Econometrica, Vol. 88, No. 6 (November, 2020), 2629–2661

MECHANISM DESIGN WITH AFTERMARKETS: CUTOFF MECHANISMS

PIOTR DWORCZAK
Department of Economics, Northwestern University

I study a mechanism design problem in which a designer allocates a single good to one of several agents, and the mechanism is followed by an *aftermarket*—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 design the aftermarket. However, she can influence its information structure by publicly disclosing information elicited from the agents 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 realization of a random threshold (cutoff) that she had to outbid to win the object. When there is a single agent in the mechanism, I show that the *optimal* cutoff mechanism offers full privacy to the agent. In contrast, when there are multiple agents, the optimal cutoff mechanism may disclose information about the winner's type; I provide sufficient conditions for optimality of simple designs. I also characterize aftermarkets for which restricting attention to cutoff mechanisms is without loss of generality in a subclass of all feasible mechanisms satisfying additional conditions.

KEYWORDS: Mechanism design, information design, aftermarkets, transparency.

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 that could 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 endowed with one-dimensional private information. 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).

Piotr Dworczak: piotr.dworczak@northwestern.edu
I am indebted to Andy Skrzypacz, Paul Milgrom, Darrell Duffie, and Michael Ostrovsky for their support and guidance. I would also like to thank five anonymous referees for their comments, as well as Mohammad Akbarpour, Ben Brooks, Eric Budish, Jeremy Bulow, Gabriel Carroll, Jeff Ely, Yossi Feinberg, Alex Frankel, Matthew Gentzkow, Ben Golub, Matthew Jackson, Emir Kamenica, Maciej Kotowski, David Kreps, Nicolas Lambert, Shengwu Li, Elliot Lipnowski, Giorgio Martini, Roger Myerson, Paul Oyer, Alessandro Pavan, Harry Di Pei, Andrés Perlroth, Doron Ravid, Bartosz Redlicki, Peter Reiss, Phil Reny, Ilya Segal, Vasiliki Skreta, Yang Song, Bruno Strulovici, Takuo Sugaya, Zhe Wang, Robert Wilson, Milena Wittwer, Anthony Zhang, participants of Paul Milgrom's research group, and multiple seminar audiences for helpful comments and discussions.

 https://doi.org/10.3982/ECTA15768

Although the mechanism designer is unable to design the aftermarket, she can influence its information structure by publicly releasing some of the information elicited by the mechanism. The design problem is therefore augmented with an additional variable—the *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. Formally, I model the aftermarket as a collection of payoffs (for the agents and the designer) that depend on the true type of the agent who acquired the object but also on the *beliefs* about that agent's type induced by the mechanism.

The resulting structure of the problem can be described as a combination of mechanism and information design. The mechanism elicits information from the agents to determine the allocation and transfers, and subsequently discloses some of that information to other market participants in order to induce the optimal distribution of posterior beliefs in the aftermarket. The two parts of the problem interact non-trivially because disclosure influences the incentives of agents to reveal their private information to the mechanism.

Suppose that the designer considers some allocation and disclosure rule, that is, a mapping from agents' types to a probability distribution over mechanism outcomes: which agent receives the good and what signal is sent. Together with the exogenous aftermarket, the rule determines the final outcome and payoffs. By the revelation principle, 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 optimal mechanism is sensitive to details of the environment and difficult to find.

Consider, however, the following class of allocation and disclosure rules called *cutoff rules*. To receive the object, the agent must report a type that is above some (possibly random) threshold which I refer to as the *cutoff*. Depending on the allocation rule, the cutoff could be, for example, a report (bid) of another agent or a reserve price. I show that such a cutoff representation exists for any monotone allocation rule. Cutoff rules are then defined by a joint restriction on allocation and signals: The allocation rule is monotone, and the signal distribution depends only on the realized cutoff of the winner. Formally, conditional on 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 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. Thus, a second-price auction with (full or partial) disclosure of the price paid by the winner implements a cutoff rule but a first-price auction with disclosure of the price does not.[1]

The key property of cutoff mechanisms is that the report of the winner does not directly influence the signal. Instead, the signal is pinned down by the realization of the cutoff

[1]Many mechanisms commonly used in practice implement cutoff rules. Examples include clock auctions (such as the "Incentive Auction" run by the Federal Communications Commission in 2017), English auctions (used widely for auctioning art and wine), as well as all trading mechanisms (including first-price auctions) that do not disclose the bid of the winner (e.g., privacy-preserving trading platforms in the financial over-the-counter market; see SIFMA (2016)).

which is determined independently of the winner's report. Because the winner cannot manipulate the signal, cutoff mechanisms admit a truthful equilibrium regardless of the details of the environment. Formally, as long as a single-crossing condition holds (fixing the posterior belief in the aftermarket, any agent's payoff from winning the object is non-decreasing in her type), irrespective of the aftermarket protocol and the prior distribution of types, any cutoff rule can be implemented by some transfer scheme. Moreover, this property is *only* satisfied by cutoff rules: For any non-cutoff mechanism, there exists an aftermarket and a prior distribution of types such that this mechanism is not truthful.

The paper focuses on the analysis of cutoff mechanisms and their properties. I develop methods for finding optimal cutoff mechanisms for general aftermarkets. By drawing a connection to information design, the paper contributes to the mechanism design literature by showing how the economic effects of post-mechanism interactions can be analyzed in a tractable way—with the optimal cutoff mechanism often found in closed form. Moreover, I show that cutoff mechanisms are uniquely characterized by some properties that may be desirable in practical design problems, under appropriate assumptions on the aftermarket interactions. In practical applications that are well approximated by these assumptions, the paper offers insights about the optimal transparency of allocation mechanisms; specifically, it supports the use of privacy-preserving mechanisms for single-agent allocation problems, and explains why and what form of information disclosure may be optimal when multiple agents compete for the good. In the remainder of this section, I discuss these findings in more detail.

Suppose that the mechanism designer wants to maximize some objective function, such as revenue or total surplus. In a general mechanism, disclosure of information interacts with incentive-compatibility constraints. But in a cutoff rule, regardless of what information about the cutoff is revealed, the agents want to report truthfully under appropriately chosen transfers. Therefore, finding the optimal disclosure rule 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. In this way, the design problem can be decomposed into two independent steps, where each step can be solved using existing mechanism and information design techniques, respectively.

When the designer contracts with a single 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: Regardless of the aftermarket game and the designer's objective, there always exists an optimal cutoff mechanism that sends no informative signals. Intuitively, in single-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 choosing an appropriate 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 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 often finds it optimal to induce competition between the agents, that is, condition the allocation for agent i on how other agents behave in the mechanism. This, however, means that the designer can no longer choose an arbitrary distribution of the cutoff for agent i. For example, if the designer decides to run an efficient auction, the distribution of the cutoff for agent i—the highest competing bid—is exogenous and cannot be chosen. More generally, when the allocation depends on the ranking of agents' types, the designer is constrained in the choice of prior distributions of the cutoffs. As a result, it may be beneficial to send signals to induce posterior beliefs that differ from the prior. I illustrate this possibility with several examples and sufficient conditions for optimality of popular mechanisms, such as a second price (or English) auction with a reserve price and revelation of the price paid by the winner. I also describe cases in which the designer prefers no disclosure despite the presence of multiple agents, and relies instead on the allocation rule (e.g., reserve prices or randomization) to optimally influence posterior beliefs in the aftermarket.

The class of cutoff mechanisms often excludes the optimal mechanism. While the question of unconstrained optimal design is beyond the scope of this paper, in Section 5, working with a single-agent model, I attempt to cast light on the question of when restricting attention to cutoff mechanisms can be justified. The main result provides conditions on the aftermarket under which restricting attention to cutoff mechanisms is without loss of generality (and hence optimality) within a subclass of all feasible mechanisms. To define the subclass on which the characterization result holds, I first strengthen the notion of implementability to what I call *ex post deterministic (ExD) implementation*. ExD implementation requires truthful reporting regardless of what beliefs the agent holds about the outcome of randomization devices used by the designer. Cutoff rules can always be implemented in this stronger sense. The main result shows that among mechanisms that are ExD implementable and satisfy a regularity condition, only cutoff mechanisms are feasible when the aftermarket is *submodular*. Informally, an aftermarket is submodular if lower types benefit more than high types from a change in posterior beliefs that shifts more probability mass toward higher types. For example, resale aftermarkets are submodular because low-value agents benefit more (relative to high-value agents) from beliefs that induce higher resale prices. I also give examples of *supermodular* aftermarkets for which cutoff mechanisms are suboptimal.[2]

The remainder of the paper is organized as follows. The next subsection discusses related literature. Section 2 introduces the model. Section 3 defines cutoff mechanisms and Section 4 derives the optimal cutoff mechanism. In Section 5, I discuss when looking at cutoff mechanisms may be justified. Section 6 concludes. Some less relevant proofs are relegated to Appendix B. Sections referred to as "Appendices" can be found in the Supplemental Material (Dworczak (2020)).

1.1. *Literature Review*

This paper combines mechanism design with information design. In a seminal paper, Myerson (1981) solved the problem of allocating a single asset in a mechanism design framework, where the designer is allowed to choose an arbitrary mechanism. 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 my model, the principal designs the mechanism and the information structure jointly. 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.[3]

[2]In a companion paper (Dworczak (2020)), I study a similar design problem in a more restricted setting in which a single third party chooses between two actions, and show that a version of submodularity of the aftermarket implies optimality of cutoff mechanisms among all feasible mechanisms.

[3]Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) combined 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.

With regard to the structure of the problem, a closely related literature is a series of papers by Calzolari and Pavan (2006a, 2006b, 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). Calzolari and Pavan (2006b) showed 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 single-agent problems, but the results are not related otherwise: Calzolari and Pavan did not restrict attention to cutoff mechanisms; I do not impose any of the three economic assumptions of the main theorem of Calzolari and Pavan. For example, the upstream principal in Calzolari and Pavan has no direct preferences over the outcome of the second stage; I focus on exactly opposite cases when the principal is concerned with the final allocation (e.g., because she maximizes total surplus). Calzolari and Pavan (2006a) considered a model of a revenue-maximizing monopolist selling an object to an agent who can later resell to a third party. They studied a simple setting with binary types which allowed them to derive a closed-form solution. My model is more general in that it allows an arbitrary objective function, multiple agents, general second-stage game, and general type spaces. I discuss the relationship in more detail in Section 5.

A number of papers analyze the consequences of post-auction interactions between the bidders and third parties. Goeree (2003), Das Varma (2003), Katzman and Rhodes-Kropf (2008), and Hu and Zhang (2017) examined the effect of different bid announcement policies on revenue in standard auctions followed by Bertrand, Cournot, or other forms of competition. Molnár and Virág (2008), assuming the post-auction payoff is type-independent and additively separable, provided sufficient conditions under which a revenue-maximizing mechanism should reveal all or no information about bidders' types. Similarly, Giovannoni and Makris (2014) modeled the aftermarket as an additive component of the objective function that depends on posterior beliefs, and they interpreted it as capturing reduced-form reputational concerns. Back, Liu, and Teguia (2020) studied the effects of transparency on welfare and dealers' profits in financial over-the-counter markets.[4] In all these papers, sufficiently strong assumptions are imposed on the aftermarket payoffs to guarantee existence of a revealing (monotone) equilibrium in the first stage, even when agents' reports (bids) are fully disclosed. Roughly, these assumptions require that higher types of agents have a (weakly) higher willingness to pay for more favorable beliefs (additive separability of the aftermarket payoff is an even stronger assumption). In the terminology introduced by this paper, this is a feature of *supermodular* aftermarkets—such aftermarkets make information disclosure "easy" for the designer. In contrast, the focus of this paper is on *submodular* aftermarkets (such as resale aftermarkets) that make information disclosure "difficult." The precise meaning of these statements is explained in the paper. Engelbrecht-Wiggans and Kahn (1991) and Dworczak (2015) explicitly constructed non-monotone equilibria using a discrete type space in auctions followed by resale games (an example of a submodular aftermarket).

Overall, previous literature made progress on studying the consequences of aftermarket interactions with third parties in two cases: when the aftermarket has a special structure (such as supermodularity) under which full disclosure (and hence any intermediate disclosure) is feasible; or in the opposite case but under restrictive conditions on the type

[4]With the exception of Molnár and Virág (2008) and Hu and Zhang (2017), these papers 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).

space, objective function, and the aftermarket interaction. By introducing cutoff mechanisms, this paper allows a tractable analysis of general aftermarkets, and moreover shows that the restriction to cutoff mechanisms has a justification precisely in the cases where progress has been hindered by lack of tractability (namely, with submodular aftermarkets).

A closely related problem is when bidders interact *with each other* after the mechanism. In general, such problems are significantly more complicated and yield different economic insights; this is primarily because agents in the first-stage mechanism consider not only the signaling effect of their behavior, but also how much they learn about others. A special case of such problems is auction design with inter-bidder resale (e.g., Gupta and Lebrun (1999), Zheng (2002), Haile (2003), Hafalir and Krishna (2008, 2009), Zhang and Wang (2013)). In this literature, to circumvent the difficulty mentioned above, the disclosure 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)), (ii) fixed for the purpose of the analysis (as in Haile (2003), who assumed 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. Carroll and Segal (2019) considered 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).

Balzer and Schneider (2019) analyzed 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.

The paper considers information disclosure *after* the auction, where outsiders learn about bidders' values. This complements a large literature on information disclosure *before* and *during* the auction, where information is controlled by the seller and refines bidders' estimates of their own values, as in Milgrom and Weber (1982), Eső and Szentes (2007), Bergemann and Wambach (2015), Li and Shi (2017), among many others. In these papers, there is no aftermarket. Lauermann and Virág (2012) considered a model where losing bidders exercise a common outside option after the auction, and the auctioneer can disclose information about the value of the outside option either before or after the auction.

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).

2. MODEL

A mechanism designer allocates an indivisible good to one of N agents. The designer chooses an allocation mechanism that specifies the probabilities with which agents receive the object, monetary transfers, and a signal distribution, as a function of agents' messages sent to the mechanism. The signal is publicly revealed after the mechanism. The agent who acquires the object in the mechanism (the "winner") participates in a post-mechanism game with third-party players. The mechanism designer cannot directly influence the post-mechanism game, and cannot contract with the third-party players.

However, the signal revealed by the mechanism may be used to influence the payoffs from the post-mechanism game by changing the beliefs over the winner's type.

Let $\mathcal{N}$ denote the set of agents. Agent $i \in \mathcal{N}$ has a type $\theta_i \in \Theta_i$, where Θ_i is a finite subset of $\mathbb{R}_+$. (The discrete type space is assumed to simplify exposition; Appendix C extends the results to continuous distributions.) Types are independent and distributed according to a prior joint distribution with probability mass function (pmf) $\boldsymbol{f}$ on $\boldsymbol{\Theta} \equiv \bigtimes_{i\in\mathcal{N}}\Theta_i$, with f_i and F_i denoting the pmf and the cdf (cumulative distribution function) of the marginal distribution of agent i's type, respectively. Throughout, bold symbols denote vectors and products, in particular $\boldsymbol{\theta} \equiv (\theta_1, \theta_2, \dots, \theta_N)$, $\boldsymbol{\theta}_{-i} \equiv (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_N)$, and $\boldsymbol{f}(\boldsymbol{\theta}) = \prod_{i\in\mathcal{N}} f_i(\theta_i)$, and a tilde is used to differentiate between random variables and their realizations, for example, $\tilde{\theta}_i$ denotes a random variable and θ_i its realization.

Assuming that the mechanism designer has commitment power and is satisfied with partial implementation, the Revelation Principle will apply; thus, I restrict attention to direct mechanisms. I assume that the mechanism can send an arbitrary public signal once the good is allocated; thus, 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 $\sum_{i\in\mathcal{N}} x_i(\boldsymbol{\theta}) \le 1$, for all $\boldsymbol{\theta}$; $\boldsymbol{\pi} : \boldsymbol{\Theta} \to \bigtimes_{i\in\mathcal{N}}\Delta(\mathcal{S}_i)$ is a signal function with a signal space $\mathcal{S}_i$ for each agent i, and $\boldsymbol{t} : \boldsymbol{\Theta} \to \mathbb{R}^N$ is a transfer function. I assume that each $\mathcal{S}_i$ is finite.[5] 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 publicly announces a signal realization $s \in \mathcal{S}_i$ drawn from the distribution $\pi_i(\cdot|\hat{\theta}_i, \boldsymbol{\theta}_{-i})$ (no other signals are sent). The identity of the winner is assumed to be observable. Thus, the posterior belief over the winner's type $\tilde{\theta}_i$ induced by a truthful mechanism $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t})$ conditional on signal realization s is given by (whenever defined)

$$f_i^s(\tau) = \frac{\displaystyle\sum_{\boldsymbol{\theta}_{-i}} \pi_i(s|\tau, \boldsymbol{\theta}_{-i}) x_i(\tau, \boldsymbol{\theta}_{-i}) \boldsymbol{f}(\tau, \boldsymbol{\theta}_{-i})}{\displaystyle\sum_{\boldsymbol{\theta}} \pi_i(s|\boldsymbol{\theta}) x_i(\boldsymbol{\theta}) \boldsymbol{f}(\boldsymbol{\theta})}, \quad \forall \tau \in \Theta_i. \tag{2.1}$$

I do not explicitly model the third-party players in the aftermarket. Instead, the post-mechanism game is described in reduced form by the conditional expected payoffs it generates for the winner given the information revealed by the mechanism. Formally, an aftermarket A is a collection of payoff functions $A \equiv \{u_i(\theta; \bar{f}) : \theta \in \Theta_i, \bar{f} \in \Delta(\Theta_i), i \in \mathcal{N}\}$, where $u_i(\theta; \bar{f})$ denotes the conditional expected payoff to agent i with type $\theta \in \Theta_i$, when the posterior belief over the type $\tilde{\theta}_i$ is $\bar{f}$, conditional on agent i holding the good. Importantly, the aftermarket is a primitive of the model in that its definition is independent of the mechanism chosen by the designer.

In the truthful equilibrium of the direct mechanism $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t})$, the expected payoff to agent i with type θ_i who deviates to reporting $\hat{\theta}_i$ conditional on other agents reporting $\boldsymbol{\theta}_{-i}$ is $\sum_{s\in\mathcal{S}_i} u_i(\theta_i; f_i^s)\pi_i(s|\hat{\theta}_i, \boldsymbol{\theta}_{-i}) x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i}) - t_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})$. The objective of the mechanism designer is to maximize

$$\sum_{i\in\mathcal{N}} \sum_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \sum_{s\in\mathcal{S}_i} V_i(\theta_i; f_i^s) \pi_i(s|\theta_i, \boldsymbol{\theta}_{-i}) x_i(\theta_i, \boldsymbol{\theta}_{-i}) \boldsymbol{f}(\boldsymbol{\theta}), \tag{2.2}$$

[5]Because of a finite type space, looking at finite signal spaces $\mathcal{S}_i$ is without loss of optimality within the class of cutoff mechanisms.

where each $V_i : \Theta_i \times \Delta(\Theta_i) \to \mathbb{R}$ is assumed to be upper semi-continuous in the second argument. Thus, the payoff of the mechanism designer is normalized to zero when the good is not allocated, and is equal to $V_i(\theta_i; f_i^s)$ otherwise, where $V_i(\theta_i; f_i^s)$ is the payoff conditional on agent i winning the object and belief f_i^s being induced in the aftermarket. The fact that the designer's payoff does not explicitly depend on transfers is essentially without loss of generality given that feasible mechanisms are required to be incentive-compatible and individually-rational: While there may be many transfer rules implementing any given allocation and disclosure rule, the set of implementing transfer rules is a complete sublattice with a largest element.[6] For example, revenue-maximizing expected transfers are uniquely pinned down by the allocation and disclosure rule, and thus formulation (2.2) allows for revenue maximization as the objective of the designer.

2.1. *The Aftermarket*

I model the aftermarket as a "black-box" without explicitly defining the underlying post-mechanism game. This approach 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. Third-party players share the common prior f over the agents' types, and observe the identity i of the winner and the public signal realization s disclosed by the mechanism. This leads to a posterior belief f_i^s over the winner's type. Given belief f_i^s and an aftermarket A, the corresponding game has a set of equilibria $EQ_i^A(f_i^s)$, where $EQ_i^A(\cdot)$ is an upper hemi-continuous correspondence mapping beliefs over the winner's type into equilibrium outcomes, where the equilibrium notion can be specified by the modeler. Then, fixing an equilibrium selection from EQ_i^A (e.g., the designer-preferred equilibrium), $u_i(\theta; f_i^s)$ is the expected equilibrium payoff to type θ of agent i conditional on s.[7]

By assumption, the signal sent by the mechanism influences the aftermarket only through the posterior belief over the winner's type. Other roles of the signal (e.g., 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.

The following single-crossing property will be needed throughout my analysis.

ASSUMPTION 1—Monotonicity: *An aftermarket A is monotone if, for any $i \in \mathcal{N}$ and any $\bar{f} \in \Delta(\Theta_i)$, the aftermarket payoff $u_i(\theta; \bar{f})$ is non-decreasing in θ.*

If there is no aftermarket and the type is equal to the value, $u_i(\theta; \bar{f}) = \theta$, the assumption is trivially satisfied. With an aftermarket, the assumption says that types can be ranked by willingness to pay for the object *irrespective* of the posterior beliefs in the aftermarket.

[6]See, for example, Kos and Messner (2013) and Dworczak and Zhang (2017) for an intuitive proof. With a continuous type space, expected transfers would be pinned down up to a constant, by the payoff equivalence theorem (see, e.g., Milgrom (2004)).

[7]The black-box approach to the aftermarket is without loss of optimality compared to a revelation-principle approach of Myerson (1982) in which the designer would send recommendations to all players in the aftermarket, under the assumption of public communication. This approach can also accommodate exogenous private information of third-party players since the realizations of their private signals (that, without loss of generality, they observe after observing the public signal s) can be integrated out in the expected payoff functions of the agent and the designer.

This is true in most applications where the type is interpreted as a value of the object to the agent. An analogous assumption is made in all papers studying aftermarkets that are surveyed in Section 1.1. I conclude with two examples of aftermarkets that will be used for illustration throughout.

EXAMPLE 1—Resale: [8] Suppose that with probability $\lambda > 0$, there is a single third-party buyer in the aftermarket with some (potentially random) value $\tilde{v}$ for the object. (With the remaining probability, there is no aftermarket and the agent keeps the good obtaining her value θ.) The third party bargains with the winner to repurchase the object. If the equilibrium price is equal to $p(\bar{f}; v)$ when the belief over the winner's type is $\bar{f}$ and $\tilde{v} = v$, then we have $u_i(\theta; \bar{f}) = \lambda \mathbb{E}[\max\{\theta, p(\bar{f}; \tilde{v})\}] + (1 - \lambda)\theta$ which is monotone in θ. If the third party has full bargaining power, and the agent-preferred equilibrium is selected, then

$$p(\bar{f}; v) = \max\left\{\operatorname*{argmax}_p (v - p) \sum_{\theta \leq p} \bar{f}(\theta)\right\}.$$

EXAMPLE 2—Ex Post binary types: Unlike the previous example, this example is a class of simple aftermarkets that capture different economic applications in a tractable manner. Suppose that $\Theta_i \subset [0, 1]$, and let θ_i be the probability with which agent i has an ex post high type h. With complementary probability $1 - \theta_i$, the agent has a low type l. The agent learns the ex post type only after acquiring the object (but before the aftermarket game). If the payoffs of all players in the aftermarket only depend on the winner's ex post type, the utility of the winner i depends on the belief $\bar{f}$ over her ex ante type $\tilde{\theta}_i$ only through its expectation $m(\bar{f}) \equiv \mathbb{E}_{\tilde{\theta}_i \sim \bar{f}}[\tilde{\theta}_i]$. Thus, denoting agent i's aftermarket payoff by $\underline{u}_i(m)$ and $\bar{u}_i(m)$ when her ex post type is high and low, respectively, we have $u_i(\theta; \bar{f}) = \theta \bar{u}_i(m(\bar{f})) + (1 - \theta)\underline{u}_i(m(\bar{f}))$. The aftermarket is monotone if $\bar{u}_i(m) \geq \underline{u}_i(m)$ for all posterior means m.

(a) [Cournot competition][9] Suppose that the mechanism allocates a patent that allows an entrant to enter a market with an incumbent (the third party). Upon acquiring the patent, the winner learns her marginal cost of production which is $c < 1$ for the low type l and $c - \Delta$ for the high type h, where $\Delta > 0$. The incumbent has cost c. Market demand is given by $Q(P) = 1 - P$, and the two firms compete à la Cournot. The equilibrium payoff for the agent in the aftermarket is given by $\underline{u}_i(m) = \frac{1}{9}(1 - c + \frac{m\Delta}{2})^2$ and $\bar{u}_i(m) = \frac{1}{9}(1 - c + \frac{m\Delta}{2} + \frac{3\Delta}{2})^2$ when the incumbent believes that the entrant's type is high with probability m. The aftermarket is monotone.

(b) [Investment game] Consider again an aftermarket where an entrant interacts with an incumbent.[10] The type θ_i of the entrant is the probability that her business model succeeds, in which case a value $v = 1$ is generated (otherwise, the entrant gets a zero payoff). Before observing whether the entrant succeeds, the incumbent takes a costly investment k that allows her to capture a fraction $\alpha(k)$ of the entrant's value in case the entrant is successful (and is a sunk cost otherwise): $k^\star(\bar{f}) \in \operatorname{argmax}_k \mathbb{E}_{\tilde{\theta}_i \sim \bar{f}}[\tilde{\theta}_i \alpha(k) - k]$. Assume that $\alpha : \mathbb{R}_+ \to [0, 1]$ is strictly increasing and concave, differentiable, $\alpha'(0) = \infty$, and $\alpha'(1) \leq 1$

[8] This example generalizes the baseline model of Calzolari and Pavan (2006a).

[9] This economic application—without the restriction to ex post binary types—was considered by Goeree (2003), Katzman and Rhodes-Kropf (2008), and Hu and Zhang (2017).

[10] The first stage could be any mechanism that equips the entrant with something necessary to run her business, for example, a license, patent, or funding.

to guarantee a solution pinned down by the first-order condition. Then, the entrant's aftermarket payoff is given by $\bar{u}_i(m) = 1 - \alpha((\alpha')^{-1}(1/m))$ and $\underline{u}_i(m) = 0$.

2.2. *Implementability*

I will refer to $(\boldsymbol{x}, \boldsymbol{\pi})$, the allocation and disclosure rule, as the *mechanism frame*.

DEFINITION 1: 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:

$$\sum_{s\in\mathcal{S}_i} u_i\big(\theta_i; f_i^s\big)\pi_i(s|\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) - t_i(\theta_i, \boldsymbol{\theta}_{-i}) \geq 0, \tag{$\boldsymbol{IR}$}$$

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

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

To appreciate the difficulty associated with adding the aftermarket, recall first that when there is no aftermarket, that is, $u_i(\theta_i; f_i^s) \equiv \theta_i$, then $(\boldsymbol{x}, \boldsymbol{\pi})$ is DS implementable if and only if $x_i(\theta, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ for any $\boldsymbol{\theta}_{-i}$, that is, an *ex post monotonicity* condition holds (of course, in this case the signal function $\boldsymbol{\pi}$ is irrelevant). In particular, the characterization of implementable outcomes is invariant to the details of the environment such as the distribution of types.

With the aftermarket, this is no longer the case. As is clear from previous work referenced in Section 1.1, the details of the aftermarket matter for how much information can be disclosed. The set of implementable mechanism frames is also sensitive to the prior distribution of types. The following simple example illustrates.

EXAMPLE 3—Resale: Consider Example 1 with $N = 1$ (subscript i can be dropped), $\Theta = \{l, h\}$, $\lambda < 1$, and a third party with a constant value $v > h$ and full bargaining power. The third party offers price h in the aftermarket when she believes the type of the agent to be h with probability at least $\kappa \equiv (h - l)/(v - l)$ (and offers l otherwise). Consider mechanism frames with binary signals, $\mathcal{S} = \{s_L, s_H\}$, with $\pi_h \equiv \pi(s_H|h) \geq 1/2$ and $\pi_l \equiv \pi(s_L|l) \geq 1/2$. By a direct calculation, (x, π) is implementable if and only if

$$x(h) - x(l) \geq \lambda\big[\pi_h x(h) - (1 - \pi_l)x(l)\big]\mathbf{1}_{\{\frac{\pi_h x(h)f(h)}{\pi_h x(h)f(h)+(1-\pi_l)x(l)f(l)} \geq \kappa\}}\mathbf{1}_{\{\frac{(1-\pi_h)x(h)f(h)}{(1-\pi_h)x(h)f(h)+\pi_l x(l)f(l)} < \kappa\}}.$$

Thus, the set of implementable mechanism frames depends on the probability of the aftermarket λ, the prior distribution f, and the value v of the third party. When the two signals induce different prices in the aftermarket (both indicator functions on the right-hand side are equal to 1), there is a trade-off between the choice of x and the choice of π. Full disclosure of the agent's type can be implemented only when x satisfies $x(h) - x(l) \geq \lambda x(h)$.

The problem of finding the optimal mechanism is intractable in the absence of restrictions on the prior distribution (as in Calzolari and Pavan (2006a), who imposed binary types as in the example above) or the aftermarket (as in Goeree (2003), Molnár and Virág (2008), Katzman and Rhodes-Kropf (2008) or Hu and Zhang (2017), who studied a restricted set of aftermarkets that permit arbitrary disclosure in the mechanism). In the next section, I instead introduce a restriction on the class of mechanisms: I study a class of allocation and disclosure rules (cutoff rules) that can always be implemented, hence circumventing the above difficulty.

3. CUTOFF MECHANISMS

For each agent $i \in \mathcal{N}$, let $\bar{c}_i$ be any number greater than $\max \Theta_i$. Then, $C_i = \Theta_i \cup \{\bar{c}_i\}$ is the space of *cutoffs* for agent i. The key observation I will explore is that monotone allocation rules define distributions on the space of cutoffs. Let $\boldsymbol{x}$ be an ex post monotone allocation rule, that is, suppose that $x_i(\theta_i, \boldsymbol{\theta}_{-i})$ is non-decreasing in θ_i for all $\boldsymbol{\theta}_{-i}$. If we extend the interim allocation rule $x_i(\cdot, \boldsymbol{\theta}_{-i})$ by assigning $x_i(\bar{c}_i, \boldsymbol{\theta}_{-i}) = 1$, then it can be treated as a cdf on C_i (I will abuse notation slightly by using the same symbol both when $x_i(\cdot, \boldsymbol{\theta}_{-i})$ is an allocation rule on Θ_i and when it is treated as a cdf on C_i). Furthermore, I let $\Delta x_i(c; \boldsymbol{\theta}_{-i})$ denote the probability that cutoff c is realized (i.e., $\Delta x_i(\cdot; \boldsymbol{\theta}_{-i})$ is the pmf corresponding to the cdf $x_i(\cdot, \boldsymbol{\theta}_{-i})$).

DEFINITION 2—Cutoff Rules: A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is a *cutoff rule* if $\boldsymbol{x}$ is an ex post monotone allocation rule, and there exists a signal function $\gamma_i : C_i \times \boldsymbol{\Theta}_{-i} \to \Delta(\mathcal{S}_i)$ such that, for all $i \in \mathcal{N}$, $\theta_i \in \Theta_i$, $\boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}$, and $s \in \mathcal{S}_i$,

$$\pi_i(s|\theta_i, \boldsymbol{\theta}_{-i}) x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{c \le \theta_i} \gamma_i(s|c, \boldsymbol{\theta}_{-i}) \Delta x_i(c; \boldsymbol{\theta}_{-i}). \tag{3.1}$$

When $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule, I will call $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t})$ a *cutoff mechanism*.

To understand the idea behind cutoff mechanisms intuitively, assume first that there is one agent, $N = 1$, so that an allocation rule is a one-dimensional function $x(\theta)$ (I drop the subscripts). Consider a random variable $\tilde{c}$ (which I will call a random cutoff) with realizations in the space of cutoffs C. I say that $\tilde{c}$ is a *random-cutoff representation* of the allocation rule x if $x(\theta) = \mathbb{P}(\theta \ge \tilde{c})$. The interpretation is that the allocation rule $x(\theta)$ can be achieved by drawing a cutoff c from the distribution of $\tilde{c}$, and giving the good to the agent if and only if the reported type θ is greater than the realized cutoff c. The observation preceding Definition 2 implies that any monotone allocation rule x admits a random-cutoff representation: It is enough to take a random variable $\tilde{c}$ on C with distribution given by cdf x.

The converse is also true: Any random variable $\tilde{c}$ on C represents some monotone allocation rule. Indeed, if y is the cdf of $\tilde{c}$, then $y(\theta)$ (restricted to Θ) is a monotone allocation rule represented by the random cutoff $\tilde{c}$. Thus, there is a one-to-one correspondence between non-decreasing allocation rules on Θ and (distributions of) random cutoffs on C.

In the general model with N agents, given an ex post monotone allocation rule $\boldsymbol{x}$, and fixing the reports $\boldsymbol{\theta}_{-i}$ of other agents, agent i's allocation can be achieved by drawing cutoffs from the conditional distribution $x_i(\cdot, \boldsymbol{\theta}_{-i})$. For example, consider the allocation rule $x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i \ge \theta_{-i}^{(1)}\}}$, where $\theta_{-i}^{(1)} = \max_{j \ne i} \theta_j$. Then, the cutoff for agent i is equal to the highest competing type.[11]

[11] In this case, the cutoff has a degenerate distribution conditional on $\boldsymbol{\theta}_{-i}$. If ties are instead broken uniformly at random, then $x_i(\theta_i, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{\theta_i \ge \theta_{-i}^{(1)}\}} / (|\{j \in \mathcal{N} : \theta_j \ge \theta_{-j}^{(1)}\}|)$ is a two-step function, and the cutoff conditional on $\boldsymbol{\theta}_{-i}$ has a binary distribution, where the lower realization has probability equal to the probability that i wins the tie-breaker.

An intuitive interpretation of Definition 2 is thus as follows: In a cutoff rule $(\boldsymbol{x}, \boldsymbol{\pi})$, each agent i reports $\hat{\theta}_i$. Conditional on other agents' reports $\boldsymbol{\theta}_{-i}$, the seller draws a cutoff c_i from the distribution with pmf $\Delta x_i(\cdot; \boldsymbol{\theta}_{-i})$. If $\hat{\theta}_i \geq c_i$, agent i gets the good, and the designer draws and announces a signal from the distribution with pmf $\gamma_i(\cdot|c_i, \boldsymbol{\theta}_{-i})$. Crucially, conditional on the cutoff realization c_i and $\boldsymbol{\theta}_{-i}$, the signal distribution is independent of i's report $\hat{\theta}_i$. If $\hat{\theta}_i < c_i$, agent i does not receive the good.[12]

Both assumptions imposed by cutoff mechanisms—that (i) the allocation rule is non-decreasing, and (ii) the signal distribution is determined by the realization of the cutoff—are restrictive. Regarding (i), when there is an aftermarket, dominant-strategy implementability does *not* imply that the allocation rule is ex post monotone (intuitively, a higher type might receive the good with lower probability if this is offset by a higher probability of a favorable signal realization). Regarding (ii), cutoff mechanisms preclude disclosure rules that reveal information about the winner directly, for example, by fully revealing her type. Nevertheless, a signal that depends on the realized cutoff c_i is informative about the type of the winner i because third-party players condition on the event $\tilde{\theta}_i \geq \tilde{c}_i$. Conditional on i winning, a cutoff rule can also disclose information about $\boldsymbol{\theta}_{-i}$. For example, full disclosure of the losing agents' reports is allowed: It is enough to set $\mathcal{S}_i = \boldsymbol{\Theta}_{-i}$ and $\gamma_i(s|c, \boldsymbol{\theta}_{-i}) = \mathbf{1}_{\{s=\boldsymbol{\theta}_{-i}\}}$ for any $s \in \mathcal{S}_i$.

Note that the the allocation rule determines how much information can be disclosed by a cutoff rule. If all types of agent i receive the good with the same probability (the allocation rule is constant), the cutoff for agent i is degenerate and hence uninformative about her type. The "steeper" the allocation rule, that is, the larger the differences in probabilities of acquiring the good between high and low types, the more informative the cutoff is about the type of the winner.

While Definition 2 provides an intuitive interpretation of cutoff mechanisms, its conditions are difficult to verify for a generic mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$. Thus, I give an equivalent definition below.

PROPOSITION 1: *A mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is a cutoff rule if and only if, for all i,*

$$\pi_i(s|\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) \quad \text{is non-decreasing in } \theta_i \text{ for all } s \in \mathcal{S}_i \text{ and } \boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}. \tag{M}$$

PROOF: That cutoff rules satisfy condition (M) is immediate from Definition 2. To prove the converse, let $(\boldsymbol{x}, \boldsymbol{\pi})$ be a mechanism frame satisfying (M). Because both the property (M) and the definition of a cutoff rule are checked for every $i \in \mathcal{N}$ separately, I fix an agent i and a profile $\boldsymbol{\theta}_{-i}$, and suppress these symbols from the notation ($x(\theta)$ stands for $x_i(\theta, \boldsymbol{\theta}_{-i})$ etc.). 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 let θ^- be the largest type in Θ smaller than θ, for any $\theta > \underline{\theta}$. Because $\beta_s(\theta)$ is non-decreasing, it induces a positive additive (not necessarily probabilistic) measure with pmf μ_s on C defined by $\mu_s(\underline{\theta}) = \beta_s(\underline{\theta})$, and $\mu_s(\theta) = \beta_s(\theta) - \beta_s(\theta^-)$ for any $\theta > \underline{\theta}$. The pmf μ_s is absolutely continuous with respect to the pmf Δx of the cutoff representing the allocation rule x:

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

[12] In order to implement a cutoff rule when $N > 1$, the designer must properly correlate the cutoffs for different agents to make sure the good is allocated to at most one agent ex post. However, the joint distribution of cutoffs is irrelevant for payoffs (because only one agent interacts in the aftermarket) and implementability (all that matters is the marginal distribution for any agent), and thus the joint distribution need not be specified.

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

$$\pi(s|\theta)x(\theta) = \beta_s(\theta) = \mu_s(\{\tau : \tau \leq \theta\}) = \sum_{c \leq \theta} g_s(c)\Delta x(c), \tag{3.2}$$

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

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

It follows that $\sum_s g_s(c) = 1$, for all c with $\Delta x(c) > 0$. I can now define the measure $\gamma : C \to \Delta(\mathcal{S})$ by $\gamma(s|c) = g_s(c)$, for all c with $\Delta x(c) > 0$ (and in an arbitrary way for c which have probability zero under Δx). Because $\sum_s g_s(c) = 1$, γ is a well-defined signal function. Moreover, equation (3.2) implies that the equality (3.1) from Definition 2 of cutoff rules holds for all s and θ. *Q.E.D.*

Proposition 1 allows me to interpret cutoff rules as mechanism frames that satisfy a strengthening of the ex post monotonicity condition—they are monotone in the type θ_i *for every signal realization* $s \in \mathcal{S}_i$. This property plays a key role in the analysis of implementability in the next subsection.

3.1. *Implementability of Cutoff Rules*

THEOREM 1: *A mechanism frame is DS implementable for any prior distribution* $\boldsymbol{f}$ *and any monotone aftermarket* A *if and only if it is a cutoff rule*.

PROOF: To prove that any cutoff rule is DS implementable, I argue that condition (M) implies implementability for any prior distribution and any monotone aftermarket (this is enough due to Proposition 1; alternatively, one could directly verify that Rochet (1987)'s cyclic monotonicity condition holds). We can formally think of signal realizations as defining distinct goods allocated by the seller. Then, condition (M) says that for each of these goods, indexed by s, the allocation rule is non-decreasing. Moreover, a monotone aftermarket guarantees that a single-crossing property holds between the types of each agent and allocations of each of the goods. Thus, for every $s \in \mathcal{S}_i$ and every fixed $\boldsymbol{\theta}_{-i}$, there exists a transfer rule $t_i^s(\theta_i, \boldsymbol{\theta}_{-i})$ that implements the allocation rule $\pi_i(s|\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i})$ of good s. Defining $t_i(\theta_i, \boldsymbol{\theta}_{-i}) = \sum_{s \in \mathcal{S}_i} t_i^s(\theta_i, \boldsymbol{\theta}_{-i})$ finishes the first part of the proof.

To prove the converse, again by Proposition 1, it is enough to show that if a mechanism frame $(\boldsymbol{x}, \boldsymbol{\pi})$ is DS implementable for any prior distribution $\boldsymbol{f}$ and any monotone aftermarket A, then it must satisfy condition (M). Fix any $(\boldsymbol{x}, \boldsymbol{\pi})$, $i \in \mathcal{N}$, $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$. Since $(\boldsymbol{x}, \boldsymbol{\pi})$ is assumed DS implementable, condition (***IC***) has to hold for θ_i and $\hat{\theta}_i$ with some transfers t_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

$$\sum_{s \in \mathcal{S}_i}\left[u_i(\theta_i; f_i^s) - u_i(\hat{\theta}_i; f_i^s)\right]\left[\pi_i(s|\theta_i, \boldsymbol{\theta}_{-i})x_i(\theta_i, \boldsymbol{\theta}_{-i}) - \pi_i(s|\hat{\theta}_i, \boldsymbol{\theta}_{-i})x_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})\right] \geq 0. \tag{3.3}$$

Denote $\beta_s(\tau) \equiv \pi_i(s|\tau, \boldsymbol{\theta}_{-i})x_i(\tau, \boldsymbol{\theta}_{-i})$. Since condition (3.3) must hold for any monotone aftermarket and any prior, it must hold in particular for aftermarkets with

$u_i(\theta_i; f_i^s) = u_i(\hat{\theta}_i; f_i^s)$ for all s with $\beta_s(\theta_i) \geq \beta_s(\hat{\theta}_i)$, and $u_i(\theta_i; f_i^s) > u_i(\hat{\theta}_i; f_i^s)$ otherwise.[13] Under such u_i, inequality (3.3) becomes

$$\sum_{\{s \in \mathcal{S}_i : \beta_s(\theta_i) < \beta_s(\hat{\theta}_i)\}} \left[u_i\left(\theta_i; f_i^s\right) - u_i\left(\hat{\theta}_i; f_i^s\right)\right]\left[\beta_s(\theta_i) - \beta_s(\hat{\theta}_i)\right] \geq 0, \tag{3.4}$$

with $u_i(\theta_i; f_i^s) > u_i(\hat{\theta}_i; f_i^s)$ for each s in the summation. We have thus obtained that a sum of strictly negative terms is nonnegative. This is only possible when the set of indices in the sum is empty: $\{s \in \mathcal{S}_i : \beta_s(\theta_i) < \beta_s(\hat{\theta}_i)\} = \emptyset$. Because $\theta_i > \hat{\theta}_i$ and $\boldsymbol{\theta}_{-i}$ were arbitrary, this shows that condition (M) holds, finishing the proof. *Q.E.D.*

The economic intuition for Theorem 1 is straightforward: Under a cutoff rule, the report of the winner does not directly influence the signal sent by the mechanism, and thus the winner cannot manipulate the aftermarket belief over her type. Losing agents can manipulate posterior beliefs, but this is irrelevant since they do not participate in the aftermarket. This is reminiscent of why VCG mechanisms (such as second-price auctions) are truthful. In a VCG mechanism, the report of an agent does not influence the transfer the agent pays, except when it changes the allocation. In a cutoff mechanism, the report does not influence the signal, except when it changes the allocation. While the agent can change the outcome by affecting the probability with which she acquires the good, monotonicity of the aftermarket implies that such a deviation can be deterred by appropriately chosen transfers.

To gain intuition for the converse part of Theorem 1, it is again helpful to think of different signal realizations $s \in \mathcal{S}$ as different goods allocated by the designer. For any fixed prior distribution and aftermarket, incentive-compatibility requires that these goods are allocated with probability that is non-decreasing in the agent's type *on average* across s. However, as we consider all possible priors and aftermarkets, the allocation probability must be monotone in *each* good s separately—this is the only way to guarantee that the average allocation probability is monotone regardless of the (endogenous) valuations $u_i(\theta_i; f_i^s)$ for different goods s. By Proposition 1, this is exactly what defines cutoff mechanisms.

In subsequent analysis, I will only use the part of Theorem 1 that guarantees that cutoff mechanisms can always be made incentive-compatible by an appropriate choice of transfers. However, the converse part has economic consequences as well. Implementability for all prior distributions and aftermarkets implies that cutoff mechanisms are a natural benchmark that can be used to establish a lower bound on the value of the objective function in *any* design problem with a monotone aftermarket. Moreover, cutoff rules are the *largest* class that can serve this purpose: Any rule outside of the class cannot be implemented in at least some cases, and hence cannot serve as a universal lower bound. Furthermore, I show in Appendix A that for some aftermarkets, such as resale, requiring implementability for all prior distributions $\boldsymbol{f}$ is already enough to rule out all but cutoff rules. This property can be useful in practical design problems due to its connection to robustness. In general, the transfer function implementing a cutoff rule will depend on the prior $\boldsymbol{f}$, and hence cutoff mechanisms are not a detail-free design.[14] However, if the

[13]Note that a payoff function u_i satisfying these properties exists because—by choosing $\boldsymbol{f}$ appropriately—we can ensure that $f_i^s \neq f_i^{s'}$ for any $s \neq s'$.

[14]This is a consequence of the setting rather than a feature of cutoff rules: With the aftermarket, the prior $\boldsymbol{f}$ and the aftermarket A directly influence the values $u_i(\theta_i; f_i^s)$ that agent i has for winning. In the analogy

designer hopes to implement a mechanism frame robustly, that is, without knowing the details of the environment, it is certainly necessary that there exist transfers that implement that frame in each possible case. Thus, Theorem 1 implies that a designer interested in robust implementation of a mechanism frame has no reason to look beyond the class of cutoff mechanisms.[15]

4. OPTIMAL CUTOFF MECHANISMS

In this section, I consider optimization in the class of cutoff mechanisms. I first focus on the single-agent case which produces a particularly sharp result and simplifies exposition. Then, I show how to generalize the solution to multi-agent mechanisms.

4.1. *Optimal Cutoff Mechanisms With a Single Agent*

In this subsection, I assume $N = 1$ (and omit the subscript i in the notation).

I say that a disclosure rule π *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 $\theta, \hat{\theta} \in \Theta, s \in \mathcal{S}$. Importantly, even when π reveals no information, the posterior belief in the aftermarket may differ from the prior because the fact that the agent participates in the aftermarket is informative of her type when the allocation rule x is non-constant. The following result establishes a strong conclusion about optimal cutoff mechanisms in the single-agent model.

THEOREM 2: *With $N = 1$, the problem of maximizing* (2.2) *subject to* (x, π) *being a cutoff rule has an optimal solution* $(x^\star, \pi^\star)$ *where* $\pi^\star$ *reveals no information.*

The conclusion of Theorem 2 holds regardless of the objective function. The type of the objective may influence the shape of the optimal allocation rule $x^\star$ but never requires the designer to make explicit announcements via $\pi^\star$.[16]

I prove the theorem in two steps: First, I consider optimization over disclosure rules for any fixed allocation rule x, and then I show that at the optimal allocation rule $x^\star$, the corresponding optimal $\pi^\star$ reveals no information. The first step provides an important building block for the multi-agent model, while the second step is specific to the case of a single agent.

Proof of Theorem 2. Step 1: Optimization Over Disclosure Rules. I fix a non-decreasing allocation rule x, and optimize over disclosure rules π subject to (x, π) being a cutoff rule. The proof strategy is as follows: As discussed in Section 3, any non-decreasing allocation rule x can be represented by a random cutoff, which I will denote $\tilde{c}_x$. In a cutoff mechanism, the signal only depends on the realization of $\tilde{c}_x$. By Theorem 1, any disclosure of the cutoff is compatible with both (***IR***) and (***IC***) constraints. Thus, the mechanism design problem becomes a pure information design problem in which the designer chooses

developed by the proof of Theorem 1, prices of goods indexed by $s \in \mathcal{S}_i$ must depend on how valuable these goods are to agent i.

[15]That being said, when the designer does not know the distribution of types and the aftermarket, it is no longer without loss of generality to restrict attention to direct mechanisms. The designer might instead fix an indirect mechanism, allowing the allocation and disclosure rule to be determined endogenously in equilibrium as the distribution and the aftermarket vary.

[16]As discussed in Section 2.1, Theorem 2 does not preclude the possibility that the designer could send uninformative signals to coordinate play in the aftermarket, if the aftermarket game were modeled explicitly instead of using the black-box approach.

a disclosure policy of the random cutoff $\tilde{c}_x$ in order to induce the optimal distribution of posterior beliefs—this is the Bayesian persuasion problem of Kamenica and Gentzkow (2011) where the relevant state is the cutoff.

The prior distribution of the cutoff $\tilde{c}_x$ (the state variable) is given by the cdf x. Given signal function γ and a signal realization s, the conditional distribution of $\tilde{c}_x$ has a cdf $x^s(c) = [\sum_{\hat{c}\le c} \gamma(s|\hat{c})\Delta x(\hat{c})]/[\sum_{\hat{c}} \gamma(s|\hat{c})\Delta x(\hat{c})]$. I will be using y as a generic symbol for a cdf of a conditional distribution of the cutoff. The aftermarket belief over the winner's type can be derived in two steps in a cutoff rule: (i) given a signal realization, the conditional cdf of the cutoff is y, (ii) conditional on the agent acquiring the object, the corresponding posterior belief over that agent's type is

$$f^y(\theta) \equiv \mathbb{P}_{\tilde{c}\sim y}(\tilde{\theta} = \theta | \tilde{\theta} \ge \tilde{c}) = \frac{y(\theta) f(\theta)}{\sum_{\tau} y(\tau) f(\tau)}. \tag{4.1}$$

The above derivation uses the fact that the order of conditioning does not matter, and that, in a cutoff rule, the signal is independent of the winner's type conditional on the cutoff, so that in step (ii), the belief over the winner's type depends on the signal only indirectly through the conditional distribution of the cutoff. Using the equivalence between non-decreasing allocation rules and distributions over cutoffs, f^y can also be interpreted as the aftermarket belief over the type of the agent who acquired the good that would arise if the designer implemented the allocation rule $y(\theta)$ (and disclosed no further information). Next, let

$$\mathcal{V}(y) = \sum_{\theta\in\Theta} V\big(\theta; f^y\big) y(\theta) f(\theta) \tag{4.2}$$

be the expected payoff to the mechanism designer conditional on the signal inducing a cdf y of the cutoff and the agent acquiring the object in the mechanism. Equivalently, $\mathcal{V}(y)$ is the expected payoff to the mechanism designer that would arise if the allocation function were y (instead of the actual x) and the mechanism revealed no additional information to the third party. It now follows from Theorem 1 and the results of Kamenica and Gentzkow (2011) that we can optimize over distributions of conditional distributions over the cutoff subject to a Bayes-plausibility constraint (this is immediate, but I provide a formal proof in Appendix B.1).

LEMMA 1: *With $N = 1$, for any non-decreasing allocation rule x, the problem of maximizing (2.2) over π subject to (x, π) being a cutoff rule is equivalent to*

$$\max_{\varrho\in\Delta(\Delta(C))} \mathbb{E}_{y\sim\varrho}\mathcal{V}(y), \tag{4.3}$$

$$s.t. \quad \mathbb{E}_{y\sim\varrho} y(\theta) = x(\theta), \quad \forall\theta\in\Theta. \tag{4.4}$$

The mechanism designer seeks to maximize her expected payoff over distributions ϱ of beliefs over the cutoff (equation (4.3)). Condition (4.4) is the Bayes-plausibility constraint—the induced distributions over the cutoff must average out to the prior (with distributions represented by cdfs).

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

COROLLARY 1: *With $N = 1$, for a fixed allocation rule x, the maximal expected payoff to the mechanism designer is equal to the concave closure of $\mathcal{V}$ at x*: $\mathrm{co}\mathcal{V}(x) \equiv \sup\{\nu : (x, \nu) \in \mathrm{CH}(\mathit{graph}(\mathcal{V}))\}$, *where* CH *denotes the convex hull*, *and* $\mathit{graph}(\mathcal{V}) \equiv \{(\hat{x}, \hat{\nu}) \in \mathcal{X} \times \mathbb{R} : \hat{\nu} = \mathcal{V}(\hat{x})\}$.

Step 2: Optimization Over Allocation Rules. By Corollary 1, the value to the designer at an optimal solution, now involving optimizing over x as well, is $\sup_{x\in\mathcal{X}} \mathrm{co}\mathcal{V}(x)$. By definition of the concave closure, $\sup_{x\in\mathcal{X}} \mathrm{co}\mathcal{V}(x) = \sup_{x\in\mathcal{X}} \mathcal{V}(x)$, that is, the value of the function and its concave closure coincide at the supremum. An optimal solution $x^\star$ exists because $\mathcal{V}$ is upper semi-continuous on a compact set. This finishes the proof of Theorem 2: $\mathcal{V}(x^\star)$ is the expected payoff to the mechanism designer when $x^\star$ is the allocation rule and the disclosure rule reveals no information. *Q.E.D.*

The proof provides a simple intuition for Theorem 2: When choosing an optimal cutoff rule, the problem of the designer is to choose a prior distribution of cutoffs (the allocation rule), and then optimally disclose information about the realized cutoff. Thus, the designer is a Sender who can *choose* the prior distribution of the state. When the posterior belief can be chosen directly by choosing the prior, there is no need to reveal additional information about the state. In the design of the optimal cutoff mechanism, there is no need to reveal additional information about the cutoff because the optimal posterior distribution can be induced directly by choosing the prior (the allocation rule). To illustrate the above results, I apply them to solve an example.

EXAMPLE 4—Resale: Consider Example 1 with $N = 1$, $\lambda = 1$, and $\Theta = \{l, h\}$ with $f(l) = f(h)$. There is a single third party with a constant value $v \in (h, 2h - l)$ that makes a take-it-or-leave-it offer to the agent in the aftermarket. The designer maximizes total surplus: $V(\theta; \bar{f}) = \theta \mathbf{1}_{\{\theta > p(\bar{f})\}} + v\mathbf{1}_{\{\theta \le p(\bar{f})\}}$, where $p(\bar{f})$ denotes the optimal offer made by the third party under posterior belief $\bar{f}$ over the agent's type (with ties broken in the designer's favor).

It is clear that $x(h) = 1$ in the optimal solution. Hence, the set of allocation rules is a one-dimensional family indexed by the probability $x(l)$ with which the object is allocated to the low type l. The cutoff representation $\tilde{c}_x$ of x is a binary random variable on $C = \{l, h\}$ with pmf Δx given by $\Delta x(l) = x(l)$ and $\Delta x(h) = 1 - x(l)$.

Fix an arbitrary distribution of the cutoff with cdf y. This distribution corresponds to a belief $f^y(h) = 1/(1 + y(l))$ that the type of the agent is high conditional on participation in the aftermarket (see equation (4.1)). The third party offers a high price h when she believes that the probability of the high *cutoff* is at least $\alpha^\star \equiv 2 - (v - l)/(h - l)$. With a binary cutoff distribution, the function $\mathcal{V}(y)$, defined by (4.2), only depends on the probability $1 - y(l)$ that the cutoff is high, and thus can be represented as

$$\mathcal{V}(y) = \mathcal{V}_1\big(1 - y(l)\big) = \begin{cases} vf(l) + hf(h) - v\big(1 - y(l)\big)f(l) & \text{if } 1 - y(l) < \alpha^\star, \\ v - v\big(1 - y(l)\big)f(l) & \text{if } 1 - y(l) \ge \alpha^\star. \end{cases} \tag{4.5}$$

(The subscript 1 in $\mathcal{V}_1$ is introduced to formally distinguish between the one-dimensional function $\mathcal{V}_1 : [0, 1] \to \mathbb{R}$ and $\mathcal{V}$ that is a function of the entire cdf; this representation will allow a simple graphical analysis.)

By Corollary 1, optimal disclosure for any fixed allocation rule x yields the concave closure of $\mathcal{V}_1$. The function $\mathcal{V}_1$ and its concave closure are depicted in Figure 1. When $1 - x(l) < \alpha^\star$, so that the third party would offer a low price when no signals are sent, it is optimal to disclose information about the cutoff in the form of a binary signal: $s \in$

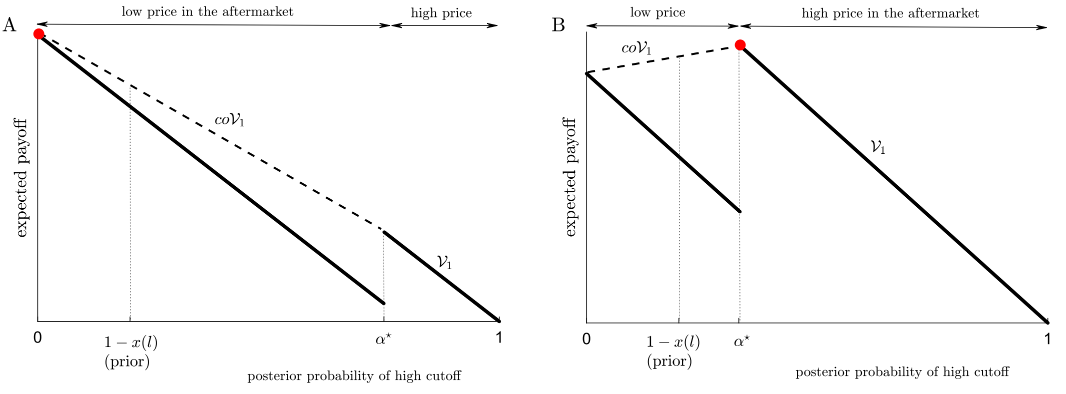

FIGURE 1.—Function $\mathcal{V}_1$ (solid line) and its concave closure (dashed line) when $v(v-h) < h(h-l)$ (panel A) or $v(v-h) > h(h-l)$ (panel B).

$\{s_L, s_H\}$. The designer sends s_L when the cutoff is low with probability η and sends s_H in all other cases. The probability η is chosen so that conditional on s_H, the third party is indifferent between offering the high and the low price (and offers the high price). When $1 - x(l) \geq \alpha^\star$, the third party already offers a high price under the prior; $\mathcal{V}_1$ coincides with its concave closure, and the designer makes no announcement in the optimal mechanism.

Next, suppose that the designer can optimize over both the allocation and the disclosure rule, that is, she can additionally choose $x(l)$. In Figure 1, since any $x(l) \in [0, 1]$ is feasible, the designer can choose an arbitrary point on the x-axis to maximize the concave closure of $\mathcal{V}_1$. There are two cases to consider depending on which of $\mathcal{V}_1(0)$ and $\mathcal{V}_1(\alpha^\star)$ is greater: $\mathcal{V}_1(0)$ is the expected surplus when the designer lets all types trade in the mechanism (which results in a low price in the aftermarket), while $\mathcal{V}_1(\alpha^\star)$ is the expected surplus when the designer excludes (exactly) enough low types from trading to always induce a high price in the aftermarket:

(A) $\mathcal{V}_1(0) > \mathcal{V}_1(\alpha^\star)$ (holds when $v(v-h) < h(h-l)$; see panel A in Figure 1): In this case, intuitively, it is difficult to induce a high price in the aftermarket. The optimal mechanism corresponds to choosing $1 - x(l) = 0$ because the concave closure of $\mathcal{V}_1$ is maximized at 0. All types trade in the mechanism, no information is disclosed, and the price in the aftermarket is low.

(B) $\mathcal{V}_1(0) < \mathcal{V}_1(\alpha^\star)$ (holds when $v(v-h) > h(h-l)$; see panel B in Figure 1): In this case, it is relatively easy to induce a high price in the aftermarket. The optimal mechanism corresponds to choosing $1 - x(l) = \alpha^\star$ because the concave closure of $\mathcal{V}_1$ is maximized at $\alpha^\star$. Low types trade with probability $1 - \alpha^\star$, no information is disclosed, and the price in the aftermarket is high.

In both cases, no information is disclosed in the optimal mechanism.

I conclude with a few remarks based on the above example. First, since no information disclosure is always optimal with one agent, the allocation rule (equivalently, the distribution of cutoffs) is chosen to optimally trade-off its direct allocative effect against the quality of beliefs in the aftermarket. The trade-off can be seen in Figure 1: A higher distribution of the cutoff induces higher beliefs in the aftermarket (the optimal price jumps up at $\alpha^\star$) at the cost of excluding the low type from trading with higher probability (the function $\mathcal{V}_1$ is decreasing for a fixed price in the aftermarket). Second, Theorem 2 implies that no information disclosure is optimal only at the *optimal* allocation rule. As the example shows, information disclosure can be optimal when the allocation rule is chosen suboptimally. Third, there could be two reasons why no information is disclosed in the optimal mechanism: (i) In case (A), information disclosure would be beneficial (since the aftermarket price is low otherwise) but no information can be disclosed because the optimal prior distribution of the cutoff is degenerate (the allocation rule is constant); (ii) in case (B), the cutoff has a non-degenerate distribution but information disclosure would lower the expected surplus.

4.2. *Optimal Cutoff Mechanisms With Multiple Agents*

I now consider the model with N agents. I show that the general problem can be reduced to one-dimensional optimization over disclosure rules, allowing the application of methods derived for the single-agent case. This is accomplished by working with reduced forms of cutoff mechanisms.[17]

Let $\bar{x}_i : \Theta_i \to [0, 1]$ denote a generic (interim expected) allocation rule for agent i. Definitions (4.1) and (4.2) are directly generalized to the multi-agent setting by putting back the subscripts i. $\mathcal{V}_i(y_i)$ is interpreted as the designer's expected payoff from interacting with agent i that would arise if y_i were the interim expected allocation rule for agent i and the mechanism revealed no additional information. Let $\mathcal{X}_i$ denote the set of one-dimensional non-decreasing allocation rules on Θ_i.

THEOREM 3: *The problem of maximizing* (2.2) *over cutoff rules has the same value as*

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

subject to the Matthews–Border condition:

$$\sum_{i \in \mathcal{N}} \sum_{\theta_i > \tau_i} \bar{x}_i(\theta_i) f_i(\theta_i) \leq 1 - \prod_{i \in \mathcal{N}} F_i(\tau_i), \quad \forall \boldsymbol{\tau} \in \mathbb{R}^N. \tag{M-B}$$

Any cutoff mechanism that maximizes (2.2) *induces a reduced form that solves the problem* (4.6) *subject to* (M-B). *Conversely*, *any solution to problem* (4.6) *subject to* (M-B) *is a reduced form of a cutoff mechanism that maximizes* (2.2).

Theorem 3 implies that to solve the general problem, it is enough to solve N one-dimensional persuasion problems—corresponding to finding the concave closure of each $\mathcal{V}_i$—and then maximize over interim expected allocation rules subject to condition (M-B). The proof of the theorem (found in Appendix B.2) follows the same steps as the derivation of the optimal mechanism in Section 4.1 for the single-agent case. However, there are two complications associated with working with interim expected allocation rules instead of a single-agent allocation rule. First, one must guarantee that the N-tuple of interim expected allocation rules $(\bar{x}_1, \dots, \bar{x}_N)$ is feasible, that is, induced by some joint allocation rule $\boldsymbol{x}$ under $\boldsymbol{f}$. This is ensured by the Matthews–Border condition (M-B) that has been derived in the literature on reduced-form auctions as a necessary and sufficient condition for feasibility (see Matthews (1984), and Border (1991)). Second, interim expected allocation rules are not sufficient to express dominant-strategy implementability—the reduced

[17]A reduced form of a mechanism $(\boldsymbol{x}, \boldsymbol{\pi}, \boldsymbol{t})$ is derived by taking expectations over $\tilde{\boldsymbol{\theta}}_{-i}$, thereby obtaining the (interim expected) allocations, signals, and transfers as a function of θ_i only, for each $i \in \mathcal{N}$; see Appendix B.2 for details.

form of a mechanism can only be used to establish Bayesian implementability. However, I show that in the class of cutoff mechanisms, there is no gap between Bayesian and dominant-strategy implementation; the argument relies on a proof technique developed by Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) in the literature on BIC-DIC equivalence (see also Manelli and Vincent (2010)).

When agents are ex ante identical, it is without loss of optimality to look at symmetric mechanisms, and the maximization problem in Theorem 3 takes a simpler form:

$$N \max_{\bar{x}\in\mathcal{X}} \mathrm{co}\mathcal{V}(\bar{x}) \quad \text{subject to} \sum_{\theta>\tau} \bar{x}(\theta)f(\theta) \le \frac{1-F^N(\tau)}{N}, \quad \forall \tau \in \mathbb{R}. \tag{4.7}$$

When $N = 1$, the Matthews–Border condition (M-B) holds vacuously; hence, unconstrained maximization of the concavified objective, $\max_{x\in\mathcal{X}} \mathrm{co}\mathcal{V}(x) = \max_{x\in\mathcal{X}} \mathcal{V}(x)$, implies the optimality of no disclosure. In contrast, when $N \ge 2$, the Matthews–Border condition is not redundant, and the optimal cutoff mechanism may disclose information. To see why, consider the symmetric case (4.7). The concave closure of $\mathcal{V}$ is taken in the space of *all* non-decreasing interim allocation rules (equivalently, all distributions over the cutoff), while the actual rule $\bar{x}$ must be chosen from a strictly smaller subset of rules that satisfy the condition (M-B), so that it is possible that

$$\max_{\bar{x}\in\mathcal{X},\bar{x}\text{ satisfies (M-B)}} \mathrm{co}\mathcal{V}(\bar{x}) > \max_{\bar{x}\in\mathcal{X},\bar{x}\text{ satisfies (M-B)}} \mathcal{V}(\bar{x}).$$

Whenever the above inequality holds, it is optimal to induce conditional distributions over the cutoff that do not correspond to an interim allocation rule satisfying (M-B), and that can only be achieved by sending informative signals in the mechanism. To illustrate this point, I revisit Example 4.

EXAMPLE 5—Resale: Consider the same problem as in Example 4 but with $N > 1$. Because agents are ex ante identical, I can look at the symmetric optimization problem (4.7). Because the high type of any agent need not receive the good with probability 1, the space of cutoffs is $C = \{l, h, \bar{c}\}$ with $\bar{c} > h$. The problem is to find $\bar{x}(l)$ and $\bar{x}(h)$ to maximize $\mathrm{co}\mathcal{V}(\bar{x})$ subject to $\bar{x}(h) \le (2/N)(1 - 1/2^N)$, $\bar{x}(l) + \bar{x}(h) \le 2/N$ (the M-B condition), and $\bar{x}(l) \le \bar{x}(h)$ (non-decreasing allocation rule).

By using the reduced-form representation of mechanisms, I solve the joint optimization problem by only looking at the interim expected allocation and disclosure from the perspective of a single agent. Because agents are symmetric, I will refer to "the" agent and "the" cutoff by fixing one (any) of the N agents.

It is intuitive that the optimal mechanism should maximize the probability of trade for high types by setting $\bar{x}(h)$ to its maximal feasible level $(2/N)(1 - 1/2^N)$ (I prove that this is indeed the case in Appendix B.3). Thus, the allocation rule can be again parameterized by a single number $\bar{x}(l)$, aiding the comparison to the case $N = 1$ in Example 4. For any cdf y on C, we have

$$\mathcal{V}(y) = \begin{cases} vy(l)f(l) + hy(h)f(h) & \text{if } \dfrac{y(h)-y(l)}{y(h)} < \alpha^\star \\ vy(l)f(l) + vy(h)f(h) & \text{if } \dfrac{y(h)-y(l)}{y(h)} \ge \alpha^\star \end{cases} = y(h)\mathcal{V}_1\left(1 - \frac{y(l)}{y(h)}\right), \tag{4.8}$$

where $\mathcal{V}_1$ is defined by (4.5). Equation (4.8) implies that the key properties of $\mathcal{V}$ can be understood by conditioning on the event that the cutoff is strictly less than $\bar{c}$ (as in the

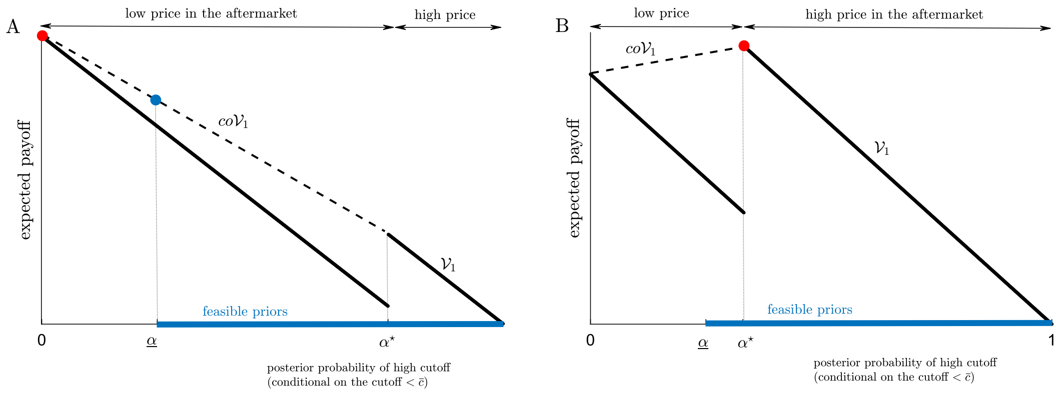

FIGURE 2.—Function $\mathcal{V}_1$ (solid line) and its concave closure (dashed line) when $v(v-h) < h(h-l)$ (panel A) or $v(v-h) > h(h-l)$ (panel B).

opposite case the agent does not receive the good). The expected contribution to total surplus is obtained by multiplying the expected payoff conditional on the cutoff being strictly below $\bar{c}$ (which is given by the same function $\mathcal{V}_1$ as in the case $N = 1$) by the probability $y(h)$ that the cutoff is strictly below $\bar{c}$. Moreover, it can be easily verified that

$$\mathrm{co}\mathcal{V}(y) = y(h)\mathrm{co}\mathcal{V}_1\left(1 - \frac{y(l)}{y(h)}\right),$$

which allows me to use the same graphical illustration as in Example 4: The x-axis will now represent the probability that the cutoff is high conditional on the cutoff being strictly below $\bar{c}$ (see Figure 2).

For any fixed (feasible) allocation rule $\bar{x}$, the optimal disclosure rule can be derived in the same way as in Example 4: Depending on whether $1 - \bar{x}(l)/\bar{x}(h)$ is below or above $\alpha^\star$, the optimal disclosure rule will either feature a binary signal or no announcement. Instead, I focus on joint optimization over allocation and disclosure rules. Since $\bar{x}(h) = (2/N)(1 - 1/2^N)$, the problem becomes

$$\max_{\bar{x}(l)} \mathrm{co}\mathcal{V}_1\left(1 - \frac{\bar{x}(l)}{(2/N)\big(1 - 1/2^N\big)}\right) \quad \text{subject to } \bar{x}(l) \le \frac{1}{N}\frac{1}{2^{N-1}}.$$

Unlike in the case $N = 1$ when the choice of $\bar{x}(l)$ was unconstrained, the designer can only choose from a subset of all prior distributions over the cutoff. Indeed, the constraint on $\bar{x}(l)$ implies that the prior belief $1 - \bar{x}(l)/\bar{x}(h)$ that the cutoff is high (conditional on the cutoff being less than $\bar{c}$) must be at least $\underline{\alpha} \equiv (2^N - 2)/(2^N - 1)$ (see Figure 2). For intuition, recall that (i) the good is always allocated to a high type when a high type is present, and (ii) the allocation rule is symmetric. Thus, before any additional information is disclosed, the third party must believe that the probability $\bar{x}(l)$ with which any given agent faces a low cutoff is bounded above by the probability that no other agent has a high type ($1/2^{N-1}$) times the probability that an agent is selected uniformly at random from N agents ($1/N$). Consequently, she must place sufficiently high probability on the cutoff being high. This is an example of a more general observation that the presence of multiple agents imposes constraints on the interim allocation rule through the Matthews–Border condition (M-B).

Panel A in Figure 2 (the case $v(v-h) < h(h-l)$) illustrates the possibility that information disclosure is optimal: As long as $\underline{\alpha} < \alpha^\star$ (a low price would be quoted if no information was disclosed), it is optimal to choose $\bar{x}(l)$ that induces the lower-bound belief $\underline{\alpha}$ over the cutoff (since $\mathrm{co}\mathcal{V}_1$ is decreasing), and disclose a binary signal that pushes the third party's posterior belief to either 0 or $\alpha^\star$. This mechanism can be implemented as a second price auction in which a low type places a low bid b_l, a high type places a high bid b_h, and the auctioneer sends a low signal s_L with probability η when the price is low, and sends a high signal s_H in all other cases. The parameter η is chosen so that conditional on s_H, the third party is indifferent between a low and a high price in the aftermarket (and offers a high price). It is easy to show that *any* optimal mechanism must disclose information in this case.

In panel B of Figure 2 (the case $v(v-h) > h(h-l)$), the unconstrained optimal $\bar{x}(l)$ corresponding to inducing belief $\alpha^\star$ is feasible when $\underline{\alpha} < \alpha^\star$. In this case, there is no information disclosure but $\bar{x}(l)$ is chosen to be $\bar{x}(h)(1-\alpha^\star)$ that is strictly lower than $(1/N)(1/2^{N-1})$. Thus, the optimal cutoff mechanism can be implemented as a second price auction with "inefficient" allocation at the low bid: When all agents place a low bid, the good is not allocated with probability $z < 1$, where z is sufficiently high so that, conditional on the good being allocated, the third party is indifferent between offering a low and high price in the aftermarket (and offers a high price).

Contrasting Example 4 with Example 5 highlights the difference between the case $N = 1$ and $N > 1$. When $N = 1$, the designer can select any prior for the persuasion problem by choosing the corresponding allocation rule. But when $N > 1$, the designer is constrained in the choice of interim expected allocation rules due to the feasibility constraint (M-B). Effectively, the designer solves N persuasion problems and chooses a prior for each problem, but the priors are jointly constrained. Therefore, it might be optimal to send signals that induce posterior beliefs that do not correspond to a feasible collection of prior distributions (i.e., a feasible collection of interim expected allocation rules). For instance, in case (A) of Example 5, when the mechanism allocates to high types with maximal probability (which is optimal), there is no feasible allocation rule that induces the belief that the cutoff is low with probability 1. However, this belief can be induced with positive probability by sending a signal revealing the low realization of the cutoff.

4.3. *Optimality of Simple Disclosure Rules*

In this subsection, I provide sufficient conditions for optimality of full and no disclosure of the cutoff. Importantly, these conditions are expressed in terms of how the designer's payoff depends on the beliefs over the winner's *type*. In Appendix B.4, 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) a Bayes-plausibility condition holds, and (ii) each posterior belief over the winner's type likelihood-ratio (LR) dominates the prior belief.[18] Condition (i) is natural in the context of information design (see Kamenica and Gentzkow (2011)), while condition (ii) is a consequence of monotonicity (M) 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. Define

$$\mathcal{W}_i(\bar{f}) = \sum_{\theta \in \Theta_i} V_i(\theta; \bar{f})\bar{f}(\theta) \tag{4.9}$$

[18] A pmf g *likelihood-ratio dominates* a full-support pmf f if $g(\theta)/f(\theta)$ is non-decreasing.

as the expected payoff to the designer conditional on agent i winning and posterior belief $\bar{f}$ over i's type. Let M_{f_i} be the set of distributions over Θ_i that likelihood-ratio dominate the prior f_i, and let $f_i^{\bar{x}_i}$, defined by (4.1), be the posterior belief over i's type given the interim expected allocation rule $\bar{x}_i$, when i is the winner and no other information is revealed.

PROPOSITION 2: *The problem of maximizing* (2.2) *over cutoff rules for a fixed* (*interim expected*) *allocation rule* $\bar{\mathbf{x}}$ *has the same value as*

$$\sum_{i\in\mathcal{N}}\left(\sum_{\theta_i\in\Theta_i}\bar{x}_i(\theta_i)f_i(\theta_i)\right)\mathrm{co}^{M_{f_i}}\mathcal{W}_i\big(f_i^{\bar{x}_i}\big), \tag{4.10}$$

where $\mathrm{co}^{M_{f_i}}\mathcal{W}_i(f_i^{\bar{x}_i}) \equiv \sup\{\nu : (f_i^{\bar{x}_i}, \nu) \in \mathrm{CH}(graph(\mathcal{W}_i)|_{M_{f_i}})\}$, *and* $graph(\mathcal{W}_i)|_{M_{f_i}}$ *is the graph of* $\mathcal{W}_i$ *restricted to domain* M_{f_i}.

Objectives (4.6) and (4.10) are analogous except for two important details. First, in (4.10), $\mathcal{W}_i$ is concavified in the subspace $M_{f_i} \subsetneq \Delta(\Theta_i)$, while in (4.6), the concave closure of $\mathcal{V}_i$ is taken in the entire space $\Delta(C_i)$. This is because a cutoff rule can induce an arbitrary belief over the cutoff but can only induce beliefs over the winner's type that LR dominate the prior. Second, in (4.10), the concavified objective is multiplied by an additional term $\sum_{\theta_i\in\Theta_i}\bar{x}_i(\theta_i)f_i(\theta_i)$—the ex ante probability of allocating the good to agent i. This is because the distribution of beliefs over the winner's type is a conditional distribution (conditional on allocating the good to agent i), so the conditional expected payoff must be converted into an ex ante expected payoff. As a corollary of Proposition 2, I obtain the following result.

COROLLARY 2: *If $\mathcal{W}_i$ is convex on its domain, the optimal cutoff mechanism fully discloses i's cutoff when i is the winner. If $\mathcal{W}_i$ is concave, the optimal cutoff mechanism reveals no information when i is the winner.*

Corollary 2 is related to a result by Molnár and Virág (2008) who derived conditions under which full and no disclosure of the type of the winner is part of a revenue-maximizing mechanism followed by a post-auction market. These results are complementary: Molnár and Virág (2008) allowed all feasible mechanisms but restricted attention to settings where the aftermarket payoff is an additively separable component that does not depend on the type of the agent (precluding all examples considered in this paper). I allow more general aftermarkets but restrict attention to the class of cutoff mechanisms.

Corollary 2 can be used to provide simple examples showing that full disclosure of the cutoff is uniquely optimal.[19] I conclude with such an example.

EXAMPLE 6—Cournot Competition: Consider the Cournot competition model (case (a) of Example 2). Suppose that there are N ex ante symmetric potential entrants competing for a single patent, and the mechanism designer chooses a disclosure rule in an auction to maximize total surplus (defined as the area under the demand curve minus

[19]When $\mathcal{W}$ is strictly convex and $N = 1$, a consequence of Corollary 2 and Theorem 2 is that the optimal cutoff distribution must be degenerate, or, equivalently, the optimal allocation rule takes the form $x(\theta) = \mathbf{1}_{\{\theta\geq r\}}$ for some r (then, and only then, full disclosure and no disclosure of the cutoff coincide).

the costs of production). Dropping subscripts (due to symmetry), we obtain $V(\theta;\bar{f}) = \theta\overline{V}(\mathbb{E}_{\bar{f}}[\tilde{\theta}]) + (1-\theta)\underline{V}(\mathbb{E}_{\bar{f}}[\tilde{\theta}])$, where $\overline{V}(m)$ and $\underline{V}(m)$ denote the total surplus conditional on the winner's type being high or low, respectively, when the aftermarket belief about the winner's type has expectation m. From this, we get that

$$\mathcal{W}(\bar{f}) = \sum_{\theta\in\Theta} V(\theta;\bar{f})\bar{f}(\theta) = \mathbb{E}_{\bar{f}}[\tilde{\theta}]\overline{V}\big(\mathbb{E}_{\bar{f}}[\tilde{\theta}]\big) + \big(1-\mathbb{E}_{\bar{f}}[\tilde{\theta}]\big)\underline{V}\big(\mathbb{E}_{\bar{f}}[\tilde{\theta}]\big) \equiv W\big(\mathbb{E}_{\bar{f}}[\tilde{\theta}]\big).$$

The objective function $\mathcal{W}(\bar{f})$ depends on the posterior belief over the winner's type only through its expectation. By direct calculation, $W(m)$ is a quadratic function of m with a strictly positive coefficient on m^2. It follows that $\mathcal{W}(\bar{f})$ is a convex function of $\bar{f}$. By Corollary 2, full disclosure of the cutoff is uniquely optimal in the class of cutoff rules.[20] For example, if the designer uses a second-price auction to allocate the patent, then disclosure of the price after the auction is optimal.

4.4. *Optimality of Simple Mechanisms Under Continuous Type Spaces*

To simplify exposition, I focused on the case of a discrete type space. However, as Appendix C formally demonstrates, all the results established so far continue to hold when the distribution of types is continuous on any compact, convex $\boldsymbol{\Theta} \subset \mathbb{R}_+^N$. While this case does not add any new economic insights, it is sometimes more tractable by permitting the use of calculus. For this section only, I adopt a continuous type space to characterize optimal cutoff mechanisms under the assumption that the payoff in the aftermarket depends on the posterior belief only through its mean. This assumption is satisfied in Example 2 which I will use for illustration of the results derived below. To streamline exposition, I further assume that agents are symmetric (hence drop the subscript i) and I normalize $\Theta \equiv [0, 1]$. I let f denote the density of the prior distribution of an agent's type, and F be the cdf.

For any density function $\bar{f}$ on $[0, 1]$, 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 measurable function $W : [0, 1] \to \mathbb{R}_+$, where $\mathcal{W}$ is defined as in (4.9): $\mathcal{W}(\bar{f}) = \int_0^1 V(\theta, \bar{f})\bar{f}(\theta)\,d\theta$. I also let $m(c) \equiv \int_c^1 \theta f(\theta)\,d\theta/(1-F(c))$ denote the expected value of $\tilde{\theta}$ conditional on $\tilde{\theta} \geq c$, and let $w(c) \equiv W(m(c))$, for any $c \in [0, 1]$. Thus, $w(c)$ is the expected payoff to the designer conditional on allocating the good and inducing a belief that the type of the winner is above c.

PROPOSITION 3: *Suppose that f is continuous and 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*

$$r^\star \in \operatorname*{argmax}_{r\in[0,1]} \big(1-F^N(r)\big) W\left(\frac{\int_r^1 \theta\, dF^N(\theta)}{1-F^N(r)}\right). \tag{4.11}$$

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

[20]This follows from strict convexity of $W(m)$ in the mean m: If any information about the cutoff was pooled, it would be possible to reveal additional information and raise total surplus.

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 maximum of the second highest type and $\underline{r}$. A sufficient condition for* (ii) *is that W is non-decreasing and log-concave.*

The proof of Proposition 3 can be found in Appendix C.1.

If W is concave and non-decreasing, it is optimal not to disclose any information, and the allocation rule is designed to maximize the posterior expected type of the winner by allocating to the highest bidder. The mechanism can additionally raise the expectation by excluding types below r from trading. However, this incurs a cost for the designer because the good is not always allocated. The $r^\star$ that solves equation (4.11) optimally trades-off these two effects.

Second, if W is concave and non-increasing, 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, a cutoff mechanism cannot 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 $J_w(c)$ that reflects the local trade-off between the allocation in the mechanism (as captured by the term $w(c)$) and the information structure induced in the aftermarket (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 the posterior belief over the winner's type conditional on allocating. The function $J_w(c)$ is similar to the virtual surplus function that reflects the trade-off between allocative efficiency and information rents in the revenue-maximization problem. In the regular case, the virtual surplus function is non-decreasing, and the seller does not use randomization to maximize revenue. Analogously, if $J_w(c)$ is non-decreasing, the designer does not use randomization to optimally influence beliefs in the aftermarket.

I apply Proposition 3 to solve two examples based on the model of Example 2.

EXAMPLE 7—Ex Post Binary Types: Consider first the investment game (case (b) of Example 2). Let $k(m) = \text{argmax}_k\{m\alpha(k) - k\}$ be the optimal investment for the incumbent when the expected type of the entrant is m. Consider a designer who maximizes total surplus in the aftermarket. Because the value generated if the entrant is successful is split between the entrant and the incumbent, maximizing surplus is equivalent to minimizing the cost of the (socially wasteful) investment $k(m)$: $W(m) = m - k(m)$. By the envelope formula, we have $W'(m) = 1 - \alpha(k(m)) \geq 0$, and $W''(m) = -\alpha'(k(m))k'(m) \leq 0$. Thus, W is non-decreasing and concave. By Proposition 3 point (1), the optimal mechanism in the first stage is an auction with a reserve price and no information disclosure.

Next, consider the Cournot model (case (a) of Example 2). By Example 6, for any fixed allocation rule, full disclosure of the cutoff maximizes total surplus. Using Proposition 3, we can also pin down the optimal allocation rule. First, $W(m)$ is a convex, non-decreasing function on $[0, 1]$. Moreover, because $W(m)$ is a quadratic function, it is log-concave. It follows from point (3) of Proposition 3 that the optimal mechanism is to run a second-price auction with some reserve price r and reveal the price paid by the winner.

The analysis so far has been silent about *when* using a cutoff mechanism is optimal when arbitrary mechanisms are allowed. It is easy to show that the optimal cutoff mechanism for the investment-game aftermarket (case (b)) is optimal overall. However, it follows from the analysis of Goeree (2003) and Hu and Zhang (2017) that with the Cournot aftermarket (case (a)), the optimal cutoff mechanism from Example 7 is not optimal overall—the optimal mechanism discloses the *type* of the agent rather than the cutoff.[21] It turns out that optimality of cutoff mechanisms depends primarily on the structure of the aftermarket. This is the subject of the next section.

5. WHEN IS THE RESTRICTION TO CUTOFF MECHANISMS JUSTIFIED?

The goal of this section is to derive conditions on the aftermarket under which restricting attention to cutoff rules is without loss of generality. The main result characterizes cutoff rules as the unique feasible class within mechanism frames that (i) satisfy a strong notion of implementability, and (ii) induce posterior beliefs that can be ranked in a certain way. While mathematically limiting, conditions (i) and (ii) have no bite when there is no aftermarket. Under (i) and (ii), I show that cutoff rules are without loss of generality under *submodular aftermarkets*. The next subsection strengthens the notion of implementability; Section 5.2 defines submodular aftermarkets; Section 5.3 contains the main result.

To simplify exposition, and because the ideas presented here are orthogonal to the complications associated with multi-agent mechanisms, I assume that there is a single agent ($N = 1$), and hence drop the subscript i.[22]

5.1. *Ex Post Deterministic Implementation*

I say that a mechanism frame (x, π) is deterministic if, for all θ and s, $x(\theta) \in \{0, 1\}$ and $\pi(s|\theta) \in \{0, 1\}$. Randomization in the mechanism can be captured by allowing the designer to have a type θ_0 drawn (independently of the agent's type) from some auxiliary type space Θ_0, and letting the mechanism depend deterministically on the extended type profile $(\theta, \theta_0) \in \Theta \times \Theta_0$.

OBSERVATION 1: *For any mechanism frame (x, π), there exists a measurable space Θ_0 and a distribution over Θ_0 with cdf F_0 such that*

$$x(\theta) = \int_{\Theta_0} \hat{x}(\theta; \theta_0)\, dF_0(\theta_0), \tag{5.1}$$

$$\pi(s|\theta) = \int_{\Theta_0} \hat{\pi}(s|\theta; \theta_0)\, dF_0(\theta_0), \tag{5.2}$$

where $(\hat{x}(\cdot; \theta_0), \hat{\pi}(\cdot; \theta_0))$ is a deterministic mechanism frame for any $\theta_0 \in \Theta_0$. I call $(\hat{x}, \hat{\pi})$ the deterministic decomposition of (x, π).

When randomization is modeled as an endogenous type of the designer, it is natural to extend the notion of implementability by requiring truthful reporting regardless of the beliefs held by the agent over the designer's type (as in dominant-strategy implementation).

[21] Goeree (2003) and Hu and Zhang (2017) focused on revenue (not total surplus) and did not assume ex post binary types, but their results can be easily modified to apply to this setting.

[22] As shown in a previous working version of the paper, the results can be extended to the general case.

DEFINITION 3: A mechanism frame (x, π) is *ex post deterministically* (ExD) implementable if there exists a deterministic decomposition $(\hat{x}, \hat{\pi})$ of (x, π) such that $(\hat{x}(\cdot; \theta_0), \hat{\pi}(\cdot; \theta_0))$ is implementable for all θ_0.

ExD implementation requires that a mechanism can be represented as an ex ante randomization over deterministic and incentive-compatible mechanisms. Thus, a mechanism is ExD-IC if the designer could disclose the outcome of any randomization *prior* to the agent reporting her type and still satisfy the IC constraints. Consequently, the agent should report truthfully regardless of what beliefs she holds about how the designer is randomizing.

ExD implementation is desirable in contexts where agents do not fully trust the mechanism designer. Arguably, as long as the designer implements an outcome that lies within the support of the distribution, it is difficult to prove that randomization was not correctly conducted. All other deviations by the designer, such as changing the payments or the allocation as a function of reports, can be directly detected. If the agent thinks that the designer has limited commitment in that she might disobey the rules of the mechanism as long as this cannot be detected, she might not want to report truthfully even if the mechanism was IC. However, the agent would want to report truthfully if the mechanism was ExD-IC.[23]

With a monotone aftermarket, ExD-IC mechanisms include all cutoff mechanisms.

PROPOSITION 4: *When the aftermarket is monotone, any cutoff rule is ExD implementable.*

PROOF: By Definition 2 of a cutoff rule, we have, for any s and θ,

$$\pi(s|\theta)x(\theta) = \sum_{c \leq \theta} \gamma(s|c)\Delta x(c) = \sum_{c \in C} \sum_{s' \in \mathcal{S}} (\mathbf{1}_{\{s'=s\}}\mathbf{1}_{\{\theta \geq c\}})\gamma(s'|c)\Delta x(c).$$

This, however, is a representation of a cutoff rule as randomization over deterministic and implementable mechanism frames, where implementability follows from the monotonicity of the allocation in θ for any $c \in C$ and $s' \in \mathcal{S}$ (and monotonicity of the aftermarket). *Q.E.D.*

The intuition is as simple as the proof: In a cutoff rule, the cutoff captures randomization in the mechanism from the perspective of the agent. Moreover, the designer can reveal the cutoff realization and the signal realization to the agent *before* asking her to report her type. This is because the allocation remains monotone in the type conditional on a cutoff and signal realization (recall property (M)).

Proposition 4 implies that ExD implementation has no additional bite without the aftermarket—any monotone allocation rule is ExD implementable because any monotone allocation rule admits a cutoff representation. However, with the aftermarket, there exist IC mechanisms that are not ExD-IC, as the following simple example illustrates.

EXAMPLE 8—Resale: Consider the setting of Example 3. As shown there, the mechanism frame $x(l) = 1 - \lambda$, $x(h) = 1$ with full disclosure of the agent's type, $\pi(s_H|h) =$

[23]This discussion is inspired by Akbarpour and Li (2020) who used a similar concern to motivate their class of credible mechanisms (see also Dequiedt and Martimort (2015)).

$\pi(s_L|l) = 1$, is implementable. However, it is not ExD implementable. To see why, note that, in any deterministic decomposition, there must be a $\theta_0 \in \Theta_0$ such that $x(h; \theta_0) = x(l; \theta_0) = 1$. Example 3 shows that full disclosure is not incentive-compatible when coupled with this allocation rule, and thus there cannot exist a deterministic decomposition of (x, π) into implementable mechanism frames.

In the next section, I identify aftermarkets for which the concept of ExD implementability has the most bite, and prove a partial converse to Proposition 4.

5.2. *Submodular Aftermarkets*

To avoid cases where players care about the "label" of a belief, I make the following assumption which is automatically satisfied when the payoffs in the aftermarket are derived from optimal choices of Bayesian agents: If for some $\bar{f}, \bar{g} \in \Delta(\Theta)$, $u(\theta; \bar{f}) = u(\theta; \bar{g})$ for all $\theta \in \Theta$ (beliefs $\bar{f}$ and $\bar{g}$ have the same payoff consequences), then also $u(\theta; \bar{f}) = u(\theta; \lambda\bar{f} + (1-\lambda)\bar{g})$ for any $\lambda \in (0, 1)$ (their convex combination has the same payoff consequences). The same property is assumed about the designer's payoff V. More substantially, I impose submodularity of the agent's payoff in her type and beliefs—implying that the willingness to pay for "high" beliefs is decreasing in the type of the agent. I use the likelihood-ratio order on beliefs defined in Section 4.3,[24] which I will denote by $\succeq^{\mathrm{LR}}$.

DEFINITION 4: An aftermarket A is *submodular* if for any $\bar{f}, \bar{g} \in \Delta(\Theta)$,

$$\bar{f} \succeq^{\mathrm{LR}} \bar{g} \quad \Longrightarrow \quad u(\theta; \bar{f}) - u(\theta; \bar{g}) \quad \text{is non-increasing in } \theta.$$

An aftermarket is *strictly submodular* if additionally

$$\bar{f} \succeq^{\mathrm{LR}} \bar{g} \quad \Longrightarrow \quad u(\theta; \bar{f}) - u(\theta; \bar{g}) \quad \text{is strictly decreasing in } \theta$$

whenever $u(\theta; \bar{f}) \neq u(\theta; \bar{g})$ for some type $\theta \in \Theta$.

An aftermarket is submodular if lower types have a higher willingness to pay for an upward shift in beliefs. For example, if all types of the agent prefer to be perceived as a high type, this means that any improvement in posterior beliefs is valued more by lower types. This is the case in resale aftermarkets because lower types benefit more (relative to keeping the good) from a high resale price. In particular, the resale aftermarket from Example 1 satisfies submodularity because beliefs higher in the LR order lead to (weakly) higher resale prices. Simple resale aftermarkets are typically not *strictly* submodular—this is because two types $\theta > \hat{\theta}$ differ in their willingness to pay for a resale price p only if that price is accepted by $\hat{\theta}$ but rejected by θ. It can be shown that a resale market becomes strictly submodular if every price happens with positive probability conditional on any given signal realization (e.g., because the value of the third party is stochastic).

Submodularity may also be consistent with agents preferring to be perceived as low types. In Example 2, submodularity requires that $\bar{u}(m) - \underline{u}(m)$ is non-increasing in the posterior mean m. The investment game from Example 2(b) induces a strictly submodular

[24]To allow for the possibility of disjoint supports, I say that $\bar{g} \succeq^{\mathrm{LR}} \bar{f}$ if there exist full-support $\bar{g}_\epsilon$ and $\bar{f}_\epsilon$ such that $\bar{f}_\epsilon \to \bar{f}$, $\bar{g}_\epsilon \to \bar{g}$, and for small enough $\epsilon > 0$, $\bar{g}_\epsilon(\theta)/\bar{f}_\epsilon(\theta)$ is non-decreasing in θ.

aftermarket because $\bar{u}(m)$ is strictly decreasing and $\underline{u}(m) = 0$: Each type benefits from being perceived as a low type but lower types are hurt (strictly) less by an increase in the posterior mean.

An example of an aftermarket that does *not* satisfy submodularity is the Cournot model from Example 2(a). Here, the agent wants to be perceived as a high type (i.e., as having a low cost), and high types benefit more from more favorable beliefs—the aftermarket is in fact *supermodular* (where a supermodular aftermarket is defined by reversing the monotonicity condition in Definition 4).

A key observation is that it is *difficult* to disclose information about the agent's type under a submodular aftermarket: Indeed, submodularity implies that the direction of single-crossing is opposite to the one dictated by Bayesian updating. If beliefs are thought of as goods allocated by the mechanism, then submodularity of the aftermarket implies that high beliefs (beliefs that put more mass on higher types) must be "allocated" to lower types. Bayesian updating requires the opposite: On average, high beliefs must be associated with high types. This tension implies that an incentive-compatible mechanism can only disclose coarse information when the aftermarket is submodular.

5.3. *The Characterization*

DEFINITION 5: A mechanism frame (x, π) is *regular* if the posterior beliefs $\{f^s\}_{s\in\mathcal{S}}$ over the agent's type can be completely ranked in the likelihood-ratio order.

The regularity condition is mathematically restrictive. However, regularity does not in itself rule out signals that directly reveal the agent's type, which is important in the context of upcoming results. It holds, for example, when the disclosure rule has a monotone partitional structure in either the type of the agent or the cutoff. Finally, regularity automatically holds when the type space is binary, which is assumed in many papers studying optimal information design, including Calzolari and Pavan (2006a, 2006b).

THEOREM 4: *Suppose that the aftermarket is strictly monotone*[25] *and strictly submodular. Then, any regular ExD implementable mechanism frame is payoff-equivalent to a cutoff rule.*

The assumption of *strict* monotonicity and submodularity can be dropped if the mechanism is robust to how the agent breaks ties between reports.

DEFINITION 6: A mechanism frame is *strictly* implementable if there exists a transfer rule t such that for each θ, the agent *strictly* prefers her allocation to any distinct allocation received by a different type:

$$\operatorname*{argmax}_{\hat{\theta}} \sum_{s\in\mathcal{S}} u\big(\theta; f^s\big)\pi(s|\hat{\theta})x(\hat{\theta}) - t(\hat{\theta}) = \mathcal{I}(\theta),$$

where $\mathcal{I}(\theta) = \{\hat{\theta} : \forall s,\ \pi(s|\hat{\theta})x(\hat{\theta}) = \pi(s|\theta)x(\theta)\}$. A mechanism frame is *strictly ex post deterministically* (SExD) implementable if it has a deterministic decomposition into strictly implementable mechanism frames.

[25] An aftermarket is *strictly monotone* if $u(\theta; \bar{f})$ is strictly increasing in θ for any $\bar{f} \in \Theta$. Strict monotonicity is satisfied in Example 2(a) and (b) and holds for any original monotone aftermarket if we add an arbitrarily small probability that the aftermarket does not take place (in which case the agent keeps the good and receives a value equal to her type), as in Example 1 for any $\lambda < 1$.

Strict implementability requires that, for some transfer rule, the agent *strictly* prefers to receive the outcome that she obtains by reporting truthfully. This has no bite without the aftermarket: If different types receive the good with different probabilities, there exists a transfer rule that makes truthful reporting a *unique* optimal strategy.[26] Failure of this property implies that the mechanism relies on all types breaking the indifference in the direction preferred by the designer.

THEOREM 4': *Suppose that the aftermarket is monotone and submodular. Then, any regular SExD implementable mechanism frame is payoff-equivalent to a cutoff rule.*

COROLLARY 3: *If* $|\Theta| = 2$, *then any SExD-IC mechanism followed by a monotone and submodular aftermarket is payoff-equivalent to a cutoff mechanism.*

The proofs of Theorems 4 and 4' can be found in Appendices B.5 and B.6. I show that incentive-compatibility of the mechanism implies that if two types receive the same allocation, then lower types must be assigned to signals that lead to higher posterior beliefs (in the likelihood-ratio order). This is a consequence of the single-crossing property in types and beliefs induced by a submodular aftermarket. On the other hand, Bayesian updating implies the opposite relationship between types and beliefs. The resulting conflict between incentive-compatibility and Bayes plausibility limits the informativeness of signals that can be sent in a feasible mechanism. Information about the cutoff can always be disclosed (Theorem 1), and the proof demonstrates that this lower bound on informativeness is achieved.

ExD implementation plays an important role in the proof because it allows me to apply the above reasoning for every endogenous type of the designer separately. Under weaker solution concepts, it would be possible to use randomization in the mechanism to disclose additional information about the agent's type. For example, in the single-agent binary-types model of Calzolari and Pavan (2006a), the aftermarket is a resale game and is therefore submodular. Nevertheless, Calzolari and Pavan showed that in one of four cases, it is optimal to use a non-cutoff mechanism (in particular, a non-cutoff mechanism is feasible). Corollary 3 implies that the incentives to report truthfully in their optimal mechanism (which is analogous to the mechanism considered in Example 8) crucially rely on providing a random outcome to the low type. If the agent did not trust the designer to correctly randomize, she would not report truthfully.

The assumption of a submodular aftermarket is crucial for the result. Under the opposite case of supermodularity (which is satisfied by the Cournot aftermarket—see case (a) of Example 2), it is easier to disclose information. In that case, the relationship between types and beliefs implied by incentive-compatibility and Bayes plausibility is aligned: Higher types are associated with higher beliefs. It is thus possible to support truthful disclosure of the type by using transfers, even when all types receive the same allocation. Indeed, Goeree (2003) and Hu and Zhang (2017) showed that the optimal mechanism for a Cournot aftermarket is to fully disclose the *type* of the winner. In those cases, as seen in Examples 6 and 7, an optimal cutoff mechanism fully discloses the *cutoff*. Thus, restricting attention to cutoff mechanism is likely to be suboptimal when the aftermarket is supermodular.

[26]Such a transfer rule would not in general guarantee the same payoff to the designer. However, it could guarantee an arbitrarily close approximation of that payoff.

6. CONCLUDING REMARKS

In this paper, I studied mechanism design in a setting where the mechanism is followed by an aftermarket, that is, 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. Under a strong notion of implementability and regularity, cutoff rules coincide with the set of feasible outcomes in cases when the aftermarket satisfies a submodularity condition.

Although the results of this paper are established under relatively strong assumptions, many of them continue to hold under much weaker conditions. For instance, cutoff rules remain dominant-strategy implementable even if types are allowed to be correlated. The assumption of a public signal and of irrelevance of beliefs over losing agents' types allowed me to characterize the payoffs in the aftermarket as a function of a single posterior belief. However, this assumption could be relaxed as well: The aftermarket payoffs would then depend on the *vector* of beliefs (one belief for each realization of a different private signal) about the entire type *profile*. In a cutoff rule, the distributions of these signals would be required to only depend on the cutoff and the losing agents' reports. With an analogous definition of monotonicity of the aftermarket, cutoff rules would remain the unique class satisfying implementability for all aftermarkets and prior distributions. However, characterizing *optimal* cutoff mechanisms with private signals would require more advanced tools such as the Bayes correlated equilibrium of Bergemann and Morris (2016a) and is left for future work.

The approach taken to mechanism design in this paper is non-standard. Instead of looking for the optimal mechanism that can depend on fine details of the model, I proposed a class of allocation and disclosure rules with a certain robustness property (implementability in the "worst case"). Within the class, the designer maximizes a Bayesian objective—distinguishing this approach from models that look for the mechanism with the highest payoff guarantee (optimality in the "worst case"). An interesting direction for future research is to apply this approach to other design problems.

REFERENCES

AKBARPOUR, M., AND S. LI (2020): "Credible Auctions: A Trilemma," *Econometrica*, 88 (2), 425–467. [2655]

AUMANN, R. J., AND M. B. MASCHLER (1995): *Repeated Games With Incomplete Information*. MIT Press. [2632, 2644]

BACK, K., R. LIU, AND A. TEGUIA (2020): "Signaling in otc Markets: Benefits and Costs of Transparency," *Journal of Financial and Quantitative Analysis*, 55 (1), 47–75. [2633]

BALZER, B., AND J. SCHNEIDER (2019): "Belief Management and Optimal Arbitration," Working Paper. [2634]

BERGEMANN, D., AND S. MORRIS (2016a): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11 (2), 487–522. [2636,2659]

——— (2016b): "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium," *American Economic Review*, 106 (5), 586–591. [2632]

BERGEMANN, D., AND A. WAMBACH (2015): "Sequential Information Disclosure in Auctions," *Journal of Economic Theory*, 159, 1074–1095. [2634]

BORDER, K. C. (1991): "Implementation of Reduced Form Auctions: A Geometric Approach," *Econometrica*, 59 (4), 1175–1187. [2647]

CALZOLARI, G., AND A. PAVAN (2006a): "Monopoly With Resale," *RAND Journal of Economics*, 37 (2), 362–375. [2633,2637,2638,2657,2658]

——— (2006b): "On the Optimality of Privacy in Sequential Contracting," *Journal of Economic Theory*, 130 (1), 168–204. [2633,2657]

——— (2009): "Sequential Contracting With Multiple Principals," *Journal of Economic Theory*, 144 (2), 503–531. [2633]

CARROLL, G., AND I. SEGAL (2019): "Robustly Optimal Auctions With Unknown Resale Opportunities," *Review of Economic Studies*, 86 (4), 1527–1555. [2634]

DAS VARMA, G. (2003): "Bidding for a Process Innovation Under Alternative Modes of Competition," *International Journal of Industrial Organization*, 21 (1), 15–37. [2633]

DEQUIEDT, V., AND D. MARTIMORT (2015): "Vertical Contracting With Informational Opportunism," *American Economic Review*, 105 (7), 2141–2182. [2655]

DWORCZAK, P. (2015): "The Effects of Post-Auction Bargaining Between Bidders," Available at SSRN, https://ssrn.com/abstract=2550653. [2633]

——— (2020): "Mechanism Design With Aftermarkets: Optimal Mechanisms Under Binary Actions," Available at SSRN, https://ssrn.com/abstract=2892083. [2632]

——— (2020): "Supplement to 'Mechanism Design With Aftermarkets: Cutoff Mechanisms'," *Econometrica Supplemental Material*, 88, https://doi.org/10.3982/ECTA15768. [2632]

DWORCZAK, P., AND A. L. ZHANG (2017): "Implementability, Walrasian Equilibria, and Efficient Matchings," *Economics Letters*, 153 (Supplement C), 57–60. [2636]

ENGELBRECHT-WIGGANS, R., AND C. M. KAHN (1991): "Protecting the Winner. Second-Price versus Oral Auctions," *Economics Letters*, 35, 243–248. [2633]

ESŐ, P., AND B. SZENTES (2007): "Optimal Information Disclosure in Auctions and the Handicap Auction," *Review of Economic Studies*, 74 (3), 705–731. [2634]

GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): "On the Equivalence of Bayesian and Dominant Strategy Implementation," *Econometrica*, 81 (1), 197–220. [2648]

GIOVANNONI, F., AND M. MAKRIS (2014): "Reputational Bidding," *International Economic Review*, 55 (3), 693–710. [2633]

GOEREE, J. (2003): "Bidding for the Future: Signaling in Auctions With an Aftermarket," *Journal of Economic Theory*, 108 (2), 345–364. [2633,2637,2638,2654,2658]

GUPTA, M., AND B. LEBRUN (1999): "First Price Auctions With Resale," *Economics Letters*, 64, 181–185. [2634]

HAFALIR, I., AND V. KRISHNA (2008): "Asymmetric Auctions With Resale," *American Economic Review*, 98 (1), 87–112. [2634]

——— (2009): "Revenue and Efficiency Effects of Resale in First-Price Auctions," *Journal of Mathematical Economics*, 45 (9–10), 589–602. [2634]

HAILE, P. A. (2003): "Auctions With Private Uncertainty and Resale Opportunities," *Journal of Economic Theory*, 108 (1), 72–110. [2634]

HU, A. X., AND J. ZHANG (2017): "Optimal Mechanism Design With Aftermarket Competitions," Working Paper. [2633,2637,2638,2654,2658]

JEHIEL, P., AND B. MOLDOVANU (2001): "Efficient Design With Interdependent Valuations," *Econometrica*, 69 (5), 1237–1259. [2634]

JEHIEL, P., B. MOLDOVANU, AND E. STACCHETTI (1996): "How (not) to Sell Nuclear Weapons," *The American Economic Review*, 86 (4), 814–829. [2634]

KAMENICA, E., AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101 (6), 2590–2615. [2632,2644,2650]

KATZMAN, B. E., AND M. RHODES-KROPF (2008): "The Consequences of Information Revealed in Auctions," *Applied Economics Research Bulletin*, 2, 53–87. [2633,2637,2638]

KOLOTILIN, A., T. MYLOVANOV, A. ZAPECHELNYUK, AND M. LI (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85 (6), 1949–1964. [2632]

KOS, N., AND M. MESSNER (2013): "Extremal Incentive Compatible Transfers," *Journal of Economic Theory*, 148 (1), 134–164. [2636]

LAUERMANN, S., AND G. VIRÁG (2012): "Auctions in Markets: Common Outside Options and the Continuation Value Effect," *American Economic Journal: Microeconomics*, 4 (4). [2634]

LI, H., AND X. SHI (2017): "Discriminatory Information Disclosure," *American Economic Review*, 107 (11), 3363–3385. [2634]

MANELLI, A. M., AND D. R. VINCENT (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model," *Econometrica*, 78 (6), 1905–1938. [2648]

MATTHEWS, S. A. (1984): "On the Implementability of Reduced Form Auctions," *Econometrica*, 52 (6), 1519–1522. [2647]

MILGROM, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press. [2636]

MILGROM, P., AND R. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122. [2634]
MOLNÁR, J., AND G. VIRÁG (2008): "Revenue Maximizing Auctions With Market Interaction and Signaling," *Economics Letters*, 99 (2), 360–363. [2633,2638,2651]
MYERSON, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6 (1), 58–73. [2632]
MYERSON, R. B. (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10 (1), 67–81. [2636]
ROCHET, J.-C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16 (2), 191–200. [2641]
SIFMA (2016): "SIFMA Electronic Bond Trading Report: US Corporate and Municipal Securities," Report. [2630]
ZHANG, J., AND R. WANG (2013): "Optimal Mechanism Design With Resale via Bargaining," *Journal of Economic Theory*, 148 (5), 2096–2123. [2634]
ZHENG, C. Z. (2002): "Optimal Auctions With Resale," *Econometrica*, 70 (6). [2634]

Co-editor Dirk Bergemann handled this manuscript.

Manuscript received 16 October, 2017; final version accepted 3 June, 2020; available online 10 June, 2020.