining protocols in the aftermarket. For example, let $\mathcal{P} : \Delta(\Theta) \to \mathbb{R}$ be an arbitrary mapping from posterior beliefs of the third party into a deterministic price proposed to the agent in the aftermarket. Then, $u(\theta;\, f^s) \equiv \max(\theta,\, \mathcal{P}(f^s))$ is non-decreasing in θ.

Corollary 1. *A cutoff rule is implementable for any mapping $\mathcal{P}$.*

3.4 Optimal cutoff mechanisms

In this section, I study optimal mechanisms. Under the assumptions of Section 3.1, the designer maximizes

$$\sum_{\theta \in \Theta} \sum_{s \in \mathcal{S}} V(\theta;\, f^s)\pi(s|\,\theta)x(\theta)f(\theta), \tag{3.4}$$

over cutoff rules $(x,\, \pi)$.

I say that a mechanism frame $(x,\, \pi)$ reveals no information if every signal realization s is uninformative about the type of the agent: $\pi(s|\,\theta) = \pi(s|\,\hat{\theta})$ for all $\theta,\, \hat{\theta} \in \Theta,\, s \in \mathcal{S}$. The main result of this subsection establishes a strong conclusion about optimal cutoff mechanisms in the one-agent model.

Proposition 3. *The problem of maximizing (3.4) subject to $(x,\, \pi)$ being a cutoff rule has an optimal solution that reveals no information.*

The conclusion of Proposition 3 holds regardless of the objective function. The type of the objective may influence the shape of the optimal allocation rule x (which also influences beliefs in the aftermarket) but never requires the designer to make explicit announcements via π. The remainder of the section is devoted to building intuition and proving Proposition 3. An essential step in the proof is a characterization of optimal disclosure for a fixed allocation rule.

3.4.1 Optimization over disclosure rules

This subsection studies an auxiliary problem: I treat the allocation rule x as given, and optimize over disclosure rules π subject to $(x,\, \pi)$ being a cutoff rule.

An allocation function x can be represented by a random cutoff c_x. In a cutoff mechanism, the signal only depends on the realization of the random variable c_x. By Proposition 1, any cutoff rule is implementable. Thus, because the objective function does not depend explicitly on transfers, we can ignore constraints (IC) and (IR) in the optimization problem. The mechanism design problem becomes a pure communication problem in which the designer chooses a disclosure policy of the random cutoff in order to induce the optimal distribution of posterior beliefs. This problem is formally equivalent to the Bayesian persuasion problem formulated by Kamenica and Gentzkow (2011) and first analyzed by Aumann and Maschler (1995).

For any posterior belief G of the cutoff held by the third party (after observing some signal), let

$$f^G(\theta) \equiv \mathbb{P}_{c\sim G}(\theta|\,\theta \geq c) = \frac{G(\theta)f(\theta)}{\sum_\tau G(\tau)f(\tau)}, \tag{3.5}$$

be the corresponding posterior belief over the type of the agent (conditional on the agent acquiring the good). Equivalently, f^G is the belief over the type of the agent held by the third party who believes that the mechanism designer implemented the allocation rule $G(\theta)$. Next, let

$$\mathcal{V}(G) = \sum_{\theta\in\Theta} V(\theta;\, f^G)G(\theta)f(\theta) \tag{3.6}$$

be the conditional expected payoff to the mechanism designer conditional on inducing a posterior belief G over the cutoff. Equivalently, $\mathcal{V}(G)$ is the expected payoff to the mechanism designer that would arise if the allocation function were G (instead of the actual x) and the mechanism revealed no additional information to the third party.

Proposition 4. *For every non-decreasing allocation function x, the problem of maximizing (3.4) over π subject to $(x,\,\pi)$ being a cutoff rule is equivalent to*

$$\max_{\varrho\in\Delta(\Delta(C))} \mathbb{E}_{G\sim\varrho}\mathcal{V}(G) \tag{3.7}$$

subject to

$$\mathbb{E}_{G\sim\varrho}G(\theta) = x(\theta),\ \forall\theta\in\Theta. \tag{3.8}$$

Proposition 4 follows directly from the results of Kamenica and Gentzkow (2011) but I nevertheless provide a discussion of the proof in Appendix A.3.

Equation (3.7) means that the mechanism designer seeks to maximize her expected payoff over distributions ϱ over posterior beliefs G of the third party. Condition (3.8) is the so-called Bayes-plausibility constraint which states that the induced posterior beliefs over cutoffs have to average out to the prior belief (beliefs are represented by cdfs). Using the random-cutoff representation, the prior belief over cutoffs is simply the allocation rule x. Proposition 4 is further illustrated by Example 1.

Equations (3.7) - (3.8), together with the random-cutoff representation, yield an alternative interpretation of one-agent cutoff mechanisms. Any cutoff rule (x, π) can be represented as a probability distribution ϱ over implementable mechanism frames $(G, \emptyset)$ that reveal no information. Here, G is treated as a non-decreasing allocation rule, and $\emptyset$ represents no announcement. The designer draws a mechanism $(G, \emptyset)$ from the distribution ϱ, and announces which mechanism is used, but the mechanism itself reveals no information. Due to (3.8), the allocation rule x is implemented in expectation.

Proposition 4 implies that the concavification result from Bayesian persuasion can be applied to my setting. Let $\mathcal{X}$ be the set of all non-decreasing allocation functions on Θ.

Corollary 2. *The optimal expected payoff to the mechanism designer in the problem* (3.7)-(3.8) *is equal to the concave closure of* $\mathcal{V}$*,*

$$co\mathcal{V}(x) \equiv \sup\{y : (x, y) \in CH(graph(\mathcal{V}))\},$$

where CH denotes the convex hull, and $graph(\mathcal{V}) \equiv \{(\tilde{x}, \tilde{y}) \in \mathcal{X} \times \mathbb{R} : \tilde{y} = \mathcal{V}(\tilde{x})\}$*.*

I illustrate the above results with a simple numerical example.[23]

Example 1. Suppose that $\eta = 1$, $v(\theta) = 1$, for all θ, and f is the uniform distribution on $[0, 1]$. Let $x(\theta) = \beta$ for $\theta < 2/3$, and $x(\theta) = 1$ for $\theta \geq 2/3$. I consider a binary allocation rule because it leads to a binary distribution of the cutoff allowing a graphical analysis. The designer maximizes total surplus over disclosure rules using a cutoff mechanism.

The random-cutoff representation c_x of the allocation function x is a binary random variable with realizations in $\{0, 2/3\}$ and distribution $dx(0) = 1 - dx(2/3) = \beta$. Feasible posterior beliefs over cutoffs form a one-dimensional family indexed by the posterior probability that the cutoff is equal to 0 which I denote by α. For each α, let $x^\alpha(\theta) = \alpha$ for $\theta < 2/3$, and $x^\alpha(\theta) = 1$ for $\theta \geq 2/3$, be the corresponding cdf (which is also an allocation rule). The resulting posterior belief f^{x^α} over the type is given by (3.5) with $G(\theta) = x^\alpha(\theta)$, for all θ. Since the third party always makes the offer, the equilibrium in the aftermarket is summarized by $p^\star(x^\alpha)$ – the optimal price quoted by the third party given belief x^α over

[23] The example uses a continuous type space in order to simplify calculations. Section 4 formally shows that the results of this subsection hold with a continuous type space.

cutoffs.[24] $\mathcal{V}$ takes the form

$$\mathcal{V}(x^\alpha) = \int_0^1 \left(\mathbf{1}_{\{p^\star(x^\alpha)\geq\theta\}} + \mathbf{1}_{\{p^\star(x^\alpha)<\theta\}}\,\theta\right) x^\alpha(\theta) f(\theta) d\theta.$$

The function $\mathcal{V}$ is strictly concave and increasing in α on $[0,\,1/4]$, drops discontinuously at $\alpha = 1/4$, and is linear increasing on $[1/4,\,1]$ (see Figure 3.1). The concave closure coincides with $\mathcal{V}$ on $[0,\,1/4]$, and is linear on $[1/4,\,1]$.

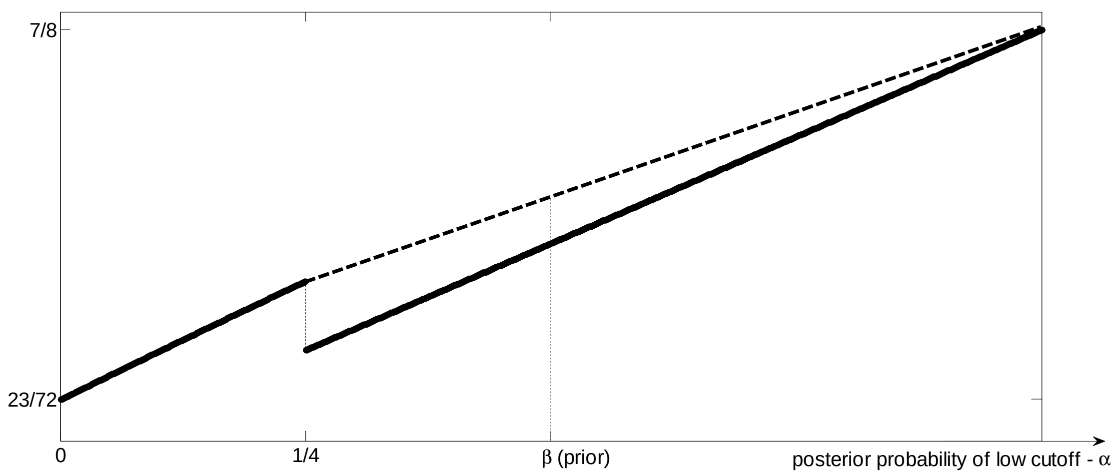

Fig. 3.1: Function $\mathcal{V}$ for Example 1 (solid line) and its concave closure (dotted line).

If $\beta \leq 1/4$, $\text{co}\mathcal{V}(x) = \mathcal{V}(x)$, and thus the optimal mechanism reveals no information. In the opposite case, $\text{co}\mathcal{V}(x) > \mathcal{V}(x)$, the unique optimal cutoff mechanism releases information in the form of a binary signal that induces a posterior belief $\alpha = 1/4$, or $\alpha = 1$, with appropriate probabilities. This mechanism frame can be implemented as follows. A cutoff c is drawn from the binary distribution dx on $\{0,\,2/3\}$. If $c = 0$, all types receive the good, and the mechanism discloses that $c = 0$ with conditional probability λ that solves $1/4 = (1-\lambda)\beta/[(1-\lambda)\beta + (1-\beta)]$. With the remaining probability $1-\lambda$, no message is sent. If $c = 3/4$, the good is awarded to all types above $3/4$, and no message is sent. The probability λ is such that conditional on no message, the posterior probability of $c = 0$ is exactly equal to $1/4$. ■

[24] By direct calculation,

$$p^\star(x^\alpha) = \begin{cases} \frac{5-2\alpha}{6} & \text{if } \alpha \leq 1/4 \\ \frac{1}{2} & \text{if } \alpha > 1/4. \end{cases}$$

3.4.2 Completion of the proof of Proposition 3 and discussion

With the characterization of the optimal payoff to the designer for a fixed allocation rule, the proof of Proposition 3 can be completed by an easy argument.

Proof of Proposition 3. By Corollary 2, the value to the designer at an optimal solution is $\sup_{x\in\mathcal{X}} \mathrm{co}\mathcal{V}(x)$. By definition of the concave closure, $\sup_{x\in\mathcal{X}} \mathrm{co}\mathcal{V}(x) = \sup_{x\in\mathcal{X}} \mathcal{V}(x)$. An optimal solution exists because $\mathrm{co}\mathcal{V}$ is an upper semi-continuous function on a compact set.[25] By definition, $\mathcal{V}(x)$ is the expected payoff to the mechanism designer when x is the allocation rule and the mechanism reveals no information. □

Recall that any cutoff mechanism can be interpreted as randomization over no-information-revealing, implementable mechanism frames $(G, \emptyset)$, with the public message disclosing which of the $(G, \emptyset)$ was used. This fact can be used to provide intuition for the above proof. Each $(G, \emptyset)$ induces a certain posterior belief in the aftermarket, and a conditional expected payoff to the mechanism designer. One of these mechanisms, denoted $(G^\star, \emptyset)$, must yield the highest conditional expected payoff. The designer can weakly increase her ex-ante expected payoff by choosing $(x, \pi) = (G^\star, \emptyset)$.

The fundamental logic behind Proposition 3 is that the mechanism design problem is a Bayesian persuasion problem in which the Sender (mechanism designer) chooses both the prior and the posterior beliefs (over cutoffs). Because the prior can be chosen, there is no need to induce a distribution over posteriors. Example 1 further illustrates this point.

Example 2. Consider the setting of Example 1. Suppose that the designer can additionally optimize over β, the probability that types below $2/3$ receive the good. This corresponds to optimizing over binary prior distribution of cutoffs on $\{0, 2/3\}$, and means that the designer can choose an arbitrary point on the concave closure of $\mathcal{V}$ in Figure 3.1. The expected payoff to the mechanism designer is maximized at $\beta = 1$. At $\beta = 1$, the function $\mathcal{V}$ coincides with its concave closure, so no-revelation is optimal.

Now suppose that the designer can choose an arbitrary allocation function x. By Proposition 3, no-revelation is optimal. It is enough to solve an unconstrained problem

$$\max_{x\in\mathcal{X}} \mathcal{V}(x) = \int_0^1 \left(\mathbf{1}_{\{p^\star(x)\geq\theta\}} + \mathbf{1}_{\{p^\star(x)<\theta\}}\,\theta\right) x(\theta) f(\theta) d\theta, \tag{3.9}$$

where

$$p^\star(x) \in \operatorname*{argmax}_p (1-p)\int_0^p x(\theta) f(\theta) d\theta. \tag{3.10}$$

[25] $\mathcal{X}$ is compact because it is a closed subset of $[0, 1]^\Theta$ which is compact by the Tychonoff's theorem.

The solution is $x(\theta) = 1$, for all θ, i.e. it is optimal to give the good to all types and reveal no information. The proof is non-trivial because (3.9) is an infinite-dimensional non-linear program due to the impact of x on the price $p^{\star}$. In Appendix A.4, I prove that giving the good to all types is optimal for a large class of distributions which includes the uniform distribution as a special case. ∎

Proposition 3 does not imply that the designer ignores the effect on the aftermarket when choosing the optimal mechanism. The choice of the allocation function, even in the absence of explicit signals, has a non-trivial impact on the information structure in the resale game because the third party conditions on the event that the agent acquired the good. In Example 2, it is optimal to give the good to all types which leads to a resale price of $1/2$ (when f is uniform). However, if the designer allocated only to types above a threshold r, the resale price would be $(1 + r)/2$. By not allocating to types $[0, r]$, the designer changes the information structure in the aftermarket, and induces beneficial trade between the third party and types $\theta \in (1/2, (1 + r)/2]$. I construct a distribution F for which a strictly positive r is optimal in Appendix A.4.

In Example 2, with the optimal allocation rule $x(\theta) = 1$, full privacy is not only optimal but in fact the only feasible information structure for the designer. This is because the cutoff representing a constant allocation rule is degenerate. However, in general, it need not be the case that no information can be revealed with the optimal allocation rule. In the Online Appendix, I provide an example (by changing the objective function of the mechanism designer) in which the optimal x leads to a non-degenerate cutoff distribution. Proposition 3 implies that it is nevertheless optimal not to disclose any information about the cutoff.

The above results are illustrated with an application to the optimal post-transaction transparency in a financial over-the-counter market in Section 5.1.[26]

3.5 Information structures induced by cutoff mechanisms

As demonstrated in the previous section, the designer can induce an arbitrary distribution of posterior beliefs over the cutoff, as long as posterior beliefs average out to the prior belief according to the the Bayes-plausibility constraint (3.8). However, the payoffs, and ultimately the form of the optimal mechanism, depend on the posterior beliefs over the *type* of the agent, rather than the cutoff.

[26] The formal model of Section 5.1 is a special case of the simple model analyzed in this section, so the details of the general model of Section 4 are not necessary to understand it.

For several reasons discussed below, it is interesting to know what restrictions are imposed on the distribution of posterior beliefs over the type when a cutoff mechanism is used. A Bayes-plausibility constraint, appropriately modified, is obviously necessary. Moreover, each posterior belief has to likelihood-ratio dominate the prior belief f. This can be seen from formula (3.5) – a posterior belief f^G is obtained by conditioning on the event that the type exceeded a cutoff which corresponds to multiplying the prior by a non-decreasing function (namely, the cdf G of the cutoff). I show in Appendix A.5 that these two conditions are also sufficient. If a distribution of posterior beliefs over the type satisfies the Bayes-plausibility constraint, and each posterior belief likelihood-ratio dominates the prior, then this distribution can be induced by a cutoff mechanism.

The above result implies an alternative characterization of the optimal expected payoff to the mechanism designer, expressed directly in terms of beliefs over the type of the agent (see Appendix A.5). In Section 4.3, I use it to provide sufficient conditions for optimality of no and full revelation of the cutoff in the multi-agent model. In Section 5.2, an application of the model to post-auction bargaining is solved by applying these results. In the Online Appendix, I show that the characterization facilitates comparing my model to the Bayesian persuasion model (in which the designer is a sender who observes the type of the agent).

3.6 Discussion

In this subsection, I discuss some practical issues related to cutoff mechanisms and their properties. Although multiple agents are introduced formally in the next section, the discussion also pertains to this case.

Definition 3 (Flexibility)**.** A mechanism frame (x, π) is *flexible* if (x, π) is implementable for every distribution of types.

Flexibility of a mechanism frame means that for every distribution of types, there exist transfers (which may depend on the distribution and other details) that implement it. Propositions 1, $2a$, and $2b$ establish conditions under which flexibility is the defining property of cutoff rules in the simple model of resale.

Three examples demonstrate the usefulness of this property in practical mechanism design problems. First, unlike in the traditional theoretical approach, many real-life situations require one mechanism to handle multiple instances of the problem. Consider designing informational requirements for a financial over-the-counter (OTC) market. The regulator cannot condition the design on the distribution of types and other details which might vary across different dealer-customer interactions. In a particular instance of the problem, a dealer in the OTC market (seller) might have a good estimate of the distribution of values of a visiting buyer (e.g. observes whether the buyer is an individual customer or a large hedge fund) but the regulator does not have access to that information. Flexibility means that the dealer can find prices that implement the recommended policy in every instance of the problem. Many auction houses use the same design across thousands of auctions for diverse items, despite the fact that the distribution of values clearly depends on the characteristics of a particular item.

Second, even if the mechanism is intended for a particular one-time problem, practical considerations often force the designer to design the mechanism in steps. In the design of large spectrum auctions[27], major parts of the mechanism have to be determined and fixed long before implementation to give time to regulators to approve it, and participants to understand it and voice concerns. Closer to implementation, as more information about the problem may arrive, minor adjustments are possible but the designer is committed to the major part of the design. In this context, flexibility can be seen as a modeling approach in which the mechanism frame is the major part, and transfers can be adjusted. Flexibility guarantees that the mechanism remains incentive-compatible even if new information about the agents' types arises after the mechanism frame is fixed.

Third, a desirable property of mechanisms is that truthful equilibria not only exist but are easily seen as being optimal. With this respect, failure of flexibility means that proving that a mechanism is truthful necessarily requires some knowledge of the underlying distribution of types. To the extent that the mechanism designer cannot be certain of some restrictions on possible distributions, mechanisms that are not in the class of cutoff mechanisms may be of limited usefulness in practice.

3.6.1 Robustness

Flexible implementation allows transfers to be a function of the distribution of types. For the sake of discussion in this paper, I informally define a stronger notion of *robust implementation*.

A mechanism frame (x, π) is *robustly implementable* if there exists an indirect mechanism, whose description does not depend on the distribution of types,[28] that implements (x, π) for every distribution of types.

[27] For example, the Incentive Auction in the US, see Milgrom and Segal (2017).

[28] In particular, I rule out mechanisms that ask agents to report the whole distribution (see Bergemann and Morris (2013) and Section 6.1 for a discussion). If such mechanisms are allowed, then flexible and robust implementation are equivalent.

Under robust implementation, the designer does not need to know the distribution of types to implement (x, π).[29] For a given distribution, by the Revelation Principle, any outcome of an indirect mechanism can be implemented by the corresponding direct mechanism. If an indirect mechanism implements (x, π) for every distribution, it has to be that there always exist transfers that implement (x, π) in a direct mechanism. Thus, flexibility is a necessary condition for robust implementation.

Robust implementability is a desirable property for practical purposes but may be cumbersome to work with in optimization. Flexibility is a much easier condition to work with, and it often turns out that the optimal cutoff mechanism can be implemented in a robust way (see Section 5.2 for a practical example, and Section 6.1 for a general discussion). In any case, robust mechanisms are a subclass of flexible (cutoff) mechanisms, so a designer who does not know the distribution has no reason to look beyond the class of cutoff mechanisms (but may need to restrict the class even further in some cases).

3.6.2 Direct versus indirect mechanisms

In discussing flexibility and robustness, I have focused on the case when the mechanism designer wants to implement the same mechanism frame (x, π) regardless of the distribution of types. For a fixed distribution f, focusing on mechanism frames as a primitive description of mechanisms is without loss of generality (by the Revelation Principle). However, when the mechanism is fixed but the distribution f varies, it might be natural to look at other representations of mechanisms. For example, in the case when the aftermarket is a resale game, the designer might want to implement the same final (ex-post) allocation regardless of the distribution of types. This approach could lead to a different theory of flexible mechanisms. The advantage of considering mechanism frames is that their description is independent of the description of the aftermarket. In Section 4, I allow for an arbitrary game in the aftermarket and the final allocation of the good may not be enough to describe the final outcome. Mechanism frames remain a valid description of the first-stage mechanism regardless of the form of the aftermarket.

Alternatively, a designer might want to fix an indirect mechanism, allowing the allocation and disclosure rule to be endogenously determined by the varying distribution f. Optimization in the class of all indirect mechanisms appears intractable, partly because it is not typically feasible to describe all their equilibria (especially that the aftermarket induces a signaling game). Traction can be gained by restricting the set of feasible indirect

[29] The notion of robustness used in this paper is still weaker than some other notions of robustness from the literature which might additionally require that the equilibrium is independent of the higher-order beliefs of agents, see for example Bergemann and Morris (2005).

mechanisms to a small class. See Dworczak (2015) for an example of such analysis in the context of post-auction bargaining.

4 General model

In this section, I extend the simple model of Section 3 by allowing *(i)* a general aftermarket, *(ii)* multiple agents in the mechanism, *(iii)* continuous distributions of types, and *(iv)* a more general definition of flexibility. The key assumptions maintained in the general model are that *(a)* the designer allocates a single object, and *(b)* only the agent who acquired the good participates in the aftermarket.

The mechanism designer owns an indivisible object that she can allocate to one of N agents. N also denotes the set of agents. If agent i acquires the object, she participates in the post-mechanism game described below. Agent $i \in N$ has a type $\theta_i \in [0,\, 1]$. Types are distributed according to a prior joint distribution with density $\boldsymbol{f}$ on $[0,\, 1]^N$, with marginals f_i. Let $\Theta_i = \mathrm{supp}(f_i)$, and $\boldsymbol{\Theta} \equiv \times_{i \in N} \Theta_i$. Throughout, bold symbols denote vectors, in particular $\boldsymbol{\theta} \equiv (\theta_1,\, \theta_2,\, ...,\, \theta_N)$ and $\boldsymbol{\theta}_{-i} \equiv (\theta_1,\, ...,\, \theta_{i-1},\, \theta_{i+1},\, ...,\, \theta_N)$. I consider two cases:[30]

1. Continuous distribution. Each Θ_i is a closed interval, and $\boldsymbol{f}$ is a density with respect to the Lebesgue measure on the Borel $\sigma-$field of $\boldsymbol{\Theta}$. I denote the set of all such distributions by $\mathcal{F}_c$.

2. Discrete distribution. Each Θ_i is finite, and $\boldsymbol{f}$ is a density with respect to discrete uniform measure on $\boldsymbol{\Theta}$. The set of all such distributions is $\mathcal{F}_d$.

I adopt the convention that for a set X, function g on X, and discrete density f on X,

$$\int_X g(x) f(x) dx \equiv \sum_{x \in \mathrm{supp}(f)} g(x) f(x).$$

With this notation, I will not distinguish between the case of continuous and discrete distributions, and unless explicitly stated, all claims pertain to both cases.

A direct mechanism is a tuple $(\boldsymbol{x},\, \boldsymbol{\pi},\, \boldsymbol{t})$, where $\boldsymbol{x} : \boldsymbol{\Theta} \to [0,\, 1]^N$ is an allocation function with $\sum_{i \in N} x_i(\boldsymbol{\theta}) \leq 1$, for all $\boldsymbol{\theta}$, $\boldsymbol{\pi} : \boldsymbol{\Theta} \to \Delta(\mathcal{S})^N$ is a signal function with signal space $\mathcal{S}$ (endowed with a respective $\sigma-$field) and $\boldsymbol{t} : \boldsymbol{\Theta} \to \mathbb{R}^N$ is a transfer function. If agent i reports $\hat{\theta}_i$, and other agents report truthfully, she receives the good with probability $x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})$ and pays $t_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})$. Conditional on allocating the good to agent i, the designer

[30] The theory to be developed works for more general distributions but the two cases seem to be sufficient for practical purposes and allow me to simplify exposition.

draws and publicly announces a signal $s \in \mathcal{S}$ according to distribution $\pi_i(\cdot|\,\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})$. No other signal is sent. To make sure that integrals are well-defined in the continuous case, I assume that $x_i(\boldsymbol{\theta})$ and $\pi_i(S|\,\boldsymbol{\theta})$ are continuous almost everywhere in $\boldsymbol{\theta}$, for any measurable set $S \subseteq \mathcal{S}$, for all i.[31]

If the distribution $\boldsymbol{f}$ is continuous, it is convenient to equate mechanisms that differ on a measure-zero set of type profiles. I will not distinguish between mechanisms $(\boldsymbol{x},\, \boldsymbol{\pi},\, \boldsymbol{t})$ and $(\boldsymbol{x}',\, \boldsymbol{\pi}',\, \boldsymbol{t}')$ if $\boldsymbol{x}(\boldsymbol{\theta}) = \boldsymbol{x}'(\boldsymbol{\theta})$, $\boldsymbol{\pi}(\cdot\,|\boldsymbol{\theta}) = \boldsymbol{\pi}'(\cdot\,|\boldsymbol{\theta})$, and $\boldsymbol{t}(\boldsymbol{\theta}) = \boldsymbol{t}'(\boldsymbol{\theta})$, for almost all $\boldsymbol{\theta}$ with respect to Lebesgue measure. Such mechanisms are identical from an ex-ante perspective for a Bayesian agent. Consequently, "for all types" should be interpreted as "for almost all types" when the distribution $\boldsymbol{f}$ is continuous. Profitable deviations are allowed for a measure-zero set of types of any agent.

I call a mechanism $\mathcal{S}$-*finite* if the signal space $\mathcal{S}$ is finite. Looking at $\mathcal{S}$-finite mechanisms is with loss of generality when $\boldsymbol{\Theta}$ is infinite but it greatly simplifies exposition and intuition. I consider mechanisms with infinite signal spaces in Section 4.1.4. For $\mathcal{S}$-finite mechanisms, $\pi_i(s|\,\boldsymbol{\theta})$ is well-defined as the probability of sending signal s conditional on agent i winning and report profile $\boldsymbol{\theta}$.

For sake of generality, I do not explicitly assume that there is a third party in the aftermarket. Instead, the post-mechanism game is described in reduced form by the conditional expected payoffs it generates given the information revealed by the mechanism. Formally, an aftermarket A is a collection of payoff functions

$$A \equiv \{u_i(\theta;\, \bar{\boldsymbol{f}}\,) :\; \theta \in \Theta_i,\; \bar{\boldsymbol{f}} \in \Delta(\boldsymbol{\Theta}),\, i \in N\},$$

where $u_i(\theta;\, \bar{\boldsymbol{f}}\,)$ denotes the conditional expected payoff to agent i with type $\theta \in \Theta_i$, when the posterior belief over the type profile $\boldsymbol{\theta}$ is $\bar{\boldsymbol{f}}$, conditional on agent i holding the good. For each i, $u_i(\theta;\, \bar{\boldsymbol{f}}\,)$ is assumed to be upper semi-continuous in $\bar{\boldsymbol{f}}$ (in the weak* topology on the space of distributions).[32]

I let $\mathcal{A}$ denote the set of possible aftermarkets. For example, $\mathcal{A}$ may include various versions of some post-mechanism game, differing in parameters governing the bargaining protocol or characteristics of the third-party players.

The "black-box" approach to modeling the aftermarket implicitly entails the following assumptions. A game is played after the mechanism between agent i who acquired the good (whose identity becomes known), and some number of third-party players. Third-party

[31] Continuity almost everywhere, i.e. except at a set of points of Lebesgue measure zero, implies measurability.

[32] Subsequent results continue to hold for a utility function u_i that depends on the entire type profile $\boldsymbol{\theta}$ rather than only on θ_i if dominant-strategy implementability is replaced with ex-post implementability.

players have the same prior belief $\boldsymbol{f}$ of the agents' types, and observe the public signal s sent by the mechanism which leads to a posterior belief over types denoted by $\boldsymbol{f}^s$. Given belief $\boldsymbol{f}^s$ and an aftermarket A, the corresponding game has a set of equilibria $EQ^A(\boldsymbol{f}^s)$, where $EQ^A(\bar{\boldsymbol{f}})$ is a hemi-continuous correspondence mapping beliefs $\bar{\boldsymbol{f}}$ into equilibrium outcomes, where the equilibrium notion can be specified by the modeler. Then, fixing an equilibrium selection from EQ^A, $u_i(\theta;\, \boldsymbol{f}^s)$ is the expected equilibrium payoff to type θ of agent i. A standard assumption in the mechanism design literature in such contexts is that the designer's preferred equilibrium is selected (the objective function of the mechanism designer will be defined in Section 4.3.) However, the theory works for any selection, as long as it generates payoff functions that satisfy the required regularity assumptions (such as upper semi-continuity of $u_i(\theta,\, \bar{\boldsymbol{f}})$ in $\bar{\boldsymbol{f}}$). By assumption, the signal s sent by the mechanism influences the aftermarket only through the posterior belief $\boldsymbol{f}^s$. Other roles of the signal (for example, as a coordination device) can be incorporated into the model by considering an appropriate equilibrium notion (e.g. a version of correlated equilibrium, see Bergemann and Morris, 2016a).

4.1 Implementability and cutoff mechanisms

In the multi-agent model, I first consider dominant-strategy implementation. Bayesian implementation is discussed in Section 4.2.

Fixing a mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi})$, let $\boldsymbol{f}^{i,s}$ denote the posterior joint belief over the profile of agents' types when agent i acquired the good, conditional on signal s, given prior joint belief $\boldsymbol{f}$, assuming truthful reporting. In an $\mathcal{S}$-finite mechanism, the marginal belief over the type of agent j when agent i is the winner and signal s was sent is

$$f_j^{i,s}(\tau) = \frac{\int_{\boldsymbol{\Theta}_{-j}} \pi_i(s|\,\tau,\, \boldsymbol{\theta}_{-j})x_i(\tau,\, \boldsymbol{\theta}_{-j})\boldsymbol{f}_{-j}(\boldsymbol{\theta}_{-j}|\,\tau)d\boldsymbol{\theta}_{-j}}{\int_{\boldsymbol{\Theta}} \pi_i(s|\,\boldsymbol{\theta})x_i(\boldsymbol{\theta})\boldsymbol{f}(\boldsymbol{\theta})d\boldsymbol{\theta}},\ \forall \tau \in \Theta_j, \tag{4.1}$$

where $\boldsymbol{f}_{-j}(\boldsymbol{\theta}_{-j}|\,\tau)$ is the conditional density of profile $\boldsymbol{\theta}_{-j}$ conditional on $\theta_j = \tau$, and $(\tau,\, \boldsymbol{\theta}_{-j})$ denotes the vector $\boldsymbol{\theta}$ with the jth coordinate θ_j replaced by τ. In particular, equation (4.1) defines the marginal belief $f_i^{i,s}$ over the type of the winner.

Definition 4. A mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi})$ is *dominant-strategy (DS) implementable* if there exist transfers $\boldsymbol{t}$ such that agents participate and report truthfully in the first-stage mechanism, taking into account the continuation payoff from the aftermarket:

$$\int_{\mathcal{S}} u_i(\theta_i;\, \boldsymbol{f}^{i,s})d\pi_i(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - t_i(\theta_i,\, \boldsymbol{\theta}_{-i}) \geq 0, \tag{\textbf{\textit{IR}}}$$

$$\theta_i \in \operatorname*{argmax}_{\hat{\theta}_i \in \Theta_i} \int_{\mathcal{S}} u_i(\theta_i;\, \boldsymbol{f}^{i,s}) d\pi_i(s|\, \hat{\theta}_i,\, \boldsymbol{\theta}_{-i}) x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i}) - t_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i}), \tag{$\boldsymbol{IC}$}$$

for all $i \in N$, $\theta_i \in \Theta_i$, and $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$.[33]

To define cutoff mechanisms for the general model, I let $C_i \equiv \Theta_i \cup \{\bar{\theta}_i\}$ be the space of cutoffs for agent i. The additional element $\bar{\theta}_i$ is an artificial type larger than any $\theta_i \in \Theta_i$. It is included to allow the possibility that the highest type does not receive the good with probability one. Suppose that the interim allocation rule $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i for any $\boldsymbol{\theta}_{-i}$, a property that is necessary and sufficient for implementability in the absence of the aftermarket (Myerson, 1981). A non-decreasing function is continuous almost everywhere, and thus there exists a non-decreasing, right-continuous $x_i'(\theta_i,\, \boldsymbol{\theta}_{-i})$ which differs from $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ on a measure-zero set of types θ_i. Because I equate mechanisms that differ on measure-zero set of types, I can without loss of generality assume that $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is right-continuous. Thus, $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ can be extended to a cumulative distribution function on C_i by defining $x_i(\bar{\theta}_i,\, \boldsymbol{\theta}_{-i}) = 1$. The random variable defined by this cdf is the random-cutoff representation of the marginal (interim) allocation rule $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ (see Section 3.3.1). I will denote the distribution of the random cutoff by $dx_i(\cdot\,,\, \boldsymbol{\theta}_{-i})$.

For any measurable function g on C_i, $\int g(c) dx_i(c,\, \boldsymbol{\theta}_{-i})$ denotes the Lebesgue integral of g with respect to the distribution of cutoffs induced by the interim allocation rule $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ on C_i. Because the allocation for agent i depends on the reports of other agents, the distribution of cutoffs depends on $\boldsymbol{\theta}_{-i}$.

Definition 5 (Cutoff rule and cutoff mechanism)**.** A mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi})$ is an $\mathcal{S}$-*finite cutoff rule* if $\mathcal{S}$ is finite, $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i for all $\boldsymbol{\theta}_{-i}$, and the signal function π_i can be represented as

$$\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}), \tag{4.2}$$

for each $\theta_i \in \Theta_i$, $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$, and $s \in \mathcal{S}$, for some measurable signal function $\gamma_i : C_i \times \boldsymbol{\Theta}_{-i} \to \Delta(\mathcal{S})$, for all $i \in N$.

$(\boldsymbol{x},\, \boldsymbol{\pi},\, \boldsymbol{t})$ is an $\mathcal{S}$-*finite cutoff mechanism* if $(\boldsymbol{x},\, \boldsymbol{\pi})$ is an $\mathcal{S}$-finite cutoff rule.

In the multi-agent setting, a cutoff mechanism reveals information about the random cutoff for every fixed report profile of other agents. Thus, the signal sent when agent i is the winner can depend on the reports of all other agents. If the random cutoff is degenerate conditional on $\boldsymbol{\theta}_{-i}$, for example $x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i \geq y(\boldsymbol{\theta}_{-i})\}}$ for some function

[33] For $\mathcal{S}$-finite mechanisms, $\int_{\mathcal{S}} u_i(\theta_i;\, \boldsymbol{f}^{i,s}) d\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) = \sum_{s \in \mathcal{S}} u_i(\theta_i;\, \boldsymbol{f}^{i,s}) \pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i})$.

y, then the signal distribution is determined solely by the profile $\boldsymbol{\theta}_{-i}$, i.e. $\pi_i(s|\,\theta_i,\,\boldsymbol{\theta}_{-i}) = \gamma_i(s|\,y(\boldsymbol{\theta}_{-i}),\,\boldsymbol{\theta}_{-i})$, for all $\theta_i \geq y(\boldsymbol{\theta}_{-i})$.

The next definition generalizes the flexibility property of cutoff rules (see Section 3.6). I use $\mathcal{F}$ to denote a generic subset of distributions.

Definition 6 (Flexibility)**.** A mechanism frame $(\boldsymbol{x},\,\boldsymbol{\pi})$ is flexible with respect to $(\mathcal{F},\,\mathcal{A})$, if $(\boldsymbol{x},\,\boldsymbol{\pi})$ is DS implementable for any prior distribution $\boldsymbol{f} \in \mathcal{F}$ and any aftermarket $A \in \mathcal{A}$.

My goal is to show that flexibility characterizes cutoff mechanisms in the general model. To do this, I need two definitions that generalize the structural properties of the simple aftermarket from Section 3.

Definition 7 (Monotonicity)**.** An aftermarket A is *monotone*, if for any agent $i \in N$, any belief $\bar{\boldsymbol{f}} \in \Delta(\boldsymbol{\Theta})$, the expected utility function $u_i(\theta;\,\bar{\boldsymbol{f}}\,)$ is non-decreasing in θ. The set of aftermarkets $\mathcal{A}$ is monotone if each $A \in \mathcal{A}$ is monotone.

The second definition is a richness condition which is an assumption about sets of prior distributions and aftermarkets.

Definition 8 (Richness)**.** The pair $(\mathcal{F},\,\mathcal{A})$ satisfies *Richness* if for any $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x},\,\boldsymbol{\pi})$, all $i \in N$, all types $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$, there exists a prior distribution $\boldsymbol{f} \in \mathcal{F}$ and an aftermarket $A \in \mathcal{A}$ such that

$$\pi_i(s|\theta_i,\,\boldsymbol{\theta}_{-i})x_i(\theta_i,\,\boldsymbol{\theta}_{-i}) < \pi_i(s|\hat{\theta}_i,\,\boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\,\boldsymbol{\theta}_{-i}) \implies u_i(\theta_i;\,\boldsymbol{f}^{i,s}) > u_i(\hat{\theta}_i;\,\boldsymbol{f}^{i,s}), \tag{4.3}$$

$$\pi_i(s|\theta_i,\,\boldsymbol{\theta}_{-i})x_i(\theta_i,\,\boldsymbol{\theta}_{-i}) > \pi_i(s|\hat{\theta}_i,\,\boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\,\boldsymbol{\theta}_{-i}) \implies u_i(\theta_i;\,\boldsymbol{f}^{i,s}) = u_i(\hat{\theta}_i;\,\boldsymbol{f}^{i,s}). \tag{4.4}$$

The meaning of Monotonicity and Richness is discussed in the next subsection.

4.1.1 Main results and discussion

The first theorem asserts that cutoff rules are flexible.

Theorem 1. *An $\mathcal{S}$-finite cutoff rule is DS implementable for any prior distribution $\boldsymbol{f}$ and any monotone aftermarket A.*

Under the Richness condition, the converse to Theorem 1 is also true.

Theorem 2. *Suppose that an $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x},\,\boldsymbol{\pi})$ is flexible with respect to $(\mathcal{F},\,\mathcal{A})$ which satisfies the Richness condition. Then, $(\boldsymbol{x},\,\boldsymbol{\pi})$ is a cutoff rule.*

Theorems 1 and 2 imply that flexibility is a defining property of cutoff mechanisms in Monotone and Rich settings. Their proofs follow from two lemmas (analogous to Lemma 1 and 2 from Section 3) which provide an alternative characterization of flexibility and cutoff mechanisms via a monotonicity condition.

Lemma 3. *An $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule if and only if*

$$\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) \text{ is non-decreasing in } \theta_i, \tag{$\boldsymbol{M}$}$$

for all $s \in \mathcal{S}$, $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$.

Lemma 4. *If $\mathcal{A}$ is monotone, and an $\mathcal{S}$-finite $(\boldsymbol{x}, \boldsymbol{\pi})$ satisfies condition ($\boldsymbol{M}$), then $(\boldsymbol{x}, \boldsymbol{\pi})$ is flexible with respect to $(\Delta(\boldsymbol{\Theta}), \mathcal{A})$. Conversely, if $(\mathcal{F}, \mathcal{A})$ satisfies Richness, and an $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is flexible with respect to $(\mathcal{F}, \mathcal{A})$, then $(\boldsymbol{x}, \boldsymbol{\pi})$ satisfies condition ($\boldsymbol{M}$).*

The proofs of the above results can be found in Appendix B.

I conclude with a discussion of the two assumptions, Monotonicity and Richness, in the context of Theorem 1 and Theorem 2. Monotonicity is a natural condition that ensures that the aftermarket does not revert the direction of the single-crossing property, i.e. that the willingness to pay for the object in the first stage is higher for high types than for low types. If Monotonicity fails, it is impossible to implement certain non-decreasing allocation rules even if no information is revealed.

Payoffs in the aftermarket have to be sensitive to information revealed by the mechanism if the existence of a post-mechanism stage is to put any restrictions on the set of implementable disclosure rules. However, this in itself is not enough. If the agent and a third party have aligned preferences, it is possible to support schemes that disclose the type of the agent exactly. The Richness condition requires a particular response of payoffs to information, for at least some distributions and aftermarkets. It implicitly assumes that there are conflicting interests between the agent and the third party. High types enjoy strictly higher payoffs than low types only if the third party believes low types to be more likely. The premise in condition (4.3) can be interpreted as "bad news" about the agent's type – after observing a signal s that satisfies the left-hand side inequality (assuming that $\boldsymbol{\theta}_{-i}$ is known), the posterior probability of the lower type $\hat{\theta}_i$ increases. Under some prior distribution $\boldsymbol{f}$ and aftermarket A, the expected payoff of the higher type θ_i has to strictly exceed the expected payoff of the lower type $\hat{\theta}_i$ when the mechanism sends "bad news". On the other hand, when the mechanism sends "good news" (condition 4.4), the expected payoffs of the two types should be equal.

For illustration, consider the resale game of Section 3 in the one-agent model, and assume that the third party makes the offer. The resale game satisfies the Richness condition (with $\mathcal{F} = \Delta(\Theta)$ and $|\mathcal{A}| = 1$) because for any $\theta > \hat{\theta}$, we can find a prior distribution f such that "bad news" induces price $\hat{\theta}$ in the aftermarket, and "good news" induces price θ in the aftermarket (see the proof of Lemma 2 in Appendix A.2). The payoff of type θ strictly exceeds the payoff of type $\hat{\theta}$ conditional on signal s if and only if type θ does not resell the good which happens exactly when the price is $\hat{\theta}$.

Remark 1. The Richness condition can be relaxed by allowing $\boldsymbol{f}$ to belong to the closure of $\mathcal{F}$, as long as u_i is continuous along some sequence of distributions $\boldsymbol{f}_n \in \mathcal{F}$ converging to $\boldsymbol{f}$. This often simplifies demonstrating Richness for continuous distributions.

4.1.2 Examples

I present two examples of settings that satisfy the above assumptions. The first example shows that the resale game of Section 3 can be extended to the multi-agent setting, and that it satisfies the Richness condition when $\mathcal{F}$ is the class of all independent distributions.

Example 3 (Resale)**.** Suppose that agent i who acquires the good in the mechanism plays a resale game with a third party, under the assumptions of Section 3 and $\eta > 0$. Assume (for now) that the value of the third party v is constant, and larger than $\max_i \max(\Theta_i)$. Let $\mathcal{F}$ be the set of all joint distributions $\boldsymbol{f} \in \Delta(\boldsymbol{\Theta})$ with independent marginals. Then, for any equilibrium selection in the resale game giving rise to expected payoff functions u_i, the aftermarket A is monotone, and with $\mathcal{A} = \{A\}$, $(\mathcal{F}, \mathcal{A})$ satisfies Richness. See Appendix B.3 for a proof.

If $(\mathcal{F}, \mathcal{A})$ satisfies Richness, and $\mathcal{F} \subseteq \mathcal{F}'$, $\mathcal{A} \subseteq \mathcal{A}'$, then $(\mathcal{F}', \mathcal{A}')$ also satisfies Richness. Consider the space $\mathbb{V} \times [0, 1] \times [0, 1]$, with typical element (v, η, λ). The function $v : \boldsymbol{\Theta} \to \mathbb{R}$ outputs the value of the third party as a function of the profile of types $\boldsymbol{\theta}$, and $\mathbb{V}$ is the set of functions that are non-decreasing in each variable. Parameter η is the probability that the third party makes the offer, and λ is the probability that the third party is present. Define $\mathcal{A}'$ as the set of aftermarkets indexed by $(v, \eta, \lambda) \in \mathbb{V} \times [0, 1] \times [0, 1]$. Then, $(\Delta(\boldsymbol{\Theta}), \mathcal{A}')$ satisfies Monotonicity and Richness. By Theorems 1 and 2, cutoff rules are DS implementable for any distribution $\boldsymbol{f}$ and any aftermarket $A \in \mathcal{A}'$, and no other mechanism frame has this property. ■

Example 4 (Buying a complementary good)**.** The mechanism designer sells an item to one of N bidders. For simplicity, assume that types of all agents come from the same type space Θ. The winner buys a second (complementary) good in the aftermarket (examples include buying infrastructure after winning a spectrum license, or subcontracting in order to complete a project after winning a procurement auction). A third-party seller quotes a monopoly price that the agent can accept or reject. If agent i with type θ_i acquires both goods, she obtains her full value θ_i. If she doesn't acquire the second good, she enjoys a reservation value $r(\theta_i)$, for some function $r : \Theta \to \mathbb{R}$ that is non-decreasing and satisfies $r(\theta) < \theta$ for all $\theta \in \Theta$. Let $\mathcal{R}$ be the set of all such functions r.

Suppose $\mathcal{F}$ is the set of all joint product distributions on Θ^N, and $\mathcal{A}$ includes aftermarkets corresponding to all $r \in \mathcal{R}$. Then, $\mathcal{A}$ is monotone (because for any price p quoted by the third party, $\max\{r(\theta), \theta - p\}$ is non-decreasing in θ), and $(\mathcal{F}, \mathcal{A})$ satisfies Richness. The proof, similar to the proof of the analogous claim in Example 3, is omitted. ■

4.1.3 Ex post incentives in cutoff mechanisms

The definition of DS implementability (Definition 4) requires that it is optimal for any agent to report truthfully for any fixed report profile of other agents, with expectation taken over the realizations of the signal. Instead, one can ask for a stronger notion of implementability where agents find it optimal to report truthully even *after* observing the signal realization. Because signals are determined endogenously, this requires a formal definition which I provide in Appendix B.4.

The consequence of Lemma 3 is that cutoff rules are implementable in this stronger sense. Moreover, under an intuitive condition on aftermarket payoffs, only cutoff rules have this property. Details are provided in Appendix B.4 which also discusses related implementability concepts from the literature.

4.1.4 Infinite signal spaces

In this subsection, I define and characterize cutoff mechanisms with infinite signal spaces. This case adds technical complications while offering few or no economic insights but is needed to formally consider full revelation in models with continuous type spaces.

Because the signal matters only to the extent that it influences posterior beliefs, it is without loss of generality to assume $\mathcal{S} = \boldsymbol{\Theta}$. Accordingly, $\mathcal{S}$ is endowed with the same Borel $\sigma-$field as the type space.

Definition 9 (Cutoff rule)**.** A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is a *cutoff rule* if $x_i(\theta_i, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i for all $\boldsymbol{\theta}_{-i}$, and the signal function π_i can be represented as

$$\pi_i(S|\,\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \gamma_i(S|\,c, \boldsymbol{\theta}_{-i})dx_i(c, \boldsymbol{\theta}_{-i}), \tag{4.5}$$

for each θ_i, $\boldsymbol{\theta}_{-i}$, and measurable $S \subseteq \mathcal{S}$, for some signal function $\gamma_i : C_i \times \boldsymbol{\Theta}_{-i} \to \Delta(\mathcal{S})$.

Cutoff mechanisms are defined accordingly. The only difference relative to an $\mathcal{S}$-finite cutoff mechanism is that the signal space $\mathcal{S}$ is allowed to be infinite, and condition (4.5) is required to hold for all measurable subsets S of $\mathcal{S}$.

In Appendix B.5, I show that a cutoff rule is DS implementable for all monotone aftermarkets and all prior distributions of types. Under a suitable richness condition, the converse conclusion holds. I also prove an approximation result: A mechanism frame is a cutoff rule if and only if it is the limit of $\mathcal{S}$-finite cutoff rules with the same allocation function. This result provides a formal tool to connect the analysis of $\mathcal{S}$-finite cutoff mechanisms to the general case.

4.2 The symmetric independent model

Having analyzed implementability in a general setting, I specialize to a symmetric independent model to present results on Bayesian implementation and optimal mechanisms. Independence and symmetry are not essential in studying the properties of cutoff mechanisms but they greatly simplify the analysis if the goal is to find a closed-form solution for an optimal design problem.

I assume that $\Theta_i = \Theta$, for all i, and the prior distribution $\boldsymbol{f}$ is a product distribution with identical marginals f, with cdf F. I assume that the aftermarket payoffs are independent of the identity of the winner, and only depend on the beliefs about the type of the winner, which I now denote by f^s (the subscript i is dropped because the belief is always over the type of the winner). With slight abuse of notation, I will thus write

$$u_i(\theta;\, \boldsymbol{f}^{i,s}) = u(\theta;\, f^s),$$

for all i and s. Because agents are ex-ante identical, it is without loss of generality to focus on symmetric mechanisms. A symmetric mechanism $(x,\, \pi,\, t)$ consists of mappings $x : \Theta^N \to [0,\, 1]$, $\pi : \Theta^N \to \Delta(\mathcal{S})$, and $t : \Theta^N \to \mathbb{R}$. For each i, $x(\theta_i,\, \boldsymbol{\theta}_{-i})$ is the probability that agent i with type θ_i receives the good,[34] $d\pi(\cdot \,|\, \theta_i,\, \boldsymbol{\theta}_{-i})$ is the probability distribution over signals conditional on agent i winning the object, and $t(\theta_i,\, \boldsymbol{\theta}_{-i})$ is the transfer paid by agent i, given reports $\boldsymbol{\theta}_{-i}$ of other players. It is without loss of generality to take $\mathcal{S} = \Theta$.

Given a symmetric mechanism frame $(x,\, \pi)$, its *reduced form* under distribution f is

[34] Since x is an allocation function, we have $\sum_{i \in N} x(\theta_i,\, \boldsymbol{\theta}_{-i}) \leq 1$, for all $\theta \in \Theta$.

defined by

$$x_f(\theta) = \int_{\mathbf{\Theta}_{-i}} x(\theta,\, \boldsymbol{\theta}_{-i}) \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) d\boldsymbol{\theta}_{-i},$$

and

$$\pi_f(S|\,\theta) x_f(\theta) = \int_{\mathbf{\Theta}_{-i}} \pi(S|\,\theta,\, \boldsymbol{\theta}_{-i}) x(\theta,\, \boldsymbol{\theta}_{-i}) \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) d\boldsymbol{\theta}_{-i},$$

for all measurable $S \subseteq \mathcal{S}$ and $\theta \in \Theta$.[35]

Definition 10. A symmetric mechanism frame $(x,\, \pi)$ is *Bayesian implementable* under distribution f if there exists a transfer function $t : \Theta \to \mathbb{R}$ such that its reduced form $(x_f,\, \pi_f)$ satisfies

$$\int_{\mathcal{S}} u(\theta;\, f^s) d\pi_f(s|\,\theta) x_f(\theta) - t(\theta) \geq 0, \tag{BIR}$$

$$\theta \in \operatorname*{argmax}_{\hat{\theta} \in \Theta} \int_{\mathcal{S}} u(\theta;\, f^s) d\pi_f(s|\,\hat{\theta}) x_f(\hat{\theta}) - t(\hat{\theta}), \tag{BIC}$$

for all $\theta \in \Theta$, where f^s is the posterior belief over the winner's type conditional on signal s.

Under Bayesian implementation, only interim expected allocation and signal functions matter to agents. The posterior belief f^s depends solely on the interim expected functions. If the designer's preferences (to be specified) do not depend on the posterior beliefs over types of agents who did not acquire the good, it is without loss of generality to represent mechanisms by their reduced forms. This is captured by the following definition of (Bayesian) equivalence of mechanism frames.

Definition 11 (Bayesian equivalence)**.** The symmetric mechanism frames $(x,\, \pi)$ and $(x',\, \pi')$ are (Bayesian) equivalent under f if they induce the same reduced form (up to relabeling of signals).

Instead of starting from a mechanism frame and deriving its reduced form, it is often convenient to work directly with reduced forms.

Definition 12 (Reduced-form mechanism)**.** A pair $(\bar{x},\, \bar{\pi})$ is a *reduced-form mechanism frame* under prior f if $\bar{x} : \Theta \to [0,\, 1]$, $\bar{\pi} : \Theta \to \Delta(\mathcal{S})$ for some signal space $\mathcal{S}$, and

$$\bar{x}(\theta) = \int_{\mathbf{\Theta}_{-i}} x(\theta,\, \boldsymbol{\theta}_{-i}) \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) d\boldsymbol{\theta}_{-i}, \tag{4.6}$$

for some joint allocation rule x.

[35] It is irrelevant how we define $\pi_f(S|\,\theta)$ for θ such that $x_f(\theta) = 0$.

A reduced-form mechanism frame $(\bar{x}, \bar{\pi})$ is only meaningful when we fix a prior distribution f. Condition (4.6) ensures that the interim expected allocation rule is feasible under prior f, i.e. induced by some joint allocation rule x.

Definition 13. A reduced-form mechanism frame $(\bar{x}, \bar{\pi})$ is a *reduced-form cutoff rule* if $\bar{x}(\theta)$ is non-decreasing in θ, and the signal function $\bar{\pi}$ can be represented as

$$\bar{\pi}(S|\,\theta)\bar{x}(\theta) = \int_0^\theta \gamma(S|\,c)d\bar{x}(c), \tag{4.7}$$

for each θ, measurable $S \subseteq \mathcal{S}$, for some signal function $\gamma : C \to \Delta(\mathcal{S})$.

A reduced-form cutoff mechanism is defined accordingly. Definition 13 is analogous to the definition of one-agent cutoff rules from Section 3.

Every N-agent cutoff rule induces a reduced-form cutoff rule. The converse result is also true.

Theorem 3. *For any symmetric cutoff rule (x, π), and any prior distribution f, (x_f, π_f) is a reduced-form cutoff rule. Conversely, for any reduced-form cutoff rule $(\bar{x}, \bar{\pi})$ under prior f, there exists a symmetric cutoff rule (x, π) such that $(x_f, \pi_f) = (\bar{x}, \bar{\pi})$.*

Theorem 3 has the flavor of a BIC-DIC equivalence result which states that in some settings Bayesian and dominant-strategy implementation are equivalent if mechanisms are identified by their interim expected allocations.[36] In my setting, the result says that if there exists a mechanism frame that induces a reduced-form cutoff rule (and hence is Bayesian implementable), then there exists an equivalent cutoff rule (which is hence dominant-strategy implementable). Compared to a standard BIC-DIC equivalence result, Theorem 3 strengthens the conclusion at the cost of strengthening the premise.

Using the famous characterization of reduced-form auctions developed by Matthews (1984) and Border (1991), I obtain the following corollary.

Corollary 3. *A pair $(\bar{x}, \bar{\pi})$ is a reduced form of some cutoff rule under f if and only if $\bar{x}(\theta)$ is non-decreasing in θ, the signal function $\bar{\pi}$ can be represented as in (4.7), and*

$$\int_\tau^1 \bar{x}(\theta)f(\theta)d\theta \leq \frac{1 - F^N(\tau)}{N}, \; \forall \tau \in \Theta. \tag{M-B}$$

[36] The classical reference is Manelli and Vincent (2010). Gershkov, Goeree, Kushnir, Moldovanu and Shi (2013) generalize the results to any one-dimensional setting. The main theorem of Gershkov et al. (2013) is not directly applicable in my model due to non-linear utilities but their proof technique can be used to prove Theorem 3.

Equation (M-B) is the so-called Matthews-Border condition that ensures that the interim expected allocation rule $\bar{x}$ can be induced by some joint allocation rule x under f.

A benefit of working with reduced forms is that the signal function becomes one-dimensional. In Section 4.1, I showed that a cutoff mechanism may reveal information about the entire vector $\boldsymbol{\theta}_{-i}$. The analysis of reduced-form mechanisms indicates that some of this information is redundant under the assumptions of Section 4.2. As an example, let the allocation rule be $x(\theta_i, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i > \boldsymbol{\theta}_{-i}^{(1)}\}}$, where $\boldsymbol{\theta}_{-i}^{(1)}$ denotes the first order statistic of $\boldsymbol{\theta}_{-i}$. Consider the reduced form of x under prior f, $x_f(\theta) = F^{N-1}(\theta)$. Under x_f, the distribution of the cutoff in a reduced-form cutoff mechanism is simply the distribution of the first order statistic of $\boldsymbol{\theta}_{-i}$, conditional on agent i acquiring the object.

Corollary 4. *If the allocation rule is given by $x(\theta_i, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i > \boldsymbol{\theta}_{-i}^{(1)}\}}$, then, up to Bayesian equivalence, any cutoff mechanism implementing x only reveals information about the second highest reported type.*

Corollary 4 could be proven directly by noting that $\boldsymbol{\theta}_{-i}^{(1)}$ is a sufficient statistic for the posterior belief over the type θ_i of the winner. Corollary 3 implies that such a one-dimensional sufficient statistic can be found for *any* non-decreasing allocation rule $x(\theta)$. Up to equivalence, a cutoff mechanism only reveals information about the reduced-form cutoff with distribution determined by the interim expected allocation $x_f(\theta)$.

4.3 Optimal cutoff mechanisms

In this section, I maintain the assumptions of Section 4.2, and consider optimization in the class of cutoff mechanisms. The objective of the designer is given by

$$\sum_{i \in N} \int_{\boldsymbol{\Theta}} \int_{\mathcal{S}} V(\theta_i;\, f_i^{i,s}) d\pi(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x(\theta_i,\, \boldsymbol{\theta}_{-i}) \boldsymbol{f}(\boldsymbol{\theta}) d\boldsymbol{\theta}, \tag{4.8}$$

where $V : \Theta \times \Delta(\Theta) \to \mathbb{R}$ is assumed to be bounded, measurable in the first argument, and upper-semi continuous in the second argument (in the weak* topology on $\Delta(\Theta)$). The payoff of the mechanism designer depends only on the type of the agent who acquires the good, and on the posterior belief over that type.[37] Exploiting the symmetry assumption,

[37] There are interesting situations that do not satisfy this assumption. As shown in Section 4.1, cutoff mechanisms can be always used in these cases but finding an optimal mechanism requires solving a multi-dimensional information disclosure problem which is typically intractable.

the objective function can be written equivalently as

$$N \int_{\Theta} \int_{\mathcal{S}} V(\theta;\, f^s) d\pi_f(s|\,\theta) x_f(\theta) f(\theta) d\theta, \tag{4.9}$$

where $(x_f,\, \pi_f)$ is the reduced form induced by $(x,\, \pi)$.

Due to the reduced-form representation, the results from Section 3 extend naturally to the multi-agent setting. To state the main result of this section, I generalize two definitions from Section 3. The interim allocation rule x_f, treated as a cdf, defines a prior distribution over cutoffs. Given posterior belief G over the cutoff (which can also be treated as an interim expected allocation rule), the belief over the type of the agent who acquired the good is given by the density

$$f^G(\theta) = \frac{G(\theta) f(\theta)}{\int_{\Theta} G(\tau) f(\tau) d\tau}. \tag{4.10}$$

The expected payoff to the mechanism designer conditional on inducing a posterior belief G over the cutoff is given by

$$\mathcal{V}(G) = N \int_{\Theta} V(\theta;\, f^G) G(\theta) f(\theta) d\theta. \tag{4.11}$$

Recall that $\mathcal{X}$ denotes the set of one-dimensional non-decreasing allocation rules on Θ.

Theorem 4. *The problem of maximizing (4.9) over the set of cutoff mechanisms is equivalent to solving*

$$\max_{\bar{x} \in \mathcal{X}} co\mathcal{V}(\bar{x}) \tag{4.12}$$

subject to

$$\int_{\tau}^{1} \bar{x}(\theta) f(\theta) d\theta \leq \frac{1 - F^N(\tau)}{N},\ \forall \tau \in \Theta. \tag{4.13}$$

If $N = 1$, the problem has an optimal solution that reveals no information.

The proof of Theorem 4 is established in two steps. In the first step, I consider optimization over disclosure rules for a fixed allocation function.

4.3.1 Optimization over disclosure rules

By Theorem 3, optimization over cutoff mechanisms can be performed directly in the space of reduced-form cutoff mechanisms. For a fixed allocation x and distribution f, a reduced-form cutoff mechanism is formally identical to a one-agent cutoff mechanism from

Section 3. Thus, we can use the proof of Proposition 4 to establish the following result (the fully analogous proof is omitted).[38]

Proposition 5. *For every allocation rule non-decreasing allocation rule x, the problem of maximizing (4.9) over π subject to (x, π) being a cutoff rule is equivalent to solving*

$$\max_{\varrho \in \Delta(\Delta(C))} \mathbb{E}_{G\sim\varrho}\mathcal{V}(G) \tag{4.14}$$

subject to

$$\mathbb{E}_{G\sim\varrho}G(\theta) = x_f(\theta),\ \forall\theta \in \Theta. \tag{4.15}$$

Applying the main result of Kamenica and Gentzkow (2011), I obtain the concave-closure characterization of the optimal payoff.

Corollary 5. *The optimal expected payoff to the mechanism designer in the problem (4.14)-(4.15) is equal to*

$$co\mathcal{V}(x_f) \equiv \sup\{y : (x_f, y) \in CH(graph(\mathcal{V}))\},$$

where $graph(\mathcal{V}) \equiv \{(\bar{x}, \bar{y}) \in \mathcal{X} \times \mathbb{R} : \bar{y} = \mathcal{V}(\bar{x})\}$.

4.3.2 Completion of the proof of Theorem 4 and discussion

The first part of Theorem 4 follows directly from Corollary 3, Proposition 5, and Corollary 5. Because mechanisms are represented by their reduced forms, the interim expected allocation rule $\bar{x}$ must satisfy the Matthews-Border condition (M-B) which is captured by constraint (4.13) in Theorem 4.

To establish the second part of Theorem 4, note that in the case $N = 1$, constraint (4.13) is vacuously satisfied. Thus, I can apply the proof of Proposition 3 from Subsection 3.4.2 – there always exists an optimal cutoff mechanism that releases no information.

When $N \geq 2$, constraint (4.13) alters the conclusion from the one-agent setting. It may be optimal to disclose information. This is because the concave closure of $\mathcal{V}$ is taken in the space of *all* non-decreasing interim allocation rules (equivalently, all posterior beliefs over cutoffs), while the actual rule $\bar{x}$ must be chosen from a subset of rules that satisfy the Matthews-Border condition (4.13). It might be optimal to induce posterior beliefs over cutoffs that do not correspond to an interim expected allocation that satisfies (4.13). For

[38] The only difference in the proof is that the space of cutoffs C is potentially infinite. The Online Appendix of Kamenica and Gentzkow (2011) extends their results to a continuous state space, so this technical complication does not cause any problems.

example, disclosing the second highest type $\theta^{(2)}$ in an efficient auction leads to a degenerate posterior belief over the cutoff represented by the cdf $G(\theta) = \mathbf{1}_{\{\theta \geq \theta^{(2)}\}}$. The function $G(\theta)$, treated as an interim allocation rule, does not satisfy the Matthews-Border condition. This means that some beliefs can only be induced by making explicit announcements.

4.4 Optimality of simple mechanisms

The predictions of mechanism design theory are particularly useful when the optimal mechanism takes a simple form, easy to implement in practice. In this section, I derive sufficient conditions for optimality of simple cutoff mechanisms.

The conditions build on the analysis of Section 3.5 (and Appendix A.5), where I characterized the set of feasible distributions of posterior beliefs over the type of the agent in a cutoff mechanism. Define $\mathcal{W}(\bar{f}) = \int_\Theta V(\theta;\, \bar{f})\bar{f}(\theta)d\theta$ as the conditional expected payoff to the mechanism designer given posterior belief $\bar{f}$ over the type of the winner, conditional on allocating the object. Because each posterior belief $\bar{f}$ has to likelihood-ratio dominate the prior f, the domain of $\mathcal{W}$ is the set of all distributions with this property.

In Appendix B.7, I show that regardless of the allocation function, if $\mathcal{W}$ is convex on its domain, it is optimal to disclose all information available under a cutoff mechanism. If $\mathcal{W}$ is concave, it is optimal to reveal no information.[39] In general, there exists a concave-closure characterization of the optimal payoff expressed in terms of $\mathcal{W}$.[40]

To characterize optimal allocation rules, I impose a simplifying assumption that the payoff in the aftermarket depends on the posterior belief only through its mean. Formally, let $M(\bar{f}) \equiv \int_0^1 \theta\bar{f}(\theta)d\theta$, and assume that $\mathcal{W}(\bar{f}) = W(M(\bar{f}))$ for some function $W : [0,\, 1] \to \mathbb{R}_+$. I also let $m(c) \equiv \int_c^1 \theta f(\theta)d\theta/(1 - F(c))$ denote the expected value of θ under the prior, conditional on $\theta \geq c$, and let $w(c) \equiv W(m(c))$, for any $c \in [0,\, 1]$. Thus, $w(c)$ is the conditional expected payoff to the mechanism designer conditional on allocating the good and inducing a belief that the type of the winner is above c. I also assume that f is continuous and fully-supported on $[0,\, 1]$.

Proposition 6. *1. If W is concave and non-decreasing, it is optimal to allocate the good to the highest type if it exceeds $r^\star$ (and to no one otherwise), and to reveal no*

[39] Molnar and Virág (2008) establish a similar result in a different model where full or no disclosure pertains to the type of the winner rather than the cutoff.

[40] These results do not follow directly from Kamenica and Gentzkow (2011) but they become a corollary of their proof technique thanks to the characterization from Section 3.5.

information, where

$$r^\star \in \underset{r\in[0,1]}{argmax} \left(1 - F^N(r)\right) W\left(\frac{\int_r^1 \theta dF^N(\theta)}{1 - F^N(r)}\right). \tag{4.16}$$

2. *If W is concave and decreasing, it is optimal to allocate the object uniformly at random and reveal no information.*

3. *Assume that W is differentiable, and let $J_w(c) \equiv w(c) - w'(c)\frac{1-F(c)}{f(c)}$. If (i) W is convex, and (ii) $J_w(c)$ is non-positive for $c \leq \underline{r}$, and positive non-decreasing for $c \geq \underline{r}$, then it is optimal to allocate the good to the highest type if it exceeds $\underline{r}$ (and to no one otherwise), and to disclose the second highest type (if the second highest type is below $\underline{r}$, it is enough to announce that the second highest type was below $\underline{r}$). A sufficient condition for property (ii) is that W is increasing and log-concave.*

Proposition 6 formulates three sufficient conditions for simple cutoff mechanisms to be optimal. First, if W is concave and increasing, it is optimal not to disclose any information, and the allocation rule is designed to maximize the posterior expected type of the winner. To do so, the mechanism allocates to the highest bidder. The mechanism can additionally raise the expectation by excluding types below r from trading. This incurs a utility cost because the good is not always allocated. The $r^\star$ that solves equation (4.16) optimally trades-off these two effects.

Second, if W is concave and decreasing, it is optimal to allocate the good randomly, with no disclosure. In this case, the designer wants to minimize the expectation of the type of the winner. However, it is not incentive-compatible to allocate to low types more often than to high types – hence the use of a uniform lottery.

Third, if W is convex, full disclosure of information (other than the type of the winner) is optimal.[41] It is enough to disclose the realization of a one-dimensional sufficient statistic which is the cutoff corresponding to the interim expected allocation rule. The optimal allocation rule is determined by the properties of the function $J_w(c)$ which captures the local trade-off between allocative efficiency and the induced information structure. (The function $J_w(c)$ is similar to the virtual surplus function which captures the trade-off between allocative efficiency and information rents.)

To gain intuition, suppose the designer starts from the fully efficient allocation (the good is given to the highest type), and considers not allocating the good conditional on

[41] The results in point *3* of Proposition 6 can be generalized beyond the case when the designer's objective function depends solely on the mean of the distribution. As long as it is optimal to fully disclose the cutoff, the function $w(c)$ can be treated as a primitive, and the same results apply.

cutoff realization ϵ-close to c (ignoring the monotonicity constraint on the allocation rule). If the type of an agent is above $c+\epsilon$, which happens with probability $1-F(c+\epsilon)$, then the designer is better off, because $w(c)$ jumps to $w(c+\epsilon)$ due to a higher belief over the agent's type.[42] But if the type of the agent is ϵ-close to c, which happens with approximate probability $\epsilon f(c)$, then the designer loses $w(c)$ because the good is not allocated. Thus, it is optimal to keep the original efficient allocation if (taking ϵ to zero) the loss $w(c)f(c)$ exceeds the gain $w'(c)(1-F(c))$, i.e. when $J_w(c)$ is positive. At the optimal allocation rule, the good is allocated precisely when $J_w(c)$ is positive, i.e. when the highest type exceeds $\underline{r}$.

To understand why $J_w(c)$ has to be non-decreasing for $c \geq \underline{r}$, suppose that the designer considers introducing randomization in some small region. Allocating the good randomly is beneficial from the perspective of the distribution of beliefs over the winner's type (the distribution of posterior beliefs becomes more dispersed which is desirable under a convex W) but only if the winner is above the randomization region. If the highest type happens to be in the randomization region, allocation is inefficient (which also leads to lower posterior beliefs conditional on the cutoff realization in the randomization region). By a similar analysis as above, it can be shown that when $J_w(c)$ is increasing in the neighborhood of c, the second effect dominates. The first effect is stronger when W is more convex. If W is log-concave and increasing, Proposition 6 asserts that second effect will always dominate, and hence the designer never finds it optimal to use randomization.

The above results are illustrated by an application in Section 5.2.

If W is neither convex nor concave, the problem can often be solved if W is sufficiently regular by applying the duality approach of Kolotilin (2016) or Dworczak and Martini (2016). These papers provide a general solution method for a class of Bayesian persuasion problems in which the preferences of the Sender only depend on the posterior mean. The results of Section 3.5 (and Appendix B.7) allow to formulate the design problem for a fixed allocation as a Bayesian persuasion problem that belongs to that class, with the caveat that posterior distributions have to likelihood-ratio dominate the prior.[43]

[42] Because I look at interim allocations, the analysis is done from the perspective of a single agent. The Matthews-Border condition ensures that everything can be translated to the multi-agent setting.

[43] This can be incorporated into the analysis of Dworczak and Martini (2016) by requiring that the distribution of posterior means put all mass on means above the prior mean. One can solve this relaxed problem, and check ex-post that the optimal distribution of posterior means can be induced by a distribution of posterior beliefs that likelihood-ratio dominate the prior belief.

5 Applications

In this section, I illustrate the theoretical results on cutoff mechanisms with two applications. In Section 5.1, I study the model of Section 3 (with a continuous type space) in the context of regulating a financial over-the-counter market. In Section 5.2, I consider designing an optimal auction followed by bargaining. Two additional applications are provided in the Online Appendix.

5.1 Regulating an OTC market

The transparency of financial OTC markets is an important topic in recent policy debates and in a growing body of theoretical and empirical literature (see for example Bessembinder and Maxwell, 2008, Asquith, Covert and Pathak, 2013, Duffie, Dworczak and Zhu, 2016 and Asriyan, Fuchs and Green, 2015). Transparency can be analyzed in the context of bilateral trade, but also as a design problem for trading platforms, such as Swap Execution Facilities (SEFs), or corporate bond market trade platforms.[44]

The analysis of cutoff mechanisms has practical implications for platform design. If the dealer who purchases an asset on the platform attempts to resell it in some aftermarket (for example in the inter-dealer market or a bilateral trade with an individual investor), the problem falls within the framework introduced in Section 4. A consequence of Theorems 1 and 2 is that there is a fundamental difference between disclosure of the winning offer and disclosure of the losing offers. For example, if dealers have a private cost of intermediation, and the platform is designed to allocate the asset to the most efficient dealer, then no information about the winner's offer should be revealed. However, the offers of dealers who do not trade may be revealed without distorting trade efficiency. This is true even if the information of the participating dealers is cross-sectionally correlated. For example, the incentive of dealers to trade according to their true information is not undermined by disclosing the average offer of all losing dealers. By disclosing the average offer, the regulator reveals information about the quality of the asset, which might reduce the adverse selection problem in the aftermarket.

In the remainder of this subsection, I focus on the problem of post-transaction transparency in bilateral trade settings. I solve for the welfare- and profit-maximizing mechanisms in a setting that is based on the simple model of Section 3.

[44] See SIFMA (2016) for some institutional details.

5.1.1 Model

The seller and the agent are dealers in an OTC market. If the agent acquires the asset, then with some probability $\lambda \in (0, 1]$ she has a chance to resell the asset to a third party (who is, for example, an individual investor or another dealer). The third party has a value for the asset that depends on the value of the agent. The third party does not observe the agent's private information but she can observe the signal revealed ex-post in the transaction between the dealers – the associated disclosure rule reflects the informational transparency of the market.

The seller has a cost $k \in [0, 1)$ of allocating the asset. I assume that the agent's type θ has a continuous distribution F with a regular density f on $[0, 1]$.[45] The value $v(\theta)$ of the third party is strictly increasing, with $v(0) < k$ and $v(1) > 1$ (these assumptions rule out uninteresting boundary solutions). Given the OTC market setting, it is more realistic to assume that the agent makes an offer (I assume that $\eta = 0$). The qualitative conclusions continue to hold when the third party has bargaining power. I restrict attention to pure-strategy equilibria in the post-mechanism game. Finally, I assume that when the seller allocates to all types above her cost, a price equal to 1 (which leads to trade with probability one) is not an equilibrium of the aftermarket:

$$\int_k^1 (v(\theta) - 1) f(\theta) d\theta < 0. \tag{5.1}$$

This assumption implies that the lemons problem in the aftermarket plays a non-trivial role in the design of information structures in the first-stage mechanism.

In the OTC market setting, it is useful to interpret the probability x as the quantity of a perfectly divisible asset. The initial seller has a quantity normalized to 1, and all players' values are linear in quantity.

5.1.2 Regulator's preferred mechanism

In this subsection, I assume that a regulator can choose a mechanism that the seller is obliged to implement. In the Online Appendix, I analyze a problem in which the regulator can only impose a disclosure rule. The regulator's objective is to maximize efficiency.

If the allocation rule is x, and the mechanism reveals no information, the price in the

[45] That is, the virtual surplus function $J(\theta) \equiv \theta - (1 - F(\theta))/f(\theta)$ is increasing.

aftermarket is

$$p(x) = \max\{p \in [0,\, 1] : \int_0^p (v(\theta) - p)x(\theta)f(\theta)d\theta \geq 0\}. \tag{5.2}$$

If the seller offers quantity 1 to all types above her cost, by assumption (5.1), the price is less than 1, and hence some gains from trade are lost due to the lemons problem. As long as the allocation rule is a threshold rule (allocates full quantity to all types above a threshold), no information can be revealed by a cutoff mechanism, because the random cutoff is degenerate (deterministic). Information can be revealed if the mechanism screens types by offering different quantities for sale. In this case, the distribution over cutoffs is no longer degenerate, and the mechanism sends messages of the form "quantity sold was at least x". Asquith et al. (2013) analyze consequences of introducing transaction reporting (TRACE) in the corporate bond market. TRACE forced dealers to reveal the price and exact quantity (up to a cap) immediately after each transaction (with some exceptions). The analysis of cutoff mechanisms implies that it is not always possible to reveal so much information without violating implementability of the underlying mechanism frame. At least in some cases, the dealers would be forced to use a mixed-strategy (if the IC constraint is violated), or to leave the market (if the IR constraint is violated) to protect their private information. This is consistent with the conclusions of Asquith et al. (2013), who show that the volume of trade went down in the informationally sensitive "speculative grade" segment of the market. They also provide anecdotal evidence that dealers found trading more difficult after TRACE was introduced, partly because they are worried about the impact of disclosed information on their bargaining position in the resale stage (aftermarket).

The designer faces a complicated trade-off between costly screening (reducing the quantity sold to lower types) and lower information asymmetry in the aftermarket. The trade-off would be difficult to resolve if not for Proposition 3, which guarantees the existence of an optimal mechanism that reveals no information.[46] Since no explicit announcements are made in the optimal mechanism, it is enough to solve the unconstrained problem

$$\max_{x \in \mathcal{X}} \int_0^{p(x)} [\lambda v(\theta) + (1-\lambda)\theta - k]\, x(\theta)f(\theta)d\theta + \int_{p(x)}^1 (\theta - k)x(\theta)f(\theta)d\theta, \tag{5.3}$$

where $p(x)$ is given by (5.2). The problem is still non-trivial because the objective function is non-linear in x, due to its impact on the price in the aftermarket.

[46] Formally, because the type space is continuous in this section, I apply the generalization of Proposition 3, Theorem 4 found in Section 4.3.

Claim 1. *Problem (5.3) admits a solution of the form* $x(\theta) = \mathbf{1}_{\{\theta \geq r^\star_{\mathit{eff}}\}}$ *for some* $r^\star_{\mathit{eff}} \in [0,\, 1]$.

The proof and the definition of $r^\star_{\text{eff}}$ can be found in Appendix C.1. The optimal scheme is a posted-price mechanism in which all quantity is offered for sale. The price is chosen by optimally trading off the losses caused by not allocating the asset to low types against the higher realized gains from trade in the aftermarket.[47]

5.1.3 Seller's preferred mechanism

I now analyze how the market outcome differs in the absence of regulation, that is, when the seller chooses a profit-maximizing cutoff mechanism. In Appendix C.2, I show that also in this case the objective function takes the form (3.4). Applying Proposition 3 (or Theorem 4) again, I obtain the following result.

Claim 2. *The profit-maximizing mechanism reveals no information and allocates the good to all types above a threshold:* $x(\theta) = \mathbf{1}_{\{\theta \geq r^\star_{\mathit{rev}}\}}$ *for some* $r^\star_{\mathit{rev}} \in [0,\, 1]$.

Moreover, $r^\star_{\mathit{eff}} \leq r^\star_{\mathit{rev}}$ *with equality if and only if* $p(\mathbf{1}_{\{\theta \geq r^\star_{\mathit{eff}}\}}) = 1$ *and* $\lambda \geq \lambda^\star$ *for some* $\lambda^\star < 1$*. That is, the welfare- and profit-maximizing mechanisms coincide when trade in the efficient mechanism occurs in the aftermarket with probability one conditional on the third party being present and the third party is present with sufficiently high probability.*

The proof and the calculations of $r^\star_{\text{rev}}$ and $\lambda^\star$ can be found in Appendix C.3.

5.1.4 Discussion

Both the efficient and the profit-maximizing schemes are posted-price mechanisms that release no information. The threshold (price) can be either *(i)* strictly lower for the efficient mechanism, or *(ii)* the same in both mechanisms.

In the less surprising case *(i)*, the profit-maximizing seller excludes types $\theta \in [r^\star_{\text{eff}},\, r^\star_{\text{rev}}]$ in order to reduce the information rents of the agent. This lowers total surplus. In case *(ii)*, the mechanism chosen by the seller also achieves maximal total surplus. There is no trade-off between efficiency and revenue. Such an outcome is possible when the lemons problem is not too severe, so that the probability of trade in the aftermarket is high (explaining why a high λ is needed). In this case, the agent acts primarily as an intermediary, and thus the seller can extract rents without excluding low types. In Appendix C.4, I illustrate the above discussion with a numerical example.

[47] A related trade-off is considered in the model of Leland and Pyle (1977) where a firm is privately informed about its value and decides about optimal level of debt. Suboptimal debt structure is used as a signaling device to alleviate the lemons problem in the capital market.

The above analysis implies that transparency is not necessarily desirable from an efficiency viewpoint. The result should be properly interpreted. I only considered revealing information that the seller does not initially have, namely the private information of the agent elicited by the mechanism. Duffie et al. (2016), in a different model, analyze an advance commitment to disclosure of information initially controlled by the seller, and show that transparency typically improves welfare. In Section 7.1, I extend the model to allow for exogenous information of the seller, and show that information disclosure may be optimal. However, only exogenous information is revealed in the optimal mechanism.

In the optimal mechanisms studied above the designer often excludes low types from trading to reduce the information frictions in the aftermarket. This is conceptually related to the idea that public intervention can effect the market outcome by changing the information structure. A series of papers by Philippon and Skreta (2012), Tirole (2012), Fuchs and Skrzypacz (2015) analyze the problem of overcoming adverse selection in a market by an intervention that alleviates the lemons problem.

5.2 Post-auction negotiations

The mechanism designer designs an auction followed by post-auction bargaining. There are $N \geq 2$ symmetric bidders (agents) whose types are distributed independently according to a full-support continuous density on $[0, 1]$. After the object is allocated, the winner negotiates with a third party in the aftermarket. Examples include subcontracting after winning a government contract in a procurement auction, bargaining over roaming agreements after spectrum auctions, and resale.

For the initial part of the analysis, I assume that the object must be allocated to the highest bidder in the first stage, and I focus on optimal disclosure. In some settings, such a restriction may be a consequence of institutional constraints. I give conditions for optimality of this allocation rule, as well as alternative designs, at the end of the section.

I consider a tractable reduced-form approach to the bargaining game. Having observed the signal released by the mechanism, the third party has to pay a cost k in order to start negotiations. If the cost is paid, negotiations lead to a total surplus of $\delta\theta+\Delta$, where θ is the type of the winner. The third party captures a fraction $\eta \in (0, 1)$ of the additional surplus generated by the negotiations (η is the Nash bargaining parameter). In the opposite case, negotiations do not take place, and total surplus is equal to the value of the winner θ. One interpretation of the game is that the third party incurs a due diligence cost – by paying the cost k she learns the value of the winner and the subsequent negotiation stage is a full-information Nash bargaining game. I assume that $\delta\theta + \Delta \geq \theta$, for all θ, $\delta > 0$, and that the cost k is random from the point of view of the first stage, distributed according to a cdf H on $\mathbb{R}_+$.[48]

The auctioneer maximizes the total expected value of allocating the good[49] over cutoff mechanisms that allocate to the highest type in the first stage.

Claim 3. *Let* $\bar{k} = \eta \max\{\Delta, \Delta + \delta - 1\}$. *If* $H(x)x$ *is convex on* $[0, \bar{k}]$, *it is optimal to disclose the second highest value. If* $H(x)x$ *is concave on* $[0, \bar{k}]$, *it is optimal to disclose no information.*

In the above problem, the function from beliefs over *cutoffs* into payoffs is complicated, but the function from beliefs over *types* into payoffs has a tractable structure. The belief over the type only influences the payoffs through the expected value of the winner. The proof of Claim 3 becomes an easy application of the sufficient conditions from Section 4.4.

Using a carefully constructed auction, the designer can implement the optimal mechanism robustly, in the sense defined in Section 3.6. Details of this construction and a discussion are provided in Section 6.1. Transfers are pinned down by the equilibrium bidding strategies of agents. The designer does not need to know anything about the distribution of types (or details of the aftermarket) to guarantee that the mechanism is truthful. A guarantee of optimality can be obtained if the designer knows enough about the distribution H – for example that $H(x)x$ is convex.

To gain intuition for Claim 3, note that the problem of the designer is to induce an optimal distribution of the posterior expected type of the winner by disclosing information about the cutoff that, in this case, corresponds to the second highest type. The third party is willing to pay the cost k if and only if it does not exceed the expected conditional gain from negotiating $\eta((\delta - 1)\mathbb{E}[\theta|\, s] + \Delta)$, where $\mathbb{E}[\theta|\, s]$ is the expected type of the winner conditional on signal s. The designer wants to increase the probability that negotiations take place. If H is convex, probability increases when beliefs are more dispersed. If H is concave, the highest probability of negotiating is obtained by pooling all information. However, the objective is to maximize the value. High probability of negotiations is most valuable when the social gains from negotiating are highest. This additional effect favors information revelation because disclosure introduces positive correlation between the gains from negotiating and the probability that negotiations take place. For optimality of full

[48] If $\mathcal{F}$ is the set of independent distributions, and $\mathcal{A}$ includes aftermarkets corresponding to all possible values of the parameters (Δ, δ, η, and H), then the setting satisfies Monotonicity and Richness. By Theorems 1 and 2, cutoff rules are characterized by flexibility.

[49] I exclude the private cost of the third party from the objective function of the designer. If the designer cares about the cost incurred by the third party, the analysis changes slightly, and in particular the optimal mechanism can depend on the bargaining parameter η.

disclosure of the second highest type, it is enough that H is not "too concave," so that $H(x)x$ is convex. If H remains concave after multiplying by x, the first effect dominates and no revelation is optimal.

I now comment on the optimality of allocating to the highest type in the first stage. I offer an informal discussion but all claims can be formalized by a simple application of Proposition 6.

First, if $H(x)x$ is concave, and $\delta > 1$, it is optimal to allocate to the highest type (conditional on allocating) to maximize the probability of negotiations (when $\delta > 1$, the third party is more willing to negotiate when she believes the type of the winner to be high). If Δ is large, so that there's a large social benefit from negotiating, the designer might find it optimal to set a reserve price in the first stage, and decrease the probability of allocating the object to induce a higher chance of negotiations conditional on allocating. When $\Delta = 0$, this last effect is absent, and it can be shown that a fully efficient auction is optimal in the first stage.

When $\delta < 1$, negotiations are more likely to happen when the third party believes that the type of the winner is low. To minimize the expectation of the winner's type, the designer can allocate the object randomly. However, allocating randomly decreases surplus in the event when negotiations fail (because in this case the final value is equal to the type of the winner). When Δ is large relative to δ, the first effect dominates, and it can be shown that it is optimal to allocate the object randomly.

Finally, suppose that $H(x)x$ is convex. To show that allocating to the highest type is optimal, I have to find conditions under which J_w, defined in Proposition 6, is positive non-decreasing. The properties of J_w depend on the local properties of the distributions F and H. Suppose that these distributions are "well-behaved" (it is enough if their densities are bounded and continuously differentiable). Then, two sufficient conditions can be given.

First, optimality can be established when η is sufficiently small. When the third party does not have too much bargaining power, for a fixed distribution of costs H, negotiations are not very likely to happen. The designer's preferences over allocation rules are driven mainly by the event that negotiations fail (in which case the total value is equal to the value of the winner, and it is trivially optimal to allocate to the highest-value bidder). Second, it is optimal to allocate efficiently in the first stage when δ is sufficiently close to 1. In this case, the effect that posterior beliefs have on the probability of negotiating is relatively weak, and eventually dominated by the effect that the allocation rule has on the value of the winner, as δ gets closer to 1 (see the discussion below Proposition 6 in Section 4.4).

Under weaker assumptions, it is optimal to allocate to the highest type subject to a reserve price (and then reveal the relevant cutoff). Using the last statement in Proposition 6, it is enough if $\delta > 1$ and the function $x + H(x)x$ is log-concave.

6 Implementation

I have defined and analyzed cutoff mechanisms as direct mechanisms. Although direct mechanisms are convenient for theoretical analysis, they are rarely used in practice. In Section 6.1, I derive conditions under which simple cutoff rules can be implemented robustly by a standard first- or second-price auction. In Section 6.2, I show that cutoff mechanisms can be characterized as monotone equilibria of a class of dynamic auctions (called Generalized Clock Auctions). Throughout, I consider the symmetric independent model of Section 4.2, and assume that the aftermarket is monotone.

6.1 Robust implementation by standard auctions

Suppose that the designer does not know the distribution of types and details of the aftermarket but agents do.[50] Is it still possible to design a mechanism that implements a cutoff rule? In this section, I study cases in which an answer to this question is positive. I assume $N \geq 2$, and work with a continuous distribution of types F on $\Theta = [0, 1]$.

It is known (see for example Bergemann and Morris, 2013) that using direct mechanisms is not without loss of generality when agents have more information than the designer. For example, if agents are ex-ante identical and arbitrary indirect mechanisms are allowed, the designer can elicit information about the distribution by asking agents to report it, and punishing if reports disagree. There are serious concerns about practical applicability of such schemes. I focus on simple indirect mechanisms with a one-dimensional message space, such as standard auctions.

In general, robust implementation of a cutoff rule in an auction with a one-dimensional bid space may be impossible. If the equilibrium bidding function is an injection on the set of types, these exists a mapping from bids into signals that induces the direct disclosure rule π. However, this mapping varies with f, so if the designer does not know f, she may be unable to invert the bidding function to recover it.

There are important cases in which the inversion need not be done. Suppose that x allocates the object to the highest-value agent. Section 4.4 provides sufficient conditions

[50] It is enough to assume that agents have the same Bayesian beliefs about these variables, and the same higher-order belief about the beliefs of the third party.

under which full revelation of the cutoff, or no revelation, are optimal. These two mechanisms are solutions to practical problems considered in Section 5.2 and in the Online Appendix. By Corollary 4, it is enough to consider signals that only depend on the second highest type. Define a function $v^\pi : \Theta^2 \to \mathbb{R}$ by $v^\pi(\theta, \hat{\theta}) = u(\theta; f_\pi^{\hat{\theta}})$, where $f_\pi^{\hat{\theta}}$ is the posterior belief over the type of the winner when the second highest value is $\hat{\theta}$ and the disclosure rule is given by π. The function $v^\pi(\theta, \hat{\theta})$ is the expected continuation value of the winner with type θ conditional on winning against second highest type $\hat{\theta}$.[51] To simplify proofs, I assume that $v^\pi(\theta, \theta)$ is strictly positive and differentiable in θ.

The function v^π is similar to an object studied by Milgrom and Weber (1982a) in the context of auctions with affiliated values. In the setting of Milgrom and Weber (1982a), the value of the winner depends on the value of the second highest bidder $\hat{\theta}$ due to statistical correlation of types and the assumption of non-private values. In my setting, the value of the winner depends on $\hat{\theta}$ because the bid of the second highest bidder influences the signal sent by the mechanism (and hence the continuation payoff of the winner). In Milgrom and Weber (1982a), affiliation of types implies that their analog of $v^\pi(\theta, \theta)$ is non-decreasing in θ, a property necessary for existence of a monotone equilibrium in standard auctions. In my model, $v(\theta, \hat{\theta})$ is non-decreasing in θ under the assumption of a monotone aftermarket. In general, there is no reason to expect monotonicity in $\hat{\theta}$, and hence $v^\pi(\theta, \theta)$ may fail to be increasing in θ. A sufficient condition for monotonicity of $v^\pi(\theta, \theta)$ is that a "higher" posterior belief of the third party leads to a higher payoff for the agent in the aftermarket. An example is provided at the end of the subsection.

Proposition 7. *Suppose that* $x(\theta_i, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i \geq \boldsymbol{\theta}_{-i}^{(1)}\}}$, *for all* $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. *If* π *is the full-disclosure rule,* (x, π) *can be robustly implemented by*

- *a second-price auction where the price paid by the winner is disclosed, or*
- *a first-price auction where the second highest bid is disclosed,*

if and only if $v^\pi(\theta, \theta)$ *is strictly increasing in* θ. *Moreover,* (x, π) *can always be robustly implemented by an all-pay auction where the second highest bid is disclosed.*

If π *is the no-revelation rule,* (x, π) *can be robustly implemented by any of the above auction formats (with no revelation).*

A candidate equilibrium bidding function is determined by the local (first-order) optimality condition. The bidding function has to be strictly increasing to guarantee that

[51] When π is the no-revelation rule, $f_\pi^{\hat{\theta}}$ is the distribution of the first order statistic of N draws from f (and does not depend on $\hat{\theta}$). If π is the full-disclosure rule, $f_\pi^{\hat{\theta}}$ is the truncation of f at $\hat{\theta}$.

full disclosure of the bid (or price) leads to full disclosure of the second highest type (and is also needed for equilibrium existence). In a SPA, the candidate bidding function in the full-disclosure case is $v^\pi(\theta, \theta)$, and so $v^\pi(\theta, \theta)$ must be strictly increasing in θ. In a FPA, it is enough that $(\int_0^\theta v^\pi(\tau, \tau)dF^{N-1}(\tau))/(F^{N-1}(\theta))$ is strictly increasing in θ. This condition holds for all F if and only if $v^\pi(\theta, \theta)$ is strictly increasing in θ. If robustness is required only for a subset of prior distributions (e.g. because the designer has some information about it), a weaker condition on v^π may suffice. Finally, in an all-pay auction, an agent with type θ bids $\int_0^\theta v^\pi(\tau, \tau)dF^{N-1}(\tau)$. The bidding function is always increasing, because, unlike in a first- or second-price auction, it is not obtained by conditioning on winning.

The above results can be applied to the solution of the post-auction bargaining model of Section 5.2. In the full-disclosure case,

$$v^\pi(\theta, \hat{\theta}) = \theta + (1-\eta)H\left(\eta[(\delta-1)\mathbb{E}[\theta_N^{(1)} \,|\, \theta_N^{(2)} = \hat{\theta}] + \Delta]\right)((\delta-1)\theta + \Delta),$$

where $\mathbb{E}[\theta_N^{(1)} \,|\, \theta_N^{(2)} = \hat{\theta}]$ is the expected value of the highest type when the second highest type is $\hat{\theta}$. A sufficient condition for the function $v^\pi(\theta, \theta)$ to be strictly increasing is that $\delta \geq 1$. By Proposition 7, a SPA with revelation of the price robustly implements the optimal mechanism in that case. Regardless of the monotonicity properties of v^π, an all-pay auction may be used.

6.2 A canonical indirect implementation of cutoff mechanisms

Section 6.1 identifies special cases in which a standard auction can be used to implement a cutoff rule. In this section, my goal is to provide a general characterization of cutoff mechanisms as equilibria of simple dynamic auctions, called Generalized Clock Auctions (GCAs). Informally, a GCA is a bidding procedure in which agents gradually drop out until a winner is determined. A disclosure rule for a GCA specifies what details of the bidding history are publicly revealed after the auction. For example, the auctioneer can reveal the final price, or the set of active bidders in every round. Intuitively, the history of bidding does not depend on the exact type of the winner because the auction ends once the second highest bidder drops out. Hence, any garbling of the public history corresponds to a cutoff disclosure rule. The analysis below formalizes this intuition.

I assume that the type space Θ is finite for ease of exposition. A Generalized Clock Auction (GCA) is characterized by a sequence of prices and a disclosure rule. Let $\mathcal{T} = \{0, 1, 2, ..., T\}$ be the set of rounds. In round 0, agents simultaneously decide to participate or not. In every subsequent round $t \in \mathcal{T}$:

- A price p^t is announced to bidders;
- Bidders simultaneously (and covertly) decide to stay in the auction, or to exit;
- The auctioneer observes bidders' decisions and implements the relevant outcome (to be specified);
- The auctioneer announces to bidders the outcome of the round (whether the auction continues, the set of active bidders, the winner in case the auction ends).

The outcome of a round is determined in the following way. If at least two bidders decide to stay (and $t < T$), the auction continues to the next round. Bidders who exited are declared inactive and do not participate in future rounds. Otherwise, the auction terminates. If all n active bidders drop out in the last round, the object is allocated uniformly at random among them, and all these bidders are declared inactive. If exactly one bidder stays (and $n-1$ bidders drop out), she wins the object, and is declared inactive with probability $1/n$ (which is the probability she would have won the object by dropping out in that round). At $t = T$, all bidders must exit. After the object is allocated, a signal s is released publicly according to a disclosure rule (to be specified). The distinction between the winner being active or inactive at the end of the auction is irrelevant for the final allocation but matters for the informational content of the signal.

Let H^t denote the public history of the bidding procedure described above up to and including round t, and let $\mathcal{H}^t$ be the set of all public histories.[52] Public history in this context is identified with the sequence of announcements made by the auctioneer to the bidders during the auction. A Generalized Clock Auction (GCA) is a sequence of functions $\{(Y^t, P^t)\}_{t=1}^T$, where $Y^t : \mathcal{H}^t \to \Delta(\mathcal{S})$, for some (finite) signal space $\mathcal{S}$, and $P^t : \mathcal{H}^{t-1} \to \Delta(\mathbb{R})$. In each round t, given a history H^{t-1}, a price p^t is drawn from the distribution $P^t(H^{t-1})$. If the auction ends in round t, the signal is drawn and announced according to distribution $Y^t(H^t)$. Hence, the signal s is an arbitrary garbling of the entire public history of the auction.

Prices do not have to change monotonically in a GCA.[53]. Because the signal distribution depends on the termination time of the auction, it is as if a different good were offered for sale in every round. Prices may have to decrease when the signal distribution gets less attractive for bidders. The auction nevertheless preserves monotonicity in types. High types always have a weakly higher incentive to stay in the auction then low types.

[52] Public history is defined as the largest information set contained in information sets of all bidders.
[53] This is one key difference to obviously strategy-proof auctions considered by Li (2016)

A GCA constitutes the first-stage mechanism. In the second stage, the winner interacts in the aftermarket and obtains her final payoff as a function of beliefs induced by the signal s. The informational content of the signal is determined by equilibrium behavior of bidders in the GCA. Because technical details associated with defining strategies and equilibria are not relevant for the message of this subsection, I relegate them to Appendix D.2, where I also define monotone and Markov strategies, monotone equilibria (in a monotone equilibrium, lower types exit earlier than higher types), and Markov GCAs (in a Markov GCA, prices depend only on the number of active bidders).

In the statement of the result, I restrict attention to a subclass of allocation rules. This allows me to focus on simple GCAs. Define a hierarchical allocation rule $x^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta})$ for any sequence $\kappa_1 < \dots < \kappa_k$, with $\kappa_m \in \Theta$ for all m, by

$$x^{\kappa_1 \dots \kappa_k}(\theta_i, \boldsymbol{\theta}_{-i}) = \begin{cases} \frac{1}{|\{j \in N : \kappa_m \leq \theta_j < \kappa_{m+1}\}|} & \text{if } \kappa_m \leq \theta_i < \kappa_{m+1} \text{ and } \forall j,\ \theta_j < \kappa_{m+1}, \\ 0 & \text{otherwise,} \end{cases}$$

where by convention $\kappa_{k+1} = \infty$. For example, if $\Theta = \{\theta_1, \dots, \theta_n\}$, then $x^{\theta_1 \dots \theta_n}(\boldsymbol{\theta})$ is the "efficient" allocation rule (highest type receives the good), $x^{\theta_m \dots \theta_n}(\boldsymbol{\theta})$ excludes types $\theta_1, \dots, \theta_{m-1}$, and $x^{\theta_1}(\boldsymbol{\theta})$ corresponds to a uniform lottery. A symmetric allocation rule x is called *decomposable* if it is a convex combination of hierarchical allocation rules. Decomposability is a mild restriction from a practical perspective. It rules out cases when the final allocation depends on types of agents who themselves never receive the good. For example, an allocation rule in which the good is allocated to the highest type if the third highest type is below $1/2$, and randomly among the top two bidders otherwise, is not decomposable.

Theorem 5. *If (x, π) is a mechanism frame implemented by a monotone equilibrium of a GCA, then (x, π) is a cutoff rule. Conversely, if x is decomposable, any symmetric cutoff rule (x, π) can be implemented (up to Bayesian equivalence) in a pure-strategy equilibrium of a Markov GCA in which randomization over prices may only happen in round* 0 *(subsequent prices are deterministic functions of the number of active bidders and the realization of the initial random price).*

The proof and discussion of Theorem 5 can be found in Appendix D.3. Due to decomposability of x, in order to implement a cutoff rule, it is enough to keep track of the number of active bidders in any round (the Markov property). This is because the allocation does not depend on the types of bidders who exited in previous rounds.[54] In order

[54] The main difficulty in the proof is to show that decomposability of x implies existence of a signal

to implement an arbitrary cutoff rule, I would have to allow for stochastic non-Markov prices, and the construction of the GCA would be more complicated.

To connect GCAs to the results of Section 6.1, suppose that x allocates to the highest type, and π is either the full-disclosure or the no-disclosure rule. If $v^{\pi}(\theta, \theta)$ is non-decreasing, prices in the corresponding GCA are also non-decreasing. Thus, robust implementation is possible because the auctioneer can use a continuously increasing clock, starting at a sufficiently low price.

7 Extensions

7.1 Exogenous information of the mechanism designer

The mechanism designer in my model can only disclose information to the aftermarket if she first elicits it from the agents. However, in many settings, the designer might have access to exogenous information. For example, a seller faces a varying cost of allocating the good, or knows whether the asset is of high or low quality. In this section, I extend the one-agent model by allowing for exogenous information of the designer. I then study the interplay between revelation of exogenous and endogenous information, showing that it results in a generalization of Theorem 4 – there always exists an optimal mechanisms that only reveals exogenous information of the seller. The section is concluded by applying the extended model in the context of the OTC market from Section 5.1.

7.1.1 Extended model

I assume that the mechanism designer observes the realization of a random variable z with realizations in some finite set $\mathcal{Z}$, and distributed according to a prior probability mass function α.

When the designer has more information than the agent, the analysis of incentives is greatly complicated, due to the possibility of signaling by proposing a mechanism. To avoid these issues (which are orthogonal to the goal of the extension), I assume that the agent also observes the realization of z. The assumption is realistic in some applications, including the application considered at the end of this section. Denoting by $\bar{\alpha}$ the posterior belief over z held by the third party in the aftermarket, the payoff to the agent with type θ conditional on acquiring the object is given by $u(\theta;\, z,\, \bar{f},\, \bar{\alpha})$.

distribution π', Bayesian equivalent to π, that inherits this property. The original distribution π may depend on types of bidders who exited in previous rounds, and thus it cannot be implemented with a deterministic price path.

In the general model of Section 4 with $N = 1$, a mechanism frame is a collection $\{(x_z, \pi_z)\}_{z\in\mathcal{Z}}$, indexed by realizations of the random variable z. The mechanism frame $\{(x_z, \pi_z)\}_{z\in\mathcal{Z}}$ is a cutoff rule if (x_z, π_z) is a cutoff rule for each z. A cutoff rule in the extended model is always implementable, as long as the aftermarket is monotone, that is, $u(\theta;\, z,\, \bar{f},\, \bar{\alpha})$ is non-decreasing in θ for any $z \in \mathcal{Z}$, $\bar{f} \in \Delta(\Theta)$, and $\bar{\alpha} \in \Delta(\mathcal{Z})$.[55] The converse result also holds if the set of distributions and aftermarkets for which implementability is required is sufficiently large.

The objective function of the designer is

$$\sum_{z\in\mathcal{Z}} \alpha(z) \int_{\Theta} \int_{\mathcal{S}} V(\theta;\, z,\, f^s,\, \alpha^s) d\pi_z(s|\,\theta) x_z(\theta) f(\theta) d\theta, \tag{7.1}$$

where $V(\theta;\, z,\, f^s,\, \alpha^s)$ is the conditional expected payoff (conditional on allocating the good) when the type of the agent is θ, the realization of the exogenous information is z, and the third party has beliefs f^s and α^s, respectively, over these two variables.[56]

7.1.2 Characterization of optimal mechanisms

In the extended model, the designer decides about disclosure of two types of information: her exogenous information about z and the endogenous information about cutoffs c. The prior distribution of z is given, while the prior distribution of c is a choice variable, determined by the chosen allocation rule. The distribution of the cutoff may endogenously depend on the realization of z. The problem can be formulated as inducing an optimal distribution of posterior joint distributions H over $(z,\, c)$ subject to a condition that the posterior marginal distributions over z average out to the prior α. Given a joint distribution H, let $dH_z(\cdot)$ denote the marginal distribution over z in the form of a pmf, and let $H_{c|z}(\cdot \mid z)$ denote the conditional distribution over c given z in the form of a cdf. For any distribution $\bar{\alpha}$ over $\mathcal{Z}$, let

$$\mathcal{U}(\bar{\alpha}) \equiv \max_{H\in\Delta(\mathcal{Z}\times C):\, \mathrm{marg}_{\mathcal{Z}}(H)=\bar{\alpha}} \mathcal{V}(H), \tag{7.2}$$

[55] This is true even if the agent does not observe z.

[56] For $\mathcal{S}-$finite mechanisms,

$$\alpha^s(z) = \frac{\alpha(z) \int_{\Theta} \pi_z(s|\,\theta) x_z(\theta) f(\theta) d\theta}{\sum_{\tilde{z}} \alpha(\tilde{z}) \int_{\Theta} \pi_{\tilde{z}}(s|\,\theta) x_{\tilde{z}}(\theta) f(\theta) d\theta}.$$

where

$$\mathcal{V}(H) = \sum_{\tilde{z}\in\mathcal{Z}} dH_z(\tilde{z}) \int_{\Theta} V(\theta;\, \tilde{z},\, f^H,\, dH_z) H_{c|z}(\theta|\, \tilde{z}) f(\theta) d\theta,$$

where f^H is the conditional belief over the type θ of the agent given joint belief H over $(z,\, c)$, conditional on $\theta \geq c$. Note that $dH_z(\tilde{z}) = \bar{\alpha}(\tilde{z})$ if $\text{marg}_{\mathcal{Z}}(H) = \bar{\alpha}$. The function $\mathcal{V}(H)$ is the expected payoff of the mechanism designer when the joint distribution over $(z,\, c)$ is H, and the mechanism reveals no information. Accordingly, $\mathcal{U}(\bar{\alpha})$ is the maximal expected payoff achievable by the mechanism designer for a fixed distribution $\bar{\alpha}$ over $\mathcal{Z}$, assuming that the designer chooses the optimal allocation rule for each z but no information is revealed.

I state the following result without proof (the proof uses similar arguments to the ones used multiple times in this paper).

Proposition 8. *The maximal value of the objective function* (7.1) *attained over the set of cutoff rules* $\{(x_z,\, \pi_z)\}_{z\in\mathcal{Z}}$ *is given by* co$\mathcal{U}(\alpha)$ *– the value of the concave closure of the function* $\mathcal{U}$ *at the prior* α.

Proposition 8 implies that the problem can be solved in two steps. In the first step, for every distribution $\bar{\alpha}$ over $\mathcal{Z}$, we compute the optimal mechanism assuming no communication (this yields the function $\mathcal{U}$). In the second step, the function $\mathcal{U}$ is concavified which corresponds to finding the optimal revelation policy. When the prior α is degenerate (the seller observes no information), Proposition 8 boils down to a one-agent version of Theorem 4. More generally, the following corollary holds.

Corollary 6. *The optimal mechanism in the extended model with one agent reveals information only about the random variable* z *(i.e. the realization of* z *is a sufficient statistic for the distribution of signals released by the mechanism).*

Corollary 6 follows directly from Proposition 8. The optimal mechanism concavifies the function $\mathcal{U}$ which corresponds to inducing a distribution of posterior beliefs over z. Conditional on inducing a posterior belief over z, no further information is revealed, by definition of $\mathcal{V}$. Unless the function $\mathcal{U}$ is concave everywhere, then at some prior α, co$\mathcal{U}(\alpha) > \mathcal{U}(\alpha)$ which means that some information is revealed.

The optimal mechanism may reveal information about z even if z does not directly influence the payoffs of the agent and the third party.[57] Disclosing information about z indirectly reveals information about the cutoff if the distribution of the cutoff varies with

[57] That is, if u depends on neither z nor the belief over z, and V depends on z but not on the belief over z.

z. For example, suppose that z is a private cost of the seller, and the seller allocates the good to all types of the agent above her cost z. Then, for a fixed z, the allocation function is a threshold rule, and the cutoff representation is degenerate. However, if z is unknown, the expected allocation rule is a step function with a non-degenerate cutoff representation. Hence, information disclosure may be optimal. To gain further intuition, note that the seller's cost z in the above example plays the role of a second highest bid in an auction setting from the perspective of the agent and the third party. The above considerations are illustrated with an example below.

7.1.3 Implications for the OTC market analysis

In the OTC market model of Section 5.1, both the welfare- and profit-maximizing mechanism can be implemented as a deterministic posted-price mechanism that reveals no information. Asriyan et al. (2015) and Duffie et al. (2016), in different settings, show that information revelation is sometimes optimal. To reconcile my findings with the findings of that literature, I extend the model of Section 5.1 by allowing for exogenous information of the seller.

In Duffie et al. (2016), dealers face a varying cost of providing the asset, known to them but not to final buyers (customers). Accordingly, I assume that the initial seller in my extended model faces a random cost $k \in \mathcal{K}$ of allocating the asset, observed by her and the agent, but not the third party. The set $\mathcal{K}$ is finite, and the distribution of the cost k has a probability mass function α. The model is otherwise identical to that of Section 5.1.

Within this framework, the following interpretation of the lemons problem can be given. The agent and the third party have an outside option to trade in a different market if resale fails. The values of these outside options are negatively correlated because the agent is a seller and the third party is a buyer. The value of the agent's outside option is θ, and the value of the third party's outside options is $o(\theta)$, for some decreasing function $o(\theta)$. Thus, the agent (dealer) is better informed about the value of the outside option. The third party has a fixed value v for the asset. Then, the function $v(\theta)$ is recovered as $v(\theta) = v - o(\theta)$, for all θ. For example, if the outside option of the third party is to buy the asset at a spread $\Delta > 0$ compared to the value of the agent, we have $o(\theta) = v - (1 + \Delta)\theta$, so that $v(\theta) = (1 + \Delta)\theta$.

Corollary 7. *In a welfare-maximizing mechanism of the extended model, the signal distribution is fully determined by the realization of the cost k (conditional on k, it does not depend on any other variable).*

The exact form of the optimal signal is sensitive to the parameters of the model. In Appendix E.1, I discuss the structure of the optimal allocation rule, and show, by means of an example, that depending on the severity of the lemons problem, either full revelation or no revelation of k may be optimal.

Instead of studying the dependence of the solution on the parameters of the model, I consider a robust approach to the disclosure problem. I have so far assumed that the third party receives signals only from the mechanism. In practice, the third party may observe other signals, or acquire more information, from sources not controlled by the mechanism designer.[58]

To model this in a tractable way, I assume that the mechanism designer does not know the distribution of the exogenous signal about k that is observed by the third party. Mechanisms are evaluated according to their worst-case performance. Formally, *(1)* the designer chooses the mechanism frame $\{(x_k, \pi_k)\}_{k \in \mathcal{K}}$, *(2)* Nature chooses a distribution of an additional exogenous signal about k, *(3)* the mechanism is implemented, *(4)* the third party observes the signal s released by the mechanism, and also observes the additional signal drawn from the distribution chosen by Nature, and *(5)* the aftermarket game is played. Nature tries to minimize the expected payoff of the mechanism designer. The following claim is a corollary of the above analysis (the proof can be found in Appendix E.2).

Claim 4. *The welfare-maximizing mechanism, under the worst-case criterion, is a posted-price mechanism with full disclosure of the cost k.*

The mechanism described in Claim 4 can be interpreted as a financial market benchmark such as LIBOR (which discloses information about the borrowing cost of major banks). Duffie et al. (2016) provide sufficient conditions for optimality of announcing a benchmark in their model. In my framework, although exact disclosure of the cost may not be optimal when the distribution of the exogenous signal is known, announcing a benchmark gives the highest welfare guarantee to an ambiguity averse regulator.

7.2 What if the loser also interacts in the aftermarket?

In the preceding sections, I assumed that only the agent who acquires the good interacts in the aftermarket. In many cases, the agent may also engage in post-mechanism interactions when she doesn't acquire the good. For example, a loser may try to purchase a similar

[58] For example, in financial OTC markets, customers often search for the best quote before transacting, and thus they can acquire more information by observing quotes of other dealers.

object in the aftermarket, or negotiate to gain access to an object owned by another market participant.

To allow for this possibility, I extend the single-agent model from Section 3. The aftermarket is formally a pair $(A_l,\, A_w)$ with

$$A_j \equiv \{u_j(\theta; \bar{f}) : \theta \in \Theta,\, \bar{f} \in \Delta(\Theta)\},$$

for $j \in \{l,\, w\}$, where subscript l denotes the aftermarket for a "loser" (agent does not acquire the good), and w – the aftermarket for a "winner". I assume that continuation payoffs are non-decreasing in the type, and, for ease of exposition, that Θ is finite. A mechanism frame in the extended setting is $(x,\, \pi_l,\, \pi_w)$, where π_l is the signal distribution conditional on not allocating the object, and π_w is the signal distribution conditional on allocating the object. Because there is only one agent, only one of the signal structures is used ex-post. Third-party players in the aftermarket know whether the agent acquired the good or not.

Definition 14. A mechanism frame $(x,\, \pi_l,\, \pi_w)$ is a *cutoff rule* if x is non-decreasing, and the signals π_l and π_w can be represented as

$$\pi_l(s|\,\theta)(1 - x(\theta)) = \sum_{c>\theta} \gamma_l(s|\,c)dx(c), \tag{7.3}$$

$$\pi_w(s|\,\theta)x(\theta) = \sum_{c\leq\theta} \gamma_w(s|\,c)dx(c), \tag{7.4}$$

for each $\theta \in \Theta$ and $s \in \mathcal{S}$, for some signal functions $\gamma_l : C \to \Delta(\mathcal{S})$ and $\gamma_w : C \to \Delta(\mathcal{S})$.

Cutoff mechanisms are defined accordingly. Condition (7.4) is analogous to condition (3.1) in the definition of cutoff rules from Section 3. Condition (7.3) is a mirror image of condition (7.4), with the sum indexed by all cutoff levels exceeding the type of the agent.

Instead of formulating a richness condition similar to the one defined in Section 4, I simplify the analysis by requiring the mechanism frames to be implementable for all distributions f and all monotone aftermarkets $(A_l,\, A_w)$.

Claim 5. *A mechanism frame $(x,\, \pi_l,\, \pi_w)$ is implementable for any distribution f and any monotone aftermarkets $(A_l,\, A_w)$ if and only if x is a constant allocation rule (in which case no information can be revealed).*

The conclusion continues to hold when only the loser interacts, and the winner enjoys the utility of holding the object: $u_w(\theta;\, \bar{f}) = \theta$, for all θ, and $\bar{f} \in \Delta(\Theta)$.

In the absence of any conditions on the aftermarket, no information about the agent's type can be elicited in a cutoff mechanism. Claim 5 should not come as a surprise. Suppose that in the loser's aftermarket, the same object is allocated but with utility doubled for every type. In this case, higher types have a relative preference for *not* acquiring the object in the first stage, which reverts the direction of single-crossing. To avoid the negative result, it is necessary to assume that the game played when the object is acquired is in some sense preferred to playing the game when the object is not allocated. Preference for winning the object should be expressed in relative terms (consistent with single-crossing) rather than absolute terms (absolute differences in utility can always be undone with transfers).

Definition 15 (Single-crossing separation)**.** The winner's aftermarket A_w is *single-crossing-separated* from the loser's aftermarket A_l if for any $\theta > \hat{\theta}$, there exists $d(\theta,\, \hat{\theta}) > 0$, such that for all $\bar{f}$,

$$u_w(\theta;\, \bar{f}) - u_w(\hat{\theta};\, \bar{f}) \geq d(\theta,\, \hat{\theta}) \geq u_l(\theta;\, \bar{f}) - u_l(\hat{\theta};\, \bar{f}).$$

Single-crossing separation requires that the difference in utilities between any two types in A_w can be separated from the difference in utilities between these two types in A_l, uniformly in posterior beliefs. For example, the condition is satisfied (with $d(\theta,\, \hat{\theta}) = \theta - \hat{\theta}$) when A_l corresponds to buying an identical object from a different seller, and there is no winner's aftermarket, i.e. $u_w(\theta;\, \bar{f}) = \theta$, for all θ.

Proposition 9. *A mechanism frame* $(x,\, \pi_l,\, \pi_w)$ *is implementable for any distribution* f *and any monotone aftermarket* $(A_l,\, A_w)$ *such that* A_w *is single-crossing-separated from* A_l *if and only if* $(x,\, \pi_l,\, \pi_w)$ *is a cutoff rule.*

Using Proposition 9 and arguments used in preceding sections, one can show that in the extended setting *(1)* for a fixed allocation function x, the problem of finding an optimal cutoff mechanism is a Bayesian persuasion problem, and *(2)* there always exists an optimal cutoff mechanism that reveals no information.

The extension to multiple players is straightforward in the case when only the losers interact (i.e. $u_w(\theta;\, \bar{f}) = \theta$ for all θ). In this case, it is the bid of the winner, not the losers, that can be revealed without compromising incentives. If losing and winning players both interact, the problem becomes significantly more complicated, because third-party players observe multiple signals, all of which could contain non-trivial information about the same agent.

8 Conclusions

In this paper, I studied mechanism design in a setting where the mechanism is followed by an aftermarket, i.e. a post-mechanism game played between the agent who acquired the object and third-party market participants. Existence of an exogenous aftermarket creates a new tool in the design problem – the disclosure rule. By disclosing information elicited by the mechanism, the designer influences the information structure of the aftermarket.

I introduced a tractable class of cutoff rules that are characterized by being always implementable – regardless of the aftermarket and the prior distribution of types. The theory was applied to study optimal transparency of financial over-the-counter markets, post-auction bargaining, and mechanisms with resale.

It is useful to distinguish three sources of information that a mechanism can attempt to disclose: *(1)* private information of agents who participate in the aftermarket, *(2)* private information of agents who do not participate in the aftermarket, and *(3)* private information of the designer, including outcomes of endogenous randomization in the mechanism. Although final payoffs in my model are determined by posterior beliefs about the first type of information, only the last two sources can be used robustly, i.e. irrespective of the fine details of the model. A natural conjecture is that this conclusion holds more generally, for example, for other first-stage social choice problems.

The analysis of mechanism design in this paper may be seen as a compromise between two extremes: a fully Bayesian approach on one hand, and ambiguity aversion (or worst-case analysis) on the other. The designer in my model faces additional robustness constraints on the feasible set of mechanisms, but is allowed to maximize a Bayesian objective function. An interesting direction for future research is to apply this approach to other design problems.

References

Asquith, Paul, Thomas R. Covert, and Parag Pathak, "The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market," 2013.

Asriyan, Vladimir, William Fuchs, and Brett S Green, "Information Spillovers in Asset Markets with Correlated Values," Working Paper 2015.

Aumann, Robert J. and Michael B. Maschler, *Repeated Games with Incomplete Information*, MIT Press, 1995.

Bergemann, Dirk and Stephen Morris, "Robust Mechanism Design," *Econometrica*, 2005, *73* (6), 1771–1813.

_ **and** _ , "Ex post implementation," *Games and Economic Behavior*, 2008, *63* (2), 527 – 566. Second World Congress of the Game Theory Society.

_ **and** _ , "An Introduction to Robust Mechanism Design," *Foundations and Trends in Microeconomics*, 2013, *8* (3), 169–230.

_ **and** _ , "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 2016, *11* (2), 487–522.

_ **and** _ , "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium," *American Economic Review*, May 2016, *106* (5), 586–91.

Bessembinder, Hendrik and William Maxwell, "Transparency and the Corporate Bond Market," *Journal of Economic Perspectives*, 2008, *22* (2), 217–234.

Border, Kim C., "Implementation of Reduced Form Auctions: A Geometric Approach," *Econometrica*, 1991, *59* (4), 1175–1187.

Calzolari, Giacomo and Alessandro Pavan, "Monopoly with resale," *RAND Journal of Economics*, 06 2006, *37* (2), 362–375.

_ **and** _ , "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, September 2006, *130* (1), 168–204.

_ **and** _ , "On the use of menus in sequential common agency," *Games and Economic Behavior*, September 2008, *64* (1), 329–334.

_ **and** _ , "Sequential contracting with multiple principals," *Journal of Economic Theory*, March 2009, *144* (2), 503–531.

Carroll, Gabriel and Ilya Segal, "Robustly Optimal Auctions with Unknown Resale Opportunities," Working Paper 2016.

Duffie, Darrell, Piotr Dworczak, and Haoxiang Zhu, "Benchmarks in search markets," forthcoming, Journal of Finance 2016.

Dworczak, Piotr, "The Effects of Post-auction Bargaining between Bidders," Working Paper 2015.

_ **and Anthony Lee Zhang**, "Implementability, Walrasian Equilibria, and Efficient Matchings," Working Paper 2015.

_ **and Giorgio Martini**, "The Simple Economics of Optimal Persuasion," Working Paper 2016.

Ező, Péter and Balázs Szentes, "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 2007, *74* (3), 705–731.

Fuchs, William and Andrzej Skrzypacz, "Government interventions in a dynamic market with adverse selection," *Journal of Economic Theory*, 2015, *158* (PA), 371–406.

Gershkov, Alex, Jacob K. Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi, "On the Equivalence of Bayesian and Dominant Strategy Implementation," *Econometrica*, 2013, *81* (1), 197–220.

Giovannoni, Francesco and Miltiadis Makris, "Reputational Bidding," *International Economic Review*, 2014, *55* (3), 693–710.

Goeree, Jacob, "Bidding for the future: signaling in auctions with an aftermarket," *Journal of Economic Theory*, 2003, *108* (2), 345–364.

Green, Jerry and Jean-Jacques Laffont, "Posterior Implementability in a Two-Person Decision Problem," *Econometrica*, 1987, *55* (1), 69–94.

Gupta, Madhurima and Bernard Lebrun, "First Price Auctions with Resale," *Economics Letters*, 1999, *64*, 181–185.

Hafalir, Isa and Vijay Krishna, "Asymmetric Auctions with Resale," *American Economic Review*, 2008, *98* (1), 87–112(26).

_ and _ , "Revenue and efficiency effects of resale in first-price auctions," *Journal of Mathematical Economics*, September 2009, *45* (9-10), 589–602.

Haile, Philip A., "Auctions with private uncertainty and resale opportunities," *Journal of Economic Theory*, 2003, *108* (1), 72 – 110.

Jehiel, Philippe and Benny Moldovanu, "Efficient Design with Interdependent Valuations," *Econometrica*, 2001, *69* (5), 1237–1259.

_ and _ , "Allocative and Informational Externalities in Auctions and Related Mechanisms," in Richard Blundell, Whitney Newey, and Torsten Persson, eds., *The Proceedings of the 9th World Congress of the Econometric Society*, Cambridge University Press, 2006.

_ , _ , and Ennio Stacchetti, "How (Not) to Sell Nuclear Weapons," *The American Economic Review*, 1996, *86* (4), 814–829.

Kamenica, Emir and Matthew Gentzkow, "Bayesian Persuasion," *American Economic Review*, October 2011, *101* (6), 2590–2615.

Katzman, Brett E. and Matthew Rhodes-Kropf, "The Consequences of Information Revealed in Auctions," *Applied Economics Research Bulletin*, 2008, *2*, 53–87.

Kolotilin, Anton, "Optimal Information Disclosure: A Linear Programming Approach," Working Paper 2016.

_ , Ming Li, Tymofiy Mylovanov, and Andriy Zapechelnyuk, "Persuasion of a Privately Informed Receiver," Working Paper 2015.

Lauermann, Stephan and Gábor Virág, "Auctions in Markets: Common Outside Options and the Continuation Value Effect," *American Economic Journal: Microeconomics*, 2012, *4* (4).

Leland, Hayne and David H Pyle, "Informational Asymmetries, Financial Structure, and Financial Intermediation," *Journal of Finance*, 1977, *32* (2), 371–87.

Li, Shengwu, "Obviously Strategy-Proof Mechanisms," Working Paper 2016.

Manelli, Alejandro M. and Daniel R. Vincent, "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 2010, *78* (6), 1905–1938.

Matthews, Steven A., "On the Implementability of Reduced Form Auctions," *Econometrica*, 1984, *52* (6), 1519–1522.

Milgrom, Paul, *Putting Auction Theory to Work*, Cambridge University Press, 2001.

_ **and Ilya Segal**, "A Design Proposal for the US Spectrum Auction," in Martin Bichler and Jacob Goeree, eds., *Handbook of Spectrum Auction Design*, Cambridge University Press, 2017.

_ **and Robert Weber**, "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1982, *50*, 1089–1122.

_ **and** _ , "The value of information in a sealed-bid auction," *Journal of Mathematical Economics*, 1982, *10* (1), 105–114.

Molnar, Jozsef and Gábor Virág, "Revenue maximizing auctions with market interaction and signaling," *Economics Letters*, 2008, *99* (2), 360–363.

Myerson, Roger, "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Myerson, Roger B., "Optimal coordination mechanisms in generalized principal-agent problems," *Journal of Mathematical Economics*, June 1982, *10* (1), 67–81.

Philippon, Thomas and Vasiliki Skreta, "Optimal Interventions in Markets with Adverse Selection," *American Economic Review*, February 2012, *102* (1), 1–28.

Rochet, Jean-Charles, "A necessary and sufficient condition for rationalizability in a quasi-linear context," *Journal of Mathematical Economics*, 1987, *16* (2), 191–200.

Seierstad, A. and Knut Sydsaeter, *Optimal Control Theory with Economic Applications*, Elsevier, North-Holland, 1987. Advanced Textbooks in Economics, volume 24.

SIFMA, "SIFMA Electronic Bond Trading Report: US Corporate and Municipal Securities," Report February 2016.

Tirole, Jean, "Overcoming Adverse Selection: How Public Intervention Can Restore Market Functioning," *American Economic Review*, February 2012, *102* (1), 29–59.

Zhang, Jun, "Optimal Mechanism Design with Aftermarket Interactions," Working Paper 2014.

_ **and Ruqu Wang**, "Optimal mechanism design with resale via bargaining," *Journal of Economic Theory*, 2013, *148* (5), 2096 – 2123.

Zheng, Charles Zhoucheng, "Optimal Auctions with Resale," *Econometrica*, 2002, *70* (6).

Zhong, Hongjun, "Information Revelation, Post-bid Market Interaction and Auction Choice," Working Paper 2002.

A Proofs and supplementary materials for Section 3

A.1 Proof of Lemma 1

Below, I provide an elementary proof which uses the special assumptions of Section 3. A more general proof technique (based on the Radon-Nikodym theorem) is used to prove a generalization of Lemma 1 – Lemma 3 in Section 4.

If (x, π) is a cutoff rule, property (M) follows directly from the definition of cutoff rules. I prove the converse part that property (M) implies a cutoff rule. To simplify notation, I denote $\underline{\theta} = \min\{\Theta\}$, and I let θ^+ be the smallest type larger than θ, and θ^- be the largest type smaller than θ, whenever they exist.

Fix a mechanism frame (x, π) that satisfies condition (M). By summing up over signals, for each θ,

$$\sum_{s\in\mathcal{S}} \left[\pi(s|\,\theta^+)x(\theta^+) - \pi(s|\,\theta)x(\theta)\right] = x(\theta^+) - x(\theta).$$

Since (x, π) satisfies condition (M), each term in the sum on the left hand side is non-negative. It follows that $x(\theta^+) \geq x(\theta)$, and

$$\pi(s|\,\theta^+)x(\theta^+) - \pi(s|\,\theta)x(\theta) \leq x(\theta^+) - x(\theta),\ \forall s \in \mathcal{S},\ \forall\theta. \tag{A.1}$$

Given the above properties, I prove that we can inductively construct a signal function $\gamma : C \to \Delta(S)$ such that

$$\pi(s|\,\theta)x(\theta) = \sum_{c\leq\theta} \gamma(s|\,\theta)dx(c),\ \forall s \in \mathcal{S}, \tag{A.2}$$

where recall that dx denotes the probability mass function over cutoffs induced by the allocation function x.

The induction is over Θ (types are totally ordered because Θ is a subset of the real line). For type $\underline{\theta}$, define $\gamma(s|\,\underline{\theta}) := \pi(s|\,\underline{\theta})$, for any $s \in \mathcal{S}$. Then, we have

$$\pi(s|\,\underline{\theta})x(\underline{\theta}) = \pi(s|\,\underline{\theta})dx(\underline{\theta}) = \sum_{c\leq\underline{\theta}} \gamma(s|\,c)dx(c).$$

Inductive hypothesis: Suppose we have constructed $\gamma(s|\,c)$ for $s \in \mathcal{S}$ and $c < \theta$ such that equation (A.2) holds for all $\hat{\theta} < \theta$.

Inductive step: I will construct $\gamma(s|\,\theta)$ for each $s \in \mathcal{S}$, and show that condition (A.2) holds for θ.

There are two cases. If $x(\theta) = x(\theta^-)$, then $dx(\theta) = 0$, so the claim holds trivially (it does not matter how we define γ for this θ). Otherwise, let us define, for each $s \in \mathcal{S}$,

$$\gamma(s|\,\theta) = \frac{\pi(s|\,\theta)x(\theta) - \pi(s|\,\theta^-)x(\theta^-)}{x(\theta) - x(\theta^-)}.$$

This is a well defined probability for each $s \in \mathcal{S}$ due to condition (M) and inequality (A.1). Moreover, we have

$$\pi(s|\,\theta)x(\theta) = \gamma(s|\,\theta)\left(x(\theta) - x(\theta^-)\right) + \pi(s|\,\theta^-)x(\theta^-).$$

Using the inductive hypothesis, and the fact that $x(\theta) - x(\theta^-) = dx(\theta)$, we obtain

$$\pi(s|\,\theta)x(\theta) = \gamma(s|\,\theta)dx(\theta) + \sum_{c \le \theta^-} \gamma(s|\,c)dx(c) = \sum_{c \le \theta} \gamma(s|\,c)dx(c)$$

which ends the induction and the proof.

A.2 Proof of Lemma 2

First, I complete the proof of Lemma 2 under the assumptions of Proposition 2a. Recall that I have to prove existence of a distribution f such that *(i)* when the third party makes an offer, she offers price θ after seeing signal s if and only if $\pi(s|\,\theta)x(\theta) \ge \pi(s|\,\hat{\theta})x(\hat{\theta})$; otherwise, she offers price $\hat{\theta}$, *(ii)* when the agent makes an offer, and signal s satisfies $\pi(s|\,\theta)x(\theta) \ge \pi(s|\,\hat{\theta})x(\hat{\theta})$, trade takes place with probability one at a price above θ in equilibrium. Recall the assumption that in case of multiplicity, the equilibrium that maximizes the probability of trade in the aftermarket is selected.

First, let's define

$$p^\star_{(x,\,\pi,\,f)}(s) \in \text{argmax}_p \sum_{\theta \le p} (v(\theta) - p)\pi(s|\,\theta)x(\theta)f(\theta),$$

as the optimal price quoted by the third party when she makes an offer, given mechanism frame $(x,\,\pi)$, distribution f, and conditional on signal s. If distribution f is supported on the set $\{\hat{\theta},\,\theta\}$, the optimal price is either $\hat{\theta}$ or θ. Price $\hat{\theta}$ is uniquely optimal if

$$(\theta - \hat{\theta})\pi(s|\,\hat{\theta})x(\hat{\theta})f(\hat{\theta}) > (v(\theta) - \theta)\pi(s|\,\theta)x(\theta)f(\theta).$$

Price θ is uniquely optimal if the opposite strict inequality holds. Define f as the unique distribution supported on $\{\hat{\theta},\,\theta\}$ such that $f(\hat{\theta})/f(\theta) = (v(\theta) - \theta)/(\theta - \hat{\theta})$. Then, f achieves property *(i)*.

To show that property *(ii)* holds under f as well, I prove two facts. First, when the signal s satisfies $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$, the conditional expected value of the third party in the aftermarket is above θ. Second, under the conclusion of the first fact, there exists an equilibrium where trade takes place with probability one, and in all such equilibria prices lie above θ.

To prove the first fact, note that given the definition of distribution f, for each signal s with $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$, the posterior probability $f^s(\theta)$ of the high type θ in the aftermarket is at least β, where β solves $(1-\beta)/\beta = (v(\theta)-\theta)/(\theta-\hat{\theta})$. Therefore, the conditional expected value of the third party is

$$(1 - f^s(\theta))v(\hat{\theta}) + f^s(\theta)v(\theta) \geq v(\hat{\theta}) + \beta(v(\theta) - v(\hat{\theta})).$$

Using the definition of β, we conclude that the expression on the right hand side is higher than θ if and only if $(\theta - \hat{\theta})(v(\theta) - v(\hat{\theta})) \geq (\theta - v(\hat{\theta}))(v(\theta) - \hat{\theta})$. Rearranging, we get $v(\theta)(v(\hat{\theta}) - \hat{\theta}) \geq \theta(v(\hat{\theta}) - \hat{\theta})$ which is always true by the assumption that $v(\tau) > \tau$ for all $\tau \in \Theta$.

To prove the second fact, for a fixed s satisfying $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$, consider a strategy profile in which both types, θ and $\hat{\theta}$ propose a price p equal to the conditional expected value of the third party, and the third party accepts with probability one. This strategy profile is an equilibrium, supported by the off-equilibrium belief of the third party that any price above p is offered by the low type $\hat{\theta}$. The third party is indifferent between accepting and rejecting, so it is a best response to accept. The price p lies above θ, so it is a best response for both types θ and $\hat{\theta}$ to propose p. Finally, in any equilibrium where the probability of trade is one, all equilibrium prices have to lie above θ. Otherwise, because the high type θ can only trade at prices above θ in equilibrium, the low type $\hat{\theta}$ would have a profitable deviation to imitate the price distribution proposed by the high type.

In the remainder of this appendix, I prove Lemma 2 under the assumptions of Proposition 2*b*. The proof is analogous to the previous case, except for the construction of distribution f. I show how to construct distribution f such that $u(\theta;\, f^s) = u(\hat{\theta};\, f^s)$ when $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$, and $u(\theta;\, f^s) > u(\hat{\theta};\, f^s)$ when $\pi(s|\,\theta)x(\theta) < \pi(s|\,\hat{\theta})x(\hat{\theta})$. Once this property is established, the remainder of the proof is identical.

Consider a class of distributions supported on $\{\hat{\theta}, \theta\}$, where $\hat{\theta} < \theta$ are two adjacent types. By the assumption that the lemons condition is locally severe, $v(\hat{\theta}) < \theta < v(\theta)$. Then, I define f as the unique distribution supported on $\{\hat{\theta}, \theta\}$ with $f(\theta) = (\theta - v(\hat{\theta}))/(v(\theta) - v(\hat{\theta}))$. Under f, in the absence of additional information, the conditional expected value of the third party is θ.

Suppose that $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$. After seeing signal s, the third party believes that her conditional expected value is (weakly) above θ. By the same argument as above, in any equilibrium with the maximal probability of trade, trade takes place with probability one at prices that lie weakly above θ. There exists at least one such equilibrium. Therefore, both types θ and $\hat{\theta}$ receive the same continuation payoff, that is, $u(\theta;\, f^s) = u(\hat{\theta};\, f^s)$ for all such s.

Now suppose that $\pi(s|\,\theta)x(\theta) < \pi(s|\,\hat{\theta})x(\hat{\theta})$. After seeing signal s, the third party believes that her conditional expected value is strictly below θ. By the same argument as above, there cannot exist an equilibrium in the aftermarket in which trade takes place with probability one at prices that lie above θ. Therefore, regardless of equilibrium selection, the high type θ must have a strictly higher expected payoff in the aftermarket than the low type $\hat{\theta}$, that is, $u(\hat{\theta};\, f^s) < u(\theta;\, f^s)$ for all such s.

A.3 Proof of Proposition 4

Consider the problem of maximizing

$$\sum_{\theta\in\Theta}\sum_{s\in\mathcal{S}} V(\theta;\, f^s)\pi(s|\,\theta)x(\theta)f(\theta)$$

over π subject to $(x,\,\pi)$ being a cutoff rule. For any cutoff mechanism, by definition, there exists a function γ such that $\pi(s|\,\theta)x(\theta) = \sum_{c\le\theta}\gamma(s|\,c)dx(c)$. Thus, the problem becomes

$$\max_{\gamma}\sum_{\theta\in\Theta}\sum_{s\in\mathcal{S}} V(\theta;\, f^s)\sum_{c\le\theta}\gamma(s|\,c)dx(c)f(\theta).$$

The problem can be rewritten as

$$\max_{\gamma}\sum_{s\in\mathcal{S}}\underbrace{\left(\sum_{c}\gamma(s|\,c)dx(c)\right)}_{\varsigma_s}\sum_{\theta\in\Theta} V(\theta;\, f^s)\underbrace{\left(\frac{\sum_{c\le\theta}\gamma(s|\,c)dx(c)}{\sum_{c}\gamma(s|\,c)dx(c)}\right)}_{G^s(\theta)} f(\theta).$$

In the above expression, ς_s is the unconditional probability of sending signal s, and the remaining expression is equal to $\mathcal{V}(G^s)$, as defined in (3.6), where G^s is the posterior cumulative distribution function of the cutoff conditional on signal s. Thus, the objective function can be written as

$$\mathbb{E}_{s\sim\varsigma}\mathcal{V}(G^s). \tag{A.3}$$

To confirm that $\mathcal{V}$ depends solely on the posterior belief over the cutoff, note that

$$\mathcal{V}(G^s) = \mathbb{E}_{c \sim G^s} \sum_{\theta \in \Theta} V(\theta;\, f^s) \mathbf{1}_{\{\theta \geq c\}} f(\theta).$$

Thus, the problem is formally equivalent to the Bayesian persuasion problem of Kamenica and Gentzkow (2011). Instead of optimizing over distributions ς over signals, we can optimize in the space of distributions over posterior beliefs $\varrho \in \Delta(\Delta(C))$ subject to a Bayes-plausibility constraint. This yields equations (3.7) and (3.8). Equation (3.8) is the Bayes-plausibility constraint on the beliefs about the cutoff c expressed in terms of a cdf.

A.4 Proof of the claims in Example 2

In this appendix, provide proofs of the two claims made in the context of Example 2.

Fix an arbitrary distribution F with continuous density f on $[0,\, 1]$, and consider the problem (3.9) - (3.10) from Example 2. I first analyze an auxiliary problem where the designer is constrained to induce a price of at least p, for some $p \in [0,\, 1]$. In the second step, I optimize over $p \in [0,\, 1]$.

In a problem where the price is constrained to lie above p, it is trivially optimal to set $x(\theta) = 1$ for all $\theta \geq p$ (this does not interact with any constraints). Let $X(\theta) \equiv \int_0^\theta x(\tau) d\tau$. Then, the auxiliary problem can be written as

$$\max_{X \colon [0,\, p] \to \mathbb{R}} X(p) \tag{A.4}$$

subject to

$$X(0) = 0, \tag{A.5}$$

$$0 \leq X'(\theta) \leq f(\theta), \tag{A.6}$$

$$X(\theta) \leq \frac{X(p)(1-p)}{1-\theta},\ \forall \theta \leq p, \tag{A.7}$$

where constraint (A.7) states that no price below p can be strictly optimal. I claim that a solution to the above problem is given by an X induced by $x(\theta) = \mathbf{1}_{\{\theta \geq r\}}$ for some r. To see why, suppose that $X^\star$ solves the above problem, and the corresponding $x^\star$ does not take this form. Consider a function $\underline{X}$ with $\underline{X}'(\theta) = f(\theta)$ for all $\theta \leq p$, and $\underline{X}(p) = X^\star(p)$. Then, $\underline{X}$ lies everywhere below $X^\star$, so it satisfies constraint (A.7), and by design it satisfies (A.6). However, it may violate (A.5) because it is possible that $\underline{X}(0) < 0$. There exists a largest r such that $\underline{X}(r) = 0$. Define $\underline{x}(\theta) = \mathbf{1}_{\{\theta \geq r\}}$. Then, $\underline{x}(\theta)$ has the desired structure, and also gives rise to an optimal solution (because the corresponding X achieves the value

$X^\star(p)$ in the objective function A.4).

Because the above argument holds for any p, we can conclude that the optimal allocation rule takes the form $x(\theta) = \mathbf{1}_{\{\theta \geq r\}}$ for some $r \in [0, 1]$, and thus it is enough to maximize over r in the second step. Let

$$p(r) \in \operatorname{argmax}_\rho (1-\rho)(F(\rho) - F(r)) \tag{A.8}$$

be the optimal price when the distribution is truncated at r. The problem becomes

$$\max_r (F(p(r)) - F(r)) + \int_{p(r)}^1 \theta f(\theta) d\theta.$$

I can relax constraint (A.8) by replacing it with a first-order condition (which has to hold because the solution cannot lie on the boundary):

$$(1-\rho)f(\rho) = F(\rho) - F(r),$$

which allows me to express the final problem as

$$\max_p (1-p)f(p) + \int_p^1 \theta f(\theta) d\theta.$$

subject to $(1-p)f(p) \leq F(p)$. (One needs to verify ex-post that the optimal solution satisfies the second-order condition).

Suppose that $f'(p) \leq f(p)(1+p)/(1-p)$ (which is in particular satisfied by the uniform distribution). Then, the objective function is decreasing, so the optimal solution is given by the smallest $p^\star$ that solves $(1-p)f(p) = F(p)$. Such $p^\star$ corresponds to choosing $r = 0$, or, $x(\theta) = 1$ for all $\theta \in [0, 1]$ (the second-order condition can be shown to be satisfied in this case).

For an example of a distribution that does not yield $r = 0$, consider

$$F_\epsilon(\theta) = \begin{cases} \theta(1 - \frac{\epsilon}{2}) & \theta \leq 1/2 \\ \frac{1-\epsilon\theta}{4(1-\theta)} & \theta \geq 1/2, \end{cases}$$

supported on $[0, 3/(4-\epsilon)]$, for some small ϵ. Intuitively, for $\epsilon = 0$, the third party is indifferent between all prices in $[1/2, 3/4]$ but for any positive ϵ, the price of $1/2$ is uniquely optimal. However, if the prior distribution is truncated at $r > \epsilon/2$, then a price of $3/(4-\epsilon)$ is uniquely optimal. For an arbitrarily small loss of trading probability (of order ϵ), the designer can increase the aftermarket price from $1/2$ to $3/(4-\epsilon)$. Thus, a strictly positive r is optimal for sufficiently small ϵ.

A.5 Supplementary materials for Section 3.5 – information structures induced by cutoff mechanisms

In this appendix, I characterize feasible distributions over posterior beliefs over the type of the agent induced by cutoff mechanisms. See Section 3.5 for a discussion of the significance of obtaining this characterization.

I start with some auxiliary definitions. For a fixed allocation rule x, I call $f^x(\theta)$, defined by (3.5), the *no-communication posterior.* The no-communication posterior is the belief over the type of the agent, conditional on the agent acquiring the good, held by the third party when the allocation function is x, and the mechanism reveals no information. Distribution f_1 *likelihood-ratio dominates* distribution f_2 (denoted $f_1 \succ^{MLR} f_2$) if $f_1(\theta)/f_2(\theta)$ is non-decreasing whenever it is defined.

Proposition 10. *A finite-support distribution of beliefs $\rho \in \Delta(\Delta(\Theta))$ is a conditional distribution of the posterior beliefs over the agent's type (conditional on the agent acquiring the good) induced by a cutoff mechanism with allocation x if and only if*

$$\bar{f} \succ^{MLR} f, \ \forall \bar{f} \in supp(\rho). \tag{A.9}$$

and

$$\mathbb{E}_{\bar{f}\sim\rho}\bar{f}(\theta) \equiv \sum_{\bar{f}\in supp(\rho)} \bar{f}(\theta)\rho(\bar{f}) = f^x(\theta), \tag{A.10}$$

The proof can be found at the end of this appendix. Condition (A.10) is the standard Bayes-plausibility constraint, except that the posterior beliefs have to average out to the no-communication posterior, instead of to the prior. This is because the distribution of beliefs is taken to be conditional on allocating the good. Condition (A.9) is an additional constraint on posterior belief – each posterior has to likelihood-ratio dominate the prior.

Proposition 10 provides a new way of characterizing the optimal payoffs achievable to the mechanism designer. Let

$$\mathcal{W}(\bar{f}) = \sum_{\theta\in\Theta} V(\theta;\, \bar{f})\bar{f}(\theta)$$

be the expected payoff to the mechanism designer that would arise if the prior distribution of types were $\bar{f}$, and the mechanism allocated to all types and revealed no information. Let $M_f \equiv \{\bar{f} \in \Delta(\Theta) :\ \bar{f} \succ^{MLR} f\}$ be the set of distributions that likelihood-ratio dominate the prior f. Note that $M_f = \{f^x : x \in \mathcal{X}\}$.

Proposition 11. *The optimal expected payoff to the mechanism designer in the problem (3.7)-(3.8) (for a fixed allocation x) is equal to*

$$\left(\sum_{\theta\in\Theta} x(\theta)f(\theta)\right) co^{M_f}\,\mathcal{W}(f^x) \tag{A.11}$$

where

$$co^{M_f}\,\mathcal{W}(f^x) \equiv \sup\{y:\ (f^x,\, y)\in CH\big(graph(\mathcal{W})|_{M_f}\big)\},$$

and $graph(\mathcal{W})|_{M_f}$ is the graph of $\mathcal{W}$ restricted to domain M_f.

The above characterization of the optimal payoff is analogous to the one from Corollary 2 but differs in that the concave closure is taken in the space of conditional distributions of beliefs over the type of the agent. I emphasize that the distribution of beliefs is conditional on allocating the good – in Corollary 2, I consider ex-ante distributions of beliefs (over the cutoff). In Proposition 11, I consider ex-post (conditional on allocating the good) distributions of beliefs (over the type). Hence, to obtain the expected payoff in Proposition 11, the conditional expected payoff $\text{co}^{M_f}\,\mathcal{W}(f^x)$ is multiplied by the unconditional probability of allocating the good, $\sum_{\theta\in\Theta} x(\theta)f(\theta)$.

A.5.1 Proof of Proposition 10

I first show that every distribution of beliefs over the cutoff $\varrho\in\Delta(\Delta(C))$ that is feasible under allocation x defines a distribution of beliefs over the type $\rho\in\Delta(\Delta(\Theta))$ which satisfies conditions (A.10)-(A.9). For every $G\in\text{supp}(\varrho)$, let f^G, defined by (3.5), be the corresponding posterior belief over the type. Each f^G satisfies condition (A.9) because G is a non-decreasing function. To show condition (A.10), define

$$\rho(f^G) = \frac{\sum_{\theta\in\Theta} G(\theta)f(\theta)\varrho(G)}{\sum_{\tilde{G}\in\text{supp}(\varrho)}\sum_{\theta\in\Theta}\tilde{G}(\theta)f(\theta)\varrho(\tilde{G})}. \tag{A.12}$$

The expression $\rho(f^G)$ is the probability of inducing belief f^G conditional on allocating the good (the probability distribution over cutoffs ϱ is unconditional, hence the need to transform the probabilities by conditioning on the event that the good was allocated). Because ϱ is a feasible distribution, i.e. it satisfies condition (3.8),

$$\sum_{\tilde{G}\in\text{supp}(\varrho)}\sum_{\theta\in\Theta}\tilde{G}(\theta)f(\theta)\varrho(\tilde{G}) = \sum_{\theta\in\Theta} x(\theta)f(\theta).$$

Then, we have

$$\sum_{G\in\text{supp}(\varrho)} f^G(\theta)\rho(f^G) = \sum_{G\in\text{supp}(\varrho)} \frac{G(\theta)f(\theta)}{\sum_{\tau\in\Theta} G(\tau)f(\tau)} \frac{\sum_{\tau\in\Theta} G(\tau)f(\tau)}{\sum_{\tau\in\Theta} x(\tau)f(\tau)} \varrho(G)$$
$$= \frac{\left(\sum_{G\in\text{supp}(\varrho)} G(\theta)\varrho(G)\right) f(\theta)}{\sum_{\tau\in\Theta} x(\tau)f(\tau)} = \frac{x(\theta)f(\theta)}{\sum_{\tau\in\Theta} x(\tau)f(\tau)} = f^x(\theta),$$

which is condition (A.10).

To show the opposite direction, start with a distribution of beliefs over the agent's type conditional on allocating the good, $\rho \in \Delta(\Delta(\Theta))$, satisfying conditions (A.9) and (A.10) for a non-decreasing allocation rule x. For each $\bar{f} \in \text{supp}(\rho)$, define

$$G^{\bar{f}}(\theta) := \left(x(\bar{\theta})\frac{f(\bar{\theta})}{\bar{f}(\bar{\theta})}\right) \frac{\bar{f}(\theta)}{f(\theta)}, \ \forall \theta \in \Theta,$$

where $\bar{\theta} = \max\{\Theta\}$.[59] Because $\bar{f}$ likelihood-ratio dominates f, the function $G^{\bar{f}}(\theta)$ is non-decreasing and bounded above by 1. Thus, it defines a non-decreasing allocation function, and hence also a distribution over cutoffs. Define a distribution of distributions over the cutoff $\varrho \in \Delta(\Delta(C))$ by

$$\varrho(G^{\bar{f}}) = \rho(\bar{f})\bar{f}(\bar{\theta}) \frac{\sum_{\theta\in\Theta} x(\theta)f(\theta)}{x(\bar{\theta})f(\bar{\theta})}, \ \forall \bar{f} \in \text{supp}(\rho).$$

This is a well defined distribution because, by condition (A.10),

$$\sum_{\bar{f}\in\text{supp}(\rho)} \varrho(G^{\bar{f}}) = \left(\sum_{\bar{f}\in\text{supp}(\rho)} \rho(\bar{f})\bar{f}(\bar{\theta})\right) \frac{\sum_{\theta\in\Theta} x(\theta)f(\theta)}{x(\bar{\theta})f(\bar{\theta})} = f^x(\bar{\theta}) \frac{\sum_{\theta\in\Theta} x(\theta)f(\theta)}{x(\bar{\theta})f(\bar{\theta})} = 1.$$

We can now check that condition (3.8) is satisfied:

$$\sum_{\bar{f}} \varrho(G^{\bar{f}})G^{\bar{f}}(\theta) = \sum_{\bar{f}} \rho(\bar{f})\bar{f}(\theta) \frac{\sum_{\tau\in\Theta} x(\tau)f(\tau)}{f(\theta)} = f^x(\theta) \frac{\sum_{\tau\in\Theta} x(\tau)f(\tau)}{f(\theta)} = x(\theta).$$

Therefore, ϱ is a feasible distribution of beliefs over the cutoff given allocation function x. It remains to be shown that the unconditional distribution ϱ of beliefs over the cutoff gives rise to the conditional distribution ρ of beliefs over the type (conditional on allocating the good). By direct calculation, equation (A.12) holds for ρ and ϱ defined as above.

[59] I have assumed that $x(\bar{\theta}) = 1$ but the proof doesn't make use of it, so I allow for a general $x(\bar{\theta}) \leq 1$.

A.5.2 Proof of Proposition 11

The proof follows almost directly from Proposition 10. Starting from the objective function (3.7) and a distribution ϱ of beliefs over cutoffs, we have

$$\mathbb{E}_{G\sim\varrho}\mathcal{V}(G) = \sum_{G\in\text{supp}(\varrho)}\left(\sum_{\theta\in\Theta} V(\theta;\, f^G)G(\theta)f(\theta)\right)\varrho(G)$$
$$= \left(\sum_{\theta\in\Theta} x(\theta)f(\theta)\right)\sum_{G\in\text{supp}(\varrho)}\underbrace{\sum_{\theta\in\Theta} V(\theta;\, f^G)f^G(\theta)}_{\mathcal{W}(f^G)}\underbrace{\frac{\sum_{\theta\in\Theta} G(\theta)f(\theta)}{\sum_{\theta\in\Theta} x(\theta)f(\theta)}\varrho(G)}_{\rho(f^G)}$$
$$= \left(\sum_{\theta\in\Theta} x(\theta)f(\theta)\right)\mathbb{E}_{\bar{f}\sim\rho}\mathcal{W}(\bar{f}),$$

where the last equality follows from the proof of Proposition 10 – the distribution ρ is the conditional distribution of beliefs over the type (conditional on allocating the good) corresponding to the unconditional distribution of beliefs over the cutoff ϱ.

Given the above representation of the objective function and Proposition 10, the concave closure characterization follows from the usual argument.

B Proofs and supplementary materials for Section 4

B.1 Proof of Lemma 3

First, if $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule, then condition ($\boldsymbol{M}$) follows directly from Definition 5 of cutoff rules. I prove that condition ($\boldsymbol{M}$) implies that an $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule.

Fix $i \in N$ and $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$. Let $\beta_s(\tau) \equiv \pi_i(s|\,\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i})$. By condition ($\boldsymbol{M}$), $\beta_s(\tau)$ is a non-decreasing function, for any s. Summing over $s \in \mathcal{S}$, we get that $x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i.

I now construct the signal function γ_i that satisfies equation (4.2). Each $\beta_s(\tau)$ is a measurable function of τ because both $x_i(\tau,\, \boldsymbol{\theta}_{-i})$ and $\pi_i(s|\,\tau,\, \boldsymbol{\theta}_{-i})$ are measurable in τ. Because $\beta_s(\tau)$ is non-decreasing, it has one-sided limits everywhere and is continuous almost everywhere. According to the convention that I identify mechanisms that differ on a measure-zero set of types, it is without loss of generality to assume that $\beta_s(\tau)$ is right-continuous in τ. It follows that β_s induces a positive $\sigma-$additive measure μ_s on C_i defined by

$$\mu_s((a,\, b] \cap C_i) = \beta_s\left(\max_{b'\in[a,\, b]\cap C_i} b'\right) - \beta_s\left(\min_{b'\in[a,\, b]\cap C_i} b'\right),$$

for any interval $(a,\, b]$ in $[0,\, 1]$. When $a,\, b \in C_i$, the above definition takes a much more transparent form

$$\mu_s((a,\, b] \cap C_i) = \beta_s(b) - \beta_s(a).$$

Because a $\sigma-$additive measure on the Borel $\sigma-$field is uniquely defined by the values it takes on intervals, the above definition uniquely characterizes μ_s.

I will show that the measure μ_s is absolutely continuous with respect to the distribution of the random cutoff $dx_i(\cdot\,,\, \boldsymbol{\theta}_{-i})$. For any $a,\, b \in C_i$, $a < b$, we have

$$\beta_s(b) - \beta_s(a) \leq \sum_{s \in \mathcal{S}} [\beta_s(b) - \beta_s(a)] = x_i(b,\, \boldsymbol{\theta}_{-i}) - x_i(a,\, \boldsymbol{\theta}_{-i}).$$

It follows that if $x_i(b,\, \boldsymbol{\theta}_{-i}) = x_i(a,\, \boldsymbol{\theta}_{-i})$, then $\beta_s(b) - \beta_s(a) = 0$. Because a and b were arbitrary, μ_s is absolutely continuous with respect to $dx_i(\cdot\,,\, \boldsymbol{\theta}_{-i})$.

By the Radon-Nikodym Theorem, there exists a measurable positive function g_s on C_i that is a density of μ_s with respect to $dx_i(\cdot\,,\, \boldsymbol{\theta}_{-i})$. In particular,

$$\beta_s(\theta_i) = \pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) \equiv \mu_s([0,\, \theta_i] \cap C_i) = \int_0^{\theta_i} g_s(c) dx_i(c,\, \boldsymbol{\theta}_{-i}), \tag{B.1}$$

for all θ_i and each $s \in \mathcal{S}$. Moreover, we have

$$\sum_{s \in \mathcal{S}} \pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \sum_{s \in \mathcal{S}} g_s(c) dx_i(c,\, \boldsymbol{\theta}_{-i}).$$

Thus, for all θ_i,

$$\int_0^{\theta_i} \left(\sum_{s \in \mathcal{S}} g_s(c) - 1 \right) dx_i(c,\, \boldsymbol{\theta}_{-i}) = 0.$$

The above equality states that the integral over $[0,\, \theta_i]$, where θ_i is arbitrary, is equal to zero. By a standard argument, this implies that the above integral is zero over any Borel subset of $[0,\, 1]$. It follows that $\sum_s g_s(c) = 1$, dx_i-almost everywhere. I can define the measure γ_i by

$$\gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}) = g_s(c),$$

for dx_i-almost all $c \in C_i$ (and in an arbitrary way on the remaining set of c of dx_i-measure 0). Because $\sum_s g_s(c) = 1$, γ_i is a well defined signal function. Moreover, equation (B.1) implies that the equality (4.2) from the definition of cutoff rules holds for all s, and all θ_i. Because i and $\boldsymbol{\theta}_{-i}$ were arbitrary, $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule.

B.2 Proof of Lemma 4

The proof is similar to the proof of Lemma 2, so I omit some details that are analogous.

I first prove that condition ($\boldsymbol{M}$) is necessary. Fix a mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$, $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$. Since $(\boldsymbol{x}, \boldsymbol{\pi})$ is assumed DS implementable, condition ($\boldsymbol{IC}$) has to hold for θ_i and $\hat{\theta}_i$. In particular, type θ_i cannot find it profitable to report $\hat{\theta}_i$, and vice versa. Summing up the two resulting inequalities, we can cancel out transfers, and obtain (using the fact that the mechanism is $\mathcal{S}$-finite)

$$\sum_{s \in \mathcal{S}} \left[u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s})\right] \left[\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - \pi_i(s|\, \hat{\theta}_i,\, \boldsymbol{\theta}_{-i}) x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})\right] \geq 0. \tag{B.2}$$

For the sake of simplifying the expressions, let $\alpha_s(\tau) \equiv u_i(\tau;\, \boldsymbol{f}^{i,s})$ and $\beta_s(\tau) \equiv \pi_i(s|\, \tau,\, \boldsymbol{\theta}_{-i}) x_i(\tau,\, \boldsymbol{\theta}_{-i})$. By the Richness condition, there exist $\boldsymbol{f} \in \mathcal{F}$, and $A \in \mathcal{A}$ such that conditions (4.3) and (4.4) hold. Under these $\boldsymbol{f}$ and A, inequality (B.2) becomes

$$\sum_{\{s \in \mathcal{S}:\, \beta_s(\theta_i) < \beta_s(\hat{\theta}_i)\}} \left[\alpha_s(\theta_i) - \alpha_s(\hat{\theta}_i)\right] \left[\beta_s(\theta_i) - \beta_s(\hat{\theta}_i)\right] \geq 0,$$

with $\alpha_s(\theta_i) > \alpha_s(\hat{\theta}_i)$ for each signal s in the summation, by condition (4.3). We have thus obtained that a sum of strictly negative terms is non-negative. This is only possible when the set of indices in the sum is empty: $\{s \in \mathcal{S} :\ \beta_s(\theta_i) < \beta_s(\hat{\theta}_i)\} = \emptyset$. Because $\theta_i > \hat{\theta}_i$ were arbitrary, condition ($\boldsymbol{M}$) holds for every signal s. And because i and $\boldsymbol{\theta}_{-i}$ were arbitrary, the first part of Lemma 4 is proven.

To prove that condition ($\boldsymbol{M}$) implies implementability for any $\boldsymbol{f} \in \mathcal{F}$ and $A \in \mathcal{A}$, I again use a condition for checking implementability in arbitrary type and allocation spaces from Dworczak and Zhang (2015). Although Dworczak and Zhang (2015) formally consider one-agent mechanisms, checking implementability in a model with multiple agents boils down to checking that for a fixed (arbitrary) $\boldsymbol{\theta}_{-i}$, conditions ($\boldsymbol{IR}$) and ($\boldsymbol{IC}$) hold for agent i. Because $\mathcal{A}$ is monotone, for any i, $\boldsymbol{f} \in \mathcal{F}$, and $A \in \mathcal{A}$, $u_i(\theta;\, \boldsymbol{f}^{i,s})$ is non-decreasing in θ. Similarly as in the proof of Lemma 2, to show that the sufficient condition of Dworczak and Zhang (2015) holds, it is enough to prove that condition (B.2) holds for any $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$.[60] The fact that $u_i(\theta;\, \boldsymbol{f}^{i,s})$ is non-decreasing in θ implies that the first square bracket is non-negative in each term of the sum, and condition ($\boldsymbol{M}$) implies that the second square bracket is non-negative. Because a sum of non-negative terms is non-negative, inequality

[60] When the type space Θ is continuous, I use the result from the appendix in Dworczak and Zhang (2015) which states that it is enough to check that the matching-efficiency condition holds for all finite subsets of the type space. Thus, the proof goes through without any modifications for a continuous type space.

(B.2) always holds.

B.3 Proof of the claim in Example 3

The proof that the resale game satisfies the Richness condition is very similar to the part of the proof of Lemma 2 contained in Appendix A.2, and I thus skip some details.

Fix a mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$, $i \in N$, $\theta_i > \hat{\theta}_i$, and $\boldsymbol{\theta}_{-i}$. I have to find a prior joint distribution $\boldsymbol{f} \in \mathcal{F}$ such that conditions (4.3) and (4.4) hold. First, I consider the case when $\mathcal{F}$ contains discrete distributions.

Let $\boldsymbol{f} = \times_{j \in N} f_j$ be a product probability mass function with marginals f_j, for $j \in N$. For any $j \neq i$, let $f_j(\theta_j) = 1$ (degenerate distribution), and let $\text{supp}(f_i) = \{\theta_i, \hat{\theta}_i\}$. Intuitively, under $\boldsymbol{f}$, the profile $\boldsymbol{\theta}_{-i}$ is deterministic, which allows me to impose restrictions on the signal distribution of the mechanism under this profile of reports. Define

$$p^{\star}_{(\boldsymbol{x},\boldsymbol{\pi},\boldsymbol{f})}(s) \in \operatorname*{argmax}_{p} \sum_{\theta \leq p} (v - p)\pi_i(s|\,\theta, \boldsymbol{\theta}_{-i})x(\theta, \boldsymbol{\theta}_{-i})f_i(\theta),$$

as the optimal price quoted by the third party when she makes an offer, conditional on signal s. Price $\hat{\theta}_i$ is uniquely optimal if

$$(\theta_i - \hat{\theta}_i)\pi_i(s|\,\hat{\theta}_i, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})f_i(\hat{\theta}_i) > (v - \theta_i)\pi_i(s|\,\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i})f_i(\theta_i).$$

Price θ_i is uniquely optimal if the opposite strict inequality holds. Define f_i as the unique distribution such that $f_i(\hat{\theta}_i)/f_i(\theta_i) = (v - \theta_i)/(\theta_i - \hat{\theta}_i)$. The choice of $\boldsymbol{f}$ implies that in the absence of additional information, the third party is indifferent between prices θ_i and $\hat{\theta}_i$.

Given the assumption about the aftermarket, recalling that $\eta > 0$ is the probability that the third party makes the offer, we have (for a fixed $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{f})$)

$$u_i(\theta;\, \boldsymbol{f}^{i,s}) = \eta \max\{\theta,\, p^{\star}_{(\boldsymbol{x},\boldsymbol{\pi},\boldsymbol{f})}(s)\} + (1 - \eta)v.$$

To prove that condition (4.3) holds, suppose that for some $s \in \mathcal{S}$, $\pi_i(s|\,\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) < \pi_i(s|\,\hat{\theta}_i, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})$. By choice of $\boldsymbol{f}$, we have $p^{\star}_{(\boldsymbol{x},\boldsymbol{\pi},\boldsymbol{f})}(s) = \hat{\theta}_i$ in that case, and thus

$$u_i(\theta_i;\, \boldsymbol{f}^{i,s}) = \eta\theta_i + (1 - \eta)v > \eta\hat{\theta}_i + (1 - \eta)v = u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s}).$$

On the other hand, when $\pi_i(s|\,\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) > \pi_i(s|\,\hat{\theta}_i, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})$, we have $p^{\star}_{(\boldsymbol{x},\boldsymbol{\pi},\boldsymbol{f})}(s) = \theta_i$, and thus

$$u_i(\theta_i;\, \boldsymbol{f}^{i,s}) = \eta\theta_i + (1 - \eta)v = \eta \max\{\hat{\theta}_i,\, p^{\star}_{(\boldsymbol{x},\boldsymbol{\pi},\boldsymbol{f})}(s)\} + (1 - \eta)v = u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s}),$$

which yields condition (4.4).

Now consider the case when $\mathcal{F}$ contains continuous distributions. By Remark 1, it is enough to show that the Richness condition holds for some distribution in the closure of $\mathcal{F}$, as long as there exists an approximating sequence in $\mathcal{F}$ along which the payoff u_i is continuous. The discrete distribution from the first part of the proof can be expressed as a limit of continuous distributions. Because I have assumed that $\pi_i(s|\,\boldsymbol{\theta})x_i(\boldsymbol{\theta})$ is continuous almost everywhere in $\boldsymbol{\theta}$, excluding a set of types of measure zero (which is irrelevant because I do not distinguish between two mechanisms that differ on a measure-zero set of types), the posterior belief over the type of the winner along the approximating sequence of prior distributions converges to the posterior belief for the limiting discrete distribution. Moreover, because $\mathcal{S}$ is assumed finite, we can choose the approximating sequence of distributions in such a way that optimal prices converge to either θ or $\hat{\theta}$, depending on which price is optimal in the limit for a given signal s. I omit the remaining technical details.

Finally, the aftermarket has monotone payoffs. This follows from the definition of u_i.

B.4 Supplementary materials for Subsection 4.1.3

In this section, I formalize the discussion of Subsection 4.1.3.

Fix a cutoff rule $(\boldsymbol{x}, \boldsymbol{\pi})$, an agent i, and a profile $\boldsymbol{\theta}_{-i}$. The cutoff representing the marginal allocation rule for agent i has a cdf $x_i(c, \boldsymbol{\theta}_{-i})$. By Definition 5 of a cutoff rule, we can represent the disclosure rule π_i for agent i by a signal function $\gamma_i(\cdot|\, c, \boldsymbol{\theta}_{-i})$ that only depends on the realization of the cutoff c, and on $\boldsymbol{\theta}_{-i}$. Thus, we can define the posterior belief over the cutoff conditional on any signal realization $s \in \mathcal{S}$ that has positive probability given $(\boldsymbol{x}, \boldsymbol{\pi})$ and $\boldsymbol{\theta}_{-i}$. I denote the cdf of that belief by $x_i(c, \boldsymbol{\theta}_{-i}|\, s)$. The function $x_i(c, \boldsymbol{\theta}_{-i}|\, s)$ can be treated as a conditional allocation rule. The interpretation is that when the agent is allowed to change the report after seeing the signal s, then $x_i(\hat{\theta}, \boldsymbol{\theta}_{-i}|\, s)$ is the conditional expected probability of winning the object with report $\hat{\theta}$. The signal realization is fixed but, having updated the beliefs about the cutoff, the agent might find it profitable to change the report to obtain a different allocation and transfer.

I say that a cutoff rule $(\boldsymbol{x}, \boldsymbol{\pi})$ is *implementable with ex post incentives*, if there exists a transfer function $\boldsymbol{t}^s(\boldsymbol{\theta})$, with $t_i^s : \Theta_i \to \mathbb{R}$ for all $i \in N$, $s \in \mathcal{S}$, such that for any i, $\boldsymbol{\theta}_{-i}$, and s, agent i wants to report truthfully after seeing signal s, given the conditional allocation rule $x_i(\cdot\,, \boldsymbol{\theta}_{-i}|\, s)$ and transfers $t_i^s(\cdot\,, \boldsymbol{\theta}_{-i})$.

Because the conditional allocation rule $x_i(\cdot\,, \boldsymbol{\theta}_{-i}|\, s)$ is a cdf, it is non-decreasing. Hence, by the standard argument, there exist transfers that support that allocation rule, for any i, s, and $\boldsymbol{\theta}_{-i}$.

Corollary 8. *If the aftermarket is monotone, any cutoff rule is implementable with ex post incentives.*

The definition of implementability with ex post incentives does not directly apply to an arbitrary (non-cutoff) mechanism. The cutoff representation ties together the probabilities of winning the object for different reports (the probabilities are defined on a single probability space). For a general mechanism, the notion of a conditional allocation rule (conditional on signal s) is not well defined. However, the conditional allocation rule $x_i(c, \boldsymbol{\theta}_{-i}| s)$ in a cutoff mechanism is uniquely defined (up to rescaling which is irrelevant for incentives) by the property that the posterior belief in the aftermarket conditional on signal s and agent i acquiring the good under the original mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is the same as the posterior belief induced by the allocation rule $x_i(c, \boldsymbol{\theta}_{-i}| s)$ when agent i is known to be the winner but no further information is disclosed. For a general mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$, I define $x_i(c, \boldsymbol{\theta}_{-i}| s)$ as the allocation rule with that property. Then, the definition of implementability with ex post incentives generalizes.

By its defining property, the conditional allocation rule $x_i(c, \boldsymbol{\theta}_{-i}| s)$ has to be proportional to $\pi_i(s| \theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i})$. If $\pi_i(s| \theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i})$ fails to be non-decreasing in θ_i, it is impossible to find transfers that support truthful reporting, as long as different types obtain a different continuation payoff in the aftermarket conditional on signal s. By Lemma 3, monotonicity of $\pi_i(s| \theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i})$ is equivalent to being a cutoff rule. We thus obtain the converse result.

Corollary 9. *Suppose that $u_i(\theta_i, \bar{\boldsymbol{f}}) > u_i(\hat{\theta}_i, \bar{\boldsymbol{f}})$ for all i, $\theta_i > \hat{\theta}_i$, and $\bar{\boldsymbol{f}} \in \Delta(\boldsymbol{\Theta})$. If a mechanism frame is implementable with ex post incentives, then it is a cutoff rule.*

Finally, I compare the notion of implementability with ex post incentives to other similar concepts in the literature. First, the notion of ex post implementability (see for example Bergemann and Morris, 2008) requires that any agent reports truthfully for every realization of private information of other agents. In my model, agents' ex post utility depends on reported, rather than actual, types of other players. It also depends on the realization of the signal s. If one is willing to treat the signal s as an endogenous "type" of the mechanism designer, then implementability with ex post incentives is an extension of ex post implementability that accounts for this additional "type".

Second, posterior implementability, introduced by Green and Laffont (1987), requires that it is optimal for any agent to report truthfully given the information revealed by the mechanism ex post. In the model of Green and Laffont (1987), posterior implementability is weaker than DS implementability. In contrast, I require DS implementability with respect to other agents' reports by assumption, and implementability with ex post incentives can be seen as additionally requiring posterior implementability with respect to the cutoff (or more generally, the conditional allocation rule induced by the realization of the signal).

B.5 Infinite signal spaces in cutoff mechanisms

This section presents supplementary results for Section 4.1.4. I extend Theorem 1 and Theorem 2 to mechanisms with infinite signal spaces, and prove an additional approximation result. Proofs are collected at the end of this appendix.

The conclusion of Theorem 1 extends to all cutoff rules – a cutoff rule is always implementable as long as the aftermarket is monotone.

Theorem 1′. *A cutoff rule is DS implementable for any prior distribution $\boldsymbol{f}$ and any monotone aftermarket A.*

Extending Theorem 2 is more difficult because with an infinite signal space $\mathcal{S}$, the Richness condition cannot in general be expressed via equations (4.3) and (4.4). For a general signal function $\boldsymbol{\pi}$, Bayes' rule may not be applicable. It is then difficult to associate each signal realization with a posterior belief. If the measure $d\pi_i(\cdot\,|\boldsymbol{\theta})$ is continuous for all i and $\boldsymbol{\theta}$, equations (4.3) and (4.4) are well defined when we interpret each π_i as a density.

Theorem 2′. *Suppose that a mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$, whose $\boldsymbol{\pi}$ is a continuous distribution over $\mathcal{S}$ for all i and $\boldsymbol{\theta}$, is flexible with respect to $(\mathcal{F}, \mathcal{A})$. Further, suppose that for all $i \in N$, types $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$, there exists a prior distribution $\boldsymbol{f} \in \mathcal{F}$ and an aftermarket $A \in \mathcal{A}$ such that equations (4.3) and (4.4) hold, with each π_i interpreted as a density. Then, $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule.*

The proof of Theorem **2′** is fully analogous to the proof of Theorem 2 (with sums replaced by integrals, and statements "for all s" replaced by "for almost all s"), and is thus omitted.

Theorem **2′** holds well beyond the case of continuous distributions over signals. For example, the conclusion remains true when the signal is deterministic conditional on each type profile. What is needed is a version of Bayes' rule, and a consistent interpretation of equations (4.3) and (4.4). Instead of pursuing the most general statement, I offer a result that provides a different justification for looking at cutoff mechanisms.

I say that a sequence of mechanism frames $\{(\boldsymbol{x}, \boldsymbol{\pi}^n)\}_{n=1}^{\infty}$ on the same signal space $\mathcal{S}$ converges to $(\boldsymbol{x}, \boldsymbol{\pi})$, if $\boldsymbol{\pi}^n(\cdot\,|\,\boldsymbol{\theta})\boldsymbol{x}(\boldsymbol{\theta})$ converges to $\boldsymbol{\pi}(\cdot\,|\,\boldsymbol{\theta})\boldsymbol{x}(\boldsymbol{\theta})$ in the weak* topology of measures on $\mathcal{S}$, for almost all $\boldsymbol{\theta}$. A mechanism frame with an infinite signal space but finite support of signals is considered $\mathcal{S}$-finite.

Proposition 12. *A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule if and only if it is the limit of $\mathcal{S}$-finite cutoff rules with the same allocation function $\boldsymbol{x}$.*

Proposition 12 implies that in rich settings only cutoff mechanisms can be approximated with flexible mechanisms admitting a finite signal space.

B.5.1 Proof of Theorem 1′

The proof is analogous to the part of proof of Lemma 4 demonstrating that $\mathcal{S}$-finite cutoff rules are always DS implementable.

An analog of equation (B.2) is sufficient for implementability, by the same argument as in the proof of Lemma 4:

$$\int_{\mathcal{S}} \left[u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s})\right] \left[d\pi_i(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - d\pi_i(s|\,\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})\right] \geq 0. \tag{B.3}$$

Using the definition of cutoff mechanisms,

$$\int_{\mathcal{S}} u_i(\tau;\, \boldsymbol{f}^{i,s}) d\pi_i(s|\,\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i}) = \int_{\mathcal{S}} u_i(\tau;\, \boldsymbol{f}^{i,s}) \int_0^{\tau} d\gamma_i(s|\,c,\, \boldsymbol{\theta}_{-i})dx_i(c,\, \boldsymbol{\theta}_{-i}).$$

Thus, equation (B.3) becomes

$$\int_{\mathcal{S}} \int_{\hat{\theta}_i}^{\theta_i} \left[u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s})\right] d\gamma_i(s|\,c,\, \boldsymbol{\theta}_{-i})dx_i(c,\, \boldsymbol{\theta}_{-i}) \geq 0,$$

which always holds because the integrand is positive by the assumption that the aftermarket is monotone.

B.5.2 Proof of Proposition 12

First, suppose that a sequence of $\mathcal{S}$-finite cutoff rules $\{(\boldsymbol{x},\, \boldsymbol{\pi}^n)\}_{n=1}^{\infty}$ converges to some mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi})$. I have to show that $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule.

Fix $\boldsymbol{\theta}$ and $i \in N$. Convergence in the weak* topology means that for any continuous bounded function g on $\mathcal{S}$, we have

$$\lim_n \int_{\mathcal{S}} g(s) d\pi_i^n(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_{\mathcal{S}} g(s) d\pi_i(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}).$$

Because for each n, $(\boldsymbol{x},\, \boldsymbol{\pi}^n)$ is a ($\mathcal{S}$-finite) cutoff rule, we have

$$\int_{\mathcal{S}} g(s) d\pi_i^n(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_{\mathcal{S}} g(s) \int_0^{\theta_i} d\gamma_i^n(s|\,c,\, \boldsymbol{\theta}_{-i})dx_i(c,\, \boldsymbol{\theta}_{-i}),$$

for some probability measure γ_i^n on $\mathcal{S}$. By the Banach-Alaoglu theorem, the set of probability measures is compact in the weak* topology, so (after passing to a subsequence if necessary) we can assume that γ_i^n converges to some γ_i. Thus

$$\lim_n \int_{\mathcal{S}} g(s) d\gamma_i^n(s|\, c,\, \boldsymbol{\theta}_{-i}) = \int_{\mathcal{S}} g(s) d\gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}).$$

By the Fubini's theorem, and the Lebesgue dominated convergence theorem,

$$\begin{aligned}\lim_n \int_{\mathcal{S}} g(s) \int_0^{\theta_i} d\gamma_i^n(s|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}) &= \lim_n \int_0^{\theta_i} \left(\int_{\mathcal{S}} g(s) d\gamma_i^n(s|\, c,\, \boldsymbol{\theta}_{-i}) \right) dx_i(c,\, \boldsymbol{\theta}_{-i}) \\ = \int_0^{\theta_i} \left(\int_{\mathcal{S}} g(s) d\gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}) \right) dx_i(c,\, \boldsymbol{\theta}_{-i}) &= \int_{\mathcal{S}} g(s) \int_0^{\theta_i} d\gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}).\end{aligned}$$

Combining the above equations,

$$\int_{\mathcal{S}} g(s) d\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_{\mathcal{S}} g(s) \int_0^{\theta_i} d\gamma_i(s|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}).$$

Because the above equality is true for all continuous bounded functions g, the two measures must be equal, i.e.

$$\pi_i(S|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \gamma_i(S|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}),$$

for all measurable $S \subseteq \mathcal{S}$. Thus, $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule.

Conversely, suppose that $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a cutoff rule. I have to find a sequence $\{(\boldsymbol{x},\, \boldsymbol{\pi}^n)\}_{n=1}^\infty$ of $\mathcal{S}$-finite cutoff rules that converges to $(\boldsymbol{x},\, \boldsymbol{\pi})$.

Fix $\boldsymbol{\theta}$ and $i \in N$, and consider the measure $\gamma_i(\cdot\,|\, \theta_i,\, \boldsymbol{\theta}_{-i})$ satisfying equation (4.5), defined on $\mathcal{S} = \boldsymbol{\Theta}$. Take an arbitrary discrete approximation of the probability measure $\gamma_i(\cdot\,|\, \theta_i,\, \boldsymbol{\theta}_{-i})$, i.e. a sequence $\{\gamma_i^n(\cdot\,|\, \theta_i,\, \boldsymbol{\theta}_{-i})\}_{n=1}^\infty$ of finite-support measures on $\mathcal{S}$ that converges in weak* topology to γ_i.[61] For each n, define a mechanism frame $(\boldsymbol{x},\, \boldsymbol{\pi}^n)$ by

$$\pi_i^n(S|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \gamma_i^n(S|\, c,\, \boldsymbol{\theta}_{-i}) dx_i(c,\, \boldsymbol{\theta}_{-i}),$$

for all $\boldsymbol{\theta}$, $i \in N$, and measurable $S \subseteq \mathcal{S}$. Because γ_i^n has finite support, $(\boldsymbol{x},\, \boldsymbol{\pi}^n)$ is an $\mathcal{S}$-finite cutoff rule. By the same argument as in the first part of the proof, $(\boldsymbol{x},\, \boldsymbol{\pi})$ is a limit of $\{(\boldsymbol{x},\, \boldsymbol{\pi}^n)\}_{n=1}^\infty$.

[61] Such an approximation can be constructed by discretizing the compact domain of γ_i.

B.6 Proof of Theorem 3

Proof of the direct part: Take an arbitrary symmetric cutoff rule (x, π), and a prior distribution f. Because $\pi(S|\,\theta_i, \boldsymbol{\theta}_{-i})x(\theta_i, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i for each $\boldsymbol{\theta}_{-i}$ (this follows directly from definition of cutoff rules), $\pi_f(S|\,\theta)x_f(\theta)$ is also non-decreasing in θ, for any measurable $S \subseteq \mathcal{S}$. By taking $S = \mathcal{S}$, we conclude that in particular $x_f(\theta)$ is non-decreasing.

To show existence of a measure γ satisfying equation (4.7) in the Definition 13 of a reduced-form cutoff rule, I use an argument similar to the one used in the proof of Lemma 3 in Appendix B.1. Without loss of generality, I can take $\mathcal{S} = \Theta \subseteq [0,\, 1]$. Denote $\beta_S(\tau) \equiv \pi_f(S|\,\tau)x_f(\tau)$, for any measurable $S \subseteq \mathcal{S}$. This time, β_S corresponds to the probability that the signal lies in the set S, to account for the fact that $\mathcal{S}$ can be an infinite space (in which case it doesn't make sense to talk about the probability of the event that the signal is *equal* to some s). I have shown that $\beta_S(\tau)$ is non-decreasing in τ. By the same argument as in the proof of Lemma 3, i.e. by an application of the Radon-Nikodym Theorem, we obtain

$$\pi_f(S|\,\theta)x_f(\theta) = \int_0^\theta g_S(c)dx_f(c), \tag{B.4}$$

for some density function $g_S(c)$, for any measurable $S \subseteq \mathcal{S}$. In particular, I can take $S = [0,\, s] \cap \Theta$, and define $G_c(s) \equiv g_{[0,\,s]\cap\Theta}(c)$, for any $s \in [0,\, 1]$. By direct inspection of the above equality, $G_c(0) = 0$, $G_c(1) = 1$, for dx_f-almost all $c \in C$. I will show that $G_c(s)$ is non-decreasing in s, for dx_f-almost all $c \in C$. To see that, consider $s < s'$, and note that

$$\begin{aligned}\int_0^\theta g_{[0,\,s']\cap\Theta}(c)dx_f(c) &= \pi_f([0, s'] \cap \Theta)x_f(\theta) = \pi_f([0, s] \cap \Theta)x_f(\theta) + \pi_f((s, s'] \cap \Theta)x_f(\theta) \\ &= \int_0^\theta g_{[0,\,s]\cap\Theta}(c)dx_f(c) + \int_0^\theta g_{(s,\,s']\cap\Theta}(c)dx_f(c),\end{aligned}$$

where $g_{(s,\,s']\cap\Theta}(c)$ is obtained by taking $S = (s,\, s'] \cap \Theta$ and applying formula (B.4). It follows that

$$\int_0^\theta \left[g_{[0,\,s']}(c) - g_{[0,\,s]}(c) - g_{(s,\,s']}(c)\right] dx_f(c) = 0,$$

for all $\theta \in \Theta$. Thus, $g_{[0,\,s']}(c) = g_{[0,\,s]}(c) + g_{(s,\,s']}(c)$ for dx_f-almost all c, and in particular, because $g_{(s,\,s']}(c)$ is non-negative, $g_{[0,\,s']}(c) \geq g_{[0,\,s]}(c)$, or $G_c(s') \geq G_c(s)$. Because $s < s'$ were arbitrary, G_c is non-decreasing.

Finally, by monotonicity of $G_c(s)$ and equation (B.4), $G_c(s)$ is right-continuous in s, for dx_f-almost all c.

Therefore, $G_c(s)$ is a cumulative distribution function, for dx_f-almost all c. We can thus define γ, for dx_f-almost all $c \in C$, by

$$\gamma([0,\, s] \cap \Theta |\, c) = G_c(s),$$

for any $s \in [0,\, 1]$. (It is irrelevant how we define γ on the remaining dx_f-measure zero set of points c.) Because a $\sigma-$ additive distribution γ is uniquely determined by the value it assigns to sets of the form $[0,\, s] \cap \Theta$, for all $s \in [0,\, 1]$, by equation (B.4) we get

$$\pi_f(S|\, \theta) x_f(\theta) = \int_0^\theta \gamma(S|\, c) dx_f(c),$$

for all measurable $S \subseteq \mathcal{S}$.

Proof of the converse part: Fix a reduced-form cutoff rule $(\bar{x},\, \bar{\pi})$ under prior distribution f. First, suppose that $(\bar{x},\, \bar{\pi})$ is $\mathcal{S}$-finite, i.e. the signal space $\mathcal{S}$ is finite. By definition of reduced-form mechanisms, there exists a joint (symmetric) allocation function x such that $\bar{x} = x_f$. Define $\pi :\, \Theta^N \to \Delta(\mathcal{S})$ by

$$\pi(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) = \bar{\pi}(s|\, \theta_i),$$

for all $s \in \mathcal{S}, \theta_i \in \Theta, \boldsymbol{\theta}_{-i} \in \Theta^{N-1}$. Then, $(x,\, \pi)$ is a symmetric, Bayesian implementable mechanism frame such that $(x_f,\, \pi_f) = (\bar{x},\, \bar{\pi})$.

The goal is to define a symmetric cutoff rule $(x^\star,\, \pi^\star)$ that induces the same reduced-form: $(x^\star_f,\, \pi^\star_f) = (\bar{x},\, \bar{\pi})$.

To use the proof technique of Gershkov et al. (2013), I introduce the following notation. Let $\mathcal{K} = (N \cup \{0\}) \times \mathcal{S}$ be the set of social alternatives, where an outcome $k = (i,\, s)$ is interpreted as player i getting the object ($i = 0$ denotes the mechanism designer) and signal s being sent. An allocation function in this setting is defined as an element of the set $\mathcal{Y} = \{\{y^{i,\, s}\} :\, y^{i,\, s}(\boldsymbol{\theta}) \geq 0,\, \forall (i,\, s) \in \mathcal{K},\, \sum_{i \in N,\, s \in \mathcal{S}} y^{i,\, s}(\boldsymbol{\theta}) \leq 1,\, \forall \boldsymbol{\theta}\}$. That is, $y^{i,\, s}(\boldsymbol{\theta})$ is the probability of implementing outcome $(i,\, s)$ conditional on type profile $\boldsymbol{\theta}$. Define an allocation function

$$x^{i,\, s}(\boldsymbol{\theta}) = \pi(s|\, \theta_i,\, \boldsymbol{\theta}_{-i}) x(\theta_i,\, \boldsymbol{\theta}_{-i}),$$

for all $i \in N$, and $\boldsymbol{\theta} \in \Theta^N$, as the probability that outcome $\{i,\, s\}$ is implemented in the mechanism $(x,\, \pi)$ (x^0 is defined as the residual probability). Clearly, $\{x^{i,\, s}\} \in \mathcal{Y}$. The following lemma follows directly from the results of Gershkov et al. (2013).

Lemma 5 (Gershkov, Goeree, Kushnir, Moldovanu and Shi, 2013)**.** *Suppose that for allocation* $\{x^{i,s}\}$, $\sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} x^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})$ *is non-decreasing in* θ_i*, for all* $i \in N$, $s \in \mathcal{S}$*. Define* $\{y^{i,s}\}$ *as the solution to the problem*[62]

$$\min_{\{y^{i,s}\}\in\mathcal{D}} \sum_{\boldsymbol{\theta}\in\Theta^N} \sum_{i\in N,\, s\in\mathcal{S}} (y^{i,s}(\boldsymbol{\theta}))^2, \tag{B.5}$$

where

$$\mathcal{D} = \left\{ \{y^{i,s}\} \in \mathcal{Y} : \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} y^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) = \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} x^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}),\ \forall i,\, \theta_i,\, s \right\}.$$

Then, $y^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})$ *is non-decreasing in* θ_i*, for all* $\boldsymbol{\theta}_{-i}$*, and all* $i \in N$, $s \in \mathcal{S}$*.*

The allocation function $\{x^{i,s}\}$ satisfies the assumption of Lemma 5 because

$$\sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} x^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) = \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} \bar{\pi}(s|\,\theta_i)x(\theta_i, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) = \bar{\pi}(s|\,\theta_i)\bar{x}(\theta_i),$$

and the last expression is non-decreasing in θ_i because $(\bar{x}, \bar{\pi})$ is a reduced-form cutoff rule (monotonicity follows directly from condition (4.7) in Definition 13). Given the allocation $\{y^{i,s}\}$ produced from $\{x^{i,s}\}$ by Lemma 5, I now define a mechanism $(x^\star, \pi^\star)$ by

$$x^\star(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{s\in\mathcal{S}} y^{i,s}(\theta_i, \boldsymbol{\theta}_{-i}),$$

and

$$\pi^\star(s|\,\theta_i, \boldsymbol{\theta}_{-i}) = \frac{y^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})}{x^\star(\theta_i, \boldsymbol{\theta}_{-i})},$$

with $\pi^\star(s|\,\theta_i, \boldsymbol{\theta}_{-i})$ defined in an arbitrary way for $x^\star(\theta_i, \boldsymbol{\theta}_{-i}) = 0$. The pair $(x^\star, \pi^\star)$ is a well-defined mechanism, and it is symmetric, without loss of generality.[63] To show that $(x^\star, \pi^\star)$ is a cutoff rule it is enough to invoke Lemma 3; because $\pi^\star(s|\,\theta_i, \boldsymbol{\theta}_{-i})x^\star(\theta_i, \boldsymbol{\theta}_{-i}) \equiv y^{i,s}(\theta_i, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i, for all $s \in \mathcal{S}$ and $\boldsymbol{\theta}_{-i} \in \Theta^{N-1}$, it must be a cutoff rule.

[62] The lemma also guarantees existence of solution.

[63] One can show that there always exists a symmetric solution to problem (B.5), or symmetry can be obtained by ex-ante uniformaly random permutation of identities of agents.

Finally, $(x^\star_f,\, \pi^\star_f) = (\bar{x},\, \bar{\pi})$ follows from the fact that $\{y^{i,\,s}\} \in \mathcal{D}$, and so

$$\sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} \pi^\star(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x^\star(\theta_i,\,\boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) = \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} y^{i,\,s}(\theta_i,\,\boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})$$
$$= \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} x^{i,\,s}(\theta_i,\,\boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}) = \sum_{\boldsymbol{\theta}_{-i}\in\Theta^{N-1}} x(\theta_i,\,\boldsymbol{\theta}_{-i})\pi(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i}),$$

for all s, and θ_i. The same calculation can be done for $x^\star$ by summing over s. This finishes the proof for $\mathcal{S}$-finite reduced-form mechanism frames.

Now consider a general $(\bar{x},\, \bar{\pi})$. By Proposition 12 in Appendix B.5 (which applies to reduced-form cutoff rules), $(\bar{x},\, \bar{\pi})$ can be represented as a limit of $\mathcal{S}$-finite reduced-form cutoff rules $\{(\bar{x},\, \bar{\pi}^n)\}_{n=1}^\infty$ (which can be taken to be symmetric). By the first part of the proof, for each n, there exists a ($\mathcal{S}$-finite) symmetric cutoff rule $(x,\, \pi^n)$ such that $(x_f,\, \pi^n_f) = (\bar{x},\, \bar{\pi}^n)$. Passing to a subsequence if necessary, we can assume that $(\bar{x},\, \bar{\pi}^n)$ converges to some $(x^\star,\, \pi^\star)$. Applying Proposition 12 again, we conclude that $(x^\star,\, \pi^\star)$ is a cutoff rule. Moreover, $(x^\star_f,\, \pi^\star_f) = (\bar{x},\, \bar{\pi})$ (because this equality holds along the sequence).

B.7 Supplementary materials for Section 4.4

In this appendix, I complete the analysis of Section 4.4.

The results of this section follow as a corollary from *(i)* Section 3.5 (along with Appendix A.5) which characterizes feasible distributions of beliefs over the type of the winner in a one-agent cutoff mechanism,[64] and *(ii)* Section 4.2 which shows that symmetric cutoff mechanisms can be represented in reduced forms which bear full analogy to one-agent allocation rules.

Recall that $\mathcal{W}(\bar{f}) = \int_\Theta V(\theta;\, \bar{f})\bar{f}(\theta)d\theta$ denotes the conditional expected payoff to the mechanism designer given posterior belief $\bar{f}$ over the type of the winner, and M^f denotes the set of beliefs over Θ that likelihood-ratio dominate the prior f (see Appendix A.5). Let $\mathcal{X}_{MB}$ denote the set of non-decreasing interim expected allocation functions that satisfy the Matthews-Border condition (M-B). Combining Corollary 3 with Proposition 11 from Appendix A.5, I obtain the following characterization of the optimal payoff to the mechanism designer.

Corollary 10. *The problem of maximizing* (4.9) *over the set of cutoff mechanisms is*

[64] In Appendix A.5, I worked with a finite type space but all results generalize easily to a continuous type space.

equivalent to solving

$$\max_{\bar{x}\in\mathcal{X}_{MB}} \left(N \int_{\Theta} \bar{x}(\theta)f(\theta)d\theta\right) co^{M^f}\, \mathcal{W}(f^{\bar{x}}).$$

In particular, if $\mathcal{W}$ is convex on M^f, it is optimal to disclose the cutoff and types of losers; if $\mathcal{W}$ is concave on M^f, it is optimal to reveal no information.

The only part of the Corollary that requires an additional argument is that $\text{co}^{M^f}\, \mathcal{W}(f^{\bar{x}})$ corresponds to the payoff from full disclosure whenever $\mathcal{W}$ is convex. This is not immediate because the concave closure is taken in the restricted space of beliefs M_f. The conclusion follows because the concave closure of a convex function $\mathcal{W}$ on a compact space M_f is given by decomposing each argument of the function (a belief $\bar{f}$) into a convex combination of extreme points of that space. Extreme points of the space M_f are exactly truncations of the prior f. Because each belief $\bar{f} \in M_f$ is a convex combination of truncations of the prior f, the corollary is proven.

Corollary 10 is analogous to Theorem 4 except that the concave closure is taken in the space of conditional posterior distributions over the type of the winner, instead of in the space of interim expected allocations. It is often more natural to work with the former space because posterior beliefs over types directly influence the aftermarket payoffs.

B.8 Proof of Proposition 6

Consider the first case when W is concave and non-decreasing. Because W is concave, and $M(\bar{f})$ is linear in $\bar{f}$, the functional $\mathcal{W}$ is concave. Thus, for any interim allocation function $\bar{x}$, it is optimal to disclose no information. Using Corollary 10 found in Appendix B.7, we can write the problem as

$$\max_{\bar{x}} \left(\int_0^1 \bar{x}(\theta)f(\theta)d\theta\right) W(M(f^{\bar{x}})) \tag{B.6}$$

subject to

$$\bar{x}(\theta) \text{ is non-decreasing in } \theta, \tag{B.7}$$

$$\int_\tau^1 \bar{x}(\theta)f(\theta)d\theta \leq \frac{1-F^N(\tau)}{N}, \ \forall\, \tau \in [0,\, 1]. \tag{B.8}$$

We can also write the objective function explicitly as

$$\left(\int_0^1 \bar{x}(\theta)f(\theta)d\theta\right) W\left(\frac{\int_0^1 \theta\bar{x}(\theta)f(\theta)d\theta}{\int_0^1 \bar{x}(\theta)f(\theta)d\theta}\right)$$

Consider an auxiliary problem in which we fix $\int_0^1 \bar{x}(\theta)f(\theta)d\theta = \beta$ for some $\beta \leq 1/N$. Since W is non-decreasing, the problem becomes

$$\max_{\bar{x}} \int_0^1 \theta\bar{x}(\theta)f(\theta)d\theta, \tag{B.9}$$

subject to (B.7), (B.8), and

$$\int_0^1 \bar{x}(\theta)f(\theta)d\theta = \beta \tag{B.10}$$

In the above problem, we can think of constraint (B.10) as specifying total probability mass. The structure of the problem implies that it is optimal to shift as much probability mass as possible to the right, subject to constraint (B.8), which will thus hold with equality for sufficiently large τ. Formally, I will show optimality of $\bar{x}(\theta) = F^{N-1}(\theta)\mathbf{1}_{\{\theta\geq r\}}$, where r is chosen so that condition (B.10) holds. Using integration by parts,

$$\int_0^1 \theta\bar{x}(\theta)f(\theta)d\theta = \int_0^1 \underbrace{\left(\int_\theta^1 \bar{x}(\tau)f(\tau)d\tau\right)}_{\Gamma(\theta)} d\theta$$

Ignoring constraint (B.7) for now, the problem is to maximize the above expression over Γ subject to $\Gamma(0) = \beta$, Γ is non-increasing, and $\Gamma(\theta) \leq (1 - F^N(\theta))/N$, for all θ. Clearly, this problem is solved by $\Gamma(\theta) = \min\{\beta, (1 - F^N(\theta))/N\}$. But then $\Gamma(\theta) = \int_\theta^1 F^{N-1}(\tau)\mathbf{1}_{\{\theta\geq r\}}f(\tau)d\tau$, by the definition of r. Moreover, $F^{N-1}(\theta)\mathbf{1}_{\{\theta\geq r\}}$ satisfies constraint (B.7), so it is a solution to problem (B.9).

In the second step, we optimize over $\beta \in [0, 1/N]$ in condition (B.10), which corresponds to optimizing over $r \in [0, 1]$ in the optimal solution to the auxiliary problem. By plugging in the optimal solution from the auxiliary problem to the objective function (B.6), we arrive at

$$\max_{r\in[0,1]} \left(\int_r^1 F^{N-1}(\theta)f(\theta)d\theta\right) W\left(\frac{\int_r^1 \theta F^{N-1}(\theta)f(\theta)d\theta}{\int_r^1 F^{N-1}(\theta)f(\theta)d\theta}\right).$$

This expression corresponds exactly to equation (4.16) in Proposition 6, and thus the first case is proven.

Consider now the case when W is concave and decreasing. In this case, the auxiliary problem reads

$$\min_{\bar{x}} \int_0^1 \theta \bar{x}(\theta) f(\theta) d\theta,$$

subject to (B.7), (B.8), and (B.10). This time, all the probability mass under $\bar{x}$ should be shifted to the left, subject to the monotonicity constraint (B.7). Thus, the optimal $\bar{x}$ will be constant, equal to β. Because $\beta \leq 1/N$, such $\bar{x}$ satisfies the Matthews-Border condition (B.8), and corresponds to uniform randomization over agents.

In the second step, because W was assumed non-negative, optimization over β yields $\beta = 1/N$, i.e. β should be set to the maximal feasible level. Such mechanism always allocates the good (to a randomly selected agent). This finishes the proof of the second case.

Finally, for the third case, assume that W is convex. Then, the functional $\mathcal{W}$ is convex, so it is optimal to fully disclose the cutoff representing the interim allocation rule $\bar{x}$. Full disclosure means that any posterior belief $\bar{f} \in M_f$ is decomposed into a distribution over truncations of the prior distribution f. Recall that $\bar{x}$ can be treated as a cdf of the cutoff. Therefore,

$$\mathrm{co}^{M_f} \mathcal{W}(\bar{f}) = \int_0^1 W(m(c)) \, \frac{1 - F(c)}{\int_0^1 \bar{x}(\theta) f(\theta) d\theta} d\bar{x}(c).$$

The additional term $(1 - F(c))/(\int_0^1 \bar{x}(\theta) f(\theta) d\theta)$ appears because, by definition, the payoff W is a conditional expected payoff conditional on allocating the good, so the ex-ante probability of cutoff c is transformed into a conditional probability (by conditioning on the event $\theta \geq c$). The objective function (B.6) can be written as

$$\max_{\bar{x}} \int_0^1 W(m(c)) \, (1 - F(c)) d\bar{x}(c).$$

Using integration by parts (by assumption, W is differentiable) we obtain

$$\int_0^1 W(m(c)) \, (1 - F(c)) d\bar{x}(c) = -W(m(0)) \bar{x}(0^-) - \int_0^1 \frac{d}{dc} \left[W(m(c)) \, (1 - F(c)) \right] x(c) dc.$$

Because $\bar{x}$ represents a cdf in the above equation, $\bar{x}(0^-)$, the left limit of $\bar{x}$ at 0, is equal to zero. By letting $w(c) \equiv W(m(c))$, the objective function can be written as

$$\max_{\bar{x}} \int_0^1 \frac{-\frac{d}{dc} \left[W(m(c)) \, (1 - F(c)) \right]}{f(c)} \bar{x}(c) f(c) dc = \max_{\bar{x}} \int_0^1 \underbrace{\left[w(c) - w'(c) \frac{1 - F(c)}{f(c)} \right]}_{J_w(c)} \bar{x}(c) f(c) dc.$$

The conclusion of Proposition 6 now follows from an argument analogous to the one used above. If $J_w(c)$ is non-positive for $c \leq \underline{r}$, and positive non-decreasing for $c \geq \underline{r}$, then it is optimal to set $\bar{x}(\theta) = 0$ for $\theta \in [0, \underline{r}]$, and push all the mass under $\bar{x}$ on $[\underline{r}, 1]$ to the right, subject to constraint (B.8). This gives us $\bar{x}(\theta) = F^{N-1}(\theta)\mathbf{1}_{\{\theta \geq \underline{r}\}}$. Under this $\bar{x}$, the distribution of the cutoff has a continuous part which is the distribution of a second highest type conditional on that type exceeding $\underline{r}$, and an atom at $\underline{r}$, with mass equal to the probability that the second highest type is below $\underline{r}$. Full disclosure of such a cutoff can be obtained by disclosing the realization of the second highest type but in the case when the second highest type is below $\underline{r}$, it is enough to inform that the second highest type was below $\underline{r}$ (such message leads to the same posterior belief over the type of the winner).

To finish the proof of Proposition 6, I have to show that when $W(c)$ is increasing and log-concave, then there exists $\underline{r}$ such that $J_w(c)$ is non-positive for $c \leq \underline{r}$, and positive non-decreasing for $c \geq \underline{r}$. It is enough to prove that $J_w(c) \geq 0$ implies $J_w'(c) \geq 0$.

By direct calculation, we have

$$m'(c) = (m(c) - c)\frac{f(c)}{1 - F(c)}.$$

The inequality $J_w(c) \geq 0$ implies that

$$m(c) - c \leq \frac{W(m(c))}{W'(m(c))}.$$

Using the assumption that $W'' \geq 0$, and the above inequality,

$$J_w'(c) = W'(m(c)) - W''(m(c))(m(c) - c) \geq W'(m(c)) - W''(m(c))\frac{W(m(c))}{W'(m(c))}$$

Using the fact that $W' \geq 0$, the above expression is greater than zero if and only if $(W')^2 \geq W''W$ which is equivalent to log-concavity of W.

C Proofs and supplementary materials for Section 5

C.1 Proof of Claim 1

The problem (5.3) of the regulator can be equivalently stated as

$$\max_{x \in \mathcal{X},\, p \in [0,\,1]} \int_0^p \left[\lambda v(\theta) + (1 - \lambda)\theta - k\right] x(\theta) f(\theta) d\theta + \int_p^1 (\theta - k) x(\theta) f(\theta) d\theta \tag{C.1}$$

subject to

$$\int_0^p (v(\theta) - p)x(\theta)f(\theta)d\theta \geq 0. \tag{C.2}$$

We can solve the problem in two steps, by first optimizing over x, and then over p. For a fixed $p \in [0, 1]$, the problem is linear, and we can apply optimal control techniques. I show that for any p, the optimal x is a threshold rule.

I first prove that $x^\star(\theta) = 1$ for $\theta \geq p$ at the optimal solution $x^\star$. In the case $p \geq k$, this is obvious. Suppose that $p < k$. Then, $x^\star(\theta) = 1$ for $\theta \geq k$, and $x^\star(\theta) = x^\star(p)$ for $\theta \in (p, k)$. The latter conclusion follows from the fact that the objective function is maximized by minimizing x point-wise in the interval (p, k) and x has to be non-decreasing. Because the objective function is linear in x on $[0, k)$, and the constraint is preserved when x is multiplied by a positive scalar, we must have either $x^\star(p) = 0$ or $x^\star(p) = 1$ (boundary solution). In the first case, we conclude that $x^\star(\theta) = \mathbf{1}_{\{\theta \geq k\}}$, and thus it is impossible that $p < k$. In the second case, we obtain the desired conclusion. By the above, we can ignore the term $\int_p^1 (\theta - k)x(\theta)f(\theta)d\theta$ in the optimization.

To deal with the constraint (C.2), I introduce an auxiliary state variable Γ with $\Gamma'(\theta) = (v(\theta) - p)x(\theta)f(\theta)$, $\Gamma(0) = 0$ and $\Gamma(p) \geq 0$. By the Mangasarian Sufficiency Theorem (see for example Seierstad and Sydsaeter, 1987), to prove optimality of a feasible candidate solution $x^\star$, it is enough to find a continuous and piece-wise continuously differentiable function $q(\theta)$ such that, for all $\theta \in [0, 1]$,

$$x^\star(\theta) \in \underset{x\in[0,1]}{\text{argmax}}\, H(x, \theta, q) \equiv \underset{x\in[0,1]}{\text{argmax}}\, [\lambda v(\theta) + (1-\lambda)\theta - k + q(\theta)(v(\theta) - p)]x(\theta)f(\theta),$$

$$q'(\theta) = 0,\ q(p) \geq 0\, (= 0 \text{ if } \Gamma(p) > 0)$$

$$H(x, \theta, q(\theta)) \text{ is concave in } x.$$

Define $q(\theta) \equiv q_0 \geq 0$. Then, $\lambda v(\theta) + (1-\lambda)\theta - k + q(\theta)(v(\theta) - p)$ is strictly increasing, so the function $x^\star$ that maximizes the Hamiltonian H point-wise is given by $x^\star(\theta) = \mathbf{1}_{\{\theta \geq r\}}$ for some $r \in [0, 1]$. The Hamiltonian H is linear in x, so the last condition is satisfied. There are two cases. If condition (C.2) holds with $r = \underline{r}$ defined as solution to equation

$$\lambda v(\underline{r}) + (1-\lambda)\underline{r} = k, \tag{C.3}$$

then we can define $q_0 = 0$, and $x^\star(\theta) = \mathbf{1}_{\{\theta \geq \underline{r}\}}$ is optimal. In the opposite case, suppose that (C.2) fails with $r = \underline{r}$. Because $v(1) > 1$, there must exist an $r^\star > \underline{r}$ such that (C.2) holds with equality. Then, I have to prove existence of $q_0 \geq 0$ such that $\lambda v(\theta) + (1 - \lambda)\theta - k + q_0(v(\theta) - p) = 0$ at $\theta = r^\star$. Since we must have $v(r^\star) < p$ when (C.2) holds with equality, we can define $q_0 = (\lambda v(r^\star) + (1-\lambda)r^\star - k)/(p - v(r^\star)) \geq 0$, where the inequality follows from the definition of $\underline{r}$ and the fact that $r^\star > \underline{r}$. Thus, in this case, $x^\star(\theta) = \mathbf{1}_{\{\theta \geq r^\star\}}$ is optimal.

Because an optimal $x^\star$ is a threshold rule for every p, it is without loss of generality to restrict attention to threshold rules when looking for the solution to problem (C.1).

Abusing notation slightly, let $p(r) = p(\mathbf{1}_{\{\theta \geq r\}})$, where $p(x)$ for $x \in \mathcal{X}$ is defined in (5.2). Then, the optimal allocation function is given by $x^\star(\theta) = \mathbf{1}_{\{\theta \geq r^\star_{\text{eff}}\}}$, where

$$r^\star_{\text{eff}} = \operatorname*{argmax}_r \int_r^{p(r)} \left[\lambda v(\theta) + (1-\lambda)\theta - k\right] f(\theta) d\theta + \int_{p(r)}^1 (\theta - k) f(\theta) d\theta.$$

This finishes the proof of Claim 1.

C.2 Derivation of the objective function for Subsection 5.1.3

Given (x, π), the probability $y(\theta)$ that the agent with type θ holds the asset after the second stage is given by

$$y(\theta) = \lambda \int_{\mathcal{S}} x(\theta) \mathbf{1}_{\{\theta \geq p(f^s)\}} d\pi(s|\,\theta) + (1-\lambda) x(\theta),$$

for all θ, where $d\pi(\cdot|\,\theta)$ is the distribution over signals conditional on θ, and $p(f^s)$ is the equilibrium price given posterior belief f^s (note that f^s depends on x). Using the envelope formula, we can calculate expected utility $U(\theta)$ of type θ as

$$U(\theta) = U(0) + \int_0^\theta y(\tau) d\tau = U(0) + \lambda \int_0^\theta \left(\int_{\mathcal{S}} x(\tau) \mathbf{1}_{\{\tau \geq p(f^s)\}} d\pi(s|\,\tau) \right) d\tau + (1-\lambda) \int_0^\theta x(\tau) d\tau.$$

In a profit-maximizing mechanism, $U(0) = 0$, and thus transfers are given by

$$\begin{aligned} t(\theta) = \lambda \left(\int_{\mathcal{S}} \max(\theta,\, p(f^s)) d\pi(s|\,\theta) x(\theta) - \int_0^\theta \left(\int_{\mathcal{S}} x(\tau) \mathbf{1}_{\{\tau \geq p(f^s)\}} d\pi(s|\,\tau) \right) d\tau \right) \\ + (1-\lambda) \left(\theta x(\theta) - \int_0^\theta x(\tau) d\tau \right). \end{aligned}$$

Using integration by parts, seller's expected profit can be expressed as

$$\int_0^1 \int_{\mathcal{S}} \left[\lambda \left(p(f^s) \mathbf{1}_{\{\theta \leq p(f^s)\}} + \mathbf{1}_{\{\theta > p(f^s)\}} J(\theta) \right) + (1-\lambda) J(\theta) - k \right] d\pi(s|\,\theta) x(\theta) f(\theta) d\theta,$$

where $J(\theta) \equiv \theta - (1-F(\theta))/f(\theta)$ is the virtual surplus function. The objective function takes the form (3.4) with $V(\theta;\, f^s) = \lambda\left(p(f^s)\mathbf{1}_{\{\theta\le p(f^s)\}} + \mathbf{1}_{\{\theta > p(f^s)\}}J(\theta)\right) + (1-\lambda)J(\theta) - k$, so we can apply Theorem 4 to conclude that the optimal mechanism reveals no information. Thus, the problem is to maximize

$$\int_0^{p(x)} \left(\lambda p(x) + (1-\lambda)J(\theta) - k\right) x(\theta) f(\theta) d\theta + \int_{p(x)}^1 (J(\theta)-k)x(\theta)f(\theta)d\theta.$$

C.3 Proof of Claim 2

The problem of the seller can be equivalently stated as

$$\max_{x\in\mathcal{X},\, p\in[0,\,1]} \int_0^p \left[\lambda p + (1-\lambda)J(\theta) - k\right] x(\theta) f(\theta) d\theta + \int_p^1 (J(\theta)-k)x(\theta)f(\theta)d\theta \tag{C.4}$$

subject to

$$\int_0^p (v(\theta)-p)x(\theta)f(\theta)d\theta \geq 0. \tag{C.5}$$

Because I assumed that $J(\theta)$ is non-decreasing, by the same argument as in the proof of Claim 1, the optimal $x^\star$ is a threshold rule: $x^\star(\theta) = \mathbf{1}_{\{\theta\ge r\}}$ for some $r\in[0,\,1]$. The optimal threshold level $r^\star_{\text{rev}}$ can be defined as

$$r^\star_{\text{rev}} = \text{argmax}_r \int_r^{p(r)} \left[\lambda p(r) + (1-\lambda)J(\theta) - k\right] f(\theta) d\theta + \int_{p(r)}^1 (J(\theta)-k)f(\theta)d\theta.$$

To prove the second part of Claim 2, I start by analyzing some properties of $r^\star_{\text{eff}}$. Define $\bar{r}$ by

$$\int_{\bar{r}}^1 (v(\theta)-1)f(\theta)d\theta = 0. \tag{C.6}$$

Because $v(1) > 1$, and v is strictly increasing, $\bar{r}$ is well defined, and by the assumption $\int_k^1 (v(\theta)-1)f(\theta) < 0$, we have $\bar{r} > k$. It follows that if $p(r^\star_{\text{eff}}) = 1$, then $r^\star_{\text{eff}} = \bar{r}$. On the other hand, $r^\star_{\text{eff}}$ cannot be lower than $\underline{r}$ defined by (C.3).

Suppose that we are in the regular case $\underline{r} < \bar{r}$. For $r\in[\underline{r},\,\bar{r}]$, we have

$$\int_r^{p(r)} (v(\theta)-p(r))f(\theta)d\theta = 0. \tag{C.7}$$

Using the implicit function theorem

$$p'(r) = \frac{(v(r)-p(r))f(r)}{(v(p(r))-p(r))f(p(r)) - F(p(r)) + F(r)}.$$

The denominator must be negative because, by definition of $p(r)$, the expression $\int_r^p (v(\theta) - p) f(\theta) d\theta$ changes sign from positive to negative as a function of p at $p = p(r)$. Moreover, $v(r) < p(r)$ because v is strictly increasing. Thus, $p'(r) > 0$.

In the opposite (irregular) case $\underline{r} \geq \bar{r}$, we have $r^\star_{\text{eff}} = \underline{r}$. Indeed, it is never optimal to choose an r below $\underline{r}$, and because in this case $p(r) = 1$ for all $r \geq \underline{r}$, it is also suboptimal to choose $r > \underline{r}$.

I now prove Claim 2 by considering the regular and irregular cases separately.

In the regular case $\underline{r} < \bar{r}$, because $J(\theta) \leq \theta$, we have $r^\star_{\text{rev}} \geq \underline{r}$. Denoting by $V_{\text{eff}}(r)$ and $V_{\text{rev}}(r)$ the expected value to the mechanism designer under allocation rule $\mathbf{1}_{\{\theta \geq r\}}$, in the cases of maximizing efficiency and revenue, respectively, we can write

$$V_{\text{rev}}(r) = V_{\text{eff}}(r) - \int_r^{p(r)} \left(\lambda(v(\theta) - p(r)) + (1 - \lambda) \frac{1 - F(\theta)}{f(\theta)} \right) f(\theta) d\theta - \int_{p(r)}^1 (1 - F(\theta)) d\theta.$$

For $r \in (\underline{r}, \bar{r})$, using equation (C.7), we obtain,

$$V'_{\text{rev}}(r) = V'_{\text{eff}}(r) + (1 - \lambda)(1 - F(r)) + \lambda(1 - F(p(r)) p'(r) > V'_{\text{eff}}(r).$$

Because of the assumption $v(0) < k$, $r^\star_{\text{eff}}$ is never equal to 0, and due to $v(1) > 1$, it is never equal to 1. Thus, the first-order condition must hold: $V'_{\text{eff}}(r^\star_{\text{eff}}) = 0$.[65] There are two cases. If $r^\star_{\text{eff}} < \bar{r}$, we must have $r^\star_{\text{rev}} > r^\star_{\text{eff}}$, due to $V'_{\text{rev}}(r^\star_{\text{rev}}) = 0$ and $V'_{\text{rev}}(r) > V'_{\text{eff}}(r)$. By definition of $\bar{r}$, $p(r^\star_{\text{eff}}) < 1$ in this case. In the second case, $r^\star_{\text{eff}} = \bar{r}$, $p(r^\star_{\text{eff}}) = 1$, and I need to show that $r^\star_{\text{rev}} \geq \bar{r}$ with equality for sufficiently large λ. By the above analysis, we know that $r^\star_{\text{rev}}$ cannot be strictly lower than $\bar{r}$, so I only need to prove that there exists $\lambda^\star < 1$ such that $r^\star_{\text{rev}} = \bar{r}$ for all $\lambda \geq \lambda^\star$. For $r \geq \bar{r}$, because $p(r) = 1$, we have

$$V_{\text{rev}}(r) = \int_r^1 (\lambda + (1 - \lambda) J(\theta) - k) f(\theta) d\theta.$$

Because J is increasing, it is optimal to take $r^\star_{\text{rev}} = \bar{r}$ if and only if

$$\lambda + (1 - \lambda) J(\bar{r}) - k \geq 0.$$

This allows us to define

$$\lambda^\star = \begin{cases} \frac{k - J(\bar{r})}{1 - J(\bar{r})} & \text{if } J(\bar{r}) \leq k \\ 0 & \text{if } J(\bar{r}) > k. \end{cases}$$

Clearly, $\lambda^\star < 1$, and due to $v(1) > 1$, $\bar{r} < 1$, so $\lambda^\star$ is well defined.

[65] If V_{eff} or V_{rev} is not differentiable at some r, we can use subdifferentials instead of derivatives.

Finally, I consider the irregular case $\underline{r} \geq \bar{r}$ (when $r^\star_{\text{eff}} = \underline{r}$). By definition of $\bar{r}$, $p(r^\star_{\text{eff}}) = 1$, so I have to prove that $r^\star_{\text{rev}} \geq \underline{r}$ with equality if and only if λ is sufficiently high. This follows from the same reasoning as above, where in the derivation of $\lambda^\star$, $\bar{r}$ is replaced by $\underline{r}$. By definition (C.3) and assumption $v(1) > 1$, we have $\underline{r} < 1$, so $\lambda^\star < 1$ is well defined.

C.4 Example for Section 5.1.3

Suppose that $\lambda = 1$, $v(\theta) = \delta\theta + \Delta$ for $\delta \in (0, 2)$, $\Delta \in [0, k]$, $\delta + \Delta > 1$, and F is the uniform distribution on $[0, 1]$. If the allocation function gives the good to all types above threshold r, the price in the aftermarket is

$$p(r) = \max\left\{\frac{\delta r + 2\Delta}{2 - \delta}, 1\right\}.$$

Under the above assumptions, $\underline{r} = (k - \Delta)/\delta$, $\bar{r} = (2 - \delta - 2\Delta)/\delta$ ($\bar{r} > k$ by assumption 5.1). Assume that we are in the regular case $\underline{r} < \bar{r}$.[66] Then, by direct calculation,

$$r^\star_{\text{eff}} = \begin{cases} \frac{4(\delta-1)\Delta + (2-\delta)^2 k}{\delta(4-3\delta)} & \text{if } \delta \leq \frac{4-2\Delta-2k}{3-k} \\ \bar{r} & \text{if } \delta > \frac{4-2\Delta-2k}{3-k} \end{cases},$$

and

$$r^\star_{\text{rev}} = \begin{cases} \frac{\delta - 2\Delta + (2-\delta)k}{2\delta} & \text{if } \delta \leq \frac{4-2\Delta-2k}{3-k} \\ \bar{r} & \text{if } \delta > \frac{4-2\Delta-2k}{3-k} \end{cases}.$$

For example, when $\delta = 1$, we have $r^\star_{\text{eff}} = k$, and $r^\star_{\text{rev}} = 1/2 - \Delta + k/2 > k$. The efficient and profit-maximizing allocations are closer when the cost is higher, and when the lemons problem becomes less severe (Δ rises).

C.5 Proof of Claim 3

By Corollary 10 found in Appendix B.7, optimization over disclosure rules (for any fixed allocation rule) can be performed directly in the space of posterior beliefs over the winner's type. Recall that $\mathcal{W}(\bar{f})$ is the conditional expected payoff to the mechanism designer conditional on posterior belief $\bar{f}$.

To calculate $\mathcal{W}(\bar{f})$, let $\bar{y} = \mathbb{E}_{\bar{f}}(\theta)$ denote the posterior mean. Then, negotiations happen in the aftermarket if and only if the realized cost k does not exceed $\eta[(\delta - 1)\bar{y} + \Delta]$.

[66] In the opposite case $\underline{r} \geq \bar{r}$, we would have $r^\star_{\text{eff}} = r^\star_{\text{eff}} = \underline{r}$.

Therefore,

$$\mathcal{W}(\bar{f}) = \int_0^1 V(\theta,\, \bar{y})\bar{f}(\theta)d\theta,$$

where

$$V(\theta,\, \bar{y}) = H(\eta[(\delta-1)\bar{y}+\Delta])(\delta\theta+\Delta) + (1 - H(\eta[(\delta-1)\bar{y}+\Delta]))\,\theta.$$

Thus, we get

$$\mathcal{W}(\bar{f}) = H(\eta[(\delta-1)\bar{y}+\Delta])((\delta-1)\bar{y}+\Delta) + \underbrace{\int_0^1 \theta\bar{f}(\theta)d\theta}_{\bar{y}}.$$

In particular, $\mathcal{W}$ depends on $\bar{f}$ only through its mean $\bar{y}$.

The range of $\eta[(\delta-1)\bar{y}+\Delta]$ is contained in $[0,\, \bar{k}]$ under the assumptions of Section 5.2. When $H(x)x$ is convex, $\mathcal{W}$ is convex in $\bar{f}$ (as a composition of a convex function with a linear functional). By Corollary 10 in Appendix B.7, it is optimal to disclose all available information. If the allocation rule is to give the good to the agent with the highest type, then it is enough to disclose the realization of the second highest type. When $H(x)x$ is concave, $\mathcal{W}$ is concave in $\bar{f}$. By Corollary 10 in Appendix B.7, it is optimal not to reveal any information.

D Proofs and supplementary materials for Section 6

D.1 Proof of Proposition 7

I only prove the first part of the proposition because the second part (about implementation of the no-revelation rule) follows from standard arguments ($v^{\pi}(\theta,\, \theta)$ is always nondecreasing in this case).

Suppose that π is the full-disclosure rule. In all three designs considered in Proposition 7, a necessary condition for robust implementation of $(x,\, \pi)$ is that there exists an equilibrium in strictly increasing bidding strategies – otherwise, disclosing the second bid does not correspond to disclosing the second highest value. Using the first-order condition, I can derive the unique candidate equilibrium bidding function.

In a SPA, it has to be that for the bidding function $\beta^{SPA}(\theta)$,

$$\theta \in \text{argmax}_{\hat{\theta}} \int_0^{\hat{\theta}} \left(v^{\pi}(\theta,\, \tau) - \beta^{SPA}(\tau)\right) dF^{N-1}(\tau), \tag{D.1}$$

for any $\theta \in \Theta$. From the first-order condition,

$$\beta^{SPA}(\theta) = v^{\pi}(\theta,\, \theta).$$

Therefore, the bidding function is strictly increasing if and only if $v^{\pi}(\theta, \theta)$ is strictly increasing in θ. Equation (D.1) holds with the above bidding function because $v^{\pi}(\theta,\, \tau)$ is non-decreasing in θ. Thus, it is optimal for type θ to bid $\beta^{SPA}(\theta)$, given that other players do it as well. In this equilibrium, disclosing the price paid by the winner corresponds exactly to disclosing the value of the second highest bidder. Thus, $(x,\, \pi)$ is robustly implemented.

In a FPA, for the bidding function $\beta^{FPA}(\theta)$, we must have

$$\theta \in \text{argmax}_{\hat{\theta}} \int_0^{\hat{\theta}} \left(v^{\pi}(\theta,\, \tau) - \beta^{FPA}(\hat{\theta}) \right) dF^{N-1}(\tau), \tag{D.2}$$

for any $\theta \in \Theta$. From the first-order condition,

$$\beta^{FPA}(\theta) = \frac{\int_0^{\theta} v^{\pi}(\tau,\, \tau) dF^{N-1}(\tau)}{F^{N-1}(\theta)}.$$

If $v^{\pi}(\theta, \theta)$ is not strictly increasing in θ, there exists a distribution f such that $\beta^{FPA}(\theta)$ is not strictly increasing in θ. On the other hand, if $v^{\pi}(\theta, \theta)$ is strictly increasing in θ, then the bidding function is strictly increasing for any f, and equation (D.2) holds – it is optimal for type θ to bid $\beta^{FPA}(\theta)$, given that other players do it as well.

Finally, in an all-pay auction, for the bidding function $\beta^{APA}(\theta)$, we must have

$$\theta \in \text{argmax}_{\hat{\theta}} \left\{ \int_0^{\hat{\theta}} v^{\pi}(\theta,\, \tau) dF^{N-1}(\tau) - \beta^{APA}(\hat{\theta}) \right\}, \tag{D.3}$$

for any $\theta \in \Theta$. From the first-order condition,

$$\beta^{APA}(\theta) = \int_0^{\theta} v^{\pi}(\tau,\, \tau) dF^{N-1}(\tau).$$

Because $v^{\pi}(\theta, \theta)$ is strictly positive, this bidding function is always strictly increasing. Equation (D.3) holds because $v^{\pi}(\theta,\, \tau)$ is non-decreasing in θ – it is optimal for type θ to bid $\beta^{APA}(\theta)$, given that other players do it as well. Thus, we have an equilibrium.

D.2 Additional definitions and discussion for Section 6.2

In this appendix, I complete the formal analysis of Section 6.2.

Let N^t denote the number of active bidders at the end of round t. A GCA is called Markov, if P^t depends on H^{t-1} only through N^{t-1} and Y^t depends on H^t only through (N^{t-1}, N^t), the number of active bidders at the beginning and at the end of the last round. If the auction ends at t, N^t can be either 0 or 1, depending on whether the winner was declared active or inactive.

A pure strategy for an agent participating in a GCA is a mapping $a_i : \Theta \times \mathcal{T} \times \mathcal{H} \to \{0, 1\}$, i.e. for a type $\theta \in \Theta$, in round $t \in \mathcal{T}$, given a partial history $H^{t-1} \in \mathcal{H}^{t-1}$, $a_i(\theta, t, H^{t-1})$ specifies whether type θ exits in round t or not.[67] A strategy for agent i is monotone if for any two types $\theta > \hat{\theta}$, any $t \in \mathcal{T}$ and $H^{t-1} \in \mathcal{H}^{t-1}$, we have $a_i(\theta, t, H^{t-1}) \geq a_i(\hat{\theta}, t, H^{t-1})$. A strategy is Markov if $a_i(\theta, t, H^{t-1})$ depends on H^{t-1} only through N^{t-1}. Mixed strategies σ_i are defined in the usual way. I call a mixed-strategy σ_i monotone if it is a randomization over monotone pure-strategies a_i.

An *equilibrium* is a Perfect Bayesian Equilibrium of the GCA with payoffs determined by the outcome of the auction and the aftermarket $A \equiv \{u(\theta; \bar{f}) : \theta \in \Theta, \bar{f} \in \Delta(\Theta)\}$. Given a strategy profile σ, if an agent with type θ wins the auction and signal s is released, I denote the posterior belief over the winner's type by f^s_σ. In that case, the ex-post payoff of the winner is $u(\theta; f^s_\sigma)$.

D.3 Proof and discussion of Theorem 5

Proof of the direct part: I first introduce some notation. A deterministic price path $p = (p^t)_{t \geq 1}$, a monotone pure-strategy profile a, and type profile $\boldsymbol{\theta}$, together pin down a unique time of exit for every agent. I let $\Gamma_i^{(y, p, a)}(\theta_i, \boldsymbol{\theta}_{-i})$ denote the exit time of agent i with type θ_i, when other types are $\boldsymbol{\theta}_{-i}$. Because strategies are assumed to be monotone, Γ_i is non-decreasing in θ_i. Let H^τ_0 denote the history in which the final winner becomes inactive in the last round τ, and let H^τ_1 denote the history in which the final winner remains active in the last round τ. For the tuple (y, p, a), the corresponding allocation

[67] Given the set of feasible actions and the definition of public history, it is irrelevant whether bidders condition their strategies on private or public histories.

and revelation rules are given by

$$\pi_i^{(y,\,p,\,a)}(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i^{(y,\,p,\,a)}(\theta_i,\,\boldsymbol{\theta}_{-i})$$
$$= \begin{cases} \left(\frac{1}{N^{\tau-1}}\right) Y^{\tau}(H_0^{\tau})(s) + \left(\frac{N^{\tau-1}-1}{N^{\tau-1}}\right) Y^{\tau}(H_1^{\tau})(s) & \text{if } \Gamma_i^{(y,\,p,\,a)}(\boldsymbol{\theta}) > \tau \equiv \max_{j\neq i} \Gamma_j^{(y,\,p,\,a)} \\ \left(\frac{1}{N^{\tau-1}}\right) Y^{\tau}(H_0^{\tau})(s) & \text{if } \Gamma_i^{(y,\,p,\,a)}(\boldsymbol{\theta}) = \tau \equiv \max_{j\neq i} \Gamma_j^{(y,\,p,\,a)} \\ 0 & \text{otherwise,} \end{cases} \tag{D.4}$$

for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and $s \in \mathcal{S}$. In the above definition, τ denotes the last round. If agent i is the only agent who decides to stay in round τ (first case), she wins the auction, and there is a $1/N^{\tau-1}$ probability that the all bidders will be announced inactive, in which case the signal is drawn from the distribution conditional on history H_0^{τ}. Otherwise, the winner i is active, and the signal is drawn from distribution $Y^{\tau}(H_1^{\tau})$. If all bidders, including agent i, decide to exit in round τ (second case), there is a $1/N^{(\tau-1)}$ probability that agent i receives the good, and the signal is drawn from $Y^{\tau}(H_0^{\tau})$. Finally, if agent i exits before round τ (third case), she does not win the good.

For a Generalized Clock Auction $(Y,\,P) = \{(Y^t,\,P^t)\}_{t\geq 1}$, monotone mixed strategy profile σ, and type profile $\boldsymbol{\theta}$ we can define

$$\pi_i(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i(\theta_i,\,\boldsymbol{\theta}_{-i}) = \mathbb{E}_{(y,\,p)\sim(Y,\,P),\,a\sim\sigma}\, \pi_i^{(y,\,p,\,a)}(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i^{(y,\,p,\,a)}(\theta_i,\,\boldsymbol{\theta}_{-i}).$$

Each $\pi_i^{(y,\,p,\,a)}(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i^{(y,\,p,\,a)}(\theta_i,\,\boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i, for any $s \in \mathcal{S}$, by direct inspection of equation (D.4). Therefore, $\pi_i(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i(\theta_i,\,\boldsymbol{\theta}_{-i})$ is also non-decreasing in θ_i. This means that condition ($\boldsymbol{M}$) from Lemma 3 holds. By Lemma 3, $(\boldsymbol{x},\,\boldsymbol{\pi})$ is a cutoff rule, and thus the first part of Theorem 5 is thus proven.

Discussion: The restriction to monotone equilibria in the first part of Theorem 5 is almost without loss of generality in the sense that higher types always have a weakly higher incentive to stay in the auction. Non-monotone equilibria are only possible when a set of types are indifferent between staying and exiting but higher types exit with higher probability.

The conclusion of the first part of Theorem 5 relies on the distinction of whether the winner is active or inactive at the end of the auction (this piece of information is included in the public history which determines the signal distribution). Informally, this is because the auction must protect the privacy of the winner to a sufficient degree to guarantee that the disclosure rule will be always implementable. If the auction always disclosed the decision of the winner (whether she decided to exit or to stay in the last round), the implemented mechanism frame could fail to be a cutoff mechanism. To see that, fix $(y,\, p,\, a)$ and $\boldsymbol{\theta}_{-i}$, and note that the random cutoff representing the allocation for agent i in the GCA has a binary distribution on $\{\theta^{\tau},\, \theta^{\tau+1}\}$ with probabilities $1/N^{\tau-1}$ and $(N^{\tau-1}-1)/N^{\tau-1}$, respectively, where θ^{τ} is the smallest type of agent i who exits in round τ, and $\theta^{\tau+1}$ is the smallest type of agent i who does not exit up to and including round τ. A cutoff mechanism only reveals the realization of the cutoff so the auctioneer cannot always disclose whether the type of the winner was above or below $\theta^{\tau+1}$.[68] By introducing the random determination of the status of the winner (active or inactive), I formally incorporated the cutoff into the definition of a GCA.

Proof of the converse part: In the first step of the proof, given a cutoff rule $(x,\, \pi)$ with a decomposable allocation function, I construct a Bayesian equivalent mechanism frame $(x,\, \pi')$ (see Definition 11). The equivalent disclosure rule π' will have the feature that the signal distribution only depends on the number of active bidders. In the second step, I show how to implement $(x,\, \pi')$ using a Markov GCA.

Because x is decomposable, it can be represented as a convex combination of hierarchical allocation rules (the convex combination is finite because there are finitely many hierarchical auctions when the type space is finite):

$$x(\boldsymbol{\theta}) = \sum_{\alpha} \lambda^{\alpha} x^{\kappa_1^{\alpha}\ldots\kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}),$$

for some $\lambda^{\alpha} \geq 0$, $\sum_{\alpha} \lambda^{\alpha} = 1$, and hierarchy $\kappa_1^{\alpha}\ldots\kappa_{k(\alpha)}^{\alpha}$, for each α. Because $(x,\, \pi)$ is a symmetric cutoff rule, there exists a signal function γ such that for all s, and $\boldsymbol{\theta}$,

$$\pi(s|\, \theta_i,\, \boldsymbol{\theta}_{-i})x(\theta_i,\, \boldsymbol{\theta}_{-i}) = \sum_{c\leq\theta_i} \gamma(s|\, c,\, \boldsymbol{\theta}_{-i})dx(c,\, \boldsymbol{\theta}_{-i}),$$

For a hierarchy $\kappa_1,\, ...,\, \kappa_k$, define

$$\kappa^{\kappa_1\ldots\kappa_k}(\boldsymbol{\theta}_{-i}) = \max\{\kappa_m \,:\, \kappa_m \leq max_{j\neq i}\theta_j\},$$

and

$$n^{\kappa_1\ldots\kappa_k}(\boldsymbol{\theta}_{-i}) = |\{j \in N \setminus \{i\} \,:\, \theta_j \geq \kappa^{\kappa_1\ldots\kappa_k}(\boldsymbol{\theta}_{-i})\}|.$$

In words, $\kappa^{\kappa_1\ldots\kappa_k}(\boldsymbol{\theta}_{-i})$ is the highest of the thresholds $\kappa_1, ..., \kappa_k$ that at least one type in

[68] Even if a cutoff mechanism discloses the cutoff exactly, there are cases in which the type of the winner is above $\theta^{\tau+1}$ but the mechanism only informs that the cutoff was θ^{τ}.

$\boldsymbol{\theta}_{-i}$ exceeds, and $n^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta}_{-i})$ is the number of types in $\boldsymbol{\theta}_{-i}$ that exceed $\kappa^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta}_{-i})$. The vector $\nu^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta}_{-i}) \equiv (\kappa^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta}_{-i}), n^{\kappa_1 \dots \kappa_k}(\boldsymbol{\theta}_{-i}))$ is a sufficient statistic for $\boldsymbol{\theta}_{-i}$ needed to implement $x_i^{\kappa_1 \dots \kappa_k}(\theta_i, \boldsymbol{\theta}_{-i})$.

Define a symmetric cutoff rule (x', π') by

$$\pi'(s|\,\theta_i, \boldsymbol{\theta}_{-i})x'(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{\alpha} \lambda^{\alpha} \sum_{c \leq \theta_i} \gamma'(s|\,c, \nu^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i}))dx^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(c, \boldsymbol{\theta}_{-i}),$$

for all s, and $\boldsymbol{\theta}$, where the signal function γ' is defined by,

$$\gamma'(s|\,c, \nu) = \frac{\sum_{\{\boldsymbol{\theta}_{-i}:\, \nu = \nu^{\kappa_1^{\alpha}, \dots, \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i})\}} \gamma(s|\,c, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})}{\sum_{\{\boldsymbol{\theta}_{-i}:\, \nu = \nu^{\kappa_1^{\alpha}, \dots, \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i})\}} \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})},$$

for any feasible vector ν. Note that $x' = x$, and π' averages the signal distribution under π across all $\boldsymbol{\theta}_{-i}$ that lead to the same allocation rule for agent i, i.e. to the same vector ν. The mechanism frames (x, π) and (x, π') are Bayesian equivalent. Moreover, (x, π') can be decomposed into hierarchical mechanism frames in such a way that the allocation and signal distribution depend on $\boldsymbol{\theta}_{-i}$ only through the sufficient statistic $\nu(\boldsymbol{\theta}_{-i})$.

In the second step of the proof, I show how to implement (x, π') in a GCA. By definition of (x, π'), it is enough to show that the hierarchical mechanism frame

$$\pi^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(s|\,\theta_i, \boldsymbol{\theta}_{-i})x^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{c \leq \theta_i} \gamma'(s|\,c, \nu^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i}))dx^{\kappa_1^{\alpha} \dots \kappa_{k(\alpha)}^{\alpha}}(c, \boldsymbol{\theta}_{-i}),$$

can be implemented in a GCA with a price path that only depends on the number of active bidders, for any α. The claim of the second part of Theorem 5 can then be obtained by randomizing over α according to the distribution $\{\lambda^{\alpha}\}$ in round 0 of the GCA.[69] In the remainder of the proof, I fix α and omit it from the notation – I will denote the hierarchy to be implemented by $\kappa_1, \dots, \kappa_k$. The description of the auction and the equilibrium is kept informal to avoid additional notation. Intuitively, in the equilibrium I construct, bidders with types in $[\kappa_t, \kappa_{t+1})$ exit in round t.

First, I specify the signal distribution for every possible outcome of the bidding process. Without loss of generality, I can assume that the auction ends no later than in round k in equilibrium.[70] If the auction ends in round $\tau \leq k$, there are two cases. Either *(i)* all $N^{\tau-1}$ bidders become inactive in round τ, or *(ii)* $N^{\tau-1} - 1$ bidders become inactive and exactly one bidder remains active. In case *(i)*, the signal s is drawn from

[69] I implicitly assume that the mechanism designer informs the bidders about the realization of α in round 0, i.e. discloses which hierarchical auction will be used.

[70] It is enough to set the price to a prohibitively high level in the subsequent round.

distribution $\gamma'(s|\,\kappa_\tau,\,(\kappa_\tau,\,N^{\tau-1}-1))$. In case *(ii)*, the signal s is drawn from distribution $\gamma'(s|\,\kappa_{\tau+1},\,(\kappa_\tau,\,N^{\tau-1}-1))$. In particular, the signal distribution depends on the public history of the auction only through $N^{\tau-1}$ and N^τ (the latter variable determines which case, *(i)* or *(ii)*, is used).

Second, I specify the price function P^t, for each $t \leq k$. Proceeding recursively from the last round k, one can calculate the expected continuation payoff for each type θ, conditional on the number of active bidders, under the assumption that in any round t', exactly bidders with types above $\kappa_{t'}$ are active (this assumption pins down the posterior belief over the types of active bidders in any subgame). I set the price P^t, as a function of the number of active bidders, to be such that type κ_{t+1} is indifferent between exiting and staying. In particular, P^t is only a function of N^{t-1}, the number of active bidders at the beginning or round t.

Third, I specify equilibrium strategies for bidders. In every round t, given the number of active bidders, type θ stays in the auction if the expected continuation payoff strictly exceeds the expected payoff from dropping out, and exits if the reverse strict inequality holds. In case of indifference type θ drops out in round t if and only if $\theta < \kappa_{t+1}$.

Fourth, by the specification of the signal distribution and the fact that bidders with types $\theta \in [\kappa_t,\,\kappa_{t+1})$ exit in round t, if bidders follow the above strategies, the auction implements the desired mechanism frame $(x^{\kappa_1\ldots\kappa_k},\,\pi^{\kappa_1\ldots\kappa_k})$. In particular, if a bidder considers a deviation, she faces a choice that is analogous to choosing a type to report given the mechanism frame, with the caveat that the agent might have access to some additional information about the types of other bidders. Because the mechanism frame is non-decreasing in θ_i conditional on every profile $\boldsymbol{\theta}_{-i}$, it is also non-decreasing in θ_i given any belief about the profile $\boldsymbol{\theta}_{-i}$.

Fifth, I argue why the above profile of strategies constitutes a Bayesian Perfect Equilibrium of the GCA. In every observable history of the game, bidder i's beliefs about $\boldsymbol{\theta}_{-i}$ coincide with the public belief, and therefore the expected continuation payoff of θ_i does not depend on i. Because in round t the price P^t is set in such a way that type κ_{t+1} is indifferent between exiting or not, in any history, because of monotonicity of the aftermarket, any type $\theta < \kappa_{t+1}$ finds it optimal to exit, and every type $\theta \geq \kappa_{t+1}$ finds it optimal to stay.

E Proofs and supplementary materials for Section 7

E.1 Additional analysis and examples for Section 7.1.3

In this appendix, I complete the analysis of Section 7.1.3 by deriving the form of the optimal allocation rule, and discussing two examples.

In the extended model, I define the function

$$V_{\text{eff}}(\theta;\, k,\, f^s) \equiv \mathbf{1}_{\{\theta \leq p(f^s)\}}(\lambda v(\theta) + (1-\lambda)\theta - k) + (1 - \mathbf{1}_{\{\theta \leq p(f^s)\}})(\theta - k),$$

where $p(f^s)$ is the highest price at which trade happens given posterior belief f^s,

$$p(f^s) = \max\{p \in [0,\, 1] : \int_0^p (v(\theta) - p) f^s(\theta) d\theta \geq 0\}.$$

Then, the function $\mathcal{U}$, defined by equation (7.2), is given by the value of the optimization problem

$$\mathcal{U}(\bar{\alpha}) \equiv \max_{\{y_k\} \in \mathcal{X}^{|\mathcal{K}|},\, p} \sum_{k \in \mathcal{K}} \bar{\alpha}(k) \int_0^p \left[\lambda v(\theta) + (1-\lambda)\theta - k\right] y_k(\theta) f(\theta) d\theta + \int_p^1 (\theta - k) y_k(\theta) f(\theta) d\theta$$

subject to

$$\int_0^p (v(\theta) - p) \left(\sum_{k \in \mathcal{K}} \bar{\alpha}(k) y_k(\theta) \right) f(\theta) d\theta \geq 0.$$

Fix p in the above problem, and a feasible vector of non-decreasing allocation functions $\{y_k\}$. Consider an alternative function y_k' for some $k \in \mathcal{K}$ such that y_k' "first-order stochastically dominates" y_k. That is, $\int_0^1 y_k(\theta) f(\theta) d\theta = \int_0^1 y_k'(\theta) f(\theta) d\theta$, and $\int_0^\tau y_k'(\theta) f(\theta) d\theta \leq \int_0^\tau y_k(\theta) f(\theta) d\theta$, for all $\tau \in [0,\, 1]$. Then, for any increasing function $g : \Theta \to \mathbb{R}$,

$$\int_0^1 g(\theta) y_k'(\theta) f(\theta) d\theta \geq \int_0^1 g(\theta) y_k'(\theta) f(\theta) d\theta.$$

This implies that replacing y_k with y_k' weakly raises the value of the objective function while keeping the solution feasible.

It follows that the optimal y_k is given as the maximal element in the first-order stochastic dominance order defined above. Because y_k has to be non-decreasing, the maximal element in the first-order stochastic dominance order takes the form $y_k(\theta) = \mathbf{1}_{\{\theta \geq \tau_k(\bar{\alpha})\}}$ for some $\tau_k(\bar{\alpha}) \in [0,\, 1]$, for each k. The thresholds $\tau_k(\bar{\alpha})$ typically depend non-trivially on $\bar{\alpha}$ because the optimization problem changes with $\bar{\alpha}$.

By Proposition 8, the value of the problem is given by $\text{co}\,\mathcal{U}(\alpha)$. The optimal mechanism induces some distribution ς of beliefs $\bar{\alpha}$ over k (the beliefs average out to the prior α), and conditional on posterior belief $\bar{\alpha}$, the allocation is a threshold rule for every k. The support of ς has at most $|\mathcal{K}|$ elements, without loss of generality. By Corollary 6, the realization of k is a sufficient statistic for the signal distribution. As a consequence, the unconditional allocation rule $x_k(\theta)$ is a step function for every k, $x_k(\theta) = \sum_{\bar{\alpha} \in \text{supp}(\varsigma)} \varsigma(\bar{\alpha}) \mathbf{1}_{\{\theta \geq \tau_k(\bar{\alpha})\}}$. The optimization problem defining function $\mathcal{U}(\bar{\alpha})$ can thus be expressed as

$$\max_{\{\tau_k\} \in [0,\,1]^{|\mathcal{K}|},\, p \in [0,\,1]} \sum_{k \in \mathcal{K}} \bar{\alpha}(k) \int_{\tau_k}^{p \vee \tau_k} \left[\lambda v(\theta) + (1-\lambda)\theta - k\right] f(\theta) d\theta + \int_{p \vee \tau_k}^{1} (\theta - k) f(\theta) d\theta \quad \text{(E.1)}$$

subject to

$$\sum_{k \in \mathcal{K}} \bar{\alpha}(k) \int_{\tau_k \wedge p}^{p} (v(\theta) - p) f(\theta) d\theta \geq 0.$$

The problem consists in maximizing a function over $|\mathcal{K}| + 1$ real variables subject to a constraint. It is difficult to obtain an analytical solution due to non-differentiability at the points where $\tau_k = p$ for some k, and because the solution will sometimes lie on the boundary (for example, we may have $p = 1$).

Example. I present a numerical solution for the case $\lambda = 1$, $v(\theta) = \delta\theta$, and f uniform on $[0,\,1]$. Suppose that $\mathcal{K} = \{0,\,1/2\}$, and let α be the probability that the cost k is 0, and suppose that the prior distribution is symmetric, $\alpha = 1/2$. I will consider $\delta = 1.25$ and $\delta = 1.5$ (see Figure E.1).

Consider the case of low gains from trade ($\delta = 1.25$). First, I determine the shape of the function $\mathcal{U}$ by solving the problem (E.1) for a fixed belief α that the cost is equal to zero. When α is low (expected cost is high), the optimal τ_0 and $\tau_{1/2}$ are relatively high (see the graph in bottom left corner of Figure E.1). By excluding enough low types, the mechanism induces a price $p = 1$ in the aftermarket, overcoming the lemons problem. As the expected cost decreases (α increases), it becomes less and less beneficial to keep the price p high at the cost of not allocating the asset in the first stage. At some point (roughly when α crosses $1/2$), it is no longer optimal to alleviate the lemons problem. The aftermarket collapses (price is 0), there is no resale, and the asset is allocated to all types above the cost.

The resulting function $\mathcal{U}$ is convex in α, and thus it is optimal to fully disclose the cost k, i.e. induce degenerate posterior beliefs (represented by the two white dots in the graph in the top left corner of Figure E.1). When the realization of k is 0, the mechanism discloses that $k = 0$, and allocates the good to all types. In this case, there is no positive price p at which trade can happen in the second stage, so the agent is the final owner of the asset. In the opposite case $k = 1/2$, the mechanism discloses that $k = 1/2$ and allocates only to type above 0.6 which ensures that the price is 1 in the aftermarket – resale always happens. In this example, the message sent by the mechanism plays a deciding role in shaping the outcome of the market interaction. Intuitively, full disclosure is optimal because the form of the optimal mechanism changes with the expected cost.

Now consider the case of high gains from trade ($\delta = 1.5$). Compared to the previous case, it becomes more beneficial to allocate to the third party. The graph in the bottom right corner of Figure E.1 shows that it is always optimal (regardless of α) to induce a price equal to 1 in the second stage by excluding enough low types in the first stage.

In this case, the function $\mathcal{U}$ is concave (see the top right corner of Figure E.1), so it is optimal to reveal no information. Because the form of the optimal mechanism is the same for all α, it is easier to ensure a high resale price by pooling the two realizations of k into one posterior belief.

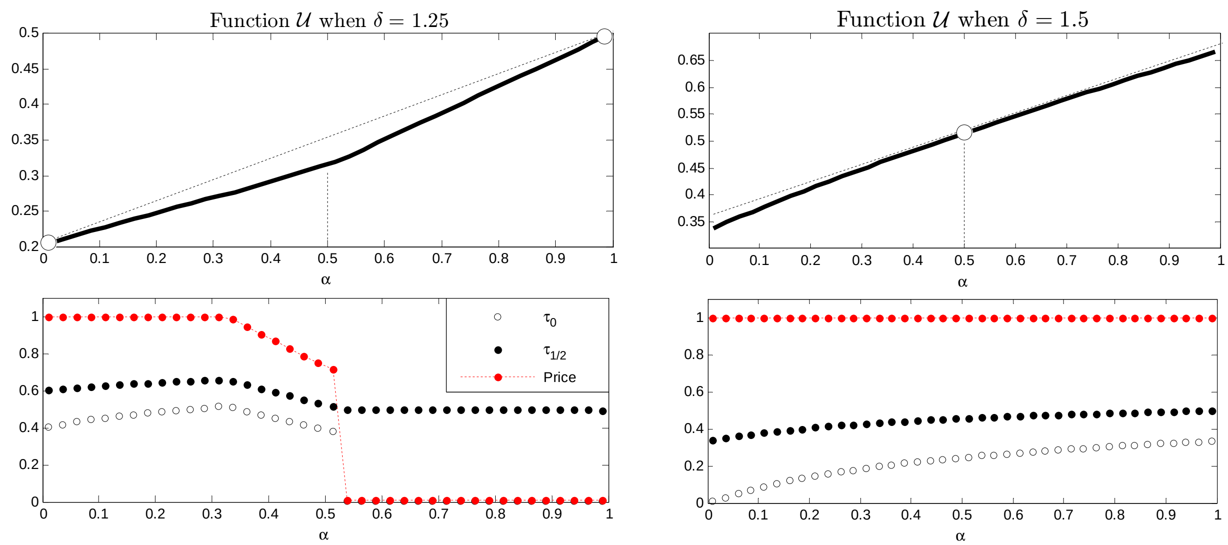

Fig. E.1: Function $\mathcal{U}$ and the optimal prices and thresholds.

E.2 Proof of Claim 4

The claim follows easily from other results in Section 7.1.

First, consider a feasible choice of Nature to send a fully revealing signal about the cost k. In this case, by the analysis of appendix E.1, the optimal mechanism is a posted-price mechanism for every realization of k. The mechanism only reveals information about k but the exact form of the disclosure policy is irrelevant because the third party observes the cost k anyway. This provides an upper bound on the expected payoff of the mechanism designer, call it $\bar{V}$.

To show that this upper bound is achieved, consider the mechanism with the same allocation functions as in the above case and full disclosure of k. Since k is fully disclosed, Nature's choice is irrelevant. Thus, this mechanism achieves the same expected payoff $\bar{V}$. Because $\bar{V}$ is the upper bound on the designer's payoff, no other mechanism gives a better welfare guarantee.

E.3 Proof of Claim 5 and Proposition 9

The proof closely resembles other proofs in this paper, so I omit some details.

In the extended setting, a necessary and sufficient condition for implementability is that for all $\theta > \hat{\theta}$, distributions f, and monotone aftermarkets A_l and A_w,

$$\sum_{s\in\mathcal{S}} \left[u_l(\theta;\, f_l^s) - u_l(\hat{\theta};\, f_l^s)\right] \left[\pi_l(s|\,\theta)(1-x(\theta)) - \pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))\right] + \sum_{s\in\mathcal{S}} \left[u_w(\theta;\, f_w^s) - u_w(\hat{\theta};\, f_w^s)\right] \left[\pi_w(s|\,\theta)x(\theta) - \pi_w(s|\,\hat{\theta})x(\hat{\theta})\right] \geq 0. \quad \text{(E.2)}$$

In equation (E.2), f_l^s denotes the posterior belief over the type of the agent conditional on not acquiring the object and signal s being sent,

$$f_l^s(\theta) = \frac{\pi_l(s|\,\theta)(1-x(\theta))}{\sum_\tau \pi_l(s|\,\tau)(1-x(\tau))},\ \theta \in \Theta,$$

and f_w^s denotes the posterior belief conditional on the agent acquiring the object and signal s being sent,

$$f_w^s(\theta) = \frac{\pi_w(s|\,\theta)x(\theta)}{\sum_\tau \pi_w(s|\,\tau)x(\tau)},\ \theta \in \Theta.$$

Proof of Claim 5. I will show that if condition (E.2) holds for all distributions and monotone aftermarkets then

$$\pi_l(s|\,\theta)(1-x(\theta)) \text{ is non-decreasing,} \quad \text{(E.3)}$$

and

$$\pi_w(s|\,\theta)x(\theta) \text{ is non-decreasing,} \quad \text{(E.4)}$$

for all $s \in \mathcal{S}$. These two conditions imply that x has to be constant (by summing up over $s \in \mathcal{S}$, we see that $x(\theta)$ has to be both non-decreasing and non-increasing).

First, set $u_w(\theta;\, \bar{f}) = \theta$, for all $\bar{f}$ and θ. For a fixed $\theta > \hat{\theta}$, let the aftermarket A_l be such that $u_l(\theta;\, f_l^s) = u_l(\hat{\theta};\, f_l^s)$ for all $s \in \mathcal{S}_1$, where

$$\mathcal{S}_1 \equiv \{s \in \mathcal{S} :\ \pi_l(s|\,\theta)(1 - x(\theta)) \geq \pi_l(s|\,\hat{\theta})(1 - x(\hat{\theta}))\},$$

and $u_l(\theta;\, f_l^s) - u_l(\hat{\theta};\, f_l^s) = \alpha(\theta - \hat{\theta})$, for all $s \notin \mathcal{S}_1$, for some $\alpha > 0$.[71] Then, equation (E.2) becomes

$$(\theta \;-\; \hat{\theta}) \left\{ \alpha \sum_{s \notin \mathcal{S}_1} \left[\pi_l(s|\,\theta)(1 - x(\theta)) - \pi_l(s|\,\hat{\theta})(1 - x(\hat{\theta})) \right] + x(\theta) - x(\hat{\theta}) \right\} \;\;\geq\;\; 0. \quad \text{(E.5)}$$

If there exists $s \notin \mathcal{S}_1$, then by taking a sufficiently high α, we obtain a contradiction in the above inequality. Thus, $\mathcal{S}_1 = \mathcal{S}$, and by definition of $\mathcal{S}_1$, $\pi_l(s|\,\theta)(1 - x(\theta)) \geq \pi_l(s|\,\hat{\theta})(1 - x(\hat{\theta}))$ for all $s \in \mathcal{S}$. Because $\theta > \hat{\theta}$ were arbitrary, condition (E.3) holds.

To show condition (E.4), take $u_l(\theta;\, \bar{f}) = 0$, for all θ and $\bar{f}$, and use the argument from Section 3 (without the loser's aftermarket, the framework is equivalent to the baseline model).

By the way I constructed the winner's aftermarket, the above argument also proves the conclusion even in the case when only the loser interacts.

The converse conclusion holds trivially: if x is constant, it is implementable for all distributions and aftermarkets.

Proof of Proposition 9. I will show that condition (E.2) is equivalent to

$$\pi_l(s|\,\theta)(1 - x(\theta)) \text{ is non-increasing,} \quad \text{(E.6)}$$

and

$$\pi_w(s|\,\theta)x(\theta) \text{ is non-decreasing,} \quad \text{(E.7)}$$

for all $s \in \mathcal{S}$. The above conditions are analogous to condition (M) in Lemma 1. By the same (or analogous in case of E.6) argument that was used to prove Lemma 1, (E.6) and (E.7) imply the cutoff representation. Conversely, if $(x,\, \pi_l,\, \pi_w)$ is a cutoff rule, conditions (E.6) and (E.7) are satisfied, by direct inspection of Definition 14.

First, assume that conditions (E.6) and (E.7) hold. I will show that (E.2) holds. Under

[71] Because implementability is required for all distributions f, it is always possible to choose an f such that distinct signals lead to distinct posteriors, making the above construction well-defined. The only exception is when two signals are indistinguishable but in this case they can be merged into one signal without impacting the analysis.

the condition that A_w is single-crossing-separated from A_l, we have

$$\sum_{s\in\mathcal{S}}\left[u_l(\theta;\, f_l^s)-u_l(\hat{\theta};\, f_l^s)\right]\left[\pi_l(s|\,\theta)(1-x(\theta))-\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))\right]$$

$$+\sum_{s\in\mathcal{S}}\left[u_w(\theta;\, f_w^s)-u_w(\hat{\theta};\, f_w^s)\right]\left[\pi_w(s|\,\theta)x(\theta)-\pi_w(s|\,\hat{\theta})x(\hat{\theta})\right]$$

$$\geq d(\theta,\,\hat{\theta})\sum_{s\in\mathcal{S}}\left[\pi_l(s|\,\theta)(1-x(\theta))-\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))\right]+d(\theta,\,\hat{\theta})\sum_{s\in\mathcal{S}}\left[\pi_w(s|\,\theta)x(\theta)-\pi_w(s|\,\hat{\theta})x(\hat{\theta})\right]$$

$$=d(\theta,\,\hat{\theta})\left[(1-x(\theta))-(1-x(\hat{\theta}))\right]+d(\theta,\,\hat{\theta})\left[x(\theta)-x(\hat{\theta})\right]=0$$

For the converse part, assume condition (E.2). To show condition (E.6), fixing $\theta>\hat{\theta}$, take A_w such that $u_w(\theta;\,\bar{f})-u_w(\hat{\theta};\,\bar{f})=d(\theta,\,\hat{\theta})$. Then, consider an aftermarket A_l that gives differences in payoffs $u_l(\theta;\, f_l^s)-u_l(\hat{\theta};\, f_l^s)=d(\theta,\,\hat{\theta})$ for $s\notin\mathcal{S}_2$ and $u_l(\theta;\, f_l^s)-u_l(\hat{\theta};\, f_l^s)=0$ otherwise, where

$$\mathcal{S}_2\equiv\{s\in\mathcal{S}:\, \pi_l(s|\,\theta)(1-x(\theta))>\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))\}.$$

Then, condition (E.2) implies

$$x(\theta)-x(\hat{\theta})\geq\sum_{s\notin\mathcal{S}_2}\left[\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))-\pi_l(s|\,\theta)(1-x(\theta))\right].$$

The above condition can be rewritten as

$$x(\theta)-x(\hat{\theta})\geq\sum_{s\notin\mathcal{S}_2}\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta}))-\sum_{s\notin\mathcal{S}_2}\pi_l(s|\,\theta)(1-x(\theta))$$

$$=\left(1-\sum_{s\in\mathcal{S}_2}\pi_l(s|\,\hat{\theta})\right)(1-x(\hat{\theta}))-\left(1-\sum_{s\in\mathcal{S}_2}\pi_l(s|\,\theta)\right)(1-x(\theta)),$$

which implies

$$0\geq\sum_{s\in\mathcal{S}_2}\left[\pi_l(s|\,\theta)(1-x(\theta))-\pi_l(s|\,\hat{\theta})(1-x(\hat{\theta})\right].$$

The definition of $\mathcal{S}_2$ together with the above equation imply that $\mathcal{S}_2=\emptyset$. Because $\theta>\hat{\theta}$ were arbitrary, condition (E.6) is proven.

To show condition (E.7), I take $u_l(\theta;\,\bar{f})=0$, for all θ and $\bar{f}$, and specify the aftermarket A_w (fixing $\theta>\hat{\theta}$) by $u_w(\theta;\, f_w^s)-u_w(\hat{\theta};\, f_w^s)=d(\theta,\,\hat{\theta})>0$ if $\pi_w(s|\,\theta)x(\theta)\geq\pi_w(s|\,\hat{\theta})x(\hat{\theta})$, and $u_w(\theta;\, f_w^s)-u_w(\hat{\theta};\, f_w^s)=\alpha d(\theta,\,\hat{\theta})$, for some $\alpha>1$, otherwise. Then, for sufficiently high α condition (E.2) is violated unless $\pi_w(s|\,\theta)x(\theta)\geq\pi_w(s|\,\hat{\theta})x(\hat{\theta})$, for all $s\in\mathcal{S}$.