Mechanism Design with Aftermarkets: Cutoff Mechanisms

Piotr Dworczak[*]

October 20, 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.

I introduce a class of allocation and disclosure rules, called cutoff rules, that disclose information about the buyer's type only by revealing information about the random threshold (cutoff) that she had to outbid to win the object. A rule is implementable regardless of the form of the aftermarket and the underlying distribution of types if and only if it is a cutoff rule. Cutoff mechanisms are tractable, and admit an indirect implementation that often makes them easy to use in practice. I provide sufficient conditions 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 an application to the design of trading rules and post-transaction transparency in financial over-the-counter markets.

Keywords: Mechanism Design, Information Design, Auctions, OTC Markets

[*]Becker Friedman Institute, University of Chicago, E-mail: dworczak@uchicago.edu, Website: http://home.uchicago.edu/∼dworczak/. I am indebted to my advisors, Andy Skrzypacz and Paul Milgrom, for their continuous support and guidance, and to Darrell Duffie and Michael Ostrovsky for comments and serving on my dissertation committee. I would like to thank Mohammad Akbarpour, Ben Brooks, Eric Budish, Jeremy Bulow, Gabriel Carroll, Yossi Feinberg, Alex Frankel, Matthew Gentzkow, Ben Golub, Matthew Jackson, Emir Kamenica, Maciej Kotowski, David Kreps, Nicolas Lambert, Shengwu Li, Elliot Lipnowski, Giorgio Martini, Paul Oyer, Alessandro Pavan, Andrés Perlroth, Doron Ravid, Bartosz Redlicki, Peter Reiss, 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, 2016 European Winter Meeting of the Econometric Society, QMUL Theory Workshop, Kellogg's Rookie Fest, and economics seminars at Princeton, MIT, Columbia, Yale, UCL, Harvard, Chicago Booth, Northwestern, UC San Diego, University of Bonn, University of Bologna, and University of Southampton for helpful comments and discussions.

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.

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, 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 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 information elicited by the mechanism determines the information available for disclosure. Conversely, disclosure influences the incentives of agents to reveal their private information to the mechanism.

On a high level, the paper addresses two questions: *(1)* How much information *can* be elicited and revealed, and *(2)* how much information *should* be elicited and revealed to maximize a designer's objective function.

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 aftermarket. Those, in turn, depend on the aftermarket protocol and the beliefs of aftermarket participants. As a result, the set of implementable allocation and disclosure rules varies with the aftermarket and the prior distribution of agents' types – the shape of the optimal mechanism is sensitive to details of the environment. Finding the optimal Bayesian mechanism may be difficult, and even when it can be found, its use may be hindered by the designer's limited knowledge of these details.

Consider, however, the following class of allocation and disclosure rules, which I call *cutoff rules*. In order to receive the object, an agent must outbid some threshold, which I refer to as the *cutoff*. Depending on the allocation rule, the cutoff could be, for example, a bid of another agent or a reserve price. As it turns out, any non-decreasing allocation rule can be represented by a (random) cutoff. What distinguishes cutoff rules is the disclosure rule: A cutoff rule only reveals information about the cutoff. Formally, conditional on the type profile and the cutoff, the signal from a cutoff mechanism does not depend on the type of the agent who acquires the good. For example, if the object is allocated to the highest bidder in an auction, the relevant cutoff is the second highest bid; Conditional on the second highest bid, the message sent after the auction should not depend on the winner's type.

The key property of cutoff mechanisms (mechanisms that implement cutoff rules) is that the report of the eventual winner of the object does not directly influence the signal. The signal is pinned down by the realization of the cutoff which is determined independently of the winner's report. Because the winner cannot manipulate the signal, cutoff mechanisms can be made truthful regardless of the details of the environment such as the distributions of agents' types – cutoff rules are always implementable.

The paper focuses on the analysis of cutoff mechanisms and their properties. My contribution is threefold. First, I develop a theory of optimal cutoff mechanisms and provide conditions for optimality of simple designs. Second, I show that the cutoff class is uniquely characterized by a number of properties that are desirable in practical design problems. Third, I apply the theoretical results to study optimal design of trading protocols and transparency in financial over-the-counter (OTC) markets.

The first major component of the paper is the analysis of *optimal* cutoff mechanisms. The analysis is based on a crucial observation that allows me to connect cutoff mechanisms to the literature on information design and Bayesian persuasion. In a general mechanism, information disclosure affects incentive-compatibility constraints. However, cutoff rules are always implementable: Regardless of what information about the cutoff is revealed, the agents want to report truthfully. Therefore, finding the optimal disclosure rule in the cutoff class reduces to a standard information design problem where the cutoff plays the role of a state variable. Choosing the allocation rule corresponds to choosing a prior distribution of the state variable. The informational content of the cutoff depends on the allocation rule. For example, if the allocation is constant in agents' types, the cutoff is deterministic and uninformative about the type of the winner. If the allocation rule is that of an efficient auction, the cutoff corresponds to the second highest bid, and its disclosure provides non-trivial information about the winner's type.

When the designer contracts with one agent, a cutoff corresponds to a random reserve price. If the allocation rule is fixed, it may benefit the designer to disclose information about the cutoff. 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 cutoff mechanism that sends no signals. Intuitively, in one-agent problems, the designer has full discretion over the choice of the prior distribution of the cutoff – any distribution of the cutoff can be induced by a non-decreasing allocation rule. Because the designer can directly *choose* the prior belief over the state variable (the cutoff), she need not send signals to induce the optimal posterior beliefs.

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. This is because the designer might want to give up full control over the distribution of the cutoff to use competition between the agents. For example, if the designer decides to run an efficient auction, the distribution of the cutoff – the second highest bid – is exogenous and cannot be chosen. More generally, whenever the allocation depends on the ranking of agents' types, the designer is constrained in the choice of prior distributions of the cutoff. As a result, it may be beneficial to send signals to induce posterior beliefs that differ from the prior. I provide sufficient conditions for optimality of simple mechanisms, such as a second price auction with a reserve price that reveals the price paid by the winner.

The above results are silent about whether it is optimal for the designer to use a cutoff mechanism. In a companion paper (Dworczak, 2017), for a class of problems in which the third party takes a binary action, I derive conditions on the aftermarket under which cutoff mechanisms are optimal. The conditions hold, for example, in certain resale games in which the initial stage is similar to a broker or wholesaler buying an object for resale.

While in general an optimal cutoff mechanism need not maximize the designer's objective among all feasible mechanisms, I provide a number of reasons why a designer may nevertheless restrict attention to this class. As argued above, cutoff rules are always implementable. The first result establishes a surprising property that *only* mechanisms that implement a cutoff rule can be made incentive-compatible regardless of the distribution of types and the form of the aftermarket. This has consequences for robust design in settings where the designer has limited knowledge of the details of the environment. Second, cutoff mechanisms are the unique class of mechanisms under which the agents want to report truthfully even if they do not fully trust the designer to implement the promised allocation and disclosure rules. 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, and are hence easy to use in practice.

The final contribution of the paper is to apply the results to the problem of post-transaction transparency in financial over-the-counter (OTC) markets. In a bilateral trade setting, a seller (mechanism designer) chooses a mechanism to sell an asset to an intermediary (agent) who resells to a customer (third party) in the aftermarket. Due to a lemons problem in the aftermarket, trade is inefficient in the absence of additional information about the value of the asset. However, neither the welfare-maximizing nor the profit-maximizing mechanisms release any signals. Instead, the designer uses the allocation rule to optimally influence beliefs in the aftermarket. Thus, the model helps to explain why regulation that aims to increase post-transaction transparency may be harmful.[1]

Next, I consider 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.[2] Suppose the dealer who buys the asset tries to resell it in the aftermarket. To guarantee allocative efficiency (regardless of the distribution of dealers' costs of intermediation, or the bargaining protocol in the aftermarket), no information about the winning dealer's offer should be disclosed. However, the offers of dealers who do not trade can be disclosed, even if private information is correlated among dealers. I derive the optimal cutoff mechanism in a model where dealers' types describe their intermediation effectiveness measured by the number of final buyers participating in the resale auction in the aftermarket.

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

[1] 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.

[2] 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 information structure jointly.[3] My analysis makes use of the concavification argument first used by Aumann and Maschler (1995), and applied to the Bayesian persuasion model by Kamenica and Gentzkow (2011). A methodological contribution of the paper is to find a connection between the mechanism design problem and the concavification result via the introduction of cutoffs.[4]

With regard to the structure of the problem, the most closely related literature is a series of papers by Calzolari and Pavan (2006a,b, 2008, 2009) on sequential agency. In a sequential agency problem, the agent contracts with multiple principals, and an upstream principal decides how much information to reveal to downstream principals (which play a role analogous to the third parties in my aftermarket). The upstream principal designs the allocation and the disclosure rule jointly. A similar structure arises in papers that study auctions followed by resale or other forms of interactions between the bidders. I provide a detailed discussion of the literature in Section 7.

Section 2 is a condensed version of the paper: I introduce the key ideas and explain the main results using a simple setting. Sections 3 – 6 then extend the results to the full-fledged model. Specifically, Section 3 introduces the general setting, and Sections 4 – 6 correspond to the three main contributions: the theory of optimal cutoff mechanisms, characterizations of the class, and applications. Section 7 discusses the literature and Section 8 – extensions. Additional results can be found in the Online Appendix. Throughout, "theorems" are results pertaining to the general model, "propositions" to special cases of the model (such as the model in Section 2), and "claims" to examples or applications.

2 Simple model of resale

In this section, I consider a simplified model: There is one agent in the mechanism, and the aftermarket is a particular resale game with a single third party.[5] This setting allows me to highlight the main ideas, and build cleaner intuition behind the proofs. I later extend the analysis to multi-agent mechanisms, continuous type spaces, and general aftermarkets.

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

[3] 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.

[4] 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.

[5] The simple model is related to the model of Calzolari and Pavan (2006a) – see Section 7.

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 a 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 signal function with some finite signal set $\mathcal{S}$, and $t : \Theta \to \mathbb{R}$ is a transfer function.[6] 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.[7] I let f^s denote the updated belief over the agent's type,

$$f^s(\theta) = \frac{\pi(s|\,\theta)x(\theta)f(\theta)}{\sum_{\tau \in \Theta} \pi(s|\,\tau)x(\tau)f(\tau)},\ \theta \in \Theta. \tag{2.1}$$

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 can be relaxed, see Section 8.2).

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.[8] In Section 3, 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,

[6] Focusing on direct mechanisms is without loss of generality, by the Revelation Principle. However, this restriction will be consequential for some practical properties of mechanisms discussed in Section 5.1.1.

[7] 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.

[8] 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. With no restriction on off-equilibrium beliefs, the lemons market may have a continuum of equilibria with inefficiently low trade.

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 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 $V(\theta;\, \bar{f}) = v(\theta)\mathbf{1}_{\{\theta \le p^\star(\bar{f})\}} + \theta\mathbf{1}_{\{\theta > p^\star(\bar{f})\}}$. 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.[9]

2.1 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: for all θ,

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

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

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. 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. In the next subsection, I introduce a class of mechanism frames whose implementability is not sensitive to the prior distribution.

[9] With a continuous type space, which is allowed by the general model of Section 3, this follows from the payoff equivalence theorem (see e.g. Milgrom, 2001). For example, with $\eta = 1$, revenue maximization corresponds to choosing

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

where $J(\theta)$ is the virtual surplus function. I work with a finite type space in Section 2 so exact payoff equivalence does not hold. However, it is well known that for a fixed $(x,\, \pi)$, the set of implementing transfers is a complete lattice with a unique highest element.

2.2 Cutoff mechanisms

To define cutoff mechanisms, I first introduce the notion of a random-cutoff representation of an allocation rule. For simplicity of exposition, I assume that $x(\max(\Theta)) = 1$, i.e. the highest type always receives the good (this assumption is relaxed in the general model).

2.2.1 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 $x(\theta)$.[10] 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.[11]

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 space of cutoffs: C contains all possible realizations of c_x as x varies, and is equal to Θ under the assumption $x(\max(\Theta)) = 1$.

2.2.2 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.

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

$$\pi(s|\,\theta)x(\theta) = \sum_{c \leq \theta} \gamma(s|\,c)dx(c), \tag{2.2}$$

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

[10] Formally, x becomes a cdf after extending it to a right-continuous step function on the real line.
[11] Using dx (which is typically used for continuous measures) allows me to keep consistent notation throughout all sections.

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 (and no signal is sent). 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.

Example 1. I will use a simple example throughout this section to illustrate the main concepts. Let $\Theta = \{0, 1\}$, $f(0) = f(1) = 1/2$. Under the assumption $x(\max(\Theta)) = 1$, the set of allocation rules is a one-dimensional family indexed by the probability that type 0 gets the object $\beta = x(0)$. The cutoff representation c_x of x is a binary random variable on $C = \{0, 1\}$ with probability mass function dx given by $dx(0) = 1 - dx(1) = \beta$. A cutoff mechanism implementing x gives the good to the low type only if the realized cutoff is low $(c_x = 0)$, and sends a signal that is measurable with respect to c_x.

Note that the amount of information that can be revealed by a cutoff mechanism depends on the allocation rule through the distribution of the cutoff 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. 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.

2.2.3 Implementability of cutoff rules

In general, implementability of a mechanism frame depends on the details of the setting, such as the distribution of types f. Proposition 1 establishes that cutoff rules can be always implemented.

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

A proof of Proposition 1 can be found in Appendix A (see also the discussion in Subsection 2.4.1). The key intuition for the result is simple: Under a cutoff rule, the report of the agent does not directly influence the signal sent by the mechanism. This is reminiscent of why VCG mechanisms (such as second price auctions) are truthful. In a VCG 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.[12] While the agent can change the outcome by manipulating the probability with which she acquires the good, it is easy to show that such a deviation can be deterred by setting up appropriate transfers. Formally, existence of such transfers follows from a single-crossing property: The continuation payoff $u(\theta;\, \bar{f})$ is non-decreasing in θ, holding $\bar{f}$ fixed, for any underlying equilibrium of the aftermarket (this is evident from a strategy-stealing argument).

It is important to emphasize that while any cutoff rule is implementable for any prior distribution f, the implementing transfer function will in general depend on f. This is a consequence of the setting rather than a feature of cutoff rules: With the aftermarket, the prior distribution f directly influences the value that the agent has for winning the object. Because the values change with f, transfers have to be rescaled accordingly. I discuss this issue, and its implication for robust implementation, in Section 5.1.1.

2.3 Optimal cutoff mechanisms

In this section, I study optimal design of cutoff mechanisms. The robustness of cutoff rules (Proposition 1) allows me to draw a connection to Bayesian persuasion which makes optimization tractable.

Under the assumptions of Section 2, the designer maximizes

$$\sum_{\theta \in \Theta} \sum_{s \in \mathcal{S}} V(\theta;\, f^s)\pi(s|\,\theta)x(\theta)f(\theta). \tag{2.3}$$

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 following result establishes a strong conclusion about optimal cutoff mechanisms in the one-agent model.

Proposition 2. *The problem of maximizing (2.3) subject to $(x,\, \pi)$ being a cutoff rule has an optimal solution that reveals no information.*

The conclusion of Proposition 2 holds regardless of the objective function. The type of the objective may influence the shape of the optimal allocation rule x but never requires the designer to make explicit announcements via π. An essential step in the proof is a characterization of optimal disclosure for a fixed allocation rule.

[12] Cutoff mechanisms are also related to the Becker-DeGroot-Marschak method used in experimental economics to elicit information about willingness to pay, see Becker, Degroot and Marschak (1964).

2.3.1 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, π) being a cutoff rule.

An allocation rule 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).

Suppose that G is a posterior distribution of the cutoff conditional on observing some signal realization, before conditioning on the event that the agent acquired the good. Then,

$$f^G(\theta) \equiv \mathbb{P}_{c\sim G}(\theta |\, \theta \geq c) = \frac{G(\theta)f(\theta)}{\sum_\tau G(\tau)f(\tau)}, \tag{2.4}$$

is the corresponding posterior belief over the type of the agent conditional on the agent acquiring the good.[13] Equivalently, f^G is the posterior belief over the type of the agent who acquired the good that would be held by the third party if she believed that the designer implemented the allocation rule $G(\theta)$. Next, let

$$\mathcal{V}(G) = \sum_{\theta\in\Theta} V(\theta;\, f^G)G(\theta)f(\theta) \tag{2.5}$$

be the conditional expected payoff to the mechanism designer conditional on the signal inducing a posterior distribution G of 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.

Lemma 1. *For every non-decreasing allocation function x, the problem of maximizing (2.3) over π subject to (x, π) being a cutoff rule is equivalent to*

$$\max_{\varrho\in\Delta(\Delta(C))} \mathbb{E}_{G\sim\varrho}\mathcal{V}(G) \tag{2.6}$$

[13] The order of conditioning does not matter, so for the inference problem we can assume that the sequence of events is as follows: (1) a cutoff is drawn, (2) a signal is drawn as a function of the cutoff, (3) allocation is determined as a function of the cutoff and the agent's report. If G is the belief over the cutoff held after step 2, then f^G is the belief over the agent's type held after step 3.

subject to

$$\mathbb{E}_{G\sim\varrho}G(\theta) = x(\theta),\ \forall\theta\in\Theta. \tag{2.7}$$

Lemma 1 follows directly from the results of Kamenica and Gentzkow (2011) but I nevertheless provide a discussion of the proof in Appendix A.1. Equation (2.6) means that the mechanism designer seeks to maximize her expected payoff over distributions ϱ of posterior beliefs G over the cutoff. Condition (2.7) is the Bayes-plausibility constraint – the induced posterior beliefs over the cutoff must average out to the prior belief (beliefs are represented by cdfs). Using the random-cutoff representation, the prior distribution of the cutoff is simply the allocation rule x.

Lemma 1 yields an alternative interpretation of one-agent cutoff mechanisms. The designer randomizes the choice of an allocation rule (subject to the constraint that on average she uses x), and only reveals information about which allocation rule she used. Formally, (2.7) implies that a cutoff rule $(x,\,\pi)$ can be represented as a probability distribution ϱ over mechanism frames $(G,\,\emptyset)$ that reveal no information: 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.

Lemma 1 implies that the concavification result of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011) can be applied to my setting. Let $\mathcal{X}$ be the set of all non-decreasing allocation rules on Θ.

Corollary 1. *The optimal expected payoff to the mechanism designer in the problem* (2.6)-(2.7) *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 an example.

Example 1. [continued] Consider an aftermarket with $\eta = 1$ and $v(\theta) = v \in (1,\,2)$ (the third party has constant value and makes an offer to the agent). For a fixed allocation rule x with $x(0) = \beta$, the designer maximizes total surplus over disclosure rules.

Each posterior belief over the cutoff is described by the probability that the cutoff is low ($c_x = 0$), denoted α. The prior belief is β. A posterior belief α corresponds to a belief $\alpha/(1+\alpha)$ that the type of the agent is low in the aftermarket. The third party offers a price of 0 when she believes that the probability of the low type in the aftermarket is at least

$(v-1)/v$ (and only the low type accepts). Otherwise, she offers a price of 1 (and both types accept). Thus, with slight abuse of notation, maximizing total surplus corresponds to

$$\mathcal{V}(\alpha) = \begin{cases} v\alpha f(0) + vf(1) = \frac{v}{2}(\alpha+1) & \text{if } \alpha \leq v-1 \\ v\alpha f(0) + f(1) = \frac{1}{2}(v\alpha+1) & \text{if } \alpha > v-1 \end{cases}.$$

The function $\mathcal{V}$ and its concave closure are depicted in Figure 2.1 for $v = 5/3$. The discontinuity at $\alpha = v - 1$ corresponds to the belief at which the third party is indifferent between quoting the two prices (0 and 1) in the aftermarket.

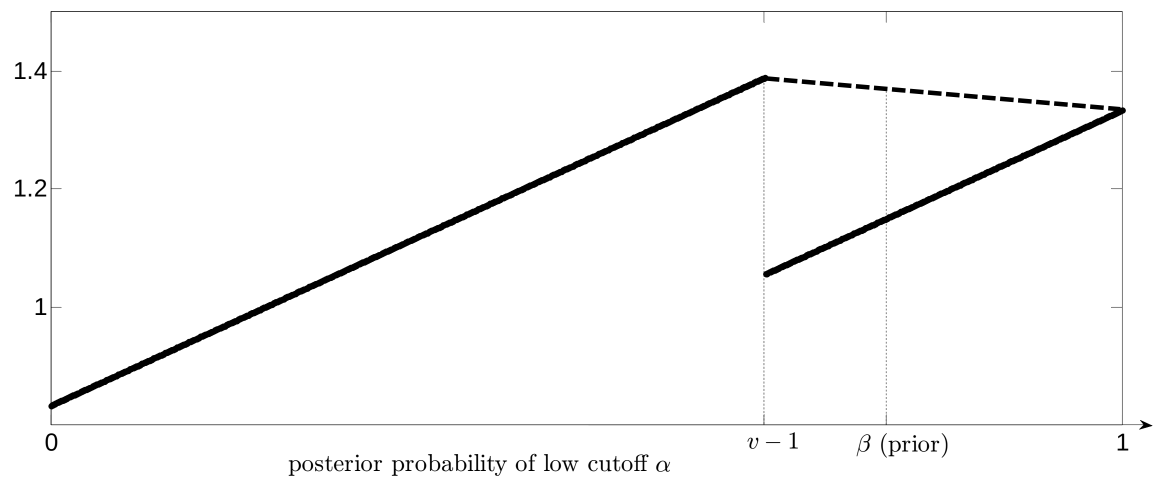

Fig. 2.1: Function $\mathcal{V}$ (solid line) and its concave closure (dotted line).

When $\beta > v - 1$ (as is the case in Figure 2.1), it is optimal to disclose some information: the designer reveals the realization of the low cutoff with probability λ and never reveals the high realization. The probability λ is chosen so that conditional on no revelation, the third party is indifferent between offering the high and the low price (and offers the high price): $v - 1 = \beta(1-\lambda)/(\beta(1-\lambda)+1-\beta)$.

When $\beta \leq v - 1$, $\mathcal{V}$ coincides with its concave closure, and the designer makes no announcement in the optimal mechanism. The third party always offers a high price in the aftermarket. Note that even when no announcement is made, the third party's beliefs in the aftermarket are different from the prior f. Because the allocation rule has $x(0) = \beta < 1$, the posterior probability of the high type conditional on acquiring the good is revised upwards to $1/(\beta + 1)$. ■

2.3.2 Completion of the proof of Proposition 2

With the characterization of the optimal payoff to the designer for a fixed allocation rule, the proof of Proposition 2 can be completed by an easy argument.

Proof of Proposition 2. By Corollary 1, 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.[14] 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. □

The intuition behind Proposition 2 is that the mechanism design problem is a Bayesian persuasion problem in which the Sender (mechanism designer) can choose the prior distribution (of the cutoff). This is because the prior distribution of the cutoff corresponds to the allocation rule. In communication problems, a Sender reveals information about a state to optimally shape the posterior beliefs of a Receiver. But when the posterior belief can be chosen directly by choosing the prior, there is no need to reveal information about the state. In the design of the optimal cutoff mechanism, there is no need to reveal information about the cutoff because the optimal posterior distribution can be induced directly by choosing the prior (allocation rule).

For an alternative intuition, recall that any cutoff rule (x, π) can be interpreted as randomization ϱ over no-information-revealing mechanism frames $(G, \emptyset)$, with the public message disclosing which of the $(G, \emptyset)$ was used. The payoff to the designer is the expectation of the payoffs from $(G, \emptyset)$ weighted by ϱ, see (2.6). Because G is a cdf, it is a non-decreasing function, so each $(G, \emptyset)$ is an implementable mechanism frame. For any cutoff rule revealing information, the designer can weakly increase the payoff by implementing the $(G, \emptyset)$ yielding the highest conditional expected payoff with probability 1.

Example 1. [continued] Suppose that the designer can optimize over both allocation and disclosure rules. This means that she can additionally optimize over β, the probability that the low type obtains the good, or equivalently, over all prior (binary) distributions of the cutoff on $\{0, 1\}$. Thus, the designer can choose an arbitrary point on the x-axis in Figure 2.1 to maximize the concave closure of $\mathcal{V}$. In Figure 2.1, the expected payoff to the mechanism designer is maximized at $\beta = v - 1$. At $\beta = v - 1$, the function $\mathcal{V}$ coincides with its concave closure, so no-revelation is optimal.

More generally, choosing $\beta = v - 1$ is optimal whenever $v \geq (1 + \sqrt{5})/2$. In the opposite case, the value of $\mathrm{co}\mathcal{V}$ is higher at $\beta = 1$ than at $\beta = v - 1$, and thus it is optimal to choose

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

a constant allocation rule $x \equiv 1$ (and reveal no information). ■

The above example illustrates two additional points. First, Proposition 2 does not imply that the designer ignores the effect on the aftermarket when choosing the mechanism. The choice of the allocation rule has an impact on the information structure in the resale game because the third party conditions on the event that the agent acquired the good. Second, when the allocation rule is a threshold rule, $x(\theta) = \mathbf{1}_{\{\theta \geq r\}}$ for some r, the cutoff is degenerate (equal to r), and hence a cutoff rules implementing x cannot reveal any additional information about the agent's type. One might suspect that Proposition 2 is true because the optimal allocation rule is always a threshold rule. Example 1 shows that this is not true: When the optimal mechanism allocates the good to the low type with interior probability, additional information could be revealed but it is not in the designer's interest to do so.[15]

2.4 A characterization of cutoff rules

Proposition 1 shows that cutoff rules are implementable for any prior distribution f. It is natural to ask what other mechanisms have this robustness property. It turns out that under mild restrictions on the structure of the aftermarket cutoff rules are the *only* mechanism frames that can be implemented regardless of the distribution f.

Proposition 3*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 payoff-irrelevant. Then, each π is trivially implementable. When $\eta = 0$, Proposition 3a remains true if the lemons problem is severe enough so 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 3*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.*

Propositions 3a and 3b can be interpreted as follows: For a given allocation rule, the cutoff is the most informative random variable with the property that its disclosure never upsets

[15] This conclusion is not an artifact of the discrete type space. In the Online Appendix OA.1.1, I construct an example with a continuous types space where the optimal allocation rule leads to a non-degenerate distribution of the cutoff (and yet no information about the cutoff is revealed).

incentive compatibility of the mechanism. If a disclosure rule induces aftermarket beliefs that are not induced by a cutoff rule, then for at least some prior distribution the mechanism is not truthful. Hence, the space of cutoffs is the largest domain on which concavification of the objective function can be performed regardless of the prior distribution.

To gain intuition for Propositions 3*a* and 3*b*, note that low types of the agent have a higher willingness to pay (relative to high types) for signals that lead to higher resale prices in the aftermarket. At the same time, high resale prices only occur for signals that are more likely conditional on high types. Informally, when the aftermarket is a resale game and signals are informative, high types get something they don't value as much as lower types do – creating problems with incentive compatibility. Take an extreme example where the mechanism attempts to disclose the type of the agent exactly (not a cutoff rule), and the third party makes an offer to the agent. The highest-value agent does not benefit from her signal because the third party extracts all the rents in the aftermarket. But a lower type has a high willingness to pay for the highest type's signal because she benefits strictly from accepting the offer in the aftermarket following the deviation. A transfer function that would deter the low type from pretending to be a high type would have to leave no information rents for the high type, contradicting incentive compatibility.

The proof of Propositions 3*a* and 3*b* exploits the above intuition to show that any non-cutoff rule must fail incentive-compatibility for some distribution of types, where that distribution is carefully constructed to amplify the sensitivity of aftermarket payoffs to the disclosed information. The reason why cutoff mechanism are not ruled out by the above reasoning is that they use the allocation rule as a "leverage". When two types receive the good with different probabilities, there is a slack in their respective incentive constraints. Cutoff rules reveal information in proportion to the slack to guarantee that disclosure will not upset incentive compatibility.

2.4.1 Discussion of the proof

Propositions 1, 3*a*, and 3*b* follow from two lemmas which are of independent interest.

Lemma 2. *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 3. *If* (x, π) *satisfies condition* (M)*, then* (x, π) *is implementable for every distribution of types* f*. Conversely, under the assumptions of Proposition* 3*a* *(or* 3*b*)*, if a mechanism frame* (x, π) *is implementable for every distribution of types* f*, then* (x, π) *satisfies condition* (M)*.*

Property (M) extends the Myerson monotonicity condition from one-stage allocation problems to settings with aftermarkets. Cutoff mechanisms satisfy the monotonicity property by definition – due to representation (2.2). Lemma 2 states that the property is a characterization of cutoff rules. Condition (M) is stronger than Myerson monotonicity in that it implies that $x(\theta)$ is non-decreasing. In problems with an aftermarket, by Lemma 3, implementability for every distribution requires this stronger monotonicity notion: The joint probability of allocating the good *and* sending any signal $s \in \mathcal{S}$ must be non-decreasing in the type. The proofs can be found in Appendix A.

3 General model

In this section, I extend the simple model of Section 2 by allowing *(i)* a general aftermarket, *(ii)* multiple agents in the mechanism, and *(iii)* continuous distributions of types. The key assumptions maintained in the general model are that *(a)* there is a single object to be allocated, and *(b)* only the agent who acquired the good participates in the aftermarket. Subsequently, Sections 4 – 5 extend the analysis of Section 2 by providing a general theory of optimal cutoff mechanisms and several characterizations of the class of cutoff mechanisms.

The mechanism designer owns an indivisible good 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 $\Theta_i = \text{supp}(f_i)$, 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:

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 $\boldsymbol{\Theta}$.

2. Discrete distribution. Each Θ_i is finite, and $\boldsymbol{f}$ is a density with respect to discrete uniform measure on $\boldsymbol{\Theta}$.

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

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

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 rule 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 σ−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 draws and publicly announces a signal $s \in \mathcal{S}$ according to distribution $\pi_i(\cdot| \hat{\theta}_i, \boldsymbol{\theta}_{-i})$.[16] 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.

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.

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).

The "black-box" approach to modeling the aftermarket implicitly entails the following assumptions. A Bayesian game is played after the mechanism between agent i who acquired the good (whose identity becomes known) and third-party players. The term "game" is understood broadly and includes cases in which, for example, the third party is another mechanism designer, and the aftermarket is a mechanism in which the winner participates. Third-party players have a common prior $\boldsymbol{f}$ over the agents' types, and observe the public signal s sent by the mechanism which leads to a posterior belief over types $\bar{\boldsymbol{f}}$. Given belief $\bar{\boldsymbol{f}}$ and an aftermarket A, the corresponding game has a set of equilibria $EQ^A(\bar{\boldsymbol{f}}\,)$, 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

[16] The restriction to public signals is motivated by practical considerations. The results on the characterization of cutoff rules go through with private signals but the results on optimality change. With private signals, the model is related to Bayes Correlated Equilibrium of Bergemann and Morris (2016a).

equilibrium selection from EQ^A, $u_i(\theta;\, \bar{\boldsymbol{f}}\,)$ is the expected equilibrium payoff to type θ of agent i.[17] 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.) 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 $\bar{\boldsymbol{f}}\,$. Other roles of the signal (for example, as a coordination device) can be incorporated by considering an appropriate equilibrium concept (e.g. a version of correlated equilibrium, see Bergemann and Morris, 2016a). Consequently, I will not distinguish between two mechanisms that induce the same distribution of posterior beliefs for any prior (relabeling of signals does not change the mechanism).

I will sometimes refer to the *symmetric case* in which agents are ex-ante identical. In particular, they have the same marginal distribution of types f on a common space Θ, and enjoy the same payoff in the aftermarket as a function of their type and beliefs.

3.1 Implementability

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 $\boldsymbol{f}$, under truthful reporting ($\boldsymbol{f}^{i,s}$ can be computed by Bayes' rule as in (2.1)).

Definition 3. 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:

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

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

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

[17] This formulation allows for exogenous private information of third-party players. The payoff u_i remains measurable with respect to the public belief $\bar{\boldsymbol{f}}$ because we can assume without loss that third-party players observe their exogenous signals after observing the public signal s – the realizations of private signals are integrated out in the payoff function u_i.

3.2 Cutoff mechanisms

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 for 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 2.2.1). 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 the cutoff 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 the cutoff for agent i depends on $\boldsymbol{\theta}_{-i}$.

Definition 4 (Cutoff rule and cutoff mechanism)**.** 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 π_i can be represented as

$$\pi_i(S|\,\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \int_0^{\theta_i} \gamma_i(S|\,c, \boldsymbol{\theta}_{-i})dx_i(c, \boldsymbol{\theta}_{-i}), \tag{3.1}$$

for some measurable signal function $\gamma_i : C_i \times \boldsymbol{\Theta}_{-i} \to \Delta(\mathcal{S})$, for all $\theta_i \in \Theta_i$, $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$, measurable $S \subset \mathcal{S}$, and $i \in N$. $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t})$ is a *cutoff mechanism* if $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule.

In the multi-agent setting, a cutoff mechanism reveals information about the random cutoff of the winner for every fixed report profile of other agents. Thus, the signal sent when agent i is the winner can depend on $i's$ cutoff but also on the reports of all other agents.

Note that $(\boldsymbol{x}, \boldsymbol{\pi})$ is characterized by N cutoffs, one for each agent i. While these cutoffs must in general be correlated to implement the allocation rule $\boldsymbol{x}$, the correlation is irrelevant in the analysis of cutoff mechanisms because ex-post only one cutoff matters for the determination of the signal (namely, the cutoff of the winner).

The next definition captures the key property of the simple aftermarket presented in Section 2 that is needed for all cutoff rules to be implementable.

Definition 5 (Monotonicity)**.** An aftermarket A is *monotone*, if for any agent $i \in N$, and any belief $\bar{\boldsymbol{f}} \in \Delta(\boldsymbol{\Theta})$, the expected utility function $u_i(\theta;\, \bar{\boldsymbol{f}}\,)$ is non-decreasing in θ.

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. Monotonicity is the only property needed to extend Proposition 1.

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

To prove Theorem 1, I show that a multi-agent cutoff rule satisfies a generalized monotonicity condition from Section 2:

$$\pi_i(S|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) \text{ is non-decreasing in } \theta_i \text{ for all measurable } S \in \mathcal{S} \text{ and all } \boldsymbol{\theta}_{-i} \in \Theta_{-i}. \tag{3.2}$$

The proof uses condition (3.2) and montonicity of the aftermarket in a way analogous to the proof of Proposition 1 (details can be found in Appendix B).

4 Optimal cutoff mechanisms

In this section, I consider optimization in the class of cutoff mechanisms. I make two additional assumptions. First, agents' types are independent. With correlated types, payoff equivalence breaks and it is well known that optimal mechanisms may have unintuitive properties (e.g. full surplus extraction is possible if the designer uses a Crémer-McLean mechanism, see Crémer and McLean, 1988). Cutoff mechanisms are still well defined and truthful with correlated types, and I comment on their properties in Section 5. Second, I assume that only the belief about the type of the winner matters in the aftermarket. This is a natural assumption in many applications, and it simplifies the search for the optimal mechanism. If beliefs about all agents' types are relevant in the aftermarket (which could be the case, for example, when agents' types are private signals about a common value of the asset), the problem involves optimal disclosure of multi-dimensional information which is not tractable in most cases. Although the designer still has access to multi-dimensional information when $N > 1$, the assumption will imply existence of a one-dimensional statistic sufficient to characterize disclosure in the optimal mechanism. This is shown in Subsection 4.1.

With slight abuse of notation, I will write agent i's payoff in the aftermarket as $u_i(\theta_i;\, f_i^s)$, where f_i^s is the posterior belief over agent i's type in the aftermarket (the superscript i is dropped because only the belief over i's type matters when i is the winner). The objective of the designer is given by

$$\sum_{i\in N}\int_{\boldsymbol{\Theta}}\int_{\mathcal{S}} V_i(\theta_i;\, f_i^s)d\pi_i(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})\boldsymbol{f}(\boldsymbol{\theta})d\boldsymbol{\theta}, \tag{4.1}$$

where each $V_i:\, \Theta_i\times\Delta(\Theta_i)\to\mathbb{R}$ is assumed to be bounded, measurable in the first argument, and upper-semi continuous in the second argument (in the weak* topology on $\Delta(\Theta_i)$). By assumption, 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.

4.1 A dimensionality reduction via reduced-form mechanisms

The goal of this subsection is to reduce the search for the optimal cutoff mechanisms in the general model to a one-dimensional problem similar to the one considered in Section 2.

Given $(\boldsymbol{x},\, \boldsymbol{\pi})$, its *reduced form* under distribution $\boldsymbol{f}$, denoted $(\boldsymbol{x}^{\boldsymbol{f}},\, \boldsymbol{\pi}^{\boldsymbol{f}})$, is defined by

$$x_i^{\boldsymbol{f}}(\theta_i) = \int_{\boldsymbol{\Theta}_{-i}} x_i(\theta_i,\, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})d\boldsymbol{\theta}_{-i},$$

and

$$\pi_i^{\boldsymbol{f}}(S|\,\theta_i)x_i^{\boldsymbol{f}}(\theta_i) = \int_{\boldsymbol{\Theta}_{-i}} \pi_i(S|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})d\boldsymbol{\theta}_{-i},$$

for all measurable $S\subseteq\mathcal{S}$, $\theta_i\in\Theta_i$, and $i\in N$.

I say that two mechanism frames $(\boldsymbol{x},\, \boldsymbol{\pi})$ and $(\boldsymbol{x}',\, \boldsymbol{\pi}')$ are *Bayesian equivalent* under $\boldsymbol{f}$ if they induce the same reduced form. I have assumed that only the belief over the winner's type is payoff-relevant in the aftermarket. The posterior belief f_i^s depends only on the reduced form of a mechanism. Thus, a reduced-form representation is sufficient to describe payoff-relevant consequences of any mechanism for the aftermarket. However, restricting attention to reduced-form mechanisms is not in general without loss when dominant-strategy implementation is required – a reduced form of a mechanism can only be used to establish Bayesian implementability. The following result shows that looking at reduced forms is without loss of generality in the class of cutoff mechanisms.

Lemma 4. *A pair $(\bar{\boldsymbol{x}},\, \bar{\boldsymbol{\pi}})$, where $\bar{x}_i:\, \Theta_i\to[0,\,1]$ and $\bar{\pi}_i:\, \Theta_i\to\Delta(\mathcal{S})$, for all i, is a reduced form of a cutoff rule under prior distribution $\boldsymbol{f}$ if and only if,*

1. *The interim allocation rule $\bar{x}_i(\theta_i)$ is non-decreasing in θ_i, for all $i\in N$;*

2. *The interim signal function* $\bar{\pi}_i$ *can be represented as*

$$\bar{\pi}_i(S|\,\theta_i)\bar{x}_i(\theta_i) = \int_0^{\theta_i} \gamma_i(S|\,c)d\bar{x}_i(c), \tag{4.2}$$

for some signal function $\gamma_i : C_i \to \Delta(\mathcal{S})$*, for all* $i \in N$*,* θ_i*, and measurable* $S \subseteq \mathcal{S}$*;*

3. *Interim expected allocation rules are jointly feasible under* $\boldsymbol{f}$*:*

$$\sum_{i\in N} \int_{\tau_i}^{1} \bar{x}_i(\theta)f_i(\theta_i)d\theta_i \leq 1 - \prod_{i\in N} F_i(\tau_i),\ \forall\tau \in \boldsymbol{\Theta}. \tag{M-B}$$

I call a reduced-form mechanism satisfying conditions *1-3* of Lemma 4 a *reduced-form cutoff rule*. Lemma 4 shows that a reduced-form cutoff rule is equivalent to a one-agent cutoff rule from Section 2 (condition (4.2) is a generalization of (2.2) to infinite type and signal spaces) with one difference – the additional constraint (M-B). Because each $\bar{x}_i$ is an interim allocation rule, there must exist a joint allocation rule that induces these marginals:

$$\bar{x}_i(\theta_i) = \int_{\boldsymbol{\Theta}_{-i}} x_i(\theta_i,\, \boldsymbol{\theta}_{-i})\boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})d\boldsymbol{\theta}_{-i}. \tag{4.3}$$

For example, the designer cannot promise to give the good to each agent regardless of her type, $\bar{x}_i(\theta_i) = 1$, for all i and θ_i, because she only has one good – there does not exist an allocation rule $\boldsymbol{x}$ with $\sum_i x_i(\boldsymbol{\theta}) \leq 1$ which induces such marginals. By the characterization of reduced-form auctions developed by Matthews (1984) and Border (1991), and extended to the asymmetric case by Border (2007) and Mierendorff (2011), condition (4.3) can be replaced by the Matthews-Border condition (M-B) in Lemma 4.

Lemma 4 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.[18] 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 DS implementable).

A consequence of Lemma 4 is that I can restrict attention to reduced-form mechanisms, and hence also to one-dimensional signals (despite the fact that a general cutoff mechanism can disclose information about the entire vector $\boldsymbol{\theta}_{-i}$). For a simple example of such a dimensionality reduction, let the allocation rule be $x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i > \boldsymbol{\theta}_{-i}^{(1)}\}}$, where $\boldsymbol{\theta}_{-i}^{(1)}$ denotes

[18] The classical reference is Manelli and Vincent (2010). Gershkov, Goeree, Kushnir, Moldovanu and Shi (2013) generalize the results to any one-dimensional setting and I build heavily on their proof technique to establish the "if" part of Lemma 4.

the first order statistic of $\boldsymbol{\theta}_{-i}$. Consider the reduced form of $\boldsymbol{x}$ under prior $\boldsymbol{f}$: $x_i^{\boldsymbol{f}}(\theta_i) = \prod_{j \neq i} F_j(\theta_i)$. Under $x_i^{\boldsymbol{f}}$, 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. For any cutoff mechanism implementing $\boldsymbol{x}$, there exists a payoff-equivalent cutoff mechanism that only reveals information about the (one-dimensional) second highest reported type. This is because $\boldsymbol{\theta}_{-i}^{(1)}$ is a sufficient statistic for the posterior belief over the type θ_i of the winner. Lemma 4 implies that such a one-dimensional sufficient statistic can be found for each agent i for *any* non-decreasing allocation rule $\boldsymbol{x}(\boldsymbol{\theta})$. In fact, this one-dimensional sufficient statistic is exactly the reduced-form cutoff defined by the interim expected allocation $x_i^{\boldsymbol{f}}(\theta_i)$.

4.2 Optimal design

Due to the dimensionality reduction via the reduced-form representation, the results from Section 2 extend naturally to the multi-agent setting. To state the main result of this section, I generalize two definitions from Section 2. The interim allocation rule $x_i^{\boldsymbol{f}}$, treated as a cdf, defines a prior distribution of the cutoff for agent i. Given posterior distribution G_i 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

$$f_i^{G_i}(\theta_i) = \frac{G_i(\theta_i) f(\theta_i)}{\int_{\Theta_i} G_i(\tau) f(\tau) d\tau}. \tag{4.4}$$

The expected payoff to the mechanism designer from agent i when the cutoff for agent i has distribution G_i is given by

$$\mathcal{V}_i(G_i) = \int_{\Theta_i} V_i(\theta_i;\, f_i^{G_i}) G_i(\theta_i) f_i(\theta_i) d\theta_i. \tag{4.5}$$

Let $\mathcal{X}_i$ denote the set of one-dimensional non-decreasing allocation rules on Θ_i.

Theorem 2. *The problem of maximizing (4.1) over the set of cutoff mechanisms is equivalent to solving*

$$\max_{\{\bar{x}_i \in \mathcal{X}_i\}_{i \in N}} \sum_{i \in N} co\mathcal{V}_i(\bar{x}_i) \tag{4.6}$$

subject to the Matthews-Border condition (M-B).

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

Because only the winner interacts in the aftermarket, solving a general design problem requires solving N one-dimensional persuasion problem independently – that is why the concave closure operator is applied to each $\mathcal{V}_i$ separately, rather than to the sum of $\mathcal{V}_i$'s. In the symmetric case, the optimization problem takes the simpler form:

$$N \max_{\bar{x} \in \mathcal{X}} \operatorname{co}\mathcal{V}(\bar{x}) \tag{4.7}$$

subject to

$$\int_{\tau}^{1} \bar{x}(\theta) f(\theta) d\theta \leq \frac{1 - F^N(\tau)}{N}, \; \forall \tau \in \Theta. \tag{4.8}$$

The proof of Theorem 2 is established in two steps, analogously as in Section 2. In the first step, I consider optimization over disclosure rules for a fixed allocation rule establishing a connection to information design and applying the concavification argument. In the second step, I optimize over allocation rules (see Appendix C.2).

For $N = 1$, the Matthews-Border condition (M-B) holds vacuously, and the argument used in the proof of Proposition 2 implies that there always exists a solution that reveals no information.

When $N \geq 2$, constraint (M-B) is binding and it alters the conclusion from the one-agent setting. It may be optimal to disclose information (examples can be found in Section 6.2). To see why, consider the symmetric case (4.7) - (4.8). The concave closure of $\mathcal{V}$ is taken in the space of *all* non-decreasing interim allocation rules (equivalently, all posterior beliefs over the cutoff), while the actual rule $\bar{x}$ must be chosen from a strictly smaller subset of rules that satisfy the Matthews-Border condition (M-B). It might be optimal to induce posterior beliefs over the cutoff that do not correspond to an interim allocation rule satisfying (M-B).

In the one-agent model, the optimal mechanism reveals no information because the seller directly chooses the distribution of the state variable (cutoff) used for persuasion. With multiple agents, the seller may choose to give up her full control over the distribution of the state variable (cutoff) in order to use competition between the agents. If the seller decides to allocate the good only to a particular agent i, then she can still choose an arbitrary cutoff distribution for that agent. But as soon as she decides to condition the allocation on reports of multiple agents (use competition), the Matthews-Border condition kicks in and introduces a constraint on the marginal distributions of the cutoffs. For example, consider an auction with no reserve price. The cutoff for the winner is the second-order statistic of the bids. The distribution of the second-order statistic is completely determined by the prior distribution $\boldsymbol{f}$ and cannot be changed by the designer. To directly affect the prior distribution of cutoffs, the seller would have to limit competition in the auction. It may instead be optimal to use signals to affect the *posterior* distribution.

4.3 Optimality of simple mechanisms

In this section, I derive sufficient conditions for optimality of simple cutoff mechanisms. To simplify notation, I restrict attention to the symmetric case (agents are ex-ante identical).

The previous section showed that optimization over cutoff rules involves optimizing over distributions of beliefs over the cutoffs. To obtain a characterization of cases in which simple designs are optimal, it is often convenient to optimize directly over posterior distributions of beliefs over the winner's type. In Appendix C.3, I prove that a conditional distribution of beliefs over the winner's type is feasible (that is, induced by some cutoff rule) if and only if *(i)* an appropriate Bayes-plausibility conditions holds, and *(ii)* each posterior belief over the winner's type monotone-likelihood-ratio dominates the prior belief f. Condition *(i)* is natural in the context of information design (see Kamenica and Gentzkow, 2011). while condition *(ii)* is a consequence of monotoniticty of cutoff rules - regardless of the signal, higher types receive the good with higher probability, so a posterior belief over the winner's type must be higher than the prior belief. Define $\mathcal{W}(\bar{f}) = \int_\Theta V(\theta;\, \bar{f})\bar{f}(\theta)d\theta$ as the conditional expected payoff to the designer, conditional on inducing posterior distribution $\bar{f}$ over the winner's type. Let M_f be the set of distributions over Θ which monotone-likelihood-ratio dominate the prior f,[19] and let $f^{\bar{x}}$, defined by (4.4), be the posterior belief over the winner's type if no information is revealed.

Lemma 5. *The optimal expected payoff to the mechanism designer in the problem* (4.7)-(4.8) *for a fixed interim allocation rule $\bar{x}$ is equal to*

$$N\left(\int_\Theta \bar{x}(\theta)f(\theta)d\theta\right) co^{M_f}\,\mathcal{W}(f^{\bar{x}}) \tag{4.9}$$

where $co^{M_f}\,\mathcal{W}(f^{\bar{x}}) \equiv \sup\{y : \,(f^{\bar{x}},\, y) \in CH\big(graph(\mathcal{W})|_{M_f}\big)\}$, and $graph(\mathcal{W})|_{M_f}$ is the graph of $\mathcal{W}$ restricted to domain M_f.

Objectives (4.7) and (4.9) are analogous except for two important details. First, in (4.9) the concavification of $\mathcal{W}$ is taken in the space $M_f \subsetneq \Delta(\Theta)$, while in (4.7) the concave closure of $\mathcal{V}$ is taken in the entire space $\Delta(C)$. This is because a cutoff rule can induce an arbitrary belief over the cutoff but can only induce certain beliefs over the winner's type (namely those that MLR-dominate the prior). Second, in (4.9) the concavified objective is multiplied by an additional term $\int_\Theta \bar{x}(\theta)f(\theta)d\theta$ – the ex-ante probability of allocating the good to any single agent. This is because the distribution of beliefs over the winner's type is a conditional distribution (conditional on allocating the good to an agent), so the conditional expected

[19] See Appendix C.3 for a formal definition.

payoff must be converted into an ex-ante expected payoff. As a corollary of Lemma 5, I obtain the following result.[20] (See Appendix C for a proof.)

Corollary 2. *If $\mathcal{W}$ is convex on its domain, the optimal cutoff mechanism fully discloses the cutoff. If $\mathcal{W}$ is concave, the optimal cutoff mechanism reveals no information.*

To give a sharper characterization of optimal cutoff mechanisms, 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.

Proposition 4. *Suppose that f is a continuous density, fully-supported on $[0, 1]$.*

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 information, where*

$$r^\star \in \underset{r \in [0,\,1]}{argmax} \left(1 - F^N(r)\right) W\left(\frac{\int_r^1 \theta dF^N(\theta)}{1 - F^N(r)}\right). \tag{4.10}$$

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.*

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

[20] 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.

good is not always allocated. The $r^\star$ that solves equation (4.10) 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 the cutoff is optimal. The optimal allocation rule is determined by the properties of the function $J_w(c)$ which captures the local trade-off between allocative efficiency (as captured by the term $w(c)$) and the induced information structure (as captured by the term $-w'(c)(1 - F(c))/f(c)$. Allocating the good with smaller probability conditional on realization c lowers surplus if $w(c)$ is positive but increases posterior beliefs over the winner's type conditional on allocating. The function $J_w(c)$ is similar to the virtual surplus function which captures the trade-off between allocative efficiency and information rents in the revenue-maximization problem. In the regular case, the virtual surplus function is increasing, and the seller does not introduce randomization to the revenue-maximizing mechanism. Analogously, if $J_w(c)$ is increasing, the designer does not use randomization in the allocation rule to optimally influence beliefs in the aftermarket. An extended discussion can be found in the Online Appendix OA.2.1.

If W is neither convex nor concave, the problem can be solved if W is sufficiently regular by applying the duality approach of Kolotilin (2016) or Dworczak and Martini (2017). 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.[21]

5 Characterizations of cutoff rules

In this section, I provide three characterization theorems for the class of cutoff mechanisms arguing that the class might be attractive from the perspective of practical mechanism design. For ease of exposition, I show the results for the case of a finite type space $\boldsymbol{\Theta}$. A continuous type space is useful for modeling purposes because it allows using calculus but I view finitiness as being essentially without loss of (economic) generality for the purposes of practical mechanism design. When the type space is finite, it is without loss to study mechanisms with a finite signal space – I call such mechanisms $\mathcal{S}-$ finite.

A key step in providing the characterizations is a generalization of Lemma 2.

[21] To use the results of Dworczak and Martini (2017), one has to modify their method to account for the fact that under a cutoff rule, each posterior belief has to MLR-dominate the prior. I show how this can be done in the proof of Claim 3 in Appendix E.4.

Lemma 6. *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 6 is of independent interest as it draws a connection between cutoff rules and Myerson monotonicity as discussed in Section 2.

5.1 Cutoff rules are implementable for all distributions and aftermarkets

The first characterization theorem is an extension of Propositions 1, 3*a*, and 3*b* from Section 2. To that end, I want to prove the converse of Theorem 1 – *only* cutoff rules are implementable regardless of the distribution and the aftermarket. 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 and characteristics of third-party players, or different equilibria of the same aftermarket game.

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 $A \in \mathcal{A}$.

Theorem 1 shows that cutoff rules are flexible with respect to all distributions and monotone aftermarkets. To prove a converse, I define an appropriate Richness condition.

Definition 7 (Richness)**.** The pair $(\mathcal{F}, \mathcal{A})$ satisfies *Richness* if for any $\mathcal{S}-$finite mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$, $i \in N$, $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$, there exists a prior distribution $\boldsymbol{f} \in \mathcal{F}$ and an aftermarket $A \in \mathcal{A}$ such that

$$\pi_i(s|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) < \pi_i(s|\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i}) \implies u_i(\theta_i;\, \boldsymbol{f}^{i,s}) > u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s}), \tag{5.1}$$

$$\pi_i(s|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) > \pi_i(s|\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i}) \implies u_i(\theta_i;\, \boldsymbol{f}^{i,s}) = u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s}). \tag{5.2}$$

Under the Richness condition, the converse to Theorem 1 holds.

Theorem 3. *Suppose that an $\mathcal{S}$-finite mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is flexible with respect to $(\mathcal{F}, \mathcal{A})$ which satisfy the Richness condition. Then, $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule.*

Thus, if $(\mathcal{F}, \mathcal{A})$ satisfy Monotonicity and Richness, flexibility is a defining property of cutoff mechanisms.

The intuition for Theorem 3 is analogous to that for Propositions $3a$ –$3b$ so I focus on the intuition for the Richness condition in the discussion below.

The existence of a post-mechanism stage imposes restrictions on the set of implementable disclosure rules only if payoffs in the aftermarket are sensitive to information revealed by the mechanism. However, this in itself is not enough. For example, if the agent and the third parties have aligned preferences, it might be 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. The premise in condition (5.1) can be interpreted as "bad news" about the agent's type – after observing a signal s that satisfies the left-hand side inequality (for a fixed $\boldsymbol{\theta}_{-i}$), 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 5.2), the expected payoffs of the two types should be equal. Conditions (5.1) and (5.2) resemble submodularity of the agent's payoff in the type and posterior belief: Higher types enjoy higher payoffs than low types for any aftermarket belief (monotonicity) but the difference in payoffs decreases (to zero) when the belief over the winner's type "increases".

This feature is a generalization of the property of the simple aftermarket from Section 2: Higher types value each allocation more than low types but lower types have a higher willingness to pay for *signals* leading to higher posterior beliefs. To see this, consider the resale game of Section 2 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 the fixed aftermarket of Section 2, i.e., a singleton $\mathcal{A}$) 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 3 in Appendix A.2.2). 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 $\hat{\theta}$.

In the multi-agent setting, an analogous resale game satisfies the Richness condition with respect to ($\mathcal{F}$, $\mathcal{A}$) as long as $\mathcal{F}$ contains independent distributions of types, and $\mathcal{A}$ contains aftermarkets corresponding to a sufficiently rich set of valuation functions for the third party (in particular, the third party must have a value higher than any agent's value in some aftermarket). Flexibility characterizes cutoff rules on every superset of such ($\mathcal{F}$, $\mathcal{A}$): There can be correlation of agents' types, and multiple third parties with private information and interdependent valuations – the designer can always use a cutoff rule, and no other mechanism frame can be implemented against all these scenarios. I give one more example (without a proof) of a setting that satisfies the Richness condition below.

Example 2 (Buying a complementary good)**.** The mechanism designer sells an item to one of N buyers. For simplicity, assume that types of all agents come from the same 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 $\mathcal{F}$ contains the set of all joint product distributions on Θ^N, and $\mathcal{A}$ includes aftermarkets corresponding to all $r \in \mathcal{R}$. Then, ($\mathcal{F}$, $\mathcal{A}$) satisfy Monotonicity and Richness.

5.1.1 Discussion

Two examples demonstrate the usefulness of flexible mechanisms in practical problems. First, many real-life situations require one mechanism to handle multiple instances of the problem. When designing informational requirements for a financial over-the-counter market, a regulator may not be able to condition the design on the distribution of types and other details which might vary across 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[22], major parts of the mechanism must be fixed long before implementation to allow enough time for administrative procedures and communication with participants. 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

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

the mechanism frame is fixed.

In general, flexible implementation allows transfers to be a function of the distribution of types (and the form of the aftermarket). In the Online Appendix OA.3.2, I show that it is sometimes possible to implement a cutoff rule in such a way that the transfers are pinned down by the equilibrium behavior of agents in an indirect mechanism. In this case, the cutoff rule can be implemented robustly, that is, even if the designer lacks any knowledge of the distribution or the aftermarket. It is easy to observe that flexibility is a necessary condition for robust implementation. This means that a designer interested in robust implementation of a mechanism frame has no reason to look beyond the class of cutoff mechanisms.

In the above discussion, I have focused on the case when the mechanism designer implements the same mechanism frame (x, π) regardless of the distribution of types and the aftermarket. For a fixed distribution and aftermarket, 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 and the aftermarket may vary, it might be natural to look at other representations of mechanisms. For example, in the case when the aftermarket is a resale game, suppose that the designer implements the same final (ex-post) allocation regardless of the distribution of types. This case could lead to a different theory of flexible mechanisms. The advantage of considering mechanism frames is that they provide a complete description of final outcomes regardless of the form of the aftermarket. In other settings, a designer might fix an indirect mechanism, allowing the allocation and disclosure rule to be determined endogenously in equilibrium as the distribution and aftermarket vary. Optimization in the class of all indirect mechanisms appears intractable, partly because it is typically not feasible to describe all their equilibria (the aftermarket induces a signaling game). Traction can be gained by restricting the set of feasible indirect mechanisms to a small class.[23]

5.2 Cutoff rules are robust to mistrust

Cutoff rules are uniquely identified by the monotonicity condition ($\boldsymbol{M}$). In this subsection, I provide an economic interpretation of monotonicity by showing that it implies desirable properties in a setting where agents have limited trust towards the mechanism designer.

Suppose that each bidder i believes that the designer might renege on the outcome of the mechanism. Let S_i be the set of signals that occur with (ex-ante) positive probability when agent i is the winner. Agent i believes that conditional on winning and s, $\boldsymbol{\theta}_{-i}$ being

[23] This approach has been adopted by a number of papers cited in the Literature Review section. These papers typically study equilibria of first- and second-price auctions for a particular aftermarket.

realized, the seller will allocate the good and send some other signal $s' \in S_i$ with probability $\mu_i(s'|\, s, \boldsymbol{\theta}_{-i})$. Moreover, $\sum_{s'} \mu_i(s'|\, s, \boldsymbol{\theta}_{-i}) \leq 1$, allowing for the possibility (when the inequality is strict) that the good is not allocated at all (in which case the signal is irrelevant and no payments are exchanged between bidder i and the designer). For simplicity, I assume that agents believe that the third parties trust the designer so that any signal $s' \in S_i$ will be interpreted by third parties as truthfully revealed by the designer given the original mechanism $(\boldsymbol{x}, \boldsymbol{\pi})$.[24] Then, when agent i holds beliefs μ_i and reports $\hat{\theta}$, her payoff is (not including transfers)

$$\sum_{s \in S_i} \sum_{s' \in S_i} u_i(\theta_i;\, \boldsymbol{f}^{i,s'}) \mu_i(s'|s, \boldsymbol{\theta}_{-i}) \pi_i(s|\hat{\theta},\, \boldsymbol{\theta}_{-i}) x_i(\hat{\theta},\, \boldsymbol{\theta}_{-i}).$$

I allow the designer to condition transfers on the revealed signal. Specifically, let $t_i^s(\theta_i,\, \boldsymbol{\theta}_{-i})$ denote the transfer paid by agent i conditional on winning the good and signal s being sent. Let $\boldsymbol{t} = \{t_i^s\}_{i=\in N}^{s \in S_i}$. The idea behind mistrust is that agents believe that they will pay t_i^s when they receive the good and signal s is revealed (it is easy to prove that the seller cheated if this is not the case) but they do not necessarily trust that the seller will generate the outcome of the mechanism according to $(\boldsymbol{x}, \boldsymbol{\pi})$ (the outcome is determined by randomization that the agent might not be able to verify). Given beliefs $\{\mu_i\}_{i \in N}$, Definition 3 of DS implementability extends naturally.

Let Λ denote an arbitrary set of belief profiles $\{\mu_i\}_{i \in N}$. A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is *weakly robust to mistrust wrt* Λ if:

$$\text{for any } \mu \in \Lambda, \text{ there exists } \boldsymbol{t} \text{ such that } (\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t}) \text{ is DS implementable under } \mu.$$

A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is *strongly robust to mistrust wrt* Λ if:

$$\text{there exists } \boldsymbol{t} \text{ such that, for any } \mu \in \Lambda, (\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t}) \text{ is DS implementable under } \mu.$$

If a mechanism is weakly robust to mistrust, the seller must know the form of mistrust to guarantee that agents will report truthfully regardless of their mistrust. Under strong robustness to mistrust, the seller can guarantee truthfulness without any knowledge of agents' mistrust. Let Λ_{all} be the set of all possible belief profiles $\mu = \{\mu_i\}_{i \in N}$, and let

$$\Lambda_{\text{alloc}} = \{\mu :\, \mu_i(s'|s, \boldsymbol{\theta}_{-i}) = \delta_i(s, \boldsymbol{\theta}_{-i}) \mathbf{1}_{\{s=s'\}},\, \forall i,\, \forall s,\, s' \in S_i, \text{ and some } \delta_i :\, S_i \times \boldsymbol{\Theta}_{-i} \to [0,\, 1]\}$$

[24] Another implicit assumption is that agents believe that their own report does not affect the probability of reneging. Without that assumption, no mechanism would be robust to mistrust towards the designer.

be the set of beliefs that the seller will report the signals truthfully but might renege on allocating the good.

Theorem 4. *Suppose that $u_i(\theta_i; \bar{\boldsymbol{f}}) > u_i(\hat{\theta}_i; \bar{\boldsymbol{f}})$ for all $\theta_i > \hat{\theta}_i$ and all posterior belief profiles $\bar{\boldsymbol{f}}$. A mechanism frame is weakly (strongly) robust to mistrust with respect to Λ_{all} (with respect to Λ_{alloc}) if and only if it is a cutoff rule.*

Remark 1. The strict inequality $u_i(\theta_i; \bar{\boldsymbol{f}}) > u_i(\hat{\theta}_i; \bar{\boldsymbol{f}})$ is only needed for the "only if" part of Theorem 4.

The intuition for the result follows from the monotonicity condition ($\boldsymbol{M}$). We can view the seller as a multi-product monopolist where every signal realization corresponds to a different product. The monopolist sells only one of these products ex-post, and commits ex-ante to a probability distribution over which product is sold. If the buyer trusts the seller, the probability of allocating some of the products can be non-monotone in the type of the buyer if it is offset by a sufficient increase in probability of allocating other products. However, when the buyer is allowed to have rich beliefs about which of these products the monopolist will eventually allocate, the only implementable schemes are the ones in which *each* product is sold with increasing probability to higher types. In the context of mechanism design with aftermarkets, these are schemes that satisfy ($\boldsymbol{M}$), i.e., cutoff rules.

5.3 Cutoff rules are implemented by generalized clock auctions

In this section, I provide a characterization of cutoff rules as equilibria of simple dynamic auctions, called Generalized Clock Auctions (GCAs). Because the details of this indirect implementation scheme are only tangentially related to the main theme of the paper, I describe the idea informally here, and formalize it in Appendix D.4.

A GCA is a dynamic bidding procedure in which agents gradually drop out until a winner is determined. In every round, the auctioneer announces the current price, and bidders simultaneously decide whether to stay or exit. At the end of the round, the auctioneer reveals to the bidders the outcome of the round. The auction ends when all active bidders drop out (in which case the object is allocated to one of them at random), or there is only one bidder remaining (in which case she becomes the winner).

A disclosure rule for a GCA specifies what part of the bidding history is publicly revealed when the auction ends. For example, the auctioneer can reveal the final price, or the set of active bidders in every round. In general, the signal can be an arbitrary garbling of the information contained in the sequence of prices and bidders' decisions at each stage.

A GCA is an indirect mechanism. The following result connects the outcomes implemented by equilibria of GCAs to cutoff rules. A formal statement of Theorem 5 is available in Appendix D.4.

Theorem 5 (Informal). *Any outcome implemented by a monotone equilibrium of any GCA is a cutoff rule. Conversely, any symmetric cutoff rule (subject to a mild restriction on allocation rules) can be implemented by some GCA.*

Intuitively, in a GCA, the history of bidding does not directly depend on the 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.

Prices do not have to change monotonically in a GCA.[25] Because the informational content of the signal in general depends on the termination time of the auction, it is as if a different good was offered for sale in every round. Prices may have to decrease when the current-round signal induces posterior beliefs that are less attractive for bidders. The auction nevertheless preserves monotonicity in types. Allocation rules that involve randomization are implemented through price jumps in the auction that induce a range of types to exit in the same round.

6 Application – Optimal design of OTC markets

In this section, I illustrate the theoretical results with an application to the design of trading protocols in the financial over-the-counter market. (Additional applications can be found in the Online Appendix OA.4.) 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). My model is well suited for studying questions of post-transaction transparency in cases where the buyer of an asset acts as a dealer and resells the asset in the aftermarket.

An immediate conclusion from the analysis of cutoff mechanisms is that certain types of information can always be disclosed in a trading mechanism, while other may be in conflict with incentive compatibility. Consider the case where dealers purchase an asset on a platform, such as a Swap Execution Facility (SEF), and attempt to resell it in an aftermarket (for example, in the inter-dealer market or in a bilateral transaction with an individual investor). Theorem 3 implies that there is a fundamental difference between disclosure of the winning offer and disclosure of the losing offers. If dealers have a private

[25] This is one key difference to obviously strategy-proof auctions considered by Li (2016).

cost of intermediation, and the platform is designed to allocate the asset efficiently, then no information about the winner's offer should be revealed. However, the offers of dealers who did not buy the asset may be revealed without distorting efficiency of the allocation. For example, the designer can always reveal the average value of losing dealers in order to disclose information about the common value component of the asset. This is true even if the information of the participating dealers is cross-sectionally correlated.

Because the implications for transparency depend on the number of participating agents, I introduce two models of trading in the OTC market. In Subsection 6.1, I focus on the problem of post-transaction transparency in bilateral trade settings, where a single dealer intermediates the trade between the seller and a final buyer. In Subsection 6.2, I consider the design of optimal trading platforms. The two models feature different aftermarkets to illustrate the generality of the methods developed in preceding sections.

6.1 Bilateral Trade Setting

The seller has an asset that she values at k (this can be though of as a deterministic cost of providing the asset). There is a single agent (dealer) with type θ, distributed according to a continuous distribution F with a regular density f on $[0, 1]$.[26] If the dealer acquires the asset, then with 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 $v(\theta)$ for the asset that depends on the value of the dealer. The value $v(\theta)$ is strictly increasing, with $v(0) < k$ and $v(1) > 1$ (these assumptions rule out uninteresting boundary solutions). In the aftermarket, the dealer makes a take-it-or-leave-it offer to the third party. I focus on this case which is more realistic in the OTC markets setting but 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.

In the above situation, the disclosure rule of the first-stage mechanism corresponds to the informational transparency of the market. Because the dealer is better informed about the asset's value for the third party, and makes an offer in the aftermarket, adverse selection may lead to inefficiently low trade volume. A frequent argument in favor of a transparent primary market (an informative disclosure rule) is that it reduces the information asymmetry in the aftermarket. To examine this claim, I assume that the lemons problem occurs in the aftermarket if no information is revealed. Formally, a price equal to 1 (which leads to trade with probability one) is not an equilibrium of the aftermarket when the seller allocates to

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

types above her cost and reveals no information:

$$\int_k^1 (v(\theta) - 1) f(\theta) d\theta < 0. \tag{6.1}$$

In the analysis of this OTC market, it will be 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.

6.1.1 Regulator's preferred mechanism

I first assume that the mechanism in the primary market is chosen by a benevolent regulator who maximizes efficiency in the market. In the Online Appendix OA.4.1, I analyze a related problem in which the regulator can only impose a disclosure rule, and the seller chooses an allocation rule.

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

$$p(x) = \max\{p \in [0, 1] : \int_0^p (v(\theta) - p) x(\theta) f(\theta) d\theta \geq 0\}. \tag{6.2}$$

If the seller offers quantity 1 to all types above her cost, by assumption (6.1), 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 of the cutoff 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, 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 were worried about the impact of disclosed information on their bargaining position in the resale stage.

The designer faces a complicated trade-off between costly screening (limiting the quantity sold to lower types) and reducing information asymmetry in the aftermarket. The trade-off would be difficult to resolve if not for the second part of Theorem 2, which, for the case $N = 1$, guarantees existence of an optimal mechanism that reveals no information (see also Proposition 2). Since no explicit signals are sent in the optimal mechanism, it is enough to solve the unconstrained problem

$$\max_{x \in \mathcal{X}} \int_0^{p(x)} \left[\lambda v(\theta) + (1-\lambda)\theta - k\right] x(\theta) f(\theta) d\theta + \int_{p(x)}^1 (\theta - k) x(\theta) f(\theta) d\theta, \tag{6.3}$$

where $p(x)$ is given by (6.2). 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.

Claim 1. *Problem (6.3) admits a solution of the form* $x(\theta) = \mathbf{1}_{\{\theta \geq r^\star_{eff}\}}$ *for some* $r^\star_{eff} \in [0, 1]$.

The proof and the definition of $r^\star_{\text{eff}}$ can be found in Appendix E.1. The optimal scheme is a posted-price mechanism in which all quantity is offered for sale. The price is chosen to optimally trade off the losses from not allocating the asset against the higher realized gains from trade in the aftermarket conditional on allocaiton.[27]

6.1.2 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 E.2, I show how to derive the objective function of the form (4.1). Applying Theorem 2 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:* $x(\theta) = \mathbf{1}_{\{\theta \geq r^\star_{rev}\}}$ *for some* $r^\star_{rev} \in [0, 1]$.

Moreover, $r^\star_{eff} \leq r^\star_{rev}$ *with equality if and only if* $p(\mathbf{1}_{\{\theta \geq r^\star_{eff}\}}) = 1$ *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 E.3.

[27] 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.

6.1.3 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 profit-maximizing mechanism 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. In this case, the agent acts primarily as an intermediary (rarely ends up holding the good), and her payoff becomes less sensitive to the type. Thus, the seller can extract rents without excluding low types.

The above analysis implies that transparency is not necessarily desirable from an efficiency viewpoint. However, 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 Online Appendix OA.5.1 (see also Section 8.1), 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 by the optimal mechanism.

In the optimal mechanisms studied above the designer 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.

6.2 Designing a trading platform

I now consider the problem of designing a trading platform with $N \geq 2$ dealers. To illustrate the generality of the framework, I consider a different model of the aftermarket.

Dealers (agents) participate in a mechanism (trading platform) which allocates a single asset. For simplicity, I treat the platform as a seller of the asset.[28] Unlike in the previous

[28] The case where the platform serves the role of matching buyers with a seller, and makes money by charging fees, could also be analyzed.

model, dealers have no intrinsic value for holding the asset but differ in intermediation effectiveness measured by each dealer i's type $\theta_i \in [0, 1]$. More effective dealers have access to a larger number of final buyers (third parties). I model the aftermarket as a first-price auction (with no reserve price) among final buyers. Each final buyer has a value $v > 0$ for the asset. Conditional on acquiring the asset, with probability $1 - \theta_i$, dealer i has only one potential buyer, called the default buyer. With probability θ_i, there is a second buyer, called the alternative buyer. Importantly, the default buyer does not know whether or not she has a competitor in the resale auction. The dealer cannot credibly reveal that there is a second buyer, so the default buyer relies on information revealed by the mechanism to form beliefs about the level of competition in the aftermarket.

To gain intuition for how the resale game depends on information revealed in the initial mechanism, suppose that the default buyer is sure that she has a competitor (she believes that $\theta_i = 1$ with probability one), and that belief is common knowledge. Then, both buyers bid v in the auction (Bertrand competition). If she assigns probability one to the event $\theta_i = 0$, she bids bid 0 (the reserve value of the dealer). For interior beliefs, final buyers randomize their offers in a mixed-strategy equilibrium.

This aftermarket should not be taken literally as a description of a resale mechanism in the OTC market. Rather, it is a tractable reduced-form framework capturing a situation in which a dealer acts as an intermediary whose value for the object depends only on the resale opportunity. By observing the outcome of the mechanism, final buyers form beliefs about the level of competition in the resale market and adjust their offers accordingly. High prices in the original mechanism can be used as credible signals of high intermediation effectiveness of the winning dealer and, when revealed, lead to more aggressive bidding in the resale stage.

I restrict attention to a class of distributions of dealers' types: $F(\theta) = 1 - (1 - \theta)^{1/\beta}$, $\theta \in [0, 1]$, $\beta > 0$. This is the largest class of distributions with the property that the virtual surplus function is affine: $\theta - (1 - F(\theta))/f(\theta) = (1 + \beta)\theta - \beta$. Uniform distribution is a special case with $\beta = 1$. This restriction allows me to apply the results of Section 4.3 and provides clean comparative statics as β increases. For low β, platform profit maximization is closely aligned with maximizing the surplus of the dealers. For high β, a profit-maximizing platform puts more weight on minimizing the information rents of dealers.

Claim 3. *The mechanism that maximizes dealers' surplus is a second-price auction with a strictly positive reserve price and full disclosure of the price paid by the winning dealer.*

Suppose instead that the objective function is to maximize the platform's profits. Then, the optimal mechanism reveals the price if it is lower than a threshold, and only reports that the price was higher than that threshold otherwise.[29] *As* β *approaches* 0*, the threshold*

[29] I prove in Online Appendix OA.4.2 that the optimal mechanism in this case is also a second price auction

converges to the highest price (the disclosure rule in the optimal mechanism converges to full revelation). Conversely, when $\beta \geq \bar{\beta} > 2$,[30] the threshold is equal to the lowest price (no information is disclosed about the price). When $\beta > \sqrt{e} - 1$, no revelation is optimal if the number of dealers N is sufficiently large.

Let m denote the expectation of the type of the winning dealer held by the default buyer. In the unique equilibrium of the resale auction, the alternative buyer (whenever present) bids b or less with probability $H(b) = (1 - m)/m \cdot b/(v - b)$ for $b \in [0, mv]$. The default buyer bids 0 with probability $1 - m$, and according to the same distribution $H(b)$ with remaining probability. Because the bid distribution depends only on the mean of the posterior belief, and the objective function of the platform is affine in the type of the dealer, I can apply the results of Section 4.3 to solve for the optimal mechanism.

If the platform maximizes dealer surplus, it is optimal to set a reserve price and fully reveal the cutoff. This corresponds to revealing the price paid by the winning dealer in a second price auction. Full disclosure of the price is in the interest of dealers because it induces (on average) more aggressive bidding by the final buyers.[31] A less intuitive feature of the dealer-optimal mechanism is that the reserve price is strictly positive even though all dealer types have strictly positive utility from intermediating the trade. The reserve price excludes less effective (low θ) dealers from trading in order to increase the resale profits of the more effective (high θ) dealers through the effect on aftermarket beliefs. The aftermarket beliefs are increased (relative to the case of a zero reserve price) whenever a dealer purchases the asset at the reserve price. The benefit of this effect outweighs the loss associated with excluding the lowest types from trading because the contribution of the lowest types to total dealer surplus under full disclosure is relatively small.

If the platform maximizes profits, it faces the usual trade-off between maximizing dealer surplus and minimizing their information rents. This can be shown to lead to an increase in the reserve price. More interestingly, the optimal disclosure policy also changes. If β is low, the platform is not too concerned about dealers' information rents, and the profit-maximizing mechanism reveals the price almost fully, similarly to the dealer-optimal mechanism. As β grows, the optimal mechanism garbles information about high prices to reduce the information rents of the winning dealer. Unlike in the standard model without the aftermarket, upward incentive compatibility constraints are not redundant in the problem with infor-

with a positive reserve price.

[30] Using a numerical calculation, $\bar{\beta} \approx 2.16$.

[31] This effect is reminiscent of the linkage principle which implies that the seller should reveal information affiliated with bidder information to maximize her revenue. It also suggests that dealer surplus could be further improved by revealing the exact type of the winning dealer - this, however, is not possible in any incentive compatible mechanism.

mation disclosure. When information about high prices is pooled, highest types enjoy less surplus but it becomes less attractive for low types to pretend that they are high types. When β is sufficiently high ($\beta \geq \bar{\beta}$) this effect dominates – the platform provides no information about the price. To gain intuition, note that as β increases, the expected type of the winner increases (for any signal realization), and the bid distribution in the resale auction converges to an atom at v. As a consequence, the difference in profits between high and low types gets smaller (the expected difference between the maximum of two draws and one draw from a distribution gets smaller as the distribution gets less dispersed). Thus, for high β, low types enjoy relatively high information rents, and the platform minimizes these rents to maximize profits. A similar effects takes place when the number of dealers N gets large.

The insights of the above example – the fact that dealers prefer more transparency than the profit-maximizing platform – rely on the structure of the aftermarket. Signaling a higher type is better for the dealers because it increases competition in the resale market. However, this is not true across all bargaining protocols. Consider a model where the type of the dealer is the value of her outside option, and a single buyer makes an offer in the aftermarket to buy the asset. In this case, the dealer would always want to preserve some private information about the value of the outside option to make non-zero profits. It is possible to show that (under certain conditions on the distribution of the value of the outside option) that the dealer-optimal mechanism in this case would not disclose any information.

7 Literature review

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.

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).[32] 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, 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)[33], *(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).[34]

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 fully revealed. In contrast, Engelbrecht-Wiggans and Kahn (1991) and Dworczak

[32] 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 (Dworczak, 2017) which establishes conditions under which cutoff mechanisms are optimal.

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

[34] The designer in Carroll and Segal (2016) does not control the information leakage to the aftermarket.

(2015) explicitly construct 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.

Balzer and Schneider (2017) analyze a model in which two players try to resolve a conflict which (if unresolved) leads to an escalation game between the two sides. Because the behavior in the conflict management mechanism is informative of the payoff-relevant types of the players, a designer can influence payoffs in the escalation game by disclosing information in the mechanism. While the two problems are related on a high level, the models, and hence the techniques used to analyze them, are different.

The presence of aftermarkets has been cited as an important motivation for studying mechanisms with allocative and informational externalities, for example in Jehiel, Moldovanu and Stacchetti (1996) and Jehiel and Moldovanu (2001, 2006). These papers focus on the case of interaction between the mechanism participants and hence do not explicitly consider the effects of ex-post information disclosure.

8 Extensions

8.1 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. In many settings, the designer has access to exogenous information. For example, the seller might have a private value for the asset or know something about its quality. In Online Appendix OA.5.1, I propose an extended model (for the case $N = 1$) in which the designer observes the realization of some exogenous variable z. To avoid problems with signaling through the choice of the mechanism, I assume that this information is also observed by the agent but not by the third party. Using a cutoff rule, which is always implementable, the designer optimizes over disclosure rules of exogenous information about z and endogenous information about the cutoff c. The marginal distribution of z is given, while the marginal distribution of c can be chosen.

The interplay between revelation of exogenous and endogenous information results in a generalization of Theorem 2 with $N = 1$: There always exists an optimal cutoff mechanism with a signal that is measurable with respect to the exogenous information z. Revealing information about z may be strictly optimal even if z itself is payoff-irrelevant for the agent and the third party. This is because z is informative about the cutoff when its distribution (allocation rule) varies with z. For example, suppose that z is the private value of the seller, and the seller only allocates to types θ above z. Then, the variable z plays a role similar to that of a second highest bid in an auction setting – its disclosure may be optimal even if it does not directly enter the utility function of the winner and the third party.

I apply this extension to the OTC market model from Section 6.1 by assuming that the cost k of the seller is random. If the cost k is interpreted as the borrowing cost in the interbank money market, and the mechanism as a bilateral transaction with a dealer, then the conclusion is that it suffices to reveal information about the borrowing costs (for example, in the form of a benchmark, such as LIBOR), and it is not necessary to reveal any information about individual transactions.

8.2 What if the loser also interacts in the aftermarket?

In 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 does not acquire the good. For example, a loser may purchase a similar object in the aftermarket, or negotiate to gain access to an object owned by another market participant.

In Online Appendix OA.5.2, I propose an extension of the model with one agent in which the winner participates in the winner's aftermarket, and the loser participates in the loser's aftermarket. A mechanism now includes a pair of signal functions. Because only one signal is sent ex-post (depending on whether the agent acquires the good), all results from the one-agent model generalize easily under an appropriate single-crossing assumption.[35] In particular, an optimal cutoff mechanism reveals no information.

The extension to multiple-players is significantly more complicated, and the conclusions will depend on the setting. If $N = 2$, and only the loser interacts in the aftermarket, it is the value of the winner, not the loser, that can be revealed without upsetting incentive compatibility. However, when $N \geq 3$, or if losing and winning players both interact (with third parties), the exact set of robustly implementable disclosure rules is difficult to characterize. In the extreme version of the model where all agents are allowed to participate in the aftermarket, and implementability is required for a large set of aftermarkets, the only type of information that can be always revealed is information about the outcome of randomization performed by the seller, e.g., a random reserve price.

[35] This assumption says, roughly, that the difference in payoffs between any two types θ and $\hat{\theta}$ is larger in the winner's aftermarket than in the loser's aftermarket.

8.3 Interdependent valuations

Even in the most general formulation of the model, I have assumed that the utility of agent i does not depend on other agents' types $\boldsymbol{\theta}_{-i}$. Such dependence is natural in models where the asset has a common-value component, or when the interaction in the aftermarket is with other agents (rather than with a third party). If u_i depends on $\boldsymbol{\theta}_{-i}$ but only the winner interacts in the aftermarket, then the results go through with one modification: DS implementability has to be changed to ex-post implementability. The latter case, when interdependence of valuations stems from the fact that agents interact with each other in the aftermarket, is outside of the scope of this paper. The problem is significantly more complicated: Except for signaling in the first-stage mechanism, agents also engage in information acquisition about the competitor's types, further complicating incentive compatibility constraints.

9 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 trading mechanisms in financial over-the-counter markets.

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 worst-case analysis on the other. The designer in my model restricts attention to subclass of mechanisms (cutoff mechanisms) with a certain robustness property (implementability in the "worst-case") but maximizes 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.

Balzer, Benjamin and Johannes Schneider, "Managing a Conflict," Working Paper 2017.

Becker, Gordon M., Morris H. Degroot, and Jacob Marschak, "Measuring utility by a single-response sequential method," *Behavioral Science*, 1964, *9* (3), 226–232.

Bergemann, Dirk and Stephen Morris, "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.

_ , "Reduced Form Auctions Revisited," *Economic Theory*, Apr 2007, *31* (1), 167–181.

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.

Crémer, Jacques and Richard P. McLean, "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1988, *56* (6), 1247–1257.

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.

— , "Mechanism Design with Aftermarkets: On the Optimality of Cutoff Mechanisms," Working Paper 2017.

— **and Anthony Lee Zhang**, "Implementability, Walrasian equilibria, and efficient matchings," *Economics Letters*, 2017, *153* (Supplement C), 57 – 60.

— **and Giorgio Martini**, "The Simple Economics of Optimal Persuasion," Working Paper 2017.

Engelbrecht-Wiggans, Richard and Charles M. Kahn, "Protecting the winner. Second-price versus oral auctions," *Economics Letters*, 1991, *35*, 243 – 248.

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.

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.

Mierendorff, Konrad, "Asymmetric reduced form Auctions," *Economics Letters*, 2011, *110* (1), 41 – 44.

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.

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 Supplementary materials for Section 2

A.1 Proof of Lemma 1

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. By definition of a cutoff rule, there exists a function $\gamma : C \to \Delta(\mathcal{S})$ such that $\pi(s|\,\theta)x(\theta) = \sum_{c\le\theta}\gamma(s|\,c)dx(c)$. Thus, the problem becomes

$$\begin{aligned}\max_{\gamma}\sum_{\theta\in\Theta}\sum_{s\in\mathcal{S}} V(\theta;\, f^s)\sum_{c\le\theta}\gamma(s|\,c)dx(c)f(\theta)&\\ = \max_{\gamma}\sum_{s\in\mathcal{S}}\underbrace{\left(\sum_{c}\gamma(s|\,c)dx(c)\right)}_{\varsigma_s}\sum_{\theta\in\Theta}V(\theta;\, f^s)\underbrace{\left(\frac{\sum_{c\le\theta}\gamma(s|\,c)dx(c)}{\sum_{c}\gamma(s|\,c)dx(c)}\right)}_{G^s(\theta)}f(\theta).&\qquad\text{(A.1)}\end{aligned}$$

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 (2.5), where G^s is the posterior cumulative distribution function of the cutoff conditional on signal s. Thus, the objective function can be written as

$$\mathbb{E}_{s\sim\varsigma}\mathcal{V}(G^s).\qquad\text{(A.2)}$$

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

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

Thus, the problem is formally equivalent to the Bayesian persuasion problem of Kamenica and Gentzkow (2011). Instead of optimizing over distributions ς of signals, we can optimize over distributions of posterior beliefs $\varrho\in\Delta(\Delta(C))$ subject to a Bayes-plausibility constraint. This yields equations (2.6) and (2.7). Equation (2.7) is the Bayes-plausibility constraint on posterior beliefs over the cutoff c expressed in terms of cdfs.

A.2 Proofs of Proposition 1, $3a$, and $3b$

First, note that to prove the three propositions, it is enough to prove Lemma 2 and 3. Indeed, any cutoff rules satisfies the monotonicity property (M) by Lemma 2, and thus, by Lemma 3, it is implementable for every distribution of types f, establishing Proposition 1. Conversely, suppose that $(x,\,\pi)$ is implementable for every distribution f, under the assumptions of

Proposition 3a (or 3b). By Lemma 3, (x, π) satisfies condition (M). Lemma 2 implies that (x, π) is a cutoff rule, yielding Proposition 3a (and 3b).

A.2.1 Proof of Lemma 2

The fact that a cutoff rule satisfies property (M) is trivial – it follows directly from Definition 2 of cutoff rules. I prove the converse part below. The proof uses the Radon-Nikodym Theorem. While an elementary proof can be provided for the case of a finite Θ, using this method shortens the proof and admits a direct extension to continuous type spaces.[36]

Fixing a cutoff rule (x, π), let $\beta_s(\theta) \equiv \pi(s|\theta)x(\theta)$. By condition (M), $\beta_s(\theta)$ is a non-decreasing function on Θ, for any s. Summing over $s \in \mathcal{S}$, we get that $x(\theta)$ is non-decreasing. Let $\underline{\theta} = \min(\Theta)$, and $\hat{\theta}$ be the largest type in Θ smaller than θ, for any $\theta > \underline{\theta}$. Because β_s is non-decreasing, it induces a positive additive (not necessarily probabilistic) measure μ_s on C defined by $\mu_s(\underline{\theta}) = \beta_s(\underline{\theta})$, and

$$\mu_s(\theta) = \beta_s(\theta) - \beta_s(\hat{\theta}),\ \theta > \underline{\theta}.$$

The measure μ_s is absolutely continuous with respect to the distribution dx of the cutoff. Indeed, because both measures are discrete, this follows from

$$\mu_s(\theta) \leq \sum_{s' \in \mathcal{S}} \mu_{s'}(\theta) = dx(\theta).$$

By the Radon-Nikodym Theorem, there exists a positive function g_s on C that is a density of μ_s with respect to dx. In particular,

$$\pi(s|\theta)x(\theta) = \beta_s(\theta) = \mu_s(\{\tau : \tau \leq \theta\}) = \sum_{c \leq \theta} g_s(c)dx(c), \tag{A.3}$$

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

$$\sum_{s \in \mathcal{S}} \pi(s|\theta)x(\theta) = x(\theta) = \sum_{c \leq \theta} \sum_{s \in \mathcal{S}} g_s(c)dx(c).$$

Thus, for all θ,

$$\sum_{c \leq \theta} \left(\sum_{s \in \mathcal{S}} g_s(c) - 1 \right) dx(c) = 0.$$

It follows that $\sum_s g_s(c) = 1$, for all c with $dx(c) > 0$. I can now define the measure $\gamma : C \to \Delta(\mathcal{S})$ by

$$\gamma(s|c) = g_s(c),$$

[36] The extension to continuous type spaces requires some care with measurability. Measures must be defined on intervals rather than on singletons.

for all c with $dx(c) > 0$ (and in an arbitrary way on the remaining set of c which have probability zero under x). Because $\sum_s g_s(c) = 1$, γ is a well defined signal function. Moreover, equation (A.3) implies that the equality (2.2) from the definition of cutoff rules holds for all s, and all θ. Thus, (x, π) is a cutoff rule.

A.2.2 Proof of Lemma 3

Part I: Necessity. I first prove that 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 $\theta > \hat{\theta}$ ($\hat{\theta}$ is the largest type smaller than θ). 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 resulting inequalities are added, transfers cancel out, and we obtain

$$\sum_{s \in \mathcal{S}} \left[u(\theta;\, f^s) - u(\hat{\theta};\, f^s)\right] \left[\pi(s|\,\theta)x(\theta) - \pi(s|\,\hat{\theta})x(\hat{\theta})\right] \geq 0. \tag{A.4}$$

Claim 4. *Under the assumptions of either Proposition 3a or Proposition 3b, there exists a 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 $\pi(s|\,\theta)x(\theta) < \pi(s|\,\hat{\theta})x(\hat{\theta})$.*

Proof of Claim 4 under the assumptions of Proposition 3a. I will 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 $\pi(s|\,\theta)x(\theta) \geq \pi(s|\,\hat{\theta})x(\hat{\theta})$; 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.

Recall the assumption that in case of multiplicity, the equilibrium that maximizes the probability of trade in the aftermarket is selected. Define

$$p^{\star}_{(x,\,\pi,\,f)}(s) = \max\left\{\text{argmax}_p \sum_{\theta \leq p} (v(\theta) - p)\pi(s|\,\theta)x(\theta)f(\theta)\right\},$$

as the optimal price quoted by the third party when she makes an offer, given mechanism frame (x, π), distribution f, and conditional on signal realization 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 $f(\hat{\theta})/f(\theta) = (v(\theta) - \theta)/(\theta - \hat{\theta})$. That is, in the absence of additional information, the third party is indifferent between offering price θ and

$\hat{\theta}$. 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 equal to

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

By the definition of β, 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 $v(\theta) > \theta$ for all $\theta \in \Theta$.

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.

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.

Proof of Claim 4 under the assumptions of Proposition *3b*. Consider a class of distributions supported on $\{\hat{\theta}, \theta\}$. By the assumption that the lemons condition is locally severe, $v(\hat{\theta}) < \theta < v(\theta)$. 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 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.

Completion of the proof of Part I. Given a distribution f with properties *(i)* and *(ii)* as in Claim 4, inequality (A.4) becomes

$$\sum_{\{s\in\mathcal{S}:\, \pi(s|\,\theta)x(\theta)<\pi(s|\,\hat{\theta})x(\hat{\theta})\}} \alpha_s \left[\pi(s|\,\theta)x(\theta) - \pi(s|\,\hat{\theta})x(\hat{\theta})\right] \geq 0, \tag{A.5}$$

where $\alpha_s \equiv u(\theta;\, f^s) - u(\hat{\theta};\, f^s)$ is strictly positive for all 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.

Part II: Sufficiency. 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 (2017).[37] 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, 2017, for a formal definition). The equilibrium payoff $u(\theta;\, f^s)$ is non-decreasing in θ for any f^s by a strategy-stealing argument, regardless of the parameters of the aftermarket and equilibrium selection (a higher type values each outcome weakly more than a lower type, so she cannot have lower utility in equilibrium). Because of monotonicity of the payoff in the type, matching efficiency is implied by pairwise stability – total surplus cannot be increased by swapping the allocations of some pair of types. Condition (A.4) 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 (A.4) holds for all $\theta > \hat{\theta}$. The fact that $u(\theta;\, f^s)$ is non-decreasing in θ implies that the first square bracket is non-negative in each term of the sum in (A.4), and condition (M) implies that the second square bracket is non-negative. Thus, inequality (A.4) always holds, regardless of the underlying prior distribution f.

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

B Supplementary materials for Section 3

B.1 Proof of Theorem 1

First, if $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule, then condition (3.2) follows directly from Definition 4 of cutoff rules. I will show that condition (3.2) implies implementability for any prior distribution and any (monotone) aftermarket. For any i and $\boldsymbol{\theta}_{-i}$, we have

$$\int_{\mathcal{S}} \left[u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s})\right] \left[d\pi_i(s|\,\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - d\pi_i(s|\,\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})\right] \geq 0. \tag{B.1}$$

Indeed, using the definition of cutoff mechanisms, for any $\tau \in \Theta_i$,

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

Thus, equation (B.1) is a consequence of

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

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

To show that condition (B.1) implies implementability, I use the condition for checking implementability in arbitrary type and allocation spaces from Dworczak and Zhang (2017), as in the proof of Lemma 3. Although Dworczak and Zhang (2017) 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. Similarly as in the proof of Lemma 3, condition (B.1) implies that the matching between types of agent i and allocations is efficient (see Dworczak and Zhang (2017) for a formal definition). Thus, Theorem 1 follows from the main result of Dworczak and Zhang (2017).[38]

C Supplementary materials for Section 4

C.1 Proof of Lemma 4

To simplify notation (drop the subscripts), I prove Lemma 4 in the symmetric case: Agents are ex-ante identical, and the mechanism frames $(\boldsymbol{x}, \boldsymbol{\pi})$ and $(\bar{\boldsymbol{x}}, \bar{\boldsymbol{\pi}})$ are symmetric. I denote them (x, π) and $(\bar{x}, \bar{\pi})$, respectively (formally, $x_i = x$ and $\pi_i = \pi$ for all i). The proof for

[38] When the type space Θ is continuous, I use the result from the appendix in Dworczak and Zhang (2017) which states that it is enough to check the matching-efficiency condition for all finite subsets of the type space. Thus, the proof goes through without any modifications for a continuous type space.

the asymmetric case is virtually identical, only notationally more complicated.[39]

Proof of the "only if" part: In this part of the proof, I show that a reduced form (x^f, π^f) of a cutoff rule satisfies conditions *1*-*3* of Lemma 4.

Take an arbitrary symmetric cutoff rule (x, π), and a prior distribution f. Condition *3* must hold by the results of Border (2007) and Mierendorff (2011). 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 $x^f(\theta)$ is non-decreasing. This shows condition *1*. Condition *2* would follow from Lemma 2 except that Lemma 2 assumes a finite type and signal spaces. To fill in this gap, I prove a generalization of Lemma 2 below.

Lemma 7. *If $x^f(\theta)$ and $\pi^f(S|\,\theta)x^f(\theta)$ are both non-decreasing in θ for any measurable $S \subseteq \mathcal{S}$, then (x^f, π^f) is a (reduced-form) cutoff rule, i.e., satisfies condition 2.*

Proof of Lemma 7. The proof is analogous to the proof of Lemma 2 but the infinite type and signal spaces require additional care. Because I do not distinguish between two mechanisms if they induce the same posterior distribution over types, it is without loss to take $\mathcal{S} = C \subseteq [0, 1]$, where C is the space of cutoffs (C contains Θ). Denote $\beta_S(\tau) \equiv \pi^f(S|\,\tau)x^f(\tau)$, for any measurable $S \subseteq \mathcal{S}$. Unlike in the proof of Lemma 2, 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. 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 defined by

$$\mu_S((a, b] \cap C) = \beta_S\left(\max_{b' \in [a, b] \cap C} b'\right) - \beta_S\left(\min_{b' \in [a, b] \cap C} b'\right),$$

for any interval $(a, b]$ contained in $[0, 1]$. To simplify notation in the proof, I will extend the measure μ_S to $[0, 1]$ by assuming that it puts no mass on points outside of C. With this, I can write

$$\mu_S((a, b]) = \beta_S(b) - \beta_S(a).$$

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.

[39] As a key step in my proof, I use a result from Gershkov et al. (2013). My proof remains the same in the asymmetric case because Gershkov et al. (2013) did the "heavy lifting" by proving their results in an asymmetric model. Similarly, the Matthews-Border condition (M-B) is more difficult to derive in the asymmetric case but Border (2007) and Mierendorff (2011) cover the asymmetric case.

I will show that the measure μ_S, for any S, is absolutely continuous with respect to the distribution dx^f. For any $a,\, b \in C$, $a < b$, we have

$$\beta_S(b) - \beta_S(a) \leq \beta_{\mathcal{S}}(b) - \beta_{\mathcal{S}}(a) = x^f(b) - x^f(a).$$

It follows that if $x^f(b) = x^f(a)$, then $\beta_S(b) - \beta_S(a) = 0$. Because a and b were arbitrary, μ_S is absolutely continuous with respect to dx^f.

By the Radon-Nikodym Theorem, for any S, there exists a measurable positive function g_S supported on C that is a density of μ_S with respect to dx^f. In particular,

$$\beta_S(\theta) = \pi^f(S|\theta)x^f(\theta) \equiv \mu_S([0,\, \theta]) = \int_0^\theta g_S(c)dx^f(c), \tag{C.1}$$

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

With $S = [0,\, s]$, I define $G_c(s) \equiv g_{[0,\, s]}(c)$, for any $s \in [0,\, 1]$. My goal is to prove that $G_c(s)$ is a cdf when treated as a function of s (for dx^f-almost all c).

Because the equality

$$\pi^f(S|\theta)x^f(\theta) = \int_0^\theta g_S(c)dx^f(c)$$

must hold for all θ, we can conclude that $G_c(0) = 0$, $G_c(1) = 1$, 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 $c \in C$. Consider $s < s'$, and note that

$$\begin{aligned}\int_0^\theta g_{[0,\, s']}(c)dx^f(c) = \pi^f([0, s'])x^f(\theta) = \pi^f([0, s])x^f(\theta) + \pi^f((s, s'])x^f(\theta) \\ = \int_0^\theta g_{[0,\, s]}(c)dx^f(c) + \int_0^\theta g_{(s,\, s']}(c)dx^f(c),\end{aligned}$$

where $g_{(s,\, s']}(c)$ is obtained by taking $S = (s,\, s']$ and applying formula (C.1). 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, $G_c(s)$ is non-decreasing in s. Finally, by monotonicity of $G_c(s)$ and equation (C.1), $G_c(s)$ 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

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

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 $[0,\, s]$, for all $s \in [0,\, 1]$, by equation (C.1) we get

$$\pi^f(S|\,\theta)x^f(\theta) = \int_0^\theta \gamma(S|\,c)dx^f(c),$$

for all measurable $S \subseteq \mathcal{S}$. Therefore, $(x^f,\, \pi^f)$ satisfies condition *2* of Lemma 4 and is thus a reduced-form cutoff rule.

□

Proof of the "if" part: Fix a reduced-form cutoff rule $(\bar{x},\, \bar{\pi})$, i.e., a rule satisfying conditions *1-3* of Lemma 4, under prior distribution f. By the results of Border (2007) and Mierendorff (2011), condition (M-B) (along with the fact that $\bar{x}$ is monotone) implies that there exists a joint (symmetric) allocation function x such that $\bar{x} = x^f$. First, suppose that $(\bar{x},\, \bar{\pi})$ is $\mathcal{S}$-finite, i.e., the signal space $\mathcal{S}$ is finite. 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) 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: $(x^f_\star,\, \pi^f_\star) = (\bar{x},\, \bar{\pi})$.

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 8 (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 (the solution always exists):*

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

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 8 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 (4.2) in Lemma 4). Given the allocation $\{y^{i,s}\}$ produced from $\{x^{i,s}\}$ by Lemma 8, 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 $x_\star(\theta_i, \boldsymbol{\theta}_{-i}) = 0$.

To show that $(x_\star, \pi_\star)$ is a cutoff rule it is enough to invoke Lemma 2 from Section 2 – cutoff rules are characterized by monotonicity in the type for every signal. Formally, an extension of Lemma 2 to multi-agent mechanisms, Lemma 6 found in Section 5, implies that 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.

Finally, $(x_\star^f, \pi_\star^f) = (\bar{x}, \bar{\pi})$ follows from the fact that $\{y^{i,s}\} \in \mathcal{D}$:

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

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 $(\bar{x}, \bar{\pi})$. I say that a sequence of reduced-form mechanism frames $\{(\bar{x}, \bar{\pi}^n)\}_{n=1}^{\infty}$ on the same signal space $\mathcal{S}$ converges to $(\bar{x}, \bar{\pi})$, if $\bar{\pi}^n(\cdot\,|\,\theta)\bar{x}(\theta)$ converges to $\bar{\pi}(\cdot\,|\,\theta)\bar{x}(\theta)$ in the weak* topology of measures on $\mathcal{S}$, for almost all θ. It is straightforward to show (a formal statement and proof of a more general statement is provided in the Online Appendix OA.3.1) that a (reduced-form) mechanism frame is a (reduced-form) cutoff rule if and only if it is a limit of (reduced-form) cutoff rules with the same allocation function and a signal with finite support. In other words, cutoff mechanisms with infinite signal spaces can be approximated by cutoff mechanisms implementing the same allocation but with finitely many signal realizations.

Thus, $(\bar{x}, \bar{\pi})$ can be represented as a limit of $\mathcal{S}$-finite reduced-form cutoff rules $\{(\bar{x}, \bar{\pi}_n)\}_{n=1}^{\infty}$. By the first part of the proof, for each n, there exists a ($\mathcal{S}$-finite) cutoff rule (x, π_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 the above characterization (a limit of cutoff rules is a cutoff rule), 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).

C.2 Proof of Theorem 2

By Lemma 4, optimization over cutoff mechanisms can be performed in the space of reduced-form cutoff mechanisms. For a fixed allocation $\boldsymbol{x}$ and distribution $\boldsymbol{f}$, a reduced-form cutoff mechanism for agent i is formally equivalent to a one-agent cutoff mechanism from Section 2. Moreover, the disclosure problem for any agent i can be solved independently from the disclosure problems for all other agents $j \neq i$ because there is only one winner (and hence one participant in the aftermarket) ex-post. Thus, we can use the proof of Lemma 1 to establish the following result (the fully analogous proof is omitted).[40]

Lemma 9. *For every non-decreasing allocation rule $\boldsymbol{x}$, the problem of maximizing (4.1) over $\boldsymbol{\pi}$ subject to $(\boldsymbol{x}, \boldsymbol{\pi})$ being a cutoff rule is equivalent to solving, for every $i \in N$,*

$$\max_{\varrho_i \in \Delta(\Delta(C_i))} \mathbb{E}_{G_i \sim \varrho_i} \mathcal{V}_i(G_i) \tag{C.3}$$

subject to

$$\mathbb{E}_{G_i \sim \varrho_i} G_i(\theta_i) = x_i^{\boldsymbol{f}}(\theta_i),\ \forall \theta_i \in \Theta_i. \tag{C.4}$$

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

Corollary 3. *The optimal expected payoff to the mechanism designer in the problem (C.3)-(C.4) is equal to*

$$\sum_{i \in N} co\mathcal{V}_i(x_i^{\boldsymbol{f}}) \equiv \sum_{i \in N} \sup\{y :\ (x_i^{\boldsymbol{f}}, y) \in CH(graph(\mathcal{V}_i))\},$$

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

[40] 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.

The first part of Theorem 2 follows directly from Lemma 4, Lemma 9, and Corollary 3.

To establish the second part of Theorem 2, note that in the case $N = 1$, constraint (M-B) is vacuously satisfied. Thus, I can apply the proof of Proposition 2 from Subsection 2.3.2 – there always exists an optimal cutoff mechanism that releases no information.

C.3 Proof of Lemma 5 and supplementary material for Section 4.3

In this appendix, I formalize the result stated in Section 4.3 about feasible distributions of posterior beliefs over the winner's type induced by cutoff mechanisms. The characterization of feasible distributions is stated in Lemma 10, and I prove Lemma 5 in Appendix C.3.1.

Because I have assumed symmetry (agents are ex-ante identical), it is without loss to consider symmetric mechanisms. An allocation rule is characterized by its one-dimensional interim allocation rule $\bar{x} : \Theta \to [0, 1]$.[41] For a fixed (interim) allocation rule $\bar{x}$, I call $f^{\bar{x}}(\theta)$, defined by (2.4), the *no-communication posterior*. The no-communication posterior is the belief over the type of the winner held by the third party when the interim allocation rule is $\bar{x}$, and the mechanism reveals no information (other than the identity of the winner). I will say that the distribution f_1 *monotone-likelihood-ratio dominates* distribution f_2 (denoted $f_1 \succ^{MLR} f_2$) if $f_1(\theta)/f_2(\theta)$ is bounded and non-decreasing.

Lemma 10. *A distribution of beliefs $\rho \in \Delta(\Delta(\Theta))$ is a conditional distribution of posterior beliefs over the winner's type induced by a cutoff mechanism with interim allocation rule $\bar{x}$ if and only if*

$$\bar{f} \succ^{MLR} f, \ \forall \bar{f} \in supp(\rho). \tag{C.5}$$

and

$$\mathbb{E}_{\bar{f}\sim\rho}\bar{f}(\theta) \equiv \int\limits_{supp(\rho)} \bar{f}(\theta)d\rho(\bar{f}) = f^{\bar{x}}(\theta), \tag{C.6}$$

Condition (C.6) 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 conditional on an agent being the winner ($\bar{x}$ is the ex-ante probability that any agent becomes the winner). Condition (C.5) is an additional constraint on posterior belief – each posterior has to likelihood-ratio dominate the prior.

Proof of Lemma 10. I first show that every ex-ante (unconditional) distribution $\varrho \in \Delta(\Delta(C))$ of beliefs over the cutoff for some agent i that is feasible under allocation $\bar{x}$ defines a posterior (conditional) distribution $\rho \in \Delta(\Delta(\Theta))$ of beliefs over i's type conditional on i being the winner which satisfies conditions (C.5)-(C.6).

[41] The symmetry assumption simplifies exposition but all results in this section follow for the asymmetric case by simply adding a subscript i to $\bar{x}$.

Because ϱ is a feasible distribution of beliefs over the cutoff, it satisfies the Bayes-plausibility condition (see equations (2.7) and (C.4)) which states that

$$\mathbb{E}_{G\sim\varrho}G(\theta) = \bar{x}(\theta),\ \forall\theta\in\Theta. \tag{C.7}$$

For every $G \in \text{supp}(\varrho)$, let f^G, defined by (4.4), be the corresponding posterior belief over the type of the winner. Each f^G satisfies condition (C.5) because G is a non-decreasing function. Given the ex-ante distribution ϱ for agent i, define the posterior distribution $\bar{\varrho}$ conditional on i being the winner:

$$\bar{\varrho}(\mathcal{G}) = \frac{\int_{\mathcal{G}}\left(\int_\Theta G(\theta)f(\theta)d\theta\right)d\varrho(G)}{\int_{\text{supp}(\varrho)}\left(\int_\Theta G(\theta)f(\theta)d\theta\right)d\varrho(G)},\ \text{ for any measurable } \mathcal{G}\subseteq\Delta(\Delta(C)). \tag{C.8}$$

Conditional on i becoming the winner, there is higher probability that the cutoff for i was drawn from a distribution that puts relatively more mass on low cutoff realizations. That is why the posterior distribution $\bar{\varrho}$ puts more weight on distributions G that allocate the good with higher probability. Define the corresponding posterior distribution $\rho\in\Delta(\Delta(\Theta))$ of beliefs over the type of the winner by

$$\rho(\mathcal{F}) = \bar{\varrho}\left(\{G\in\Delta(\Delta(C)) : f^G\in\mathcal{F}\}\right),\ \text{ for any measurable } \mathcal{F}\subseteq\Delta(\Delta(\Theta)).$$

To show that condition (C.6) holds, note that because ϱ is a feasible ex-ante distribution, it satisfies condition (C.7), and hence

$$\int_{\text{supp}(\varrho)}\left(\int_\Theta \tilde{G}(\theta)f(\theta)d\theta\right)d\varrho(\tilde{G}) = \int_\Theta\left(\int_{\text{supp}(\varrho)}\tilde{G}(\theta)d\varrho(\tilde{G})\right)f(\theta)d\theta = \int_\Theta \bar{x}(\theta)f(\theta)d\theta.$$

Then, we have

$$\begin{aligned}
\int_{\text{supp}(\rho)}\bar{f}(\theta)d\rho(\bar{f}) &= \int_{\text{supp}(\varrho)} f^G(\theta)d\bar{\varrho}(G)\\
= \int_{\text{supp}(\varrho)}\frac{G(\theta)f(\theta)}{\int_\Theta G(\tau)f(\tau)d\tau}&\frac{\int_\Theta G(\tau)f(\tau)d\tau}{\int_{\text{supp}(\varrho)}\left(\int_\Theta \tilde{G}(\tau)f(\tau)d\tau\right)d\varrho(\tilde{G})}d\varrho(G)\\
= \frac{\left(\int_{\text{supp}(\varrho)}G(\theta)d\varrho(G)\right)f(\theta)}{\int_\Theta\bar{x}(\tau)f(\tau)d\tau} &= \frac{\bar{x}(\theta)f(\theta)}{\int_\Theta\bar{x}(\tau)f(\tau)d\tau} = f^{\bar{x}}(\theta),
\end{aligned}$$

which is condition (C.6).

To show the opposite direction, start with a posterior distribution of beliefs over the winner's type $\rho\in\Delta(\Delta(\Theta))$, satisfying conditions (C.5) and (C.6) for a non-decreasing allocation rule $\bar{x}$. I will define a feasible ex-ante (unconditional) distribution of beliefs over the cutoff $\varrho\in\Delta(\Delta(C))$ such that ϱ induces ρ.

First, for each $\bar{f} \in \text{supp}(\rho)$, define

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

where $\bar{\theta} = \max\{\Theta\}$. Because $\bar{f}$ monotone-likelihood-ratio dominates f, the function $G^{\bar{f}}(\theta)$ is non-decreasing and bounded above by 1 on Θ.[42] Thus, it defines a non-decreasing allocation rule, and hence also a distribution of the cutoff (after extending it to C). Define a distribution $\varrho \in \Delta(\Delta(C))$ supported on $\{G^{\bar{f}} : \bar{f} \in \text{supp}(\rho)\}$ and defined by

$$\varrho(\{G^{\bar{f}} : \bar{f} \in \mathcal{F}\}) = \frac{\int_{\mathcal{F}} \bar{f}(\bar{\theta})d\rho(\bar{f})}{\int_{\text{supp}(\rho)} \bar{f}(\bar{\theta})d\rho(\bar{f})},\ \text{for any measurable } \mathcal{F} \subseteq \Delta(\Delta(\Theta)). \tag{C.9}$$

With this, I have to verify that ϱ is feasible, i.e., it satisfies (C.7), and induces ρ. We have

$$\begin{aligned}\int_{\text{supp}(\varrho)} G^{\bar{f}}(\theta)d\varrho(G^{\bar{f}}) &= \int_{\text{supp}(\rho)} \left(\bar{x}(\bar{\theta})\frac{f(\bar{\theta})}{\bar{f}(\bar{\theta})}\right)\frac{\bar{f}(\theta)}{f(\theta)}\frac{\bar{f}(\bar{\theta})}{\int_{\text{supp}(\rho)}\tilde{f}(\bar{\theta})d\rho(\tilde{f})}d\rho(\bar{f})\\ &= \int_{\text{supp}(\rho)}\frac{\bar{f}(\theta)}{f(\theta)}\left(\int_\Theta \bar{x}(\tau)f(\tau)d\tau\right)d\rho(\bar{f}) = f^{\bar{x}}(\theta)\frac{1}{f(\theta)}\left(\int_\Theta \bar{x}(\tau)f(\tau)d\tau\right) = \bar{x}(\theta),\end{aligned} \tag{C.10}$$

where I have used condition (C.6) twice.

To show that ϱ induces ρ, note that $f^{G^{\bar{f}}} = \bar{f}$. Moreover, using (C.8) and (C.9), the posterior distribution (conditional on the agent being the winner) over $G^{\bar{f}}$ is given by, for any measurable $\mathcal{F} \in \Delta(\Delta(\Theta))$,

$$\bar{\varrho}(\{G^{\bar{f}} : \bar{f} \in \mathcal{F}\}) = \frac{\int_{\mathcal{F}}\left(\int_\Theta G^{\bar{f}}(\theta)f(\theta)d\theta\right)\bar{f}(\bar{\theta})d\rho(\bar{f})}{\int_{\text{supp}(\rho)}\left(\int_\Theta G^{\bar{f}}(\theta)f(\theta)d\theta\right)\bar{f}(\bar{\theta})d\rho(\bar{f})} = \int_{\mathcal{F}} d\rho(\bar{f}) = \rho(\mathcal{F})$$

which shows that ϱ induces the posterior distribution ρ over the winner's type. □

C.3.1 Proof of Lemma 5

The proof follows from Lemma 10 and Lemma 9 found in Appendix C.2. Starting from the (symmetric version of) objective function (C.3) and a feasible ex-ante distribution ϱ of

[42] By condition (C.5), $\bar{f}(\theta)/f(\theta)$ is bounded and non-decreasing, so in the case when $f(\bar{\theta}) = 0$, the expression $\bar{f}(\bar{\theta})/f(\bar{\theta})$ should be understood as the limit of $\bar{f}(\theta)/f(\theta)$ as $\theta \to \bar{\theta}$.

beliefs over the cutoff,

$$\begin{aligned}
\mathbb{E}_{G\sim\varrho}\mathcal{V}(G) &\overset{(1)}{=} \int_{\text{supp}(\rho)} \mathcal{V}(G^{\bar{f}}) \frac{\bar{f}(\bar{\theta})}{\int_{\text{supp}(\rho)} \tilde{f}(\bar{\theta})d\rho(\tilde{f})} d\rho(\bar{f}) \\
&\overset{(2)}{=} \int_{\text{supp}(\rho)} \left(\int_{\Theta} V(\theta;\, f^{G^{\bar{f}}}) G^{\bar{f}}(\theta) f(\theta) d\theta \right) \frac{\bar{f}(\bar{\theta})}{f^{\bar{x}}(\bar{\theta})} d\rho(\bar{f}) \\
&\overset{(3)}{=} \left(\int_{\theta\in\Theta} \bar{x}(\theta) f(\theta) d\theta \right) \int_{\text{supp}(\rho)} \underbrace{\int_{\Theta} V(\theta;\, \bar{f}) \bar{f}(\theta) d\theta}_{\mathcal{W}(\bar{f})} d\rho(\bar{f}),
\end{aligned}$$

where (1) follows from (C.9), (2) uses definitions (4.4) and (4.5), and (3) uses the definition of $\mathcal{W}$ and $f^{\bar{x}}$. This proves that the objective function can be written as

$$\left(\int_{\theta\in\Theta} \bar{x}(\theta) f(\theta) d\theta \right) \mathbb{E}_{\bar{f}\sim\rho} \mathcal{W}(\bar{f}),$$

where feasible ρ satisfy conditions (C.5) and (C.6), by Lemma 10. Given this representation of the objective function and Lemma 10, the concave closure characterization follows immediately.

C.4 Proof of Corollary 2

The only part of the corollary which does not follow immediately from Lemma 5 is that $\text{co}^{M_f}\, \mathcal{W}(f^{\bar{x}})$ corresponds to the payoff from full disclosure when $\mathcal{W}$ is convex. The conclusion follows because the concave closure of a convex function $\mathcal{W}$ on a compact space M_f is given by decomposing the 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. Truncations of the prior are induced by disclosing the cutoff.

C.5 Proof of Proposition 4

Case 1. Consider the 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, by Corollary 2. Using Lemma 5 and the characterization of feasible interim allocation rules form Section 4.1, we can write the problem as

$$\max_{\bar{x}} \left(\int_0^1 \bar{x}(\theta) f(\theta) d\theta \right) W(M(f^{\bar{x}})) \tag{C.11}$$

subject to

$$\bar{x}(\theta) \text{ is non-decreasing in } \theta, \tag{C.12}$$

$$\int_\tau^1 \bar{x}(\theta) f(\theta) d\theta \leq \frac{1 - F^N(\tau)}{N}, \; \forall \tau \in [0,\, 1]. \tag{C.13}$$

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

$$\max_{\bar{x}} \int_0^1 \theta \bar{x}(\theta) f(\theta) d\theta, \tag{C.14}$$

subject to (C.12), (C.13), and

$$\int_0^1 \bar{x}(\theta) f(\theta) d\theta = \beta \tag{C.15}$$

In the above problem, we can think of constraint (C.15) as an equal mass condition. Intuitively, it is optimal to shift as much mass as possible to the right, subject to constraint (C.13), which will thus hold with equality for large enough 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 (C.15) 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 (C.12) 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 (C.12), so it is a solution to problem (C.14).

In the second step, I optimize over $\beta \in [0,\, 1/N]$ in condition (C.15), 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 (C.11), we obtain

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

This corresponds to equation (4.10) in Proposition 4, and thus the first case is proven.

Case 2. Consider the case when W is concave and decreasing. Following the same steps as previously, I consider the auxiliary problem with constraint (C.15). Because W is decreasing, the objective is

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

subject to (C.12), (C.13), and (C.15). This time, all the mass under $\bar{x}$ should be shifted to the left, subject to the monotonicity constraint (C.12). Thus, the optimal $\bar{x}$ will be constant, equal to β. Because $\beta \leq 1/N$, such $\bar{x}$ satisfies the Matthews-Border condition (C.13), and corresponds to allocating the object uniformly at random.

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

Case 3. Finally, 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 $\bar{x}$, by Corollary 2. 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,

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

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. The distribution $\bar{x}$ is the ex-ante distribution over the cutoff for agent i. Conditional on agent i being the winner, the posterior distribution over the cutoff for agent i must be adjusted (lower cutoffs are more likely). The ex-ante probability of cutoff c is transformed into a conditional probability by conditioning on the event $\theta \geq c$. The objective function (C.11) 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

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

The conclusion of Proposition 4 now follows from an argument analogous to the one used in previous cases. 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 (C.13). 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 $\underline{r}$. Notice that full disclosure of the cutoff leads to the same posterior beliefs over the winner's type as full disclosure of the second highest type. This is because, when the second highest type is below $\underline{r}$, the exact value of the second highest type does not influence the allocation for the highest type.

To finish the proof of Proposition 4, 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 $m'(c) = (m(c)-c)f(c)/(1-F(c))$. The inequality $J_w(c) \geq 0$ implies that $m(c) - c \leq W(m(c))/W'(m(c))$. 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.

D Supplementary materials for Section 5

D.1 Proof of Lemma 6

The fact that a cutoff rule satisfies property (3.2) is trivial – it follows directly from Definition 4 of cutoff rules. Because I have assumed that Θ is finite, and that the mechanism frame is $\mathcal{S}-$finite, the proof of the converse part is almost identical to the proof of Lemma 2 found in Appendix A.2.1. The only difference is that one must fix $i \in N$ and $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$, and apply the same reasoning with $\beta_s(\theta) \equiv \pi_i(s|\,\theta, \boldsymbol{\theta}_{-i})x_i(\theta, \boldsymbol{\theta}_{-i})$ to obtain a signal function $\gamma_i(s|\,c, \boldsymbol{\theta}_{-i})$ indexed by i and $\boldsymbol{\theta}_{-i}$. This signal function satisfies the equality (3.1) from the definition of cutoff rules.

D.2 Proof of Theorem 3

It is enough to prove the following lemma, analogous to Lemma 3 in Section 2.

Lemma 11. *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* ($\boldsymbol{M}$).

Indeed, once Lemma 11 is proven, the assumptions of Theorem 3 imply that condition ($\boldsymbol{M}$) holds. Then, the conclusion of Theorem 3 follows from Lemma 6.

Proof of Lemma 11. The proof is similar to the proof of the corresponding part of Lemma 3, so I omit details that are analogous.

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)

$$\sum_{s\in\mathcal{S}} \Big[u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s})\Big] \Big[\pi_i(s|\, \theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - \pi_i(s|\, \hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})\Big] \geq 0. \tag{D.1}$$

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 (5.1) and (5.2) hold. Under these $\boldsymbol{f}$ and A, inequality (D.1) becomes

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

with $\alpha_s(\theta_i) > \alpha_s(\hat{\theta}_i)$ for each signal s in the summation, by condition (5.1). 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, Lemma 11 is proven.

□

D.3 Proof of Theorem 4

I first show that cutoff rules are robust to mistrust. Fix a cutoff rule $(\boldsymbol{x}, \boldsymbol{\pi})$ and suppose that agent i holds a belief μ_i. Define the transfer function for agent i by

$$t_i^s(\theta_i,\, \boldsymbol{\theta}_{-i}) = u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - \frac{\int_0^{\theta_i} u_i'(\tau;\, \boldsymbol{f}^{i,s}) \sum_{s'\in S_i} \mu_i(s|s', \boldsymbol{\theta}_{-i})\pi_i(s'|\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i})d\tau}{\sum_{s'\in S_i} \mu_i(s|s', \boldsymbol{\theta}_{-i})\pi_i(s'|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})},$$

where the transfer can be arbitrarily defined when the denominator is equal to zero. Under this transfer function, the utility of agent i with type θ_i when reporting $\hat{\theta}$ is

$$
\begin{aligned}
&\sum_{s\in S_i}\sum_{s'\in S_i}[u_i(\theta_i;\, \boldsymbol{f}^{i,s'}) - t_i^{s'}(\hat{\theta},\, \boldsymbol{\theta}_{-i})]\mu_i(s'|s, \boldsymbol{\theta}_{-i})\pi_i(s|\hat{\theta},\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta},\, \boldsymbol{\theta}_{-i}) \\
&\qquad = \sum_{s'\in S_i}[u_i(\theta_i;\, \boldsymbol{f}^{i,s'}) - u_i(\hat{\theta};\, \boldsymbol{f}^{i,s'})]\sum_{s\in S_i}\mu_i(s'|s, \boldsymbol{\theta}_{-i})\pi_i(s|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) \\
&\qquad\qquad + \sum_{s'\in S_i}\int_0^{\hat{\theta}} u_i'(\tau;\, \boldsymbol{f}^{i,s'})\sum_{s\in S_i}\mu_i(s'|s, \boldsymbol{\theta}_{-i})\pi_i(s|\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i})d\tau \\
&\simeq \sum_{s'\in S_i}\int_{\hat{\theta}}^{\theta_i} u_i'(\tau;\, \boldsymbol{f}^{i,s'})\sum_{s\in S_i}\mu_i(s'|s, \boldsymbol{\theta}_{-i})\left[\pi_i(s|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - \pi_i(s|\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i})\right]d\tau,
\end{aligned}
\tag{D.2}
$$

where $\simeq$ means that the two expressions are equal up to a term that does not depend on $\hat{\theta}$. By the monotonicity property ($\boldsymbol{M}$) of cutoff rules, the above expression is maximized at $\hat{\theta} = \theta_i$ which shows that a cutoff rule is weakly robust to mistrust with respect to Λ_{all}.

Strong robustness to mistrust with respect to Λ_{alloc} follows because, for $\mu \in \Lambda_{\text{alloc}}$,

$$
t_i^s(\theta_i,\, \boldsymbol{\theta}_{-i}) = u_i(\theta_i;\, \boldsymbol{f}^{i,s}) - \frac{\int_0^{\theta_i} u_i'(\tau;\, \boldsymbol{f}^{i,s})\pi_i(s'|\tau,\, \boldsymbol{\theta}_{-i})x_i(\tau,\, \boldsymbol{\theta}_{-i})d\tau}{\pi_i(s'|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})}.
$$

That is, on Λ_{alloc}, the transfer rule implementing the cutoff rule does not depend on the particular form of μ, and thus we obtain strong robustness.

To show the converse part of Theorem 4, it is enough to prove that if $(\boldsymbol{x},\, \boldsymbol{\pi})$ is not a cutoff rule, then it is not weakly robust to mistrust wrt Λ_{alloc} (because strong robustness implies weak robustness, and $\Lambda_{\text{alloc}} \subset \Lambda_{\text{all}}$, this will cover both cases). If $(\boldsymbol{x},\, \boldsymbol{\pi})$ is not a cutoff rule, there exists $i \in N$, $\boldsymbol{\theta}_{-i}$, and $s_0 \in S_i$ such that $\pi_i(s_0|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is not monotone in θ_i. Take $\mu_i(s'|s, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{s=s'=s_0\}}$. That is, agent i believes that the seller will allocate the good only when the signal is s_0. Clearly, such a belief is consistent with the class Λ_{alloc}. For this belief, incentive compatibility requires that

$$
\left[u_i(\theta_i;\, \boldsymbol{f}^{i,s_0}) - u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s_0})\right]\left[\pi_i(s_0|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i}) - \pi_i(s_0|\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i,\, \boldsymbol{\theta}_{-i})\right] \geq 0
$$

for all $\theta_i > \hat{\theta}_i$. Because $u_i(\theta_i;\, \boldsymbol{f}^{i,s_0}) > u_i(\hat{\theta}_i;\, \boldsymbol{f}^{i,s_0})$ by assumption, and $\pi_i(s_0|\theta_i,\, \boldsymbol{\theta}_{-i})x_i(\theta_i,\, \boldsymbol{\theta}_{-i})$ is not monotone in θ_i, this condition fails for some pair $(\theta_i,\, \hat{\theta}_i)$.

D.4 Supplementary materials for Section 5.3

I assume symmetry for ease of exposition. A Generalized Clock Auction (GCA) is characterized by a (possibly random) 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 $t \in \mathcal{T}$: (1) A price p^t is announced to bidders; (2) Bidders simultaneously (and covertly) decide to stay in the auction, or to exit; (3) The auctioneer observes bidders' decisions and implements the relevant outcome (to be specified); (4) 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 will matter 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.[43] 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. In a GCA, the informational content of signals is determined by the equilibrium behavior of bidders.

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.[44] 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

[43] Public history is defined as the largest information set contained in information sets of all bidders.

[44] 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.

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 $A \equiv \{u(\theta; \bar{f}) : \theta \in \Theta, \bar{f} \in \Delta(\Theta)\}$. I assume that the aftermarket is monotone. 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 $u(\theta; f^s_\sigma)$.

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.

Theorem 6 (Formal statement)**.** *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 of Theorem 6 can be found in Appendix D.5. 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. The main difficulty in the proof is to show that decomposability of x implies existence of a signal 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 order 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.

D.5 Proof of Theorem 6

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_0^τ denote the history in which the final winner becomes inactive in the last round τ, and let H_1^τ denote the history in which the final winner remains active in the last round τ. For the tuple $(y,\,p,\,a)$, the corresponding allocation and revelation rules are given by

$$\pi_i^{(y,\,p,\,a)}(s|\,\theta_i,\,\boldsymbol{\theta}_{-i})x_i^{(y,\,p,\,a)}(\theta_i,\,\boldsymbol{\theta}_{-i})$$
$$= \begin{cases} \left(\frac{1}{N^{\tau-1}}\right) Y^\tau(H_0^\tau)(s) + \left(\frac{N^{\tau-1}-1}{N^{\tau-1}}\right) Y^\tau(H_1^\tau)(s) & \text{if } \Gamma_i^{(y,\,p,\,a)}(\boldsymbol{\theta}) > \tau \equiv \max_{j\neq i}\Gamma_j^{(y,\,p,\,a)} \\ \left(\frac{1}{N^{\tau-1}}\right) Y^\tau(H_0^\tau)(s) & \text{if } \Gamma_i^{(y,\,p,\,a)}(\boldsymbol{\theta}) = \tau \equiv \max_{j\neq i}\Gamma_j^{(y,\,p,\,a)} \\ 0 & \text{otherwise,} \end{cases} \tag{D.3}$$

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.3). 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 6 holds. By Lemma 6, $(\boldsymbol{x},\,\boldsymbol{\pi})$ is a cutoff rule, and thus the first part of Theorem 6 is thus proven.

Discussion: The restriction to monotone equilibria in the first part of Theorem 6 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 6 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 $\theta^{\tau+1}$.[45] 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 Section 4.1 for definition). 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

$$n^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i}) = |\{j \in N \setminus \{i\} :\, \theta_j \geq \kappa^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i})\}|.$$

[45] 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}.

In words, $\kappa^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i})$ is the highest of the thresholds $\kappa_1, \ldots, \kappa_k$ that at least one type in $\boldsymbol{\theta}_{-i}$ exceeds, and $n^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i})$ is the number of types in $\boldsymbol{\theta}_{-i}$ that exceed $\kappa^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i})$. The vector $\nu^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i}) \equiv (\kappa^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i}), n^{\kappa_1 \ldots \kappa_k}(\boldsymbol{\theta}_{-i}))$ is a sufficient statistic for $\boldsymbol{\theta}_{-i}$ needed to implement $x_i^{\kappa_1 \ldots \kappa_k}(\theta_i, \boldsymbol{\theta}_{-i})$.

Define a symmetric cutoff rule (x', π') by

$$\pi'(s| \theta_i, \boldsymbol{\theta}_{-i})x'(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{\alpha} \lambda^{\alpha} \sum_{c \leq \theta_i} \gamma'(s| c, \nu^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i}))dx^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(c, \boldsymbol{\theta}_{-i}),$$

for all s, and $\boldsymbol{\theta}$, where the signal function γ' is defined by,

$$\gamma'(s| c, \nu) = \frac{\sum_{\{\boldsymbol{\theta}_{-i}:\, \nu = \nu^{\kappa_1^{\alpha}, \ldots, \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i})\}} \gamma(s| c, \boldsymbol{\theta}_{-i}) \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})}{\sum_{\{\boldsymbol{\theta}_{-i}:\, \nu = \nu^{\kappa_1^{\alpha}, \ldots, \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i})\}} \boldsymbol{f}_{-i}(\boldsymbol{\theta}_{-i})},$$

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

$$\pi^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(s| \theta_i, \boldsymbol{\theta}_{-i})x^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{c \leq \theta_i} \gamma'(s| c, \nu^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(\boldsymbol{\theta}_{-i}))dx^{\kappa_1^{\alpha} \ldots \kappa_{k(\alpha)}^{\alpha}}(c, \boldsymbol{\theta}_{-i}),$$

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 6 can then be obtained by randomizing over α according to the distribution $\{\lambda^{\alpha}\}$ in round 0 of the GCA.[46] In the remainder of the proof, I fix α and omit it from the notation – I will denote the hierarchy to be implemented by $\kappa_1, \ldots, \kappa_k$. The description of the auction and the equilibrium is kept informal to avoid additional notation. Intuitively, in the equilibrium I will 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.[47] 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 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

[46] 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.

[47] It is enough to set the price to a prohibitively high level in the subsequent round.

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 Supplementary materials for Section 6

E.1 Proof of Claim 1

The problem (6.3) of the regulator can be equivalently stated as

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

subject to

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

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 I apply optimal control techniques to 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. Because the solution is pinned down for $\theta \geq p$, I can ignore the term $\int_p^1 (\theta - k)x(\theta)f(\theta)d\theta$ in the optimization.

To deal with the constraint (E.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 $\theta \in [0, 1]$,

$$x^\star(\theta) \in \operatorname*{argmax}_{x \in [0,\, 1]} H(x,\, \theta,\, q) \equiv \operatorname*{argmax}_{x \in [0,\, 1]} [\lambda v(\theta) + (1-\lambda)\theta - k + q(\theta)(v(\theta) - p)]x(\theta)f(\theta),$$

$$q'(\theta) = 0,\; q(p) \geq 0\, (= 0 \text{ if } \Gamma(p) > 0)$$

$$H(x,\, \theta,\, q(\theta)) \text{ is concave in } x.$$

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 (E.2) holds with $r = \underline{r}$ defined as solution to equation

$$\lambda v(\underline{r}) + (1-\lambda)\underline{r} = k, \tag{E.3}$$

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 (E.2) fails with $r = \underline{r}$. Because $v(1) > 1$, there must exist an $r^\star > \underline{r}$ such that (E.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 (E.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 (E.1).

Abusing notation slightly, let $p(r) = p(\mathbf{1}_{\{\theta \geq r\}})$, where $p(x)$ for $x \in \mathcal{X}$ is defined in (6.2).

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.

E.2 Derivation of the objective function for Subsection 6.1.2

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

$$t(\theta) = \lambda \left(\int_{\mathcal{S}} \max(\theta,\, p(f^s)) d\pi(s|\, \theta) x(\theta) - \int_0^\theta \left(\int_{\mathcal{S}} x(\tau) \mathbf{1}_{\{\tau \geq p(f^s)\}} d\pi(s|\, \tau) \right) d\tau \right) + (1-\lambda) \left(\theta x(\theta) - \int_0^\theta x(\tau) d\tau \right).$$

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 (4.1) for $N = 1$ with $V(\theta;\, f^s) = \lambda \left(p(f^s) \mathbf{1}_{\{\theta \leq p(f^s)\}} + \mathbf{1}_{\{\theta > p(f^s)\}} J(\theta) \right) + (1-\lambda) J(\theta) - k$. Theorem 2 implies that the optimal mechanism reveals no information. Thus, the problem is to maximize

$$\int_0^{p(x)} \left(\lambda p(x) + (1-\lambda) J(\theta) - k \right) x(\theta) f(\theta) d\theta + \int_{p(x)}^1 (J(\theta) - k) x(\theta) f(\theta) d\theta.$$

E.3 Proof of Claim 2

The problem of the seller can be equivalently stated as

$$\max_{x\in\mathcal{X},\, p\in[0,\,1]} \int_0^p \left[\lambda p + (1-\lambda)J(\theta) - k\right] x(\theta)f(\theta)d\theta + \int_p^1 (J(\theta)-k)x(\theta)f(\theta)d\theta \tag{E.4}$$

subject to

$$\int_0^p (v(\theta)-p)x(\theta)f(\theta)d\theta \geq 0. \tag{E.5}$$

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: $x^\star(\theta) = \mathbf{1}_{\{\theta\geq r\}}$ 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

$$\int_{\bar{r}}^1 (v(\theta)-1)f(\theta)d\theta = 0. \tag{E.6}$$

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 (E.3).

Suppose that we are in the regular case $\underline{r} < \bar{r}$. For $r \in [\underline{r},\,\bar{r}]$, we have

$$\int_r^{p(r)} (v(\theta)-p(r))f(\theta)d\theta = 0. \tag{E.7}$$

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 (E.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: $V'_{\text{eff}}(r^\star_{\text{eff}}) = 0$.[48] 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.

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 (E.3) and assumption $v(1) > 1$, we have $\underline{r} < 1$, so $\lambda^\star < 1$ is well defined.

E.4 Proof of Claim 3

First, I solve for the equilibrium in the aftermarket for any posterior expectation $m = \mathbb{E}[\theta_i]$ of the type of dealer i reselling the asset. If $m = 0$, the final buyer offers to buy the good for 0 (the reserve value of the dealer). If $m = 1$ (the alternative buyer is believed to be

[48] If V_{eff} or V_{rev} is not differentiable at some r, we can use subdifferentials instead of derivatives.

present with probability one), there is Bertrand competition between final buyers, so each final buyer offers her value v. When $m \in (0, 1)$, applying standard arguments from auction theory, it can be shown that the unique equilibrium bid distribution for the default buyer is an atom at 0 with probability $1 - m$ and a continuous distribution

$$H(b;\, m) = \frac{1-m}{m}\frac{b}{v-b}$$

supported on $[0,\, mv]$, with probability m. The alternative buyer bids b or less with probability $H(b;\, m)$.

Thus, conditional on acquiring the object and aftermarket belief m, a dealer with type θ has a resale profit equal to the expectation of the first-order statistic of two draws from $H(b;\, m)$ with probability $m\theta$, and the expectation of a draw from $H(b;\, m)$ with probability $(1 - m)\theta + m(1 - \theta)$:

$$U(\theta;\, m) = \theta m \int_0^{mv} b dH^2(b;\, m) + [(1-m)\theta + m(1-\theta)] \int_0^{mv} b dH(b;\, m).$$

Maximizing dealer surplus If $\bar{f}$ denotes the posterior belief over the winning dealer's type, and $M(\bar{f})$ denotes the mean of $\bar{f}$, we have (using the notation from Section 4.3)

$$\begin{aligned} \mathcal{W}_d(\bar{f}) &= \int_0^1 U(\theta;\, M(\bar{f}))\bar{f}(\theta)d\theta \\ &= [M(\bar{f})]^2 \int_0^{M(\bar{f})v} b dH^2(b;\, M(\bar{f})) + 2(1-M(\bar{f}))M(\bar{f}) \int_0^{M(\bar{f})v} b dH(b;\, M(\bar{f})) =: W_d(M(\bar{f})). \end{aligned} \tag{E.8}$$

Thus, the objective function depends only on the expected type of the winner $M(\bar{f})$ through the function W_d. To apply Proposition 4, I first show that W_d is convex. Indeed, by direct calculation, $W_d(m) = vm^2$.

Next, I prove that $J_d(c) = w_d(c) - w_d'(c)(1 - F(c))/f(c)$ is increasing, where $w_d(c) = W_d(m(c))$ (see Proposition 4). Because $m(c) = (c + \beta)/(1 + \beta)$ for the class of distributions assumed in Section 6.2, this is equivalent to $W_d(m) - \beta W_d'(m)(1 - m)$ being increasing on $[\beta/(1 + \beta),\, 1]$ which holds by direct calculation.

Applying Part *3* of Proposition 4 yields the first part of Claim 3. It is a routine exercise to check that the optimal allocation and disclosure rule can be implemented indirectly through the second-price auction described in the statement.

However, I still have to prove that the function $J_d(c)$ is strictly negative for small enough c to prove that the reserve price $r_d^\star$ (defined by $J_d(r_d^\star) = 0$) is strictly positive. It is enough (because of continuity) to show that $J_d(c)$ is strictly negative at $0:$

$$J_d(0) = \left(\frac{\beta}{1+\beta}\right)^2 - 2\left(\frac{\beta}{1+\beta}\right)^2 < 0.$$

Maximizing platform profits From the previous case, the value for winning the object for a dealer with type θ when the posterior mean of the dealer's type is m equals $U(\theta;\, m)$. As usually, we can represent the problem of maximizing seller's profit as the problem of maximizing the expectation of virtual surplus. In a way analogous to the derivation in Subsection E.2, I obtain that a profit-maximizing platform has an objective function

$$V(\theta;\, \bar{f}) = U(\theta;\, M(\bar{f})) - \beta(1-\theta)U'_\theta(\theta;\, M(\bar{f})),$$

where I used the special assumption about the prior F implying that $(1 - F(\theta))/f(\theta)$ is affine, equal to $\beta(1-\theta)$. Next, we set

$$\mathcal{W}_p(\bar{f}) = \int_0^1 V(\theta;\, \bar{f})\bar{f}(\theta)d\theta.$$

Because $V(\theta;\, M(\bar{f}))$ is affine in θ, we get that $\mathcal{W}_p(\bar{f}) = W_p(M(\bar{f}))$, where W_p is defined as

$$W_p(m) = W_d(m) - \beta\underbrace{(1-m)\left[\int_0^{mv} bdH^2(b;\, m) - \int_r^{mv} bdH(b;\, m)\right]}_{I_d(m)}. \tag{E.9}$$

Thus, the objective function W_p can be decomposed into dealer surplus W_d and "information rents" I_d with β determining the weight that the designer puts on information rent minimization. By direct calculation,

$$I_d(m) = -v\frac{1-m}{m}\left(2(1-m)log(1-m) + m(2-m)\right)$$

with

$$I''_d(m) = -2v\frac{2m + 2log(1-m) + m^2 + m^3}{m^3}.$$

The function I''_d is strictly increasing on $[0,\, 1]$ with values ranging from $-\infty$ to ∞. We can define $m^\star(\beta) \in (0,\, 1)$ as the unique solution to the equation $I''_d(m) = 2/\beta$. Then, W_p is convex on $[0,\, m^\star(\beta)]$ and concave on $[m^\star(\beta),\, 1]$. Because the smallest possible posterior mean in the aftermarket is induced by revealing that the cutoff was equal to zero, the distribution of posterior means over the winner's type is supported on a subset of $[\beta/(1+\beta),\, 1]$. Therefore, for a fixed β, the relevant domain of W_p is $[\beta/(1+\beta),\, 1]$. Figure E.1 illustrates the shape of the function W_p.

To prove that the optimal disclosure policy is to reveal the price below a threshold,

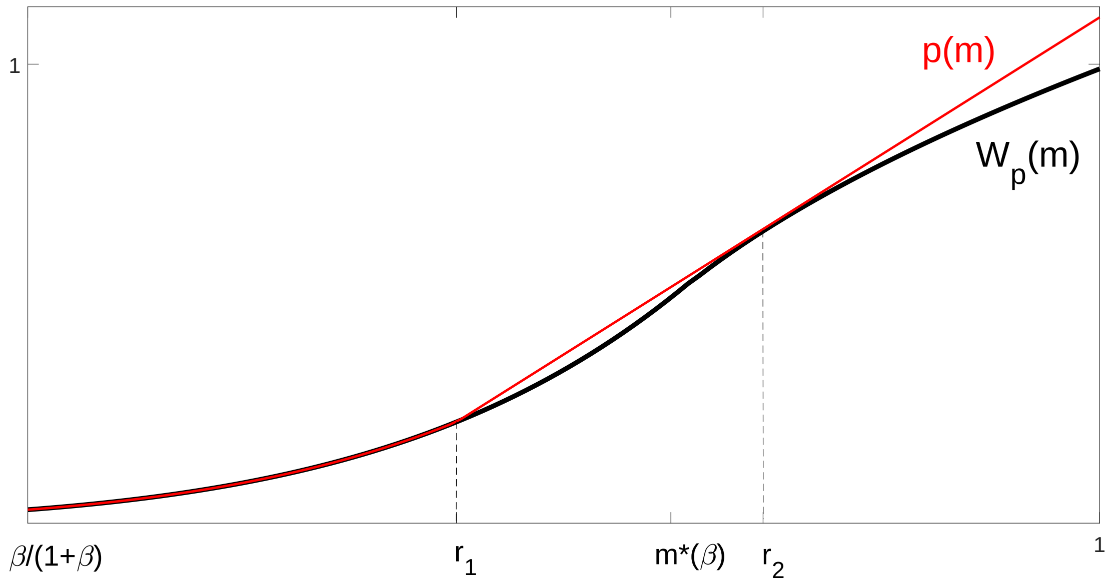

Fig. E.1: The function W_p (black line) and the convex supporting function p (red line).

and announce that the price was above that threshold otherwise, I will use the results of Appendix C.3. Fixing the allocation rule $\bar{x}$, the optimization problem is a Bayesian persuasion problem where the objective function of the designer, W_p, only depends on the posterior mean. Posterior beliefs have to average out to the no-communication posterior $f^{\bar{x}}$, and each posterior belief must likelihood-ratio dominate the prior f.

Following Dworczak and Martini (2017) and Kolotilin (2016) who provide sufficient conditions for optimality of persuasion schemes in cases where the Sender's preferences only depend on the posterior mean, the optimization problem can be represented as maximizing the expectation of W_p over distributions of posterior means on $[\beta/(1+\beta), 1]$ subject to a mean-preserving spread condition. Specifically, let $\bar{H}(m)$ be the distribution of posterior means of the type of the winner when the realization of the cutoff is fully disclosed, conditional on allocating the good:

$$\bar{H}(m) = \mathbb{P}_{c\sim\bar{x},\,\theta_i\sim F}\left(\frac{c+\beta}{1+\beta} \leq m |\, \theta_i \geq c\right) = \frac{\int_0^{m(1+\beta)-\beta}(1-F(c))d\bar{x}(c)}{\int_0^1(1-F(c))d\bar{x}(c)}.$$

When the cutoff is fully disclosed under the allocation rule $\bar{x}$, the platform receives an expected profit of $\int W_p(m)d\bar{H}(m)$. In a cutoff mechanism, the platform (sender) can garble the information about the cutoff in an arbitrary way. It follows that any feasible (induced by some garbling of the cutoff) distribution H over posterior means of the winner's type must have the feature that $\bar{H}$ is a mean-preserving spread of H. I will solve a relaxed problem where this is the only constraint imposed on distributions of posterior means, and verify ex-post that the optimal persuasion scheme can be implemented by a certain disclosure rule of the cutoff.

The relaxed problem is

$$\max_{H} \int_{\frac{\beta}{1+\beta}}^{1} W_p(x) dH(x)$$

subject to

$$\bar{H} \text{ is a mean-preserving spread of } H.$$

This is exactly the problem considered by Dworczak and Martini (2017) and Kolotilin (2016). I want to prove that the optimal H coincides with $\bar{H}$ (full disclosure) on $[\beta/(1+\beta),\, r_1]$, and pools the realizations of $\bar{H}$ into a single mass point at the conditional mean of $\bar{H}$ conditional on a realization in $[r_1,\, 1]$, for some r_1. Such H is feasible. Figure E.1 shows how Theorem 1 of Dworczak and Martini (2017) can be applied to verify optimality of H – it is enough to find a convex function p which coincides with W_p on $[\beta/(1+\beta),\, r_1]$, and is tangent to W_p at $r_2 = \mathbb{E}_{M\sim\bar{H}}[M | M \geq r_1]$. Because $W_p(m)$ is convex for $m \leq m^\star(\beta)$ and concave for $m \geq m^\star(\beta)$, such a function p can always be found for some r_1. Formally, this is true under the convention that when $r_1 = \beta/(1+\beta)$, we do not require $W_p(\beta/(1+\beta)) = p(\beta/(1+\beta))$ (in this case $-p$ is a supporting hyperplane of $-W_p$ at r_2, and H corresponds to full pooling). Therefore, some H of the form described above is optimal.

The optimal distribution H over posterior means of the winner's type can be implemented by revealing the cutoff c whenever $c \leq \bar{c}$, and pooling information about the realization of c into one signal otherwise.[49] Therefore, the optimal distribution is feasible for the non-relaxed problem. Treating the cutoff as the price faced by the winner, I have proven that the optimal disclosure rule reveals the price perfectly when it is below a threshold, and pools information about the price otherwise.

In the Online Appendix, using the characterization of optimal disclosure, I prove that the optimal mechanism is an auction with a reserve price $r_p^\star$. Because that proof is technical and tedious, it is omitted here. What remains to be shown is that, under conditions, more can be said about optimal disclosure.

First, when $\beta \to 0$, we have $m^\star(\beta) \to 1$, and thus the function W_p is convex everywhere except in the small region $[m^\star(\beta),\, 1]$. I want to prove that in this case the truncation point r_1 in Figure E.1 also converges to 1. It is enough to prove that if we draw a line tangent to W_p at 1, then it will intersect W_p at a point $r(\beta)$ that converges to 1 as $\beta \to 0$ (this implies that if p is tangent to W_p at any point $r_2 \in [m^\star(\beta),\, 1]$, then p must intersect W_p to the right of $r(\beta)$). Consider

$$W_p(1) - W_p(1)(1 - r(\beta)) = v[1 - (2-\beta)(1 - r(\beta))] = vm^2(\beta) - \beta I_d(r(\beta)).$$

[49] If there is an atom in the distribution of the cutoff at $\bar{c}$, it is possible that the optimal disclosure rule will be random at $c = \bar{c}$, i.e. the realization $\bar{c}$ is disclosed with some probability.

At $\beta = 0$, the unique solution is $m(0) = 1$, and by continuity of the solution (which follows from the implicit function theorem), this means that $r(\beta) \to 1$, as $\beta \to 0$. Therefore, I have proven that as $\beta \to 0$, the optimal solution converges to full disclosure (the threshold $\bar{c}$ converges to 1).

Second, I want to prove that for large enough β, it is optimal not to disclose any information. An immediate sufficient condition is that W_p is concave on its relevant domain $[\beta/(1+\beta), 1]$ which is equivalent to $m^\star(\beta) = \beta/(1+\beta)$. This equation has a unique solution $\bar{\beta}$ because $m^\star(\beta)$ is strictly decreasing. By a numerical calculation $\bar{\beta} \approx 2.16$.

Third, I want to prove that when $\beta > \sqrt{e} - 1$, full pooling is optimal when N is large enough. When N is sufficiently large, the posterior mean under no disclosure gets arbitrarily close to 1. Therefore, it is enough to prove that a line tangent to W_p at 1 lies strictly above W_p on $[\beta/(1+\beta), 1]$. Indeed, this implies that $r_1 = \beta/(1+\beta)$ in Figure E.1 when the function p is affine and tangent to W_p at the posterior mean r_2 (which lies sufficiently close to 1). The condition is

$$W(1) - W'(1)(1 - \beta/(1+\beta)) > W(\beta/(1+\beta)).$$

This translates to $\log(1+\beta) > 1/2$ or $\beta > \sqrt{e} - 1$ (equal to approximately 0.65). Therefore, for $\beta > \sqrt{e} - 1$, full pooling is optimal for N large enough.