Revisiting the Foundations of Dominant-Strategy Mechanisms*

Yi-Chun Chen† Jiangtao Li‡

July 28, 2016

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 where 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 grateful 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 also thank Tilman Börgers, Gabriel Carroll, Kim-Sau Chung, Takashi Kunimoto, Satoru Takahashi, Takuro Yamashita, and seminar participants at SAET 2014, CUHK, and the Summer School of the Econometric Society 2015 for their valuable comments. 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. Chen and Li 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

‡Department of Economics, National University of Singapore, jasonli1017@gmail.com

Contents

1 Introduction

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-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.

We start with the following contrast between two bilateral trade models (Section 3). In the standard bilateral trade model in which traders are ex ante identified buyers or sellers, the use of dominant-strategy mechanisms has maxmin and Bayesian foundations. We then consider a bilateral trade model with ex ante unidentified traders. In this economic environment, 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. In other words, there is neither a Bayesian foundation nor a maxmin foundation. 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.[1]

From this contract, we abstract the uniform shortest-path tree condition. 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) and Kos and Messner (2013).[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 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 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 prominent environments. First, the uniform shortest-path tree condition is satisfied in environments with linear utilities and one-dimensional types. 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)).

[1] Chung and Ely (2007, Proposition 2) construct an example in which a Bayesian foundation does not exist, but their construction is silent about the existence of a maxmin foundation. Bergemann and Morris (2005) study an implementability problem. Börgers (2013) adopts a different notion of optimality.

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

[3] We are indebted to Rakesh Vohra for bringing to our attention a closely related paper by Sher and Vohra (2015), as well as for suggestions along this direction.

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 and utilities are linear. 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 capacitated 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. If the optimal dominant-strategy mechanism exhibits certain properties, we can construct a single Bayesian mechanism that robustly achieves strictly higher expected revenue than the optimal dominant-strategy mechanism, regardless of the agents' beliefs (Theorem 2). We stress that as a no-foundation result, this is remarkably strong. In addition to bilateral trade with ex ante unidentified traders, we apply this result to auction with type-dependent outside option.

The remainder of this introduction discusses some related literature. Section 2 presents the notations, concepts, and the model. Section 3 contrasts two bilateral trade models. Section 4 formulates the notion of the uniform shortest-path tree and presents the results. Section 5 studies three applications of the results and Section 6 concludes with discussions.

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 of Bergemann and Morris (2005), Chung and Ely (2007), and more recently, Carroll (2015, 2016), Yamashita (2015, 2016), and Du (2016), among 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 (2015). Our paper differs from these in that the mechanism designer in our model does not make any assumptions about the agents' beliefs.

2 Preliminaries

2.1 Notation

There is a finite set $\mathcal{I} = \{1, 2, ..., I\}$ of risk-neutral agents and a finite set $\mathcal{K} = \{1, 2, ..., K\}$ of social alternatives. Agent i's payoff type $v_i \in \mathbb{R}^K$ represents her gross utility under the K alternatives.[4] 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 $v = (v_1, v_2, ..., v_I)$. 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 $\Omega_i = V_i$, $f_i(v_i) = v_i$, 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

[4] We may represent the agent's payoff types in different ways. For instance, when studying one-dimensional payoff types (Section 5.1), it is more convenient to represent agent i's payoff type by $v_i \in \mathbb{R}$.

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.[5]

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 alternative k is chosen.

The mechanism defines a game form, which together with the type space constitutes

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

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-relevation 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 $M(\pi)$ 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 uses a dominant-strategy mechanism, then his maximum 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 worse 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.[6]

Proposition 1. *The Bayesian foundation is a stronger notion than the maxmin foundation.*

3 Motivating examples

Before we present the results, it is instructive to contrast two bilateral trade models. In the standard bilateral trade model (see Myerson and Satterthwaite (1983)), whether an agent is the buyer or the seller is exogenously given. Either the seller sells some units to the buyer or no trade occurs. In the bilateral trade model with ex ante unidentified traders (see Cramton, Gibbons, and Klemperer (1987) and Lu and Robert (2001)), 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. The mechanism designer chooses a mechanism that maximizes the expected profit in both models.

Section 3.1 presents the basics shared by both models. Section 3.2 studies the standard bilateral trade model. In this case, the use of dominant-strategy mechanisms has maxmin and Bayesian foundations. Section 3.3 studies the bilateral trade model with ex ante unidentified traders. We show that there is neither a Bayesian foundation nor a maxmin foundation.

3.1 Setup

Consider a broker who chooses trading mechanisms that maximize the expected profit; see for example, Myerson and Satterthwaite (1983, Section 5), Lu and Robert (2001) and Börgers (2015). 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. Agent 1's valuation for the good could be either 18 or

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

38. Agent 2's valuation for the good could be either 10 or 30. The broker has the following estimate of the distribution of the agents' valuations:

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

(1)

3.2 Standard bilateral trade

In the standard bilateral trade model, agent 1 is the buyer and agent 2 is the seller. The trading mechanism is characterized by three outcome functions (p, t_1, t_2), where $p(v_1, v_2)$ is the expected trading amount, $t_1(v_1, v_2)$ is the expected payment from agent 1 to the broker and $t_2(v_1, v_2)$ is the expected 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 $0 \leq p \leq \frac{1}{2}$.

Clearly, this model belongs to the class of environments with linear utilities and one-dimensional payoff types; see Section 5.1. Following Corollary 1, the use of dominant-strategy mechanisms has maxmin and Bayesian foundations.

3.3 Bilateral trade with ex ante unidentified traders

In this section, we study the bilateral trade model with ex ante unidentified traders. Each agent may be either the buyer or the seller. The trading mechanism is characterized by three outcomes functions (p, t_B, t_S), where $p(v_1, v_2)$ is the expected trading amount, $t_B(v_1, v_2)$ is the expected payment from *the buyer* to the broker and $t_S(v_1, v_2)$ is the expected payment from *the seller* to the broker, if v_1 and v_2 are the reported valuations of agent 1 and agent 2. The buyer's utility from purchasing p units of the good and paying a transfer t_B is $pv_B - t_B$ and the seller's utility from selling p unit of the good and paying a transfer t_S is $-pv_S - t_S$, where $0 \leq p \leq \frac{1}{2}$.

In the context of this economic environment, this example illustrates that, maxmin/ Bayesian foundations might not exist. Section 3.3.1 calculates the maximum expected revenue that could be achieved by a dominant-strategy mechanism, and Section 3.3.2 explicitly constructs a single Bayesian mechanism that achieves a strictly higher expected revenue, regardless of the assumption about (the distribution of) the agents' beliefs. It should be obvious from the exposition below that this example is robust to small perturbations in the agents' valuations or the broker's estimate of the distribution of the payoff types.

3.3.1 Optimal dominant-strategy mechanism

Using a linear programming solver, we have the optimal dominant-strategy mechanism Γ as follows, where the first number in each cell indicates the amount of good agent 1 buys from agent 2, the second number is the transfer from agent 1 and the third number is the transfer from agent 2. The maximum expected revenue the mechanism designer can generate from a dominant-strategy mechanism is 3.

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

(2)

3.3.2 Neither a Bayesian foundation nor a maxmin foundation

To show that there is no maxmin foundation, it suffices to construct a single Bayesian mechanism and achieve a strictly higher expected revenue than he does using the optimal dominant-strategy mechanism, regardless of the assumption about (the distribution of) the agents' beliefs. Since Bayesian foundation is a stronger notion than maxmin foundation, this further implies that there is no Bayesian foundation.

The construction of the mechanism Γ' follows immediately from Theorem 2. We shall save the arguments in Section 4. Following Chung and Ely (2007), 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. To see that Γ' is expected revenue improving, note that Γ' achieves revenue of at least 4 everywhere and hence the expected revenue is at least 4, regardless of the agents' beliefs.

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

4 Results

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

$$\max_{p^k(\cdot)\geq 0, t_i(\cdot)} \sum_{v\in V} \pi(v) \sum_{i\in\mathcal{I}} t_i(v) \qquad (DIC-P)$$

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

$$p(v_i, v_{-i}) \cdot v_i - t_i(v_i, v_{-i}) \geq p(v_i', v_{-i}) \cdot v_i - t_i(v_i', v_{-i}), \qquad (3)$$

$$\forall v \in V, \sum_{k\in\mathcal{K}} p^k(v) = 1. \qquad (4)$$

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.

Say that a decision rule p is dsIC if there exists transfer scheme t such that the mechanism (p, t) satisfies the incentive constraints (3). We omit the proof of the following standard lemma, due to Rochet (1987).

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 every sequence of payoff types of agent $i, (v_{i,1}, v_{i,2}, ..., v_{i,k})$ with $v_{i,k} = v_{i,1}$, we have*

$$\sum_{\kappa=1}^{k-1} [p(v_{i,\kappa}, v_{-i}) \cdot v_{i,\kappa+1} - p(v_{i,\kappa}, v_{-i}) \cdot v_{i,\kappa}] \leq 0. \qquad (5)$$

4.1 Uniform shortest-path tree

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

Definition 6. *Fix a decision rule p that is dsIC and other agents' reports v_{-i}.*[7] *(1) The set of nodes for agent i is $V_i \cup \{v_0\}$; (2) For any $v_i \in V_i$ and $v_i' \in V_i \backslash \{v_i\} \cup \{v_0\}, v_i' \to v_i$ is a*

[7]In the remainder of this section, whenever we fix a decision rule p, we mean a decision rule p that is dsIC.

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* $v_{i,k} \in V_i$ *is a sequence* $P = (v_0, v_{i,1}, v_{i,2}, ..., v_{i,k})$ *where (i)* $v_{i,j} \in V_i, \forall j = 1, 2, ..., k$*; (ii)* $v_0 \to v_{i,1}$*; (iii)* $v_{i,j-1} \to v_{i,j}, \forall j = 2, ..., k$ *and (iv)* $j \neq j' \implies v_{i,j} \neq v_{i,j'}$*.*

To understand the maximization problem $(DIC - P)$ and in particular the associated incentive constraints (3), it helps to flip to its dual. The dual is a network flow problem that can be described in the following way. Fix a decision rule p and other agents' reports v_{-i}. 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 all 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 decision rule* p *and other agents' reports* v_{-i}*. A shortest-path tree is the union of all shortest-paths from the source to all nodes.*

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

When the uniform shortest-path tree condition is satisfied, we drop the dependence on p, v_{-i}. Uniform shortest-path tree induces an order on the agents' payoff types. For a typical shortest-path $(v_0, v_{i,1}, v_{i,2}, ..., v_{i,k})$ of the shortest-path tree, we write $v_{i,k} \succ_i v_{i,k-1} \succ_i \ldots \succ_i v_{i,1} \succ_i v_0$. 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, 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 write $v_i^- = v_i'$.

With the uniform shortest-path tree, the rent of any payoff type can be easily calculated and all incentive constraints can be replaced by the cyclical monotonicity constraints on the decision rule. We record this as the following proposition.

Proposition 2. *With the uniform shortest-path tree $\succ_i$, the maximization problem $(DIC-P)$ is equivalent to*

$$\max_{p(\cdot)} \sum_{i\in\mathcal{I}} \sum_{v_i\in V_i} \sum_{v_{-i}\in V_{-i}} \pi(v_i, v_{-i}) \left[p(v_i, v_{-i}) \cdot v_i - \sum_{v_i'\in V_i : v_i \succeq_i^+ v_i'} p((v_i')^-, v_{-i}) \cdot (v_i' - (v_i')^-) \right], \quad (6)$$

subject to $p(\cdot)$ satisfies the cyclical monotonicity constraint (5).

Definition 9. *Say π is regular if the cyclical monotonicity constraint (5) is automatically satisfied for p^* that maximizes the reduced objective function (6).*

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).[8] In the applications we study in Section 5.1 and Section 5.2, additional structure is imposed and we provide primitive condition for regularity.

4.2 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.*

Proof. The structure of the proof is as follows. Step 1) considers the optimal dominant-strategy mechanism design problem $(DIC-P)$ and derives its dual $(DIC-D)$. Step 2) restricts attention to a subclass of type spaces, formulates the Bayesian mechanism design problem $(BIC-P)$ and derives its dual $(BIC-D)$. 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 in Step 4) 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 $V_{DIC-P} = V_{DIC-D} \geq V_{BIC-D} \geq V_{BIC-P}$.

Step 1) First consider the optimal dominant-strategy mechanism design problem $(DIC-P)$. We derive its dual $(DIC-D)$, where $\lambda^{DIC}(v_i'; v_i, v_{-i})$ is the multiplier associated

[8] 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 p satisfies the cyclical monotonicity constraint. The regularity condition is then imposed to make sure that such optimal decision rule p automatically satisfies the cyclical monotonicity constraint.

with the incentive constraint (3) and $\mu^{DIC}(v)$ is the multiplier associated with the feasibility constraint (4).

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

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

$$\sum_{v_i' \in \{V_i \setminus \{v_i\}\} \cup \{v_0\}} \lambda^{DIC}(v_i'; v_i, v_{-i}) - \sum_{v_i' \in V_i \setminus \{v_i\}} \lambda^{DIC}(v_i; v_i', v_{-i}) = \pi(v_i, v_{-i}), \tag{7}$$

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

$$\pi(v) \sum_{i\in\mathcal{I}} v_i(k) + \sum_{i\in\mathcal{I}} \sum_{v_i' \in V_i \setminus \{v_i\}} \lambda^{DIC}(v_i; v_i', v_{-i})(v_i(k) - v_i'(k)) \leq \mu^{DIC}(v), \tag{8}$$

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

$$\lambda^{DIC}(v_i'; v_i, v_{-i}) \geq 0. \tag{9}$$

As $\lambda^{DIC}(v_i'; v_i, v_{-i})$ is the multiplier for the incentive constraint (3), by the uniform shortest-path tree and regularity, there is a dual optimum satisfying

$$\lambda^{DIC}(v_i'; v_i, v_{-i}) > 0 \text{ only if } v_i \succ_i v_i',$$

and (7) simplifies to

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

By induction,

$$\lambda^{DIC}(v_i'; v_i, v_{-i}) = \begin{cases} \sum_{\hat{v}_i : \hat{v}_i \succeq_i^+ v_i} \pi(\hat{v}_i, v_{-i}); & \text{if } v_i \succ_i v_i'; \\ 0; & \text{otherwise.} \end{cases} \tag{10}$$

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 valuation v_i. Let the set of types for agent i be equal to the set of possible valuations, i.e. $\Omega_i = V_i$. We take f_i to be the identity, and for notational ease, we will write $\tau_i(\cdot|v_i) = g_i(v_i)$ for the belief of type v_i of agent i about the types of the other agents. From now on, we restrict attention to such type spaces.

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

$$\max_{p^k(\cdot)\geq 0, t_i(\cdot)} \sum_{v\in V} \pi(v) \sum_{i\in\mathcal{I}} t_i(v) \tag{$BIC-P$}$$

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

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

$$\geq \sum_{v_{-i}\in V_{-i}} \tau_i(v_{-i}|v_i)(p(v_i', v_{-i}) \cdot v_i - t_i(v_i', v_{-i})), \tag{11}$$

$$\forall v \in V, \sum_{k\in\mathcal{K}} p^k(v) = 1. \tag{12}$$

We derive the dual minimization problem $(BIC-D)$, where $\lambda^{BIC}(v_i'; v_i)$ is the multiplier for the incentive constraint (11) and $\mu^{BIC}(v)$ is the multiplier for the feasibility constraint (12).

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

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

$$\sum_{v_i'\in\{V_i\backslash\{v_i\}\}\cup\{v_0\}} \lambda^{BIC}(v_i'; v_i)\tau_i(v_{-i}|v_i) - \sum_{v_i'\in V_i\backslash\{v_i\}} \lambda^{BIC}(v_i; v_i')\tau_i(v_{-i}|v_i') = \pi(v_i, v_{-i}), \tag{13}$$

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

$$\pi(v)\sum_{i\in\mathcal{I}} v_i(k) + \sum_{i\in\mathcal{I}} \sum_{v_i'\in V_i\backslash\{v_i\}} \lambda^{BIC}(v_i; v_i')\tau_i(v_{-i}|v_i')(v_i(k) - v_i'(k)) \leq \mu^{BIC}(v), \tag{14}$$

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

$$\lambda^{BIC}(v_i'; v_i) \geq 0. \tag{15}$$

Step 3) Now we construct a particular assumption about (the distribution of) agents' beliefs. Given π, for any $v_i \in V_i$, write

$$\pi_i(v_i) = \sum_{v_{-i}\in V_{-i}} \pi(v_i, v_{-i})$$

for the marginal probability of v_i and write

$$G_i(v_i) = \sum_{\hat{v}_i:\hat{v}_i \succeq_i^+ v_i} \pi_i(\hat{v}_i) \tag{16}$$

for the associated distribution function. We define agent i's beliefs as follows:

$$\tau_i(v_{-i}|v_i) = \frac{1}{G_i(v_i)} \sum_{\hat{v}_i:\hat{v}_i\succeq_i^+ v_i} \pi(\hat{v}_i, v_{-i}). \tag{17}$$

Step 4) Fix any feasible dual variables $\lambda^{DIC}(v_i'; v_i, v_{-i})$ and $\mu^{DIC}(v)$ of the minimization problem $(DIC - D)$ that satisfy (10), let

$$\lambda^{BIC}(v_i'; v_i) = \sum_{v_{-i} \in V_{-i}} \lambda^{DIC}(v_i'; v_i, v_{-i})$$
$$\text{and } \mu^{BIC}(v) = \mu^{DIC}(v).$$

If $v_i \succ_i v_i'$, we have

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

where the second equality follows from (10), the third equality follows from (16), and the last equality follows from (17). Otherwise,

$$\lambda^{BIC}(v_i'; v_i)\tau_i(v_{-i}|v_i) = 0.$$

In either case, we have

$$\lambda^{BIC}(v_i'; v_i)\tau_i(v_{-i}|v_i) = \lambda^{DIC}(v_i'; v_i, v_{-i}). \tag{18}$$

We now show that the dual variables $\lambda^{BIC}(v_i'; v_i)$ and $\mu^{BIC}(v)$ are feasible under the minimization problem $(BIC - D)$. (15) are trivially satisfied. It follows from (18) that (13) reduces to (7), and (14) reduces to (8). Since $\lambda^{DIC}(v_i'; v_i, v_{-i})$ and $\mu^{DIC}(v)$ are feasible under the problem $(DIC-D)$, $\lambda^{BIC}(v_i'; v_i)$ and $\mu^{BIC}(v)$ are feasible under the minimization problem $(BIC - D)$. Furthermore, the value of the objective function of the minimization problem $(BIC - D)$ is $\sum_{v \in V} \mu^{BIC}(v) = \sum_{v \in V} \mu^{DIC}(v)$. We conclude that $V_{DIC-D} \geq V_{BIC-D}$. □

4.3 No foundations of dominant-strategy mechanisms

This subsection considers violations of the uniform shortest-path tree condition. When the uniform shortest-path tree condition is not satisfied, as illustrated in the bilateral trade model with ex ante unidentified traders (Section 3.3), maxmin/ Bayesian foundations might not exist. 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.

In environments where the uniform shortest path is violated, it is difficult to find the optimal dominant-strategy mechanisms, not to mention the construction of the superior Bayesian mechanism. To have a meaningful discussion, we shall take the optimal dominant-strategy mechanisms (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). In addition to bilateral trade with ex ante unidentified traders, the result can also be applied to auction with type-dependent outside option (Section 5.3).

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

	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.

Remark 1. *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.*

Proof. 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 satisfy the incentive constraints, we have $x \geq 0, y \geq w > 0$.[9]

We show that there is no maxmin foundation. That is, 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. To do this, we first explicitly identify one such mechanism and proceed by verifying i) the mechanism is BIC for the universal type space; and ii) this mechanism achieves a strictly higher expected revenue regardless of the agents' beliefs. Since the Bayesian foundation is a stronger notion than the maxmin foundation, 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 this mechanism, the mechanism designer elicits agent 2's first-order belief about agent 1's payoff type. Consider the following Bayesian mechanism Γ':

	v_1	v_1'
$v_2, a \in [0, \frac{w}{w-z})$	$p(v_1, v_2'), t_1(v_1, v_2'), t_2(v_1, v_2')$	$p(v_1', v_2'), t_1(v_1', v_2'), t_2(v_1', v_2')$
$v_2, a \in [\frac{w}{w-z}, 1]$	$p(v_1, v_2), t_1(v_1, v_2), t_2(v_1, v_2) + x$	$p(v_1', v_2), t_1(v_1', v_2), t_2(v_1', v_2) + y$
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 $a \in [0, \frac{w}{w-z})$ will not announce v_2' as utility is unchanged;

[9]As a matter of fact, it must be that $x = 0$, and $y = w$. Otherwise, the dominant-strategy mechanism would not have been optimal. Note that the uniform shortest-path tree condition is violated.

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

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

To see that the mechanism achieves a strictly higher expected revenue than the optimal dominant-strategy mechanism, regardless of the assumption about (the distribution of) the agents' belief, 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 $x \geq 0$ and $y \geq w > 0$.

□

5 Applications

This section is devoted to the applications of the results. The uniform shortest-path tree condition holds in the standard social choice environment with linear utilities and one-dimensional payoff types as well as some multi-dimensional environments. Section 5.1 applies our result to environments with linear utilities and one-dimensional types, and Section 5.2 considers a multi-dimensional environment. For both applications, we provide primitive conditions for regularity. As we illustrated in Section 3, Theorem 2 can be applied to bilateral trade with ex ante unidentified traders. Section 5.3 applies Theorem 2 to another environment, namely, auction with type-dependent outside option.

5.1 Linear utilities and one-dimensional payoff types

In this subsection, we consider the standard social choice environment with linear utilities and one-dimensional payoff types.[10] 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)).

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

There is a finite set $\mathcal{I} = \{1, 2, ..., I\}$ of risk neutral agents and a finite set $\mathcal{K} = \{1, 2, ..., K\}$ 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 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 $A_i = (a_i^1, a_i^2, ..., a_i^K)$. For notational simplicity, we assume that each agent has M possible payoff types and that the set V_i is the same for each agent: $V_i = \{v^1, v^2, ..., v^M\}$, where $v^m - v^{m-1} = \gamma$ for each m and some $\gamma > 0$.

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

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

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

$$p(v^m, v_{-i}) \cdot A_i v^m - t_i(v^m, v_{-i}) \geq 0, \tag{19}$$

$$p(v^m, v_{-i}) \cdot A_i v^m - t_i(v^m, v_{-i}) \geq p(v^l, v_{-i}) \cdot A_i v^m - t_i(v^l, v_{-i}). \tag{20}$$

In the environment with linear utilities and 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 constraints (19) and (20).

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. *Fix any decision rule p that is dsIC, the shortest path from the source v_0 to any payoff type $v^m \in V_i$ is $P = (v_0, v^1, v^2, ..., v^m)$ and*

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

Next, we present the primitive condition for regularity. It is well known that an equivalent formulation of the problem is

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

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

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

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

$$p(v^m, v_{-i}) \cdot A_i \geq p(v^l, v_{-i}) \cdot A_i, \text{ for } m \geq l.$$

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 valuation conditional on the other agents having payoff type profile v_{-i}. Define the virtual valuation of agent i as

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

and rewrite the objective function as

$$\begin{aligned} \sum_{v\in V} \pi(v) \sum_{i\in\mathcal{I}} t_i(v) &= \sum_{v\in V} \sum_{i\in\mathcal{I}} \pi(v) p(v) \cdot A_i r_i(v) \\ &= \sum_{v\in V} \pi(v) p(v) \cdot \sum_{i\in\mathcal{I}} A_i r_i(v). \end{aligned} \tag{21}$$

For each alternative k, let $K_i^{k,\inf} = \{k' \in \mathcal{K} : a_i^{k'} < a_i^k\}$. That is, $K_i^{k,\inf}$ is the collection of alternatives that agent i considers inferior to alternative k.

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

$$k \in \arg\max_k \sum_{i\in\mathcal{I}} a_i^k r_i(v) \Rightarrow K_j^{k,\inf} \cap \arg\max_k \sum_{i\in\mathcal{I}} a_i^k r_i(\hat{v}_j, v_{-j}) = \emptyset \tag{22}$$

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

We establish the foundations of dominant-strategy mechanisms in Corollary 1.

Corollary 1. *If π satisfies the regularity condition (22), the use of dominant-strategy mechanisms has a Bayesian/ maxmin foundation.*

Proof. By (21), the objective function becomes

$$\sum_{v\in V} \pi(v) p(v) \cdot \sum_{i\in\mathcal{I}} A_i r_i(v).$$

Regularity condition (22) ensures that for any alternative k chosen with positive probability for payoff type profile (v^l, v_{-i}), when agent i's payoff type increases say 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 $m \geq l$. It is well known that this is equivalent to cyclical monotonicity in environments with linear utilities and one-dimensional payoff types. The uniform shortest-path tree condition follows from Lemma 2. The result then follows from Theorem 1. □

5.2 Multi-unit auction with capacity-constrained bidders

In addition to environments with linear utilities and one-dimensional payoff types, the uniform shortest-path tree condition is also satisfied in some multi-dimensional environments. Solving for the optimal mechanism in a multi-dimensional environment is in general a daunting task. In this section, 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.[11] 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 $\mathcal{A} = \{1, 2, ..., A\}$ and the range of b be $\mathcal{B} = \{1, 2, ..., B\}$. 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. Equivalently, we may suppose that the mechanism designer can verify that claims that exceed capacity are false.

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

Lemma 3. *Fix any decision rule p that is dsIC, the shortest-path from the source v_0 to any payoff type (a, b) is*

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

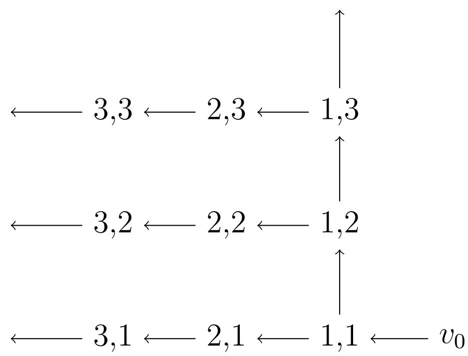

Figure 1: $(a, b) \succ_i (a-1, b) \succ_i ... \succ_i (1, b) \succ_i (1, b-1) \succ_i ... \succ_i (1, 1)$.

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

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

$$a - \frac{1 - F_{b,v_{-i}}(a)}{\pi((a, b), v_{-i})} \geq a' - \frac{1 - F_{b',v_{-i}}(a')}{\pi((a', b'), v_{-i})}, \tag{23}$$

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

The proof of Lemma 3 and the derivation of the regularity condition (23) is a straightforward extension of Malakhov and Vohra (2009) and omitted. When π is independent, the regularity condition (23) reduces to the regularity condition in Malakhov and Vohra (2009). Corollary 2 then follows from Theorem 1.

5.3 Auction with type-dependent outside option

Besides the bilateral trade model with ex ante unidentified traders (Section 3), we hereby present 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 private information could be either $(20, 0)$ or

(40, 5). Bidder 2's private information could be either (10, 0) or (30, 5). The auctioneer has the following estimate of the distribution of the agents' valuations:

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

(24)

The optimal dominant-strategy mechanism Γ is 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.

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

(25)

6 Discussion

6.1 Foundations of ex post incentive-compatible mechanisms

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. Indeed, in an independent and contemporaneous work, Yamashita and Zhu (2014) study the so-called "digital-goods" auctions in the interdependent-value setting. They show that under "ordinal invariability" (which entails that each agent has a stable preference ordering over all her payoff types, regardless of what payoff type profile the other agents have) and additional assumptions, the use of ex post incentive-compatible mechanisms has maxmin and Bayesian foundations.

6.2 Multi-dimensional environments

Theorem 1 provides a sufficient condition for the maxmin and Bayesian foundations of dominant-strategy mechanisms. The key notion is the uniform shortest-path tree; that is, the binding dominant-strategy incentive constraints are independent of the allocation rule and the reports of the other agents. This property is largely responsible for the success of mechanism design in numerical applications across various fields, since this allows us to express the payment variables in terms of the decision variables.

In multi-dimensional settings, the uniform shortest-path tree condition is typically violated. The binding dominant-strategy incentive constraints are endogenous and vary with the allocation rule and other agents' reports. Theorem 2 shows that when 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.

6.3 On the notion of the maxmin foundation

Börgers (2013) argues that the maxmin foundation requires too little of an optimal mechanism. For every dominant-strategy mechanism, Börgers constructs another mechanism which never yields lower revenue and sometimes yields strictly higher revenue. 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 i) if agents could arrange side bets without requiring the mechanism designer as an intermediary; and ii) if the mechanism designer restricts her attention to the type spaces characterized by Morris (1994), which do not allow speculative trade.

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 for every assumption about 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 may not even satisfy *the minimum requirement* of the maxmin foundation.

A Appendix

Proof of Proposition 2: With the uniform shortest-path tree $\succ_i$, for any p and v_{-i}, the rent of payoff type v_i of agent i can be calculated as follows:

$$\begin{aligned} U_i(v_i, v_{-i}) &= p(v_i, v_{-i}) \cdot v_i - t_i(v_i, v_{-i}) \\ &= p(v_i^-, v_{-i}) \cdot v_i - t_i(v_i^-, v_{-i}) \\ &= p(v_i^-, v_{-i}) \cdot v_i^- - t_i(v_i^-, v_{-i}) + p(v_i^-, v_{-i}) \cdot (v_i - v_i^-) \\ &= U_i(v_i^-, v_{-i}) + p(v_i^-, v_{-i}) \cdot (v_i - v_i^-). \end{aligned}$$

By induction,

$$U_i(v_i, v_{-i}) = \sum_{v_i' \in V_i : v_i \succeq_i^+ v_i'} p((v_i')^-, v_{-i}) \cdot (v_i' - (v_i')^-).$$

Therefore,

$$\begin{aligned} t_i(v_i, v_{-i}) &= p(v_i, v_{-i}) \cdot v_i - U_i(v_i, v_{-i}) \\ &= p(v_i, v_{-i}) \cdot v_i - \sum_{v_i' \in V_i : v_i \succeq_i^+ v_i'} p((v_i')^-, v_{-i}) \cdot (v_i' - (v_i')^-). \end{aligned}$$

The maximization problem $(DIC - P)$ is equivalent to

$$\max_{p(\cdot)} \sum_{i \in \mathcal{I}} \sum_{v_i \in V_i} \sum_{v_{-i} \in V_{-i}} \pi(v_i, v_{-i}) \left[p(v_i, v_{-i}) \cdot v_i - \sum_{v_i' \in V_i : v_i \succeq_i^+ v_i'} p((v_i')^-, v_{-i}) \cdot (v_i' - (v_i')^-) \right].$$

By Lemma 1, $p(\cdot)$ is subject to the cyclical monotonicity constraint (5).

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.

Börgers, T. (2013): "(No) Foundations of Dominant-Strategy Mechanisms: A Comment on Chung and Ely (2007)," mimeo. University of Michigan.

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

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

——— (2016): "Robustness and Separation in Multidimensional Screening," mimeo. Stanford University.

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. (2016): "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 (2015): "A Geometric Approach to Mechanism Design," mimeo. University of Technology Sydney 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.

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.

——— (2016): "Revenue Guarantee in Auction with Common Prior," mimeo, Toulouse School of Economics.

YAMASHITA, T., AND S. ZHU (2014): "Foundations of Ex Post Incentive Compatible Mechanisms," mimeo, Toulouse School of Economics.