Revisiting the Foundations of Dominant-Strategy Mechanisms*

Yi-Chun Chen† Jiangtao Li‡

November 28, 2017

Abstract

An important question in mechanism design is whether there is any theoretical foundation for the use of dominant-strategy mechanisms. This paper studies the maxmin and Bayesian foundations of dominant-strategy mechanisms in general social choice environments with quasi-linear preferences and private values. We propose a condition called the uniform shortest-path tree that, under regularity, ensures the foundations of dominant-strategy mechanisms. This exposes the underlying logic of the existence of such foundations in the single-unit auction setting, and extends the argument to cases where it was hitherto unknown. To prove this result, we adopt the linear programming approach to mechanism design. In settings in which the uniform shortest-path tree condition is violated, maxmin/ Bayesian foundations might not exist. We illustrate this by two examples: bilateral trade with ex ante unidentified traders and auction with type-dependent outside option.

*We are indebted to Jeff Ely for bringing us to this topic and Stephen Morris for encouragement. We thank Rakesh Vohra for detailed comments that substantially improved the paper. We further wish to thank Tilman Börgers, Gabriel Carroll, Kim-Sau Chung, Takashi Kunimoto, Satoru Takahashi, Takuro Yamashita, and many seminar audiences for their valuable comments. We are grateful to the Associate Editor and anonymous referees for the comments and suggestions. Part of this paper was written while Li was visiting the University of Michigan, and he would like to thank the institution for its hospitality and support. We acknowledge the financial support from the Singapore Ministry of Education Academic Research Fund Tier 1. Li acknowledges the financial support of the President's Graduate Fellowship from the National University of Singapore. All remaining errors are our own.

†Department of Economics, National University of Singapore, ecsycc@nus.edu.sg

‡School of Economics, University of New South Wales, jiangtao.li@unsw.edu.au

Contents

1 Introduction

Wilson (1987) criticizes applied game theory's reliance on common knowledge assumptions. In reaction to Wilson's critique, the literature of mechanism design has adopted the goal of finding detail-free mechanisms in order to eliminate this reliance.[1] The usual approach is to adopt stronger solution concepts, such as dominant-strategy mechanisms. A dominant-strategy mechanism does not rely on any assumptions of agents' beliefs and is robust to changes in agents' beliefs. However, dominant-strategy mechanisms constitute just one special class of detail-free mechanisms. A fundamental issue is to justify the leap from detail-free mechanisms in general to dominant-strategy mechanisms in particular.

Suppose that a revenue-maximizing mechanism designer has an estimate of the distribution of the agents' payoff types, but she does not have any reliable information about the agents' beliefs (including their beliefs about one another's payoff types, their beliefs about these beliefs, etc.), as these are arguably never observed. The mechanism designer ranks mechanisms according to their worst-case performance - the minimum expected revenue - where the minimum is taken over all possible agents' beliefs. The use of dominant-strategy mechanisms has a maxmin foundation if the mechanism designer finds it optimal to use a dominant-strategy mechanism.

A closely related notion is the Bayesian foundation. The use of dominant-strategy mechanisms is said to have a Bayesian foundation if there exists a particular assumption about (the distribution of) the agents' beliefs, against which the optimal dominant-strategy mechanism achieves the highest expected revenue among all detail-free mechanisms. Note that if there exists such an assumption, then the worst-case expected revenue of an arbitrary detail-free mechanism obviously cannot exceed its expected revenue against this particular assumption, which in turn cannot exceed the worst-case expected revenue of the optimal dominant-strategy mechanism. Therefore, the Bayesian foundation is a stronger notion than the maxmin foundation.

In the context of a revenue-maximizing auctioneer, Chung and Ely (2007) show that, under a regularity condition on the distribution of the bidders' valuations, the use of dominant-

[1]A detail-free mechanism does not rely on implicit assumptions about higher-order beliefs, and is Bayesian incentive compatible for all belief hierarchies. In other words, a detail-free mechanism is Bayesian incentive compatible relative to the universal type space; see Definition 4 for the formal definition.

strategy mechanisms has maxmin and Bayesian foundations. What has been missing thus far from the literature on mechanism design is the study of such foundations in general environments. In this paper, we study the maxmin and Bayesian foundations in general social choice environments with quasi-linear preferences and private values. This exposes the underlying logic of the existence of such foundations in the single-unit auction setting, and extends the argument to cases where it was hitherto unknown.

Our result builds on the recent literature on the network approach to mechanism design, in particular, Rochet and Stole (2003), Heydenreich, Müller, Uetz, and Vohra (2009), Vohra (2011), Kos and Messner (2013), and Sher and Vohra (2015).[2] We formulate the optimal mechanism design question as a network flow problem, and the optimization problem reduces to determining the shortest-path tree (the union of all shortest-paths from the source to all nodes) in this network. We say that there is a uniform shortest-path tree if for each agent, the shortest-path tree is the same for all dominant-strategy implementable decision rules and other agents' reports.

We show that under an additional regularity condition, the existence of a uniform shortest-path tree ensures the maxmin and Bayesian foundations of dominant-strategy mechanisms (Theorem 1). The uniform shortest-path tree is largely responsible for the success of mechanism design in numerous applications across various fields. Loosely speaking, the same features that make optimal mechanism design tractable also provide maxmin and Bayesian foundations for the use of dominant-strategy mechanisms. To prove this result, we adopt the linear programming approach to mechanism design, which exposes the underlying logic behind the existence of such foundations.[3] In particular, this gives us a recipe for constructing the assumption about (the distribution of) the agents' beliefs for the Bayesian foundation.

The uniform shortest-path tree condition is of interest because a number of resource allocation problems satisfy this condition. We examine its applicability in several prominent environments. First, the uniform shortest-path tree condition is satisfied in environments with one-dimensional types. This fits many classical applications of mechanism design, including

[2] Also see Rochet (1987), Gui, Müller, and Vohra (2004), and Müller, Perea, and Wolf (2007).

[3] We are indebted to Rakesh Vohra for suggesting the linear programming approach to us. In a recent paper, Sher and Vohra (2015) use the linear programming approach to study price discrimination when a buyer may present evidence relevant to her value.

single-unit auction (e.g., Myerson (1981)), public good (e.g., Mailath and Postlewaite (1990)), and standard bilateral trade (e.g., Myerson and Satterthwaite (1983)). The uniform shortest-path tree condition also holds in multi-unit auctions with homogeneous or heterogeneous goods, combinatorial auctions and the like, as long as the agents' private values are one-dimensional. In such a case, the payoff types are linearly ordered via a single path. Second, the uniform shortest-path tree condition can also be satisfied in some multi-dimensional environments. In particular, we consider the multi-unit auction with capacity-constrained bidders (see Malakhov and Vohra (2009)). In this case, the agent's payoff types are located on different paths and are only partially ordered. For both applications, we provide primitive conditions for regularity.

When the uniform shortest-path tree condition is violated, maxmin/ Bayesian foundations might not exist. We stress that the notion of no maxmin foundation is remarkably strong. No maxmin foundation means that there exists a *single* Bayesian mechanism that achieves *strictly* higher expected revenue than the optimal dominant-strategy mechanism, *regardless of the agents' beliefs*. Theorem 2 shows that if the optimal dominant-strategy mechanism exhibits certain properties, we could explicitly construct such a superior Bayesian mechanism. To the best of our knowledge, this is the first example of a revenue maximization setting in which the use of dominant-strategy mechanisms does not have a maxmin foundation.[4] We apply this result to bilateral trade with ex ante unidentified traders and auction with type-dependent outside option.

We hasten to emphasize that our analysis on the foundations of dominant-strategy mechanisms is constrained to the notions of maxmin and Bayesian foundations. There are, of course, other notions of optimality to examine the foundations of dominant-strategy mechanisms. We view the maxmin foundation as *the minimum requirement* that the optimal mechanism needs to satisfy. Indeed, if the use of dominant-strategy mechanisms does not have a maxmin foundation, then by definition, there exists a single Bayesian mechanism that achieves strictly higher expected revenue regardless of the agents' beliefs. Consequently, it becomes problematic to rationalize the use of dominant-strategy mechanisms. In settings in which the uniform shortest-path tree condition is violated, dominant-strategy mechanisms

[4]Chung and Ely (2007, Proposition 2) construct an example in which the Bayesian foundation does not exist, but their construction is silent about the existence of the maxmin foundation. Bergemann and Morris (2005) study an implementability problem. Börgers (2017) adopts a different notion of optimality.

may not even satisfy *the minimum requirement* of the maxmin foundation.

A natural question that our paper leaves open is whether there are other mechanisms that have a maxmin foundation, and among all mechanisms that have a maxmin foundation, also achieve other desirable properties. We don't yet know whether the optimistic view of dominant strategy mechanisms (in settings in which the use of dominant-strategy mechanisms has a maxmin foundation) will survive once the research agenda is completed. Notably, Börgers (2017) constructs another mechanism that *weakly dominates* the optimal dominant-strategy mechanism, in the sense that it never yields lower revenue and sometimes yields strictly higher revenue.[5]

The remainder of this introduction discusses some related literature. Section 2 presents the notations, concepts, and the model. Section 3 formulates the notion of the uniform shortest-path tree and presents the results. Section 4 studies applications of the results. Section 5 concludes.

1.1 Related literature

In a seminal paper, Bergemann and Morris (2005) ask whether a fixed social choice correspondence - mapping payoff type profiles to sets of possible allocations - can or cannot be robustly partially implemented. Thus they focus on a "yes or no" question. In contrast, we consider the objective of revenue maximization for the mechanism designer (under her estimate about the distribution of the agents' payoff types), allowing all possible beliefs and higher-order beliefs of the agents. The best mechanism from the point of view of the mechanism designer will in general not be separable, and thus the results of Bergemann and Morris (2005) do not apply.

This paper joins a growing literature exploring mechanism design with worst-case objectives. This includes the seminal work Chung and Ely (2007), and more recently, Carroll

[5]The construction builds on the possibility of side bets among agents, and the mechanism designer charges a small fee for each bet. As Börgers points out, the argument would not be valid if the mechanism designer restricts her attention to the type spaces characterized by Morris (1994), which do not allow speculative trades. Whether there is a mechanism that *weakly dominates* the optimal dominant-strategy mechanism when we consider only type spaces that do not allow profitable side bets is an interesting question for further research.

(2015, 2017), Yamashita (2015, 2017), Yamashita and Zhu (2017), and Du (2017), among many others.

Another recent line of literature studies the equivalence of Bayesian and dominant-strategy mechanisms; see, for example, Manelli and Vincent (2010), Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013), and Goeree and Kushnir (2017). An important difference from our work is that these papers focus on independent types.

2 Preliminaries

2.1 Notation

There is a finite set $\mathcal{I} = \{1, 2, \dots, I\}$ of risk-neutral agents and a finite set $\mathcal{K} = \{1, 2, \dots, K\}$ of alternatives. Agent i's payoff type $v_i \in \mathbb{R}^K$ represents her gross utility under the K alternatives.[6] We write $v_i(k)$ to denote agent i's gross utility under alternative k, if her payoff type is v_i. The set of possible payoff types of agent i is a finite set $V_i \subset \mathbb{R}^K$. The set of possible payoff type profiles is $V = \Pi_{i\in\mathcal{I}} V_i$ with generic payoff type profile $v = (v_1, v_2, \dots, v_I)$. We write v_{-i} for a payoff type profile of agent i's opponents, i.e., $v_{-i} \in V_{-i} = \Pi_{j\neq i} V_j$. If Y is a measurable space, then ΔY is the set of all probability measures on Y. If Y is a metric space, then we treat it as a measurable space with its Borel σ-algebra.

2.2 Types

We follow the standard approach to model agents' information using a type space. A type space, denoted $\Omega = (\Omega_i, f_i, g_i)_{i\in\mathcal{I}}$, is defined by a measurable space of types Ω_i for each agent, and a pair of measurable mappings $f_i : \Omega_i \to V_i$, defining the payoff type of each type, and $g_i : \Omega_i \to \Delta(\Omega_{-i})$, defining each type's belief about the types of the other agents.

A type space encodes in a parsimonious way the beliefs and all higher-order beliefs of the agents. One simple kind of type space is the naive type space generated by a payoff type distribution $\pi \in \Delta(V)$. In the naive type space, each agent believes that all agents' payoff types are drawn from the distribution π, and this is common knowledge. Formally, a naive type space associated with π is a type space $\Omega^\pi = (\Omega_i, f_i, g_i)_{i\in\mathcal{I}}$ such that $\Omega_i = V_i$, $f_i(v_i) = v_i$,

[6] We may represent the agent's payoff types in different ways. For instance, when studying one-dimensional payoff types (Section 4.1), it is more convenient to represent agent i's payoff type by $v_i \in \mathbb{R}$.

and $g_i(v_i)[v_{-i}] = \pi(v_{-i}|v_i)$ for every v_i and v_{-i}. The naive type space is used almost without exception in auction theory and mechanism design. The cost of this parsimonious model is that it implicitly embeds some strong assumptions about the agents' beliefs, and these assumptions are not innocuous. For example, if the agents' payoff types are independent under π, then in the naive type space, the agents' beliefs are common knowledge. On the other hand, for a generic π, it is common knowledge that there is a one-to-one correspondence between payoff types and beliefs. Myerson (1981) characterizes the optimal auction in the independent case and Crémer and McLean (1988) in the other case. Which of these cases holds makes a big difference for the structure and welfare properties of the optimal auction. The spirit of the Wilson Doctrine is to avoid making such assumptions.

To implement the Wilson Doctrine, the common approach is to maintain the naive type space, but try to diminish its adverse effect by imposing stronger solution concepts. To provide foundations for this methodology, we have to return to the fundamentals. Formally, weaker assumptions about the agents' beliefs are captured by larger type spaces. Indeed, we can remove these assumptions altogether by allowing for every conceivable hierarchy of higher-order beliefs. By the results of Mertens and Zamir (1985), there exists a universal type space, $\Omega^* = (\Omega_i^*, f_i^*, g_i^*)_{i \in \mathcal{I}}$, with the property that, for every payoff type v_i and every infinite hierarchy of beliefs $\hat{h}_i$, there is a type $\omega_i \in \Omega_i^*$ of agent i with payoff type v_i and whose hierarchy is $\hat{h}_i$. Moreover, each Ω_i^* is a compact topological space.[7]

When we start with the universal type space, we remove any implicit assumptions about the agents' beliefs. We can now explicitly model any such assumption as a probability distribution over the agents' universal types. Specifically, an assumption for the mechanism designer is a distribution μ over Ω^*.

2.3 Mechanisms

A mechanism consists of a set of messages M_i for each agent i, a decision rule $p : M \to \Delta\mathcal{K}$ and payment functions $t_i : M \to \mathbb{R}$. Each agent i selects a message from M_i. Based on the resulting profile of messages m, the decision rule p specifies the outcome from $\Delta\mathcal{K}$ (lotteries are allowed) and the payment function t_i specifies the transfer from agent i to the mechanism designer. Agent i obtains utility $p \cdot v_i - t_i$. We write p^k for the probability that

[7] Also see Heifetz and Neeman (2006).

alternative k is chosen. We find it convenient to work with direct-revelation mechanisms. In a direct-revelation mechanism for type space $\Omega = (\Omega_i, f_i, g_i)_{i\in\mathcal{I}}$, $M_i = \Omega_i$; that is, agents directly report their types to the mechanism designer.

The mechanism defines a game form, which together with the type space constitutes a game of incomplete information. The mechanism design problem is to fix a solution concept and search for the mechanism that delivers the maximum expected revenue for the mechanism designer in some outcome consistent with the solution concept. To implement the Wilson Doctrine and minimize the role of assumptions built into the naive type space, the common approach is to adopt a strong solution concept which does not rely on these assumptions. In practice, the solution concept that is often used for this purpose is dominant-strategy equilibrium. The revelation principle holds, and we can restrict attention to direct mechanisms.

Definition 1. *A direct-revelation mechanism Γ for type space Ω is dominant-strategy incentive compatible (dsIC) if for each agent i and type profile $\omega \in \Omega$,*

$$\begin{aligned} p(\omega)\cdot f_i(\omega_i) - t_i(\omega) &\geq 0, \text{ and} \\ p(\omega)\cdot f_i(\omega_i) - t_i(\omega) &\geq p(\omega_i', \omega_{-i})\cdot f_i(\omega_i) - t_i(\omega_i', \omega_{-i}), \end{aligned}$$

for any alternative type $\omega_i' \in \Omega_i$.

Definition 2. *A dominant-strategy mechanism is a dsIC direct-revelation mechanism for the naive type space Ω^π. We denote by Φ the class of all dominant-strategy mechanisms.*

To provide a foundation for using dominant-strategy mechanisms, we shall compare it to the route of completely eliminating common knowledge assumptions about beliefs. We maintain the standard solution concept of Bayesian equilibrium, but now we enlarge the type space all the way to the universal type space. By the revelation principle, we restrict attention to direct mechanisms.

Definition 3. *A direct-revelation mechanism Γ for type space $\Omega = (\Omega_i, f_i, g_i)$ is Bayesian incentive compatible (BIC) if for each agent i and type $\omega_i \in \Omega_i$,*

$$\begin{aligned} \int_{\Omega_{-i}} \left(p(\omega)\cdot f_i(\omega_i) - t_i(\omega)\right) g_i(\omega_i) d\omega_{-i} &\geq 0, \text{ and} \\ \int_{\Omega_{-i}} \left(p(\omega)\cdot f_i(\omega_i) - t_i(\omega)\right) g_i(\omega_i) d\omega_{-i} &\geq \int_{\Omega_{-i}} \left(p(\omega_i', \omega_{-i})\cdot f_i(\omega_i) - t_i(\omega_i', \omega_{-i})\right) g_i(\omega_i) d\omega_{-i} \end{aligned}$$

for any alternative type $\omega_i' \in \Omega_i$.

A mechanism, which does not rely on implicit assumptions about higher-order beliefs, should be incentive compatible for all belief hierarchies. In other words, it should be BIC relative to the universal type space.

Definition 4. *Let Ψ be the class of all BIC direct-revelation mechanism for the universal type space. We say that such a mechanism is detail-free.*

For simplicity of exposition, we add a dummy type v_0 for each agent $i \in \mathcal{I}$ and set $p(v_0, v_{-i}) \cdot v_i = t_i(v_0, v_{-i}) = 0$ for all $v_i \in V_i, v_{-i} \in V_{-i}$.

2.4 The mechanism designer as a maxmin decision maker

The mechanism designer has an estimate of the distribution of the agents' payoff types, π. Following Chung and Ely (2007), we assume that π has full support. An assumption μ about the distribution of the payoff types and beliefs of the agents is consistent with this estimate if the induced marginal distribution on V is π. Let $\mathcal{M}(\pi)$ denote the compact subset of such assumptions. For any mechanism Γ, the μ-expected revenue of Γ is

$$R_\mu(\Gamma) = \int_{\Omega^*} \sum_{i \in \mathcal{I}} t_i(\omega) d\mu(\omega).$$

We do not assume that the mechanism designer has confidence in the naive type space as his model of agents' beliefs. Rather he considers other assumptions within the set $\mathcal{M}(\pi)$ as possible as well. The mechanism designer who chooses a mechanism that maximizes the worst-case performance solves the maxmin problem of

$$\sup_{\Gamma \in \Psi} \inf_{\mu \in \mathcal{M}(\pi)} R_\mu(\Gamma).$$

If the mechanism designer used an optimal dominant-strategy mechanism, then his revenue would be

$$\Pi^D(\pi) = \sup_{\Gamma \in \Phi} R_\pi(\Gamma),$$

where

$$R_\pi(\Gamma) = \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v)$$

for any dominant-strategy mechanism $\Gamma \in \Phi$.

Definition 5. *The use of dominant-strategy mechanisms has a maxmin foundation if*

$$\Pi^D(\pi) = \sup_{\Gamma \in \Psi} \inf_{\mu \in \mathcal{M}(\pi)} R_\mu(\Gamma).$$

The use of dominant-strategy mechanisms has a Bayesian foundation if for some belief $\mu^* \in \mathcal{M}(\pi)$,

$$\Pi^D(\pi) = \sup_{\Gamma \in \Psi} R_{\mu^*}(\Gamma).$$

The Bayesian foundation is a stronger notion than the maxmin foundation. The Bayesian foundation says that there exists an assumption about (the distribution of) agents' beliefs, against which the optimal dominant-strategy mechanism achieves the highest expected revenue among all detail-free mechanisms. It follows that the worst-case expected revenue of an arbitrary detail-free mechanism cannot exceed its expected revenue against this particular assumption, which in turn cannot exceed the worst-case expected revenue of the optimal dominant-strategy mechanism. We record this observation as the following proposition.[8]

Proposition 1. *If the use of dominant-strategy mechanisms has a Bayesian foundation, then it has a maxmin foundation.*

3 Results

We can formulate the optimal dominant-strategy mechanism design problem as follows:

$$\max_{p(\cdot),\, t_i(\cdot)} \quad \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v) \tag{DIC}$$

$$\text{subject to } \forall i \in \mathcal{I},\, \forall v_i \in V_i,\, \forall v_i' \in V_i \cup \{v_0\},\, \forall v_{-i} \in V_{-i},$$

$$p(v_i, v_{-i}) \cdot v_i - t_i(v_i, v_{-i}) \geq p(v_i', v_{-i}) \cdot v_i - t_i(v_i', v_{-i}), \tag{1}$$

$$\forall v \in V,\, \forall k \in \mathcal{K},\, p^k(v) \geq 0, \tag{2}$$

$$\forall v \in V,\, \sum_{k \in \mathcal{K}} p^k(v) = 1. \tag{3}$$

In the optimal dominant-strategy mechanism design problem (DIC), the mechanism designer maximizes expected revenue by choosing (p, t), subject to incentive constraints (1) and feasibility constraints (2) - (3). We denote by V_{DIC} the value of the objective function of the program (DIC) at an optimum.

Say that a decision rule p is dsIC if there exists transfer scheme t such that the mechanism (p, t) satisfies the incentive constraints (1). We omit the proof of the following standard lemma, due to Rochet (1987).

[8] Also see Chung and Ely (2007, Section 2.5).

Lemma 1. *A necessary and sufficient condition for a decision rule p to be dsIC is the following cyclical monotonicity condition:* $\forall i \in \mathcal{I}$, $\forall v_{-i} \in V_{-i}$, *and for every sequence of payoff types of agent i,* $(v_{i,1}, v_{i,2}, \ldots, v_{i,t})$ *with* $v_{i,t} = v_{i,1}$, *we have*

$$\sum_{n=1}^{t-1} \Big[p(v_{i,n}, v_{-i}) \cdot v_{i,n+1} - p(v_{i,n}, v_{-i}) \cdot v_{i,n} \Big] \leq 0. \tag{4}$$

3.1 Uniform shortest-path tree

We first collect some graph-theoretic terminology used in the sequel.

Definition 6. *Fix a dsIC decision rule p and other agents' reports v_{-i}. (1) The set of nodes for agent i is $V_i \cup \{v_0\}$; (2) For any $v_i \in V_i$ and $v_i' \in \{V_i \setminus \{v_i\}\} \cup \{v_0\}$, $v_i' \to v_i$ is a directed edge with length $p(v_i, v_{-i}) \cdot v_i - p(v_i', v_{-i}) \cdot v_i$; and (3) A path from the dummy type v_0 to payoff type $v_{i,t} \in V_i$ is a sequence $P = (v_0, v_{i,1}, v_{i,2}, \ldots, v_{i,t})$ of nodes where (i) $v_{i,j} \in V_i$, $\forall j = 1, 2, \ldots, t$; and (ii) $j \neq j' \implies v_{i,j} \neq v_{i,j'}$.*

We can decompose the mechanism designer's problem (DIC) into two steps. In the first step, we fix a decision rule p that is dsIC, and optimize over the transfer scheme t. In the second step, we optimal over the dsIC decision rule p.

Fix a dsIC decision rule p and other agents' reports v_{-i}. The optimization problem over the the transfer scheme t has a corresponding network flow problem that can be described in the following way. Introduce one node for each type $v_i \in V_i \cup \{v_0\}$ (the node corresponding to the dummy type v_0 will be the source), and to each directed edge $v_i' \to v_i$, assign a length of $p(v_i, v_{-i}) \cdot v_i - p(v_i', v_{-i}) \cdot v_i$. The optimization problem reduces to determining the shortest-path tree (the union of shortest-paths from the source to all nodes) in this network. Edges on the shortest-path tree correspond to binding dominant-strategy incentive constraints. Readers unfamiliar with network flows may consult Ahuja, Magnanti, and Orlin (1993) and Vohra (2011).

Definition 7. *Fix a dsIC decision rule p and other agents' reports v_{-i}. A shortest-path tree is the union of shortest-paths from the source to all nodes such that if v_i' belongs to the shortest-path from the source v_0 to some $v_i \in V_i$, the truncation of the path from v_0 to v_i' defines the shortest-path from v_0 to v_i'.*

Definition 8. *There is a uniform shortest-path tree if, for each agent $i \in \mathcal{I}$, there is the same shortest-path tree for all dsIC decision rules p and other agents' reports v_{-i}.*

When the uniform shortest-path tree condition is satisfied, we can drop the dependence of the shortest-path tree on p and v_{-i}. A uniform shortest-path tree induces a partial order on agents' payoff types. For a typical shortest-path $(v_0, v_{i,1}, v_{i,2}, \ldots, v_{i,t})$, we write $v_{i,t} \succ_i v_{i,t-1} \succ_i \ldots \succ_i v_{i,1} \succ_i v_0$. It is convenient to represent the uniform shortest-path tree of agent i using $\succ_i$ and its transitive closure by $\succ_i^+$. For notational convenience, we write $v_i' \succeq_i^+ v_i$ if $v_i' \succ_i^+ v_i$ or $v_i' = v_i$. If $v_i \succ_i v_i'$, we sometimes denote v_i' by v_i^-.

3.2 Foundations of dominant-strategy mechanisms

In this subsection, we consider environments in which the uniform shortest-path tree condition holds. For any dsIC decision rule p and other agents' reports v_{-i}, it suffices to only consider constraints that correspond to edges on the shortest-path tree. In settings in which the uniform shortest-path tree condition holds, the edges on the shortest-path tree are the same, for any dsIC decision rule p and other agents' reports v_{-i}. Therefore, in the grand maximization problem (DIC), it suffices to only consider constraints that correspond to these edges, subject to that the decision rule p satisfies the cyclical monotonicity constraint (4). As is standard in the mechanism design literature, we first consider the following relaxed problem, in which we ignore the cyclical monotonicity constraint. A regularity condition on π is then imposed to ensure that some optimal decision rule p^* that solves the relaxed maximization problem automatically satisfies the cyclical monotonicity constraint.

$$\max_{p(\cdot),\, t_i(\cdot)} \quad \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v) \tag{$DIC-P$}$$

$$\text{subject to } \forall i \in \mathcal{I},\, \forall v_i \in V_i,\, \forall v_{-i} \in V_{-i},$$

$$p(v_i, v_{-i}) \cdot v_i - t_i(v_i, v_{-i}) \geq p(v_i^-, v_{-i}) \cdot v_i - t_i(v_i^-, v_{-i}), \tag{5}$$

$$\forall v \in V,\, \forall k \in \mathcal{K},\, p^k(v) \geq 0, \tag{6}$$

$$\forall v \in V, \sum_{k \in \mathcal{K}} p^k(v) = 1. \tag{7}$$

By compactness arguments, the maximization problem $(DIC - P)$ has a finite optimal value. Denote by V_{DIC-P} the value of the objective function of the program $(DIC - P)$ at an optimum.

Definition 9. *We say that π is regular if the cyclical monotonicity constraint (4) is automatically satisfied for some p^* that solves the optimization problem $(DIC - P)$.*

To the best of our knowledge, there is no formal definition of regularity in the general environments. Our definition of regularity captures how it has been used in the literature; see for example, Myerson (1981). That is, we first ask which decision rule p the mechanism designer would choose if she does not have to make sure that the decision rule satisfies the cyclical monotonicity constraint. The regularity condition is then imposed to make sure that such optimal decision rule automatically satisfies the cyclical monotonicity constraint. In the applications that we study in Section 4.1 and Section 4.2, additional structure is imposed, and we provide primitive conditions for regularity.

Theorem 1 below provides a sufficient condition for the maxmin and Bayesian foundations of dominant-strategy mechanisms.

Theorem 1. *In environments in which the uniform shortest-path tree condition holds, if π is regular, then the use of dominant-strategy mechanisms has maxmin and Bayesian foundations.*

By Proposition 1, it suffices to show the Bayesian foundation. In other words, it suffices to identify one particular assumption about (the distribution of) agents' beliefs such that the optimal dominant-strategy mechanism achieves the highest expected revenue among all detail-free mechanisms. To prove this result, we adopt the linear programming approach to mechanism design. In particular, we make use of the duality theorem in linear programming. In what follows, we first provide a sketch of the proof, and then compare our proof technique with that of Chung and Ely (2007).

The structure of the proof is as follows. Step (1) considers the optimal dominant-strategy mechanism design problem. We shall work with the maximization problem $(DIC-P)$ and derive its dual $(DIC-D)$. This is without loss of generality, as the regularity condition ensures that $V_{DIC} = V_{DIC-P}$. Step (2) restricts attention to a subclass of type spaces, formulates the optimal Bayesian mechanism design problem $(BIC-P)$, and derives its dual $(BIC-D)$. We denote by V_{DIC-D} (resp. V_{BIC-P} and V_{BIC-D}) the value of the objective function of the program $(DIC-D)$ (resp. $(BIC-P)$ and $(BIC-D)$) at an optimum. Step (3) then explicitly constructs an assumption about (the distribution of) the agents' beliefs, against which we show that, $V_{DIC-D} \geq V_{BIC-D}$. It follows from the duality theorem in linear programming (see for example, Bradley, Hax, and Magnanti (1977, Chapter 4)) that

$V_{DIC-P} = V_{DIC-D} \geq V_{BIC-D} \geq V_{BIC-P}$.[9] Therefore, we have identified an assumption about (the distribution of) the agents' beliefs, against which $V_{DIC} \geq V_{BIC-P}$. The details of the proof are relegated to the Appendix.

We now compare our duality approach with that of Chung and Ely (2007). Both approaches seek to identify an assumption against which the optimal dominant-strategy mechanism achieves the highest expected revenue among all detail-free mechanisms. Chung and Ely (2007) work with the primal maximization problems, and compares the optimal dominant-strategy mechanism design problem and the optimal Bayesian mechanism design problem. They have to identify an assumption such that *no* feasible variables in the Bayesian mechanism design problem could generate a higher revenue than the optimal value of the dominant-strategy mechanism design problem. In particular, Chung and Ely (2007) have to explicitly identify the set of binding constraints under the optimal Bayesian mechanism design problem. This approach involves a complex set of algebra which makes the construction of the assumption less transparent. Furthermore, their analysis requires the discussions of different cases, namely, the nonsingular case and the singular case. They first show the foundation result in the nonsingular case and then rely on a limiting argument to establish the foundation result in the singular case.

Whereas in the duality approach that we take, we have the benefit of directly manipulating the dual variables. It suffices to identify an assumption such that *there exists* one set of dual variables of the Bayesian minimization problem that generates the same objective value as the optimal value of the dominant-strategy minimization problem. This greatly simplifies the analysis. We do not have to explicitly identify the set of binding constraints in the optimal Bayesian mechanism design problem. Furthermore, our approach dispenses the discussion of the nonsingular versus singular cases. The key part of the proof is to construct the assumption about (the distribution) of agents' beliefs. As we shall see in Step (3) of the proof, this assumption can be easily identified using the dual variables.

[9] To apply the strong duality theorem, one has to show that the optimal value of the programming problem is finite. This follows from compactness arguments in the case of the optimal dominant-strategy mechanism design problem. Thus, by the strong duality theorem, we have $V_{DIC-P} = V_{DIC-D}$. In the case of the optimal Bayesian mechanism design problem, it suffices to invoke the weak duality theorem for our purpose. Our arguments do not require (and actually show) that the optimal value of the optimal Bayesian mechanism design problem is finite.

Remark 1. *Our paper focuses on the private-value setting. The uniform shortest-path tree condition has a natural counterpart in the interdependent-value setting that, under an additional regularity condition, ensures the maxmin and Bayesian foundations of ex post incentive compatible mechanisms. We sketch below how to modify the notion of the uniform shortest-path tree to accommodate the interdependent-value setting. The proof of Theorem 1 can be easily adopted, and we won't repeat the arguments. In interdependent-value settings, agent i's gross utility under each of the alternatives depends on the payoff types of her opponents. When agent i has payoff type v_i, we use $v_i(v_{-i}) \in \mathbb{R}^K$ to denote her gross utility under the K alternatives, when her opponents have payoff types v_{-i}. Fix a dsIC decision rule p and other agents' reports v_{-i}, (1) the set of nodes for agent i is $V_i \cup \{v_0\}$; (2) For any $v_i \in V_i$ and $v_i' \in \{V_i \setminus \{v_i\}\} \cup \{v_0\}$, $v_i' \to v_i$ is a directed edge with length $p(v_i, v_{-i}) \cdot v_i(v_{-i}) - p(v_i', v_{-i}) \cdot v_i(v_{-i})$. The path, shortest-path tree, uniform shortest-path tree, and regularity condition can be defined accordingly.*

3.3 No foundations of dominant-strategy mechanisms

In this subsection, we show that when the uniform shortest-path tree condition is violated, maxmin/ Bayesian foundations might not exist. Note that in settings in which the uniform shortest-path tree condition is violated, it is difficult to find the optimal dominant-strategy mechanism, not to mention the construction of a superior Bayesian mechanism. To have a meaningful discussion, we shall take the optimal dominant-strategy mechanism (the binding structure, and payments of the agents) as primitives. While the conditions of the theorem may be restrictive, the conditions can be verified whenever the optimal dominant-strategy mechanism can be solved (possibly by a linear programming solver). We apply Theorem 2 to bilateral trade with ex ante unidentified traders in Section 4.3, and to auction with type-dependent outside option in Section 4.4.

Theorem 2. *In environments with two agents and binary payoff types for each agent, if for the optimal dominant-strategy mechanism*

	v_1	v_1'
v_2	$p(v_1, v_2), t_1(v_1, v_2), t_2(v_1, v_2)$	$p(v_1', v_2), t_1(v_1', v_2), t_2(v_1', v_2)$
v_2'	$p(v_1, v_2'), t_1(v_1, v_2'), t_2(v_1, v_2')$	$p(v_1', v_2'), t_1(v_1', v_2'), t_2(v_1', v_2')$

1) binding structure:

$$p(v_1, v_2') \cdot v_2 - t_2(v_1, v_2') < 0,$$
$$\textit{and } p(v_1', v_2') \cdot v_2 - t_2(v_1', v_2') > 0;$$

2) payment dominance:

$$t_1(v_1, v_2') + t_2(v_1, v_2') \geq t_1(v_1, v_2) + t_2(v_1, v_2),$$
$$\textit{and } t_1(v_1', v_2') + t_2(v_1', v_2') > t_1(v_1', v_2) + t_2(v_1', v_2),$$

then there is neither a Bayesian foundation nor a maxmin foundation.

Under the binding structure in Theorem 2, the uniform shortest-path tree condition is violated. Let

$$x = p(v_1, v_2) \cdot v_2 - t_2(v_1, v_2);$$
$$y = p(v_1', v_2) \cdot v_2 - t_2(v_1', v_2);$$
$$z = p(v_1, v_2') \cdot v_2 - t_2(v_1, v_2') < 0;$$
$$w = p(v_1', v_2') \cdot v_2 - t_2(v_1', v_2') > 0.$$

Since the optimal dominant-strategy mechanism necessarily satisfies the incentive constraints, we have $x \geq 0, y \geq w > 0$. Furthermore, it must be the case that $x = 0$, and $y = w$. Otherwise, the dominant-strategy mechanism would not have been optimal. In words, when agent 1's payoff type is v_1, the participation constraint of payoff type v_2 is binding ($x = 0$), and the incentive constraint corresponding to payoff type v_2 mimicking payoff type v_2' is not binding ($z < 0$). When agent 1's payoff type is v_1', the participation constraint of payoff type v_2 is not binding, and the incentive constraint corresponding to payoff type v_2 mimicking payoff type v_2' is binding ($y = w > 0$).

We show that there is no maxmin foundation. In particular, we explicitly construct a single Bayesian mechanism that does strictly better than the optimal dominant-strategy mechanism, regardless of the assumption about (the distribution of) the agents' beliefs. Since the Bayesian foundation is a stronger notion than the maxmin foundation (Proposition 1), this further implies that there is no Bayesian foundation.

We use a to denote the first-order belief of payoff type v_2 of agent 2 that agent 1 has payoff type v_1. In the optimal dominant-strategy mechanism, agent 1 has a dominant strategy to truthfully report her payoff type. We can think of agent 2 choosing between two lotteries. The first lottery is $p(v_1, v_2)$, $t_2(v_1, v_1)$ with probability a and $p(v_1', v_2)$, $t_2(v_1', v_2)$ with probability $1 - a$.[10] The second lottery is $p(v_1, v_2')$, $t_2(v_1, v_2')$ with probability a and $p(v_1', v_2')$, $t_2(v_1', v_2')$ with probability $1-a$. Agent 2 with payoff type v_2 prefers the first lottery, and agent 2 with payoff type v_2' prefers the second lottery.

Our argument centers around increasing the payment from agent 2 at the type profile (v_1', v_2). We amend the first lottery as follows: $p(v_1, v_2)$, $t_2(v_1, v_2)$ with probability a and $p(v_1', v_2)$, $t_2(v_1', v_2) + y$ with probability $1-a$. The second lottery is left unchanged. Although the increase in payment necessarily violates the dominant-strategy incentive constraints, for agent 2 with payoff type v_2, if she believes with sufficiently high probability ($a \geq \frac{z}{w-z}$) that agent 1's valuation is v_1, she prefers the amended first lottery, since it delivers expected utility of 0 and the second lottery delivers expected utility of $az + (1-a)w$. This weakly increased the designer' expected revenue, since $y > 0$. If agent 2 with payoff type v_2 believes with sufficiently low probability ($a < \frac{z}{w-z}$) that agent 1's valuation is v_1, she would now prefer the second lottery. But following the payment dominance condition in Theorem 2, this also weakly increases the seller's expected revenue. As agent 2 with payoff type v_2' prefers the second lottery to the original first lottery, she also prefers the second lottery to the amended lottery. Agent 1 still has a dominant strategy in the new mechanism. The proof below formally shows this intuition, and verifies the incentive constraints for the universal type space.

Proof of Theorem 2. We show that the mechanism designer could employ a single Bayesian mechanism and achieve a strictly higher expected revenue than he does using the optimal dominant-strategy mechanism, regardless of the agents' beliefs. We first explicitly identify one such mechanism and proceed by verifying that (1) the mechanism is BIC for the universal type space; and (2) the mechanism achieves a strictly higher expected revenue than the optimal dominant-strategy mechanism, regardless of the assumption about (the distribution of) the agents' beliefs.

[10] We ignore the payment from agent 1 to the mechanism designer for now, since such payment does not have any effect on agent 2's incentive.

Consider the following Bayesian mechanism Γ'. In this mechanism, the designer elicits agent 2's first-order belief about agent 1's payoff type.[11]

	v_1	v_1'
$v_2, a \in [0, \frac{w}{w-z})$	$p(v_1, v_2'), t_1(v_1, v_2'), t_2(v_1, v_2')$	$p(v_1', v_2'), t_1(v_1', v_2'), t_2(v_1', v_2')$
$v_2, a \in [\frac{w}{w-z}, 1]$	$p(v_1, v_2), t_1(v_1, v_2), t_2(v_1, v_2)$	$p(v_1', v_2), t_1(v_1', v_2), t_2(v_1', v_2) + y$
v_2'	$p(v_1, v_2'), t_1(v_1, v_2'), t_2(v_1, v_2')$	$p(v_1', v_2'), t_1(v_1', v_2'), t_2(v_1', v_2')$

To see that Γ' is BIC for the universal type space, note that

i truth telling continues to be a dominant strategy for agent 1;

ii truth telling continues to be a dominant strategy for payoff type v_2' of agent 2;

iii agent 2 with payoff type v_2 and $a \in [0, \frac{w}{w-z})$ will not announce v_2' as utility is unchanged;

iv agent 2 with payoff type v_2 and $a \in [\frac{w}{w-z}, 1]$ will not announce v_2' as expected utility is lower; and

v between $v_2, a \in [0, \frac{w}{w-z})$ and $v_2, a \in [\frac{w}{w-z}, 1]$, agent 2 with payoff type v_2 will announce $v_2, a \in [\frac{w}{w-z}, 1]$ if and only if $a \in [\frac{w}{w-z}, 1]$.

To see that Γ' achieves a strictly higher expected revenue than the optimal dominant-strategy mechanism, regardless of the assumption about (the distribution of) the agents' beliefs, note that

vi $t_1(v_1, v_2') + t_2(v_1, v_2') \geq t_1(v_1, v_2) + t_2(v_1, v_2)$;

vii $t_1(v_1', v_2') + t_2(v_1', v_2') > t_1(v_1', v_2) + t_2(v_1', v_2)$;

viii $y = w > 0$;

ix π has full support.

[11] An interesting aspect of the mechanism Γ' is that it suffices for the designer to elicit agent 2's first-order belief. In other words, agent 2's optimal strategic choice can be based on her first-order belief alone. There is no need for her to form higher-order beliefs, because such beliefs are irrelevant for her optimal choice. The mechanism Γ' is strategically simple in the terminology of Börgers and Li (2017).

□

Remark 2. *For ease of exposition, we state Theorem 2 in environments with two agents and binary payoff types for each agent. The argument extends to environments with multiple agents and each agent has multiple payoff types, as long as there are two agents and two payoff types for each agent, where the structure as stated in Theorem 2 exists.*

4 Applications

This section considers various applications of our results. The uniform shortest-path tree condition holds in the standard social choice environment with one-dimensional payoff types as well as some multi-dimensional environments. Section 4.1 applies Theorem 1 to environments with one-dimensional types, and Section 4.2 to a multi-dimensional environment. For both applications, we provide primitive conditions for regularity. Section 4.3 applies Theorem 2 to the bilateral trade model with ex ante unidentified traders, and Section 4.4 to auction with type-dependent outside option.

4.1 One-dimensional payoff types

In this subsection, we consider the standard social choice environment with one-dimensional payoff types.[12] This fits many classical applications of mechanism design, including single-unit auction (e.g., Myerson (1981)), public good (e.g., Mailath and Postlewaite (1990)), and standard bilateral trade (e.g., Myerson and Satterthwaite (1983)).

There is a finite set $\mathcal{I} = \{1, 2, \dots, I\}$ of risk neutral agents and a finite set $\mathcal{K} = \{1, 2, \dots, K\}$ of social alternatives. Agent i's gross utility in alternative k equals $u_i^k(v_i) = a_i^k v_i$, where $v_i \in \mathbb{R}$ is agent i's payoff type, $a_i^k \in \mathbb{R}$ are constants and $a_i^k \geq 0$ for all k. Agent i with payoff type v_i obtains utility

$$p(v) \cdot A_i v_i - t_i(v)$$

for decision rule $p \in \Delta\mathcal{K}$ and transfer t_i, where $A_i = (a_1^1, a_i^2, \dots, a_i^K)$. For notational simplicity, we assume that each agent has M possible payoff types, and that the set V_i is the same for

[12] This set-up covers the environment studied in Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013, Section 2).

each agent: $V_i = \{v^1, v^2, \ldots, v^M\}$, where $v^m - v^{m-1} = \gamma$ for each $m = 2, 3, \ldots, M$ and some $\gamma > 0$.

We can formulate the optimal dominant-strategy mechanism design problem as follows:

$$\max_{p(\cdot),\, t(\cdot)} \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v)$$

$$\text{subject to } \forall i \in \mathcal{I},\ \forall m, l = 1, 2, \ldots, M,\ \forall v_{-i} \in V_{-i},$$

$$p(v^m, v_{-i}) \cdot A_i v^m - t_i(v^m, v_{-i}) \geq 0, \tag{8}$$

$$p(v^m, v_{-i}) \cdot A_i v^m - t_i(v^m, v_{-i}) \geq p(v^l, v_{-i}) \cdot A_i v^m - t_i(v^l, v_{-i}). \tag{9}$$

$$\forall v \in V,\ \forall k \in \mathcal{K},\ p^k(v) \geq 0,$$

$$\forall v \in V, \sum_{k \in \mathcal{K}} p^k(v) = 1.$$

In the environment with one-dimensional payoff types, we say that a decision rule p is dsIC if there exists transfer scheme t such that the mechanism (p, t) satisfies the incentive constraints (8) and (9).

Uniform shortest-path tree condition is naturally satisfied in such settings. In particular, for any agent $i \in \mathcal{I}$, the payoff types are completely ordered via a single path. We omit the proof of the following standard lemma.

Lemma 2. *For any dsIC decision rule p and other agents' reports v_{-i}, the shortest path from the source v_0 to any payoff type $v^m \in V_i$ is*

$$v^m \succ_i v^{m-1} \succ_i \ldots \succ_i v^1 \succ_i v_0.$$

In what follows, we present the primitive condition for regularity. An equivalent formulation of the optimal dominant-strategy mechanism design problem is

$$\max_{p(\cdot),\, t(\cdot)} \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v)$$

$$\text{subject to } \forall i \in \mathcal{I},\ \forall m = 2, \ldots, M,\ \forall v_{-i} \in V_{-i},$$

$$p(v^1, v_{-i}) \cdot A_i v^1 - t_i(v^1, v_{-i}) = 0, \tag{10}$$

$$p(v^m, v_{-i}) \cdot A_i v^m - t_i(v^m, v_{-i}) = p(v^{m-1}, v_{-i}) \cdot A_i v^m - t_i(v^{m-1}, v_{-i}), \tag{11}$$

$$\forall m \geq l,\ p(v^m, v_{-i}) \cdot A_i \geq p(v^l, v_{-i}) \cdot A_i, \tag{12}$$

$$\forall v \in V,\ \forall k \in \mathcal{K},\ p^k(v) \geq 0,$$

$$\forall v \in V, \sum_{k \in \mathcal{K}} p^k(v) = 1,$$

where (10) and (11) are binding incentive constraints and (12) is the monotonicity constraint.[13] From (10) and (11), by induction, we have

$$t_i(v^m, v_{-i}) = p(v^m, v_{-i}) \cdot A_i v^m - \gamma \sum_{m'=1}^{m-1} p(v^{m'}, v_{-i}) \cdot A_i. \tag{13}$$

Let $F_i(v_i, v_{-i}) = \sum_{\hat{v}_i \leq v_i} \pi(\hat{v}_i, v_{-i})$ denote the cumulative distribution function of i's payoff type conditional on the other agents having payoff type profile v_{-i}. Define the virtual valuation of agent i as

$$r_i(v) = v_i - \gamma \frac{1 - F_i(v)}{\pi(v)}.$$

Using (13), we can rewrite the objective function as follows:

$$\begin{aligned} \sum_{v \in V} \pi(v) \sum_{i \in \mathcal{I}} t_i(v) &= \sum_{v \in V} \sum_{i \in \mathcal{I}} \pi(v) p(v) \cdot A_i r_i(v) \\ &= \sum_{v \in V} \pi(v) p(v) \cdot \sum_{i \in \mathcal{I}} A_i r_i(v). \end{aligned} \tag{14}$$

First ignore the monotonicity constraint (12). The optimization problem can be solved by pointwise maximization. The mechanism designer chooses decision rule p^* such that for each $v \in V$, p^* attaches positive probability to alternative k only if $k \in \arg\max_{k'} \sum_{i \in \mathcal{I}} a_i^{k'} r_i(v)$. We now impose a regularity condition that ensures that p^* automatically satisfies the monotonicity constraint (12). For each alternative k, let $K_i^{k,\inf}$ denote the collection of alternatives that agent i considers inferior to alternative k; that is, $K_i^{k,\inf} = \{k' \in \mathcal{K} : a_i^{k'} < a_i^k\}$.

Definition 10. *We say that π is regular if the virtual valuations satisfy the following condition: for each $v \in V$, $j \in \mathcal{I}$,*

$$k \in \arg\max_{k'} \sum_{i \in \mathcal{I}} a_i^{k'} r_i(v) \Rightarrow K_j^{k,\inf} \cap \arg\max_{k'} \sum_{i \in \mathcal{I}} a_i^{k'} r_i(\hat{v}_j, v_{-j}) = \emptyset \tag{15}$$

for every $\hat{v}_j > v_j$.

Remark 3. *The regularity condition (15) is the primitive condition for the regularity in Definition 9, when we restrict attention to environments with one-dimensional types. In the single-unit auction setting, our regularity condition (15) reduces to the regularity condition in Chung and Ely (2007).*

[13] It is well known that this is equivalent to cyclical monotonicity in environments with one-dimensional payoff types.

The regularity condition (15) ensures that under p^*, for any alternative k chosen with positive probability for payoff type profile (v^l, v_{-i}), when agent i's payoff type increases from v^l to v^m, alternatives that are inferior than alternative k from agent i's point of view will not be chosen. It must be that $p^*(v^m, v_{-i}) \cdot A_i \geq p^*(v^l, v_{-i}) \cdot A_i$ for $m \geq l$. The decision rule automatically satisfies the monotonicity constraint (12). The following corollary establishes the foundations of dominant-strategy mechanisms in environments with one-dimensional types.

Corollary 1. *If π satisfies the regularity condition (15), the use of dominant-strategy mechanisms has a Bayesian/ maxmin foundation.*

4.2 Multi-unit auction with capacity-constrained bidders

In addition to environments with one-dimensional payoff types, the uniform shortest-path tree condition is also satisfied in some multi-dimensional environments. Solving for the optimal dominant-strategy mechanism in a multi-dimensional environment is in general a daunting task. In this subsection, we examine a specific case where the multi-dimensional analysis can be simplified.

Consider the problem of finding the revenue maximizing auction when bidders have constant marginal valuations as well as capacity constraints.[14] Both the marginal values and capacity constraints are private information to the bidders. Bidder i's payoff type is represented by $v_i = (a, b)$, where a is the maximum amount she is willing to pay for each unit and b is the largest number of units she seeks. Units beyond the b^{th} unit are worthless. Let the range of a be $\mathcal{A} = \{1, 2, \ldots, A\}$ and the range of b be $\mathcal{B} = \{1, 2, \ldots, B\}$. The seller has Q units to sell.

A crucial assumption is that bidders cannot inflate the capacity but can shade it down. In other words, the auctioneer can verify, partially, the claims made by a bidder. Although this assumption seems odd in the selling context, it is natural in a procurement setting. Consider a procurement auction where the auctioneer wishes to procure Q units from bidders with constant marginal costs and limited capacity. No bidder will inflate his capacity when bidding because of the huge penalties associated with not being able to fulfill the order.

[14] Malakhov and Vohra (2009) studies the optimal Bayesian mechanism in such an environment, assuming independent types.

Equivalently, we may suppose that the designer can verify that claims that exceed capacity are false. The logic of our arguments can be easily adopted to deal with this case. More explicitly, for any payoff types v_i and v'_i, we introduce a directed edge $v'_i \to v_i$ only if v_i can pretend to be v'_i.

The following lemma is an analogous result of Malakhov and Vohra (2009, Theorem 6). While they focus on the optimal Bayesian mechanism in such an environment, we study the optimal dominant-strategy mechanism. Their proof technique can be adopted to establish the shortest-path tree for any dsIC decision rule p and other agents' reports v_{-i}.

Lemma 3. *For any dsIC decision rule p and other agents' reports v_{-i}, the shortest path from the source v_0 to any payoff type (a, b) is*

$$(a, b) \succ_i (a-1, b) \succ_i \ldots \succ_i (1, b) \succ_i (1, b-1) \succ_i \ldots \succ_i (1, 1) \succ_i v_0.$$

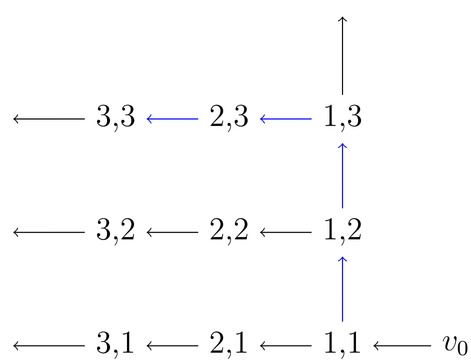

Figure 1: This figure illustrates the shortest path from the source v_0 to any payoff type (a, b). For example, the shortest-path from v_0 to $(3, 3)$ is: $(3, 3) \succ_i (2, 3) \succ_i (1, 3) \succ_i (1, 2) \succ_i (1, 1) \succ_i v_0$, highlighted in blue.

Let $F_{b,v_{-i}}(a) = \sum_{x=1}^{a} \pi((x, b), v_{-i})$.

Corollary 2. *If π satisfies the following regularity condition: $\forall v_{-i}, \forall (a, b) \geq (a', b')$,*

$$a - \frac{1 - F_{b,v_{-i}}(a)}{\pi((a, b), v_{-i})} \geq a' - \frac{1 - F_{b',v_{-i}}(a')}{\pi((a', b'), v_{-i})}, \tag{16}$$

then the use of dominant-strategy mechanisms has maxmin and Bayesian foundations.

The derivation of the regularity condition (16) is an analogous result of Malakhov and Vohra (2009, Theorem 7). When π is independent, the regularity condition (16) reduces to the regularity condition in Malakhov and Vohra (2009, Theorem 7).

4.3 Bilateral trade with ex ante unidentified traders

In this subsection, we apply Theorem 2 to the bilateral trade model with ex ante unidentified traders; see for example, Cramton, Gibbons, and Klemperer (1987) and Lu and Robert (2001). Each agent is endowed with $\frac{1}{2}$ unit of a good to be traded and has private information about her valuation for the good. Each agent may be either the buyer or the seller, depending on the realization of the privately observed information and the choice of the mechanism: the agent's role as the buyer or the seller is endogenously determined by her report and cannot be identified prior to trade. Agent 1's valuation for the good could be either 18 or 38. Agent 2's valuation for the good could be either 10 or 30.

A broker chooses trading mechanisms that maximize the expected profit; see for example, Myerson and Satterthwaite (1983, Section 5), Lu and Robert (2001), and Börgers (2015). The broker has the following estimate of the distribution of the agents' valuations:[15]

	$v_1 = 18$	$v_1 = 38$
$v_2 = 10$	$\frac{3}{8}$	$\frac{1}{8}$
$v_2 = 30$	$\frac{1}{8}$	$\frac{3}{8}$

Each trading mechanism can characterized by three outcomes functions (p, t_1, t_2), where $p(v_1, v_2)$ is the number of units agent 1 buys from agent 2, $t_1(v_1, v_2)$ is the payment from agent 1 to the broker, and $t_2(v_1, v_2)$ is the payment from agent 2 to the broker, if v_1 and v_2 are the reported valuations of agent 1 and agent 2. Agent 1's utility from purchasing p units of the good and paying a transfer t_1 is $pv_1 - t_1$, and agent 2's utility from selling p unit of the good and paying a transfer t_2 is $-pv_2 - t_2$, where $-\frac{1}{2} \leq p \leq \frac{1}{2}$. The broker chooses a trading mechanism that maximizes the expected profit.

Using a linear programming solver (MATLAB), we can solve for the optimal dominant-strategy mechanism Γ as follows, where the first number in each cell indicates the number of units agent 1 buys from agent 2, the second number is the transfer from agent 1 to the mechanism designer, and the third number is the transfer from agent 2 to the mechanism

[15]This example is robust to small perturbations in the agents' valuations or the broker's estimate of the distribution of the agents' valuations.

designer.

	$v_1 = 18$	$v_1 = 38$
$v_2 = 10$	$\frac{1}{2}, 9, -5$	$\frac{1}{2}, 9, -15$
$v_2 = 30$	$-\frac{1}{2}, -9, 15$	$\frac{1}{2}, 19, -15$

Following Theorem 2, there is neither a Bayesian foundation nor a maxmin foundation. The construction of a superior mechanism Γ' follows immediately from Theorem 2. We use a to denote the first-order belief of a low-valuation type of agent 2 that agent 1 has low valuation. In this mechanism, the mechanism designer elicits agent 2's first-order belief about agent 1's valuation.

	$v_1 = 18$	$v_1 = 38$
$v_2 = 10,\ a \in [0, \frac{1}{2})$	$-\frac{1}{2}, -9, 15$	$\frac{1}{2}, 19, -15$
$v_2 = 10,\ a \in [\frac{1}{2}, 1]$	$\frac{1}{2}, 9, -5$	$\frac{1}{2}, 9, -5$
$v_2 = 30$	$-\frac{1}{2}, -9, 15$	$\frac{1}{2}, 19, -15$

4.4 Auction with type-dependent outside option

Besides the bilateral trade model with ex ante unidentified traders (Section 4.3), we present here another environment to illustrate the usefulness of Theorem 2. A single unit of an indivisible object is up for sale. There are two risk-neutral bidders. Each bidder's payoff type is represented by $(a, b) \in \mathbb{R}^2_+$, where a is the maximum amount she is willing to pay, and b is the value of her outside option. Bidder 1's payoff type could be either $(20, 0)$ or $(40, 5)$. Bidder 2's payoff type could be either $(10, 0)$ or $(30, 5)$. The auctioneer has the following estimate of the distribution of the agents' payoff types:

	$v_1 = (20, 0)$	$v_1 = (40, 5)$
$v_2 = (10, 0)$	$\frac{3}{8}$	$\frac{1}{8}$
$v_2 = (30, 5)$	$\frac{1}{8}$	$\frac{3}{8}$

Using a linear programming solver (MATLAB), we can solve for the optimal dominant-strategy mechanism Γ as follows, where the first number in each cell indicates the probability that agent 1 gets the object, the second number is the probability that agent 2 gets the object, the third number is the transfer from agent 1 to the auctioneer, and the fourth number is the transfer from agent 2 to the auctioneer. Following Theorem 2, there is neither a Bayesian foundation nor a maxmin foundation.

	$v_1 = (20, 0)$	$v_1 = (40, 5)$
$v_2 = (10, 0)$	$1, 0, 20, 0$	$1, 0, 20, -5$
$v_2 = (30, 5)$	$0, 1, 0, 25$	$1, 0, 35, -5$

5 Conclusion

This paper revisits the maxmin and Bayesian foundations of dominant-strategy mechanisms in a general social choice environment with quasi-linear preferences and private values. We propose a notion of uniform shortest-path tree, that under regularity, ensures the maxmin and Bayesian foundations of dominant-strategy mechanisms. When the condition is violated, we show that maxmin/ Bayesian foundations might not exist, and we apply this no foundation result to bilateral trade with ex ante unidentified traders and auction with type dependent outside options.

In an independent and contemporaneous work, Yamashita and Zhu (2017) study the foundations of ex post incentive compatible mechanisms in the interdependent-value setting. They focus on a class of environments in which a single-crossing condition holds, and in which the order derived from the single-crossing condition is a total order. They show that, under "no ordinal interdependence" and certain regularity conditions, the use of ex post incentive compatible mechanisms has maxmin and Bayesian foundations. Their proof is a direct extension of Chung and Ely (2007) in the private-value setting to the interdependent-value environment. Our proof technique (duality approach) in Theorem 1 can be used to establish the foundation result in their paper; see Remark 1. Yamashita and Zhu (2017) also present a no-foundation result under certain conditions. Our no-foundation result does not imply their no-foundation result, and vice versa. In particular, our Theorem 2 and its applications show that even in private-value settings in which there is no (cardinal or ordinal) interdependence in payoffs, the maxmin and Bayesian foundations of dominant-strategy mechanisms might not exist.[16]

[16]We thank the Associate Editor for encouraging us to explore the relation between our work and that of Yamashita and Zhu (2017).

A Proof of Theorem 1

Step (1) First consider the optimal dominant-strategy mechanism design problem $(DIC-P)$. We derive its dual minimization problem $(DIC-D)$ as follows:

$$\min_{\lambda^{DIC}(v_i,v_{-i}),\,\mu^{DIC}(v)} \quad \sum_{v\in V} \mu^{DIC}(v) \tag{$DIC-D$}$$

$$\text{subject to } \forall i\in\mathcal{I},\ \forall v_i\in V_i,\ \forall v_{-i}\in V_{-i},$$

$$\lambda^{DIC}(v_i,v_{-i}) - \sum_{v_i':\, v_i'\succ_i v_i} \lambda^{DIC}(v_i',v_{-i}) = \pi(v_i,v_{-i}), \tag{17}$$

$$\forall v\in V,\ \forall k\in\mathcal{K},$$

$$\sum_{i\in\mathcal{I}} \lambda^{DIC}(v_i,v_{-i})v_i(k) - \sum_{i\in\mathcal{I}}\sum_{v_i':\, v_i'\succ_i v_i} \lambda^{DIC}(v_i',v_{-i})v_i'(k) \le \mu^{DIC}(v), \tag{18}$$

$$\forall i\in\mathcal{I},\ \forall v_i\in V_i, \forall v_{-i}\in V_{-i},$$

$$\lambda^{DIC}(v_i,v_{-i}) \ge 0, \tag{19}$$

where $\lambda^{DIC}(v_i,v_{-i})$ is the multiplier on the incentive constraint of preferring v_i over v_i^- when other agents report v_{-i} (see (5)), and $\mu^{DIC}(v)$ is the multiplier associated with the feasibility constraint of $\sum_{k\in\mathcal{K}} p^k(v)=1$ (see (7)). By induction, we can derive the following from (17):

$$\lambda^{DIC}(v_i,v_{-i}) = \sum_{\hat{v}_i:\, \hat{v}_i\succeq_i^+ v_i} \pi(\hat{v}_i,v_{-i}). \tag{20}$$

Step (2) Say that a type space is *simple* if for each agent $i\in\mathcal{I}$ and payoff type $v_i\in V_i$, there is a unique type for agent i with payoff type v_i. Let the set of types for agent i be equal to the set of possible payoff types; that is, $\Omega_i = V_i$. We take f_i to be the identity. For notational ease, we will write $\tau_i(\cdot|v_i) = g_i(v_i)[\cdot]$ for the belief of type v_i of agent i about the types of the other agents.

A simple type space can be embedded into the universal type space; see Chung and Ely (2007, Lemma 1). Indeed, the simple type space which we construct can be embedded into the universal type space by an embedding $\phi: V\to\Omega^*$. Since the designer's assumption is required to be consistent with her estimate π on V, the estimate induces a distribution/assumption $\pi\circ\phi^{-1}$ on the universal type space. From now on, we restrict attention to simple type spaces. This simplifies the analysis in the rest of the proof.

We can formulate the optimal Bayesian mechanism design problem as follows:

$$\max_{p(\cdot),\, t_i(\cdot)} \quad \sum_{v\in V} \pi(v) \sum_{i\in\mathcal{I}} t_i(v) \tag{$BIC-P$}$$

$$\text{subject to } \forall i \in \mathcal{I},\, \forall v_i \in V_i,\, \forall v_i' \in V_i \cup \{v_0\},$$

$$\sum_{v_{-i}\in V_{-i}} \tau_i(v_{-i}|v_i) \Big[p(v_i, v_{-i}) \cdot v_i - t_i(v_i, v_{-i})\Big]$$

$$\geq \sum_{v_{-i}\in V_{-i}} \tau_i(v_{-i}|v_i) \Big[p(v_i', v_{-i}) \cdot v_i - t_i(v_i', v_{-i})\Big], \tag{21}$$

$$\forall v \in V,\, \forall k \in \mathcal{K},\, p^k(v) \geq 0,$$

$$\forall v \in V,\, \sum_{k\in\mathcal{K}} p^k(v) = 1. \tag{22}$$

We derive the dual minimization problem $(BIC-D)$ as follows:

$$\min_{\lambda^{BIC}(v_i';v_i),\, \mu^{BIC}(v)} \quad \sum_{v\in V} \mu^{BIC}(v) \tag{$BIC-D$}$$

$$\text{subject to } \forall i \in \mathcal{I},\, \forall v_i \in V_i,\, \forall v_{-i} \in V_{-i},$$

$$\sum_{v_i'\in V_i\cup\{v_0\}} \lambda^{BIC}(v_i'; v_i)\tau_i(v_{-i}|v_i) - \sum_{v_i'\in V_i} \lambda^{BIC}(v_i; v_i')\tau_i(v_{-i}|v_i') = \pi(v_i, v_{-i}), \tag{23}$$

$$\forall v \in V,\, \forall k \in \mathcal{K},$$

$$\sum_{i\in\mathcal{I}} \sum_{v_i'\in V_i\cup\{v_0\}} \lambda^{BIC}(v_i'; v_i)\tau_i(v_{-i}|v_i)v_i(k)$$

$$+\sum_{i\in\mathcal{I}} \sum_{v_i'\in V_i} \lambda^{BIC}(v_i; v_i')\tau_i(v_{-i}|v_i')v_i'(k) \leq \mu^{BIC}(v), \tag{24}$$

$$\forall i \in \mathcal{I},\, \forall v_i \in V_i,\, \forall v_i' \in V_i \cup \{v_0\},$$

$$\lambda^{BIC}(v_i'; v_i) \geq 0, \tag{25}$$

where $\lambda^{BIC}(v_i'; v_i)$ is the multiplier on the incentive constraint of preferring v_i over v_i' (see (21)), and $\mu^{BIC}(v)$ is the multiplier associated with the feasibility constraint of $\sum_{k\in\mathcal{K}} p^k(v) = 1$ (see (22)).

Step (3) Our objective is to identify a simple type space, against which $V_{DIC-D} \geq V_{BIC-D}$. We claim that under simple type space (to be identified), for any feasible dual variables $\lambda^{DIC}(v_i, v_{-i})$ and $\mu^{DIC}(v)$, and its corresponding objective function value in the minimization problem $(DIC-D)$, there exist feasible dual variables $\lambda^{BIC}(v_i'; v_i)$ and $\mu^{BIC}(v)$ in the minimization problem $(BIC-D)$ such that its objective function achieves the same value.

By inspecting the constraints in the two minimization problems $(DIC - D)$ and $(BIC - D)$, in particular, by comparing (17) with (23), (18) with (24), and (19) with (25), the claim is true if there exists $\tau_i(v_{-i}|v_i)$ such that[17]

$$\lambda^{BIC}(v_i^-; v_i)\tau_i(v_{-i}|v_i) = \lambda^{DIC}(v_i, v_{-i}), \tag{26}$$

$$\lambda^{BIC}(v_i'; v_i) = 0 \text{ if } v_i' \neq v_i^-, \tag{27}$$

$$\mu^{BIC}(v) = \mu^{DIC}(v). \tag{28}$$

Summing across v_{-i} for (26), we have

$$\begin{aligned} \lambda^{BIC}(v_i^-; v_i) &= \sum_{v_{-i} \in V_{-i}} \lambda^{DIC}(v_i, v_{-i}), \\ &= \sum_{v_{-i} \in V_{-i}} \sum_{\hat{v}_i : \hat{v}_i \succeq_i^+ v_i} \pi(\hat{v}_i, v_{-i}), \end{aligned}$$

where the second line follows from (20). From (26), we can solve for $\tau_i(v_{-i}|v_i)$ as follows:

$$\begin{aligned} \tau_i(v_{-i}|v_i) &= \frac{\lambda^{DIC}(v_i, v_{-i})}{\lambda^{BIC}(v_i^-; v_i)} \\ &= \frac{\sum_{\hat{v}_i : \hat{v}_i \succeq_i^+ v_i} \pi(\hat{v}_i, v_{-i})}{\sum_{v_{-i} \in V_{-i}} \sum_{\hat{v}_i : \hat{v}_i \succeq_i^+ v_i} \pi(\hat{v}_i, v_{-i})}. \end{aligned}$$

References

Ahuja, R. K., T. L. Magnanti, and J. B. Orlin (1993): *Network Flows: Theory, Algorithms and Applications.* Prentice Hall, New Jersey.

Bergemann, D., and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73(6), 1771–1813.

Bradley, S., A. Hax, and T. Magnanti (1977): *Applied Mathematical Programming.* Addison Wesley.

[17]Suppose that $\lambda^{DIC}(v_i, v_{-i})$ and $\mu^{DIC}(v)$ are feasible under the minimization problem $(DIC - D)$. It follows from (26) and (27) that (25) is satisfied. Furthermore, (23) reduces to (17), and (24) reduces to (18). Therefore, $\lambda^{BIC}(v_i'; v_i)$ and $\mu^{BIC}(v)$ are feasible under the minimization problem $(BIC - D)$. Finally, it follows from (28) that the value of the objective function of the minimization problem $(BIC - D)$ is $\sum_{v \in V} \mu^{BIC}(v) = \sum_{v \in V} \mu^{DIC}(v)$.

Börgers, T. (2015): *An Introduction to the Theory of Mechanism Design.* Oxford University Press.

——— (2017): "(No) Foundations of Dominant-Strategy Mechanisms: A Comment on Chung and Ely (2007)," *Review of Economic Design*, 21(2), 73–82.

Börgers, T., and J. Li (2017): "Strategically Simple Mechanisms," mimeo, University of Michigan and University of New South Wales.

Carroll, G. (2015): "Robustness and Linear Contracts," *American Economic Review*, 105(2), 536–563.

——— (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85(2), 453–488.

Chung, K.-S., and J. C. Ely (2007): "Foundations of Dominant-Strategy Mechanisms," *Review of Economic Studies*, 74(2), 447–476.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a Partnership Efficiently," *Econometrica*, 55(3), 615–632.

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

Du, S. (2017): "Robust Mechanisms Under Common Valuation," mimeo, Simon Fraser University.

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.

Goeree, J., and A. Kushnir (2017): "A Geometric Approach to Mechanism Design," mimeo, University of New South Wales and Carnegie Mellon University.

Gui, H., R. Müller, and R. V. Vohra (2004): "Dominant Strategy Mechanisms with Multidimensional Types," Discussion Paper 1392, The Center for Mathematical Studies in Economics and Management Sciences, Northwestern University, Evanston, IL.

Heifetz, A., and Z. Neeman (2006): "On the Generic (Im)Possibility of Full Surplus Extraction in Mechanism Design," *Econometrica*, 74(1), 213–233.

Heydenreich, B., R. Müller, M. Uetz, and R. Vohra (2009): "Characterization of Revenue Equivalence," *Econometrica*, 77(1), 307–316.

Kos, N., and M. Messner (2013): "Extremal Incentive Compatible Transfers," *Journal of Economic Theory*, 148(1), 134–164.

Lu, H., and J. Robert (2001): "Optimal Trading Mechanisms with Ex Ante Unidentified Traders," *Journal of Economic Theory*, 97(1), 50–80.

Mailath, G., and A. Postlewaite (1990): "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57(3), 351–367.

Malakhov, A., and R. V. Vohra (2009): "An Optimal Auction for Capacity Constrained Bidders: A Network Perspective," *Economic Theory*, 39(1), 113–128.

Manelli, A., and D. Vincent (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private Values Model," *Econometrica*, 78(6), 1905–1938.

Mertens, J.-F., and S. Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 14(1), 1–29.

Morris, S. (1994): "Trade with Heterogeneous Prior Beliefs and Asymmetric Information," *Econometrica*, 62(6), 1327–1347.

Müller, R., A. Perea, and S. Wolf (2007): "Weak Monotonicity and Bayes-Nash Incentive Compatibility," *Games and Economic Behavior*, 61(2), 344–358.

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

Myerson, R., and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29(2), 265–281.

Rochet, J.-C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16(2), 191–200.

Rochet, J.-C., and L. A. Stole (2003): "The Economics of Multidimensional Screening," in *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress*. Cambridge University Press, Cambridge.

Sher, I., and R. Vohra (2015): "Price Discrimination through Communication," *Theoretical Economics*, 10(2), 597–648.

Vohra, R. V. (2011): *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, Cambridge.

Wilson, R. (1987): "Game-Theoretic Approaches to Trading Processes," in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, chap. 2, pp. 33–77. Cambridge University Press, Cambridge.

Yamashita, T. (2015): "Implementation in Weakly Undominated Strategies: Optimality of Second-Price Auction and Posted-Price Mechanism," *Review of Economic Studies*, 82(3), 1223–1246.

——— (2017): "Revenue Guarantee in Auction with a (Correlated) Common Prior and Additional Information," mimeo, Toulouse School of Economics.

Yamashita, T., and S. Zhu (2017): "On the Foundations of Ex Post Incentive Compatible Mechanisms," mimeo, Toulouse School of Economics.