Equivalence of Stochastic and Deterministic Mechanisms

Yi-Chun Chen* Wei He† Jiangtao Li‡ Yeneng Sun§

September 14, 2016

Abstract

We consider a general social choice environment that has multiple agents, a finite set of alternatives, and independent and dispersed information. We prove that for any Bayesian incentive compatible mechanism, there exists an equivalent deterministic mechanism that i) is Bayesian incentive compatible; ii) delivers the same interim expected allocation probabilities/utilities for all agents; and iii) delivers the same ex ante expected social surplus. This result holds in settings with a rich class of utility functions, multi-dimensional types, interdependent valuations, and non-transferable utilities. To prove our result, we develop a new methodology of "mutual purification", and establish its link with the mechanism design literature.

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

†Department of Economics, Chinese University of Hong Kong, hewei@cuhk.edu.hk

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

§Department of Economics, National University of Singapore, ecssuny@nus.edu.sg

1 Introduction

Myerson (1981) provides the framework that has become the paradigm for the study of optimal auction design. Under a "regularity" condition, the optimal auction allocates the object to the bidder with the highest "virtual value", provided that this virtual value is above the seller's opportunity cost. In other words, the optimal auction in Myerson's setting is deterministic.[1]

A natural conjecture is that the optimality of deterministic mechanisms generalizes beyond Myerson's setting. McAfee and McMillan (1988, Section 4) claim that under a general regularity condition on consumers' demand, stochastic delivery was not optimal for a multi-product monopolist. However, this result has been proven to be incorrect with a single agent. Several papers have shown that a multi-product monopolist may find it beneficial to include lotteries as part of the selling mechanism; see for example, Thanassoulis (2004), Manelli and Vincent (2006, 2007), Pycia (2006), Pavlov (2011), and more recently, Hart and Reny (2015), and Rochet and Thanassoulis (2015).[2] In this paper, we restore the optimality of deterministic mechanisms in remarkably general environments with multiple agents.

We consider a general social choice environment that has multiple agents, a finite set of alternatives, and independent and dispersed information.[3] We show that for any Bayesian incentive compatible mechanism, there exists an equivalent deterministic mechanism that i) is Bayesian incentive compatible; ii) delivers the same interim expected allocation probabilities and the same interim expected utilities for all agents; and iii) delivers the same ex ante expected social surplus. In addition to the standard social choice environments with linear utilities and one-dimensional, private types, our result holds in settings with a rich class of utility functions, multi-dimensional types, interdependent valuations, and non-transferable utilities.

Our result implies that any mechanism, including the optimal mechanisms (whether in terms of revenue or efficiency), can be implemented using a deterministic mechanism and

[1]Also see Riley and Zeckhauser (1983) who consider a one-good monopolist selling to a population of consumers with unit demand and show that lotteries do not help the one-good monopolist.

[2]In environments in which different types are associated with different risk attitudes, it is known that stochastic mechanisms may perform better; see for example, Laffont and Martimort (2002, p. 67) and Strausz (2006).

[3]Throughout this paper, we say that an agent has "dispersed information" if her type distribution is atomless.

nothing can be gained from designing more intricate mechanisms with possibly more complex randomization. As pointed out in Hart and Reny (2015, p. 912), Aumann commented that it is surprising that randomization can not increase revenue when there is only one good. Indeed, aforementioned papers in the screening literature establish that randomization helps when there are multiple goods. Nevertheless, we show that in general social choice environment with multiple agents, the revenue maximizing mechanism can always be deterministically implemented. This is in sharp contrast with the results in the screening literature.

Our result also has important implications beyond the revenue contrast. The mechanism design literature essentially builds on the assumption that a mechanism designer can credibly commit to any outcome of a mechanism. This requirement implies that any outcome of the mechanism must be verifiable before it can be employed. In this vein, a stochastic mechanism demands not only that a randomization device be available to the mechanism designer, but also that the outcome of the randomization device be objectively verified. As noted in Laffont and Martimort (2002, p. 67), "Ensuring this verifiability is a more difficult problem than ensuring that a deterministic mechanism is enforced, because any deviation away from a given randomization can only be statistically detected once sufficiently many realizations of the contracts have been observed. ... The enforcement of such stochastic mechanisms is thus particularly problematic. This has led scholars to give up those random mechanisms or, at least, to focus on economic settings where they are not optimal."[4] Our result implies that every mechanism can in fact be deterministically implemented, and thereby irons out the conceptual difficulties associated with stochastic mechanisms.

This paper joins the strand of literature that studies mechanism equivalence. Though motivations vary, these results show that it is without loss of generality to consider the various subclasses of mechanisms. As in the case of dominant-strategy mechanisms (see Manelli and Vincent (2010) and Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013)) and symmetric auctions (see Deb and Pai (2015)), our findings imply that the requirement of deterministic

[4]Also see Bester and Strausz (2001) and Strausz (2003).

mechanisms is not restrictive in itself.[5] In this sense, our result provides a foundation for the use of deterministic allocations in mechanism design settings such as auctions, bilateral trades, and so on.

In order to prove the existence of an equivalent deterministic mechanism, we develop a new methodology of "mutual purification", and establish its link with the literature of mechanism design.[6] The notion of mutual purification is both conceptually and technically different from the usual purification principle in the literature related to Bayesian games. We shall clarify these two different notions of purification in the next three paragraphs.

It follows from the general purification principle in Dvoretzky, Wald, and Wolfowitz (1950) that any behavioral-strategy Nash equilibrium in a finite-action Bayesian game with independent and dispersed information corresponds to some pure-strategy Bayesian Nash equilibrium with the same payoff.[7] In particular, independent and dispersed information allows the agents to replace their behavioral strategies by some equivalent pure strategies one-by-one.[8] The point is that under the independent information assumption, any agent who has dispersed information could purify her own behavioral strategy regardless whether other agents have dispersed information. Example 2 illustrates this idea of "self purification". Given a behavioral-strategy Nash equilibrium in a 2-agent Bayesian game with independent information, there is an equivalent pure strategy for the agent with dispersed information, while the other agent with an atom in her type space could not purify her behavioral strategy.

In contrast, the purification result of this paper is based on the dispersed information associated with the other agents. Example 3 partially illustrates this idea of "mutual

[5] Manelli and Vincent (2010) show that for any Bayesian incentive compatible auction, there exists an equivalent dominant-strategy incentive compatible auction that yields the same interim expected utilities for all agents. Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) extend this equivalence result to social choice environments with linear utilities and independent, one-dimensional, private types; also see Footnote 12 for related discussion. Deb and Pai (2015) show that restricting the seller to a using symmetric auction imposes virtually no restriction on her ability to achieve discriminatory outcomes.

[6] Some of our technical results extend the corresponding mathematical results in Arkin and Levin (1972); see the Supplementary Material for more detailed discussion.

[7] See Radner and Rosenthal (1982), Milgrom and Weber (1985) and Khan, Rath, and Sun (2006). Furthermore, by applying the purification idea to a sequence of Bayesian games, Harsanyi (1973) provided an interpretation of mixed-strategy equilibrium in complete information games; see Govindan, Reny, and Robson (2003) and Morris (2008) for more discussion.

[8] See the proof of Theorem 1 in Khan, Rath, and Sun (2006).

purification". For a given randomized mechanism in a 2-agent setting with independent information, the agent with an atom in her type space can achieve the same interim payoff by some deterministic mechanism, while there does not exist such a deterministic mechanism for the other agent with dispersed information. In other words, our result becomes possible because each agent relies on the dispersed information of the other agents rather than her own. This also explains why a similar result does not hold in the one-agent setting since there is no dispersed information from other agents for such a single agent to purify the relevant randomized mechanism. In addition, we emphasize that in the multiple-agent setting, the notion of "mutual purification" requires not only that each agent obtain the same interim payoff under some deterministic mechanism, but also that a *single* deterministic mechanism deliver the same interim payoffs for all the agents *simultaneously*.

From a methodological point of view, the general purification principle in Dvoretzky, Wald, and Wolfowitz (1950) is simply a version of the classical Lyapunov Theorem about the convex range of an atomless finite-dimensional vector measure. Our purification result is technically different. First, the problem we consider is infinite-dimensional because we require the same expected allocation probabilities/ utilities for the equivalent mechanism at the interim level with a continuum of types. Note that Lyapunov's Theorem fails in an infinite-dimensional setting.[9] Second, it is clearly impossible to obtain a purified deterministic mechanism that delivers the same expected allocation probabilities as the original stochastic mechanism, conditioned on the *joint* types of all the agents.[10] However, our result on mutual purification shows that such an equivalence becomes possible when the conditioning operation is imposed on the *individual* types of every agent simultaneously, although the combination of the individual types of every agent is the joint types of all the agents. To the best of our knowledge, this paper is the first to consider the purification of a randomized decision rule that retains the same expected payoffs conditioned on the individual types of every agent in an economic model.

Our paper contributes to the Bayesian mechanism design literature in relying on specific aspects of agents' private information. These information aspects are often crucial in pinning down different properties of the optimal mechanism. For instance, agents with independent

[9]See, for example, Diestel and Uhl (1977, p. 261).

[10]Since the joint types of all the agents carry the full information, the expected allocation probability of a stochastic mechanism conditioned on the joint types is simply the stochastic mechanism itself.

types retain information rents (see Myerson (1981)), whereas the mechanism designer can fully extract the surplus when the agents' types are correlated (see Crémer and McLean (1988)). Our result builds on the assumption that the agents' private information is independent and dispersed. This assumption facilitates the development of the novel methodology of "mutual purification", which lies at the core of our arguments.

The rest of the paper is organized as follows. Section 2 introduces the basics. Section 3 illustrates our equivalence notion and the idea of "mutual purification" through examples. Section 4 presents the equivalence result. Section 5 discusses the benefit of randomness, an implementation perspective of our result, and various assumptions of our result. Section 6 concludes. The appendix contains proofs omitted from the main body of the paper.

2 Preliminaries

2.1 Notation

There is a finite set $\mathcal{I} = \{1, 2, \ldots, I\}$ of risk neutral agents with $I \geq 2$ and a finite set $\mathcal{K} = \{1, 2, ..., K\}$ of social alternatives. The set of possible types V_i of agent i is a closed subset of finite dimensional Euclidean space $\mathbb{R}^l$ with generic element v_i. The set of possible type profiles is $V \equiv V_1 \times V_2 \times \cdots \times V_I$ with generic element $v = (v_1, v_2, \ldots, v_I)$. Write v_{-i} for a type profile of agent i's opponents; that is, $v_{-i} \in V_{-i} = \Pi_{j \neq i} V_j$. Denote by λ the common prior distribution on V. For each $i \in \mathcal{I}$, λ_i is the marginal distribution of λ on V_i and is assumed to be atomless. Throughout this paper, types are assumed to be independent.[11] If $(Y, \mathcal{Y})$ is a measurable space, then ΔY is the set of all probability measures on $(Y, \mathcal{Y})$. If Y is a metric space, then we treat it as a measurable space with its Borel σ-algebra.

2.2 Mechanism

The revelation principle applies, and we restrict attention to direct mechanisms characterized by $K + I$ functions, $\{q^k(v)\}_{k \in \mathcal{K}}$ and $\{t_i(v)\}_{i \in \mathcal{I}}$, where v is the profile of reports, $q^k(v) \geq 0$ is the probability that alternative k is implemented with $\sum_{k \in \mathcal{K}} q^k(v) = 1$, and $t_i(v)$ is the monetary transfer that agent i makes to the mechanism designer. Agent i's gross utility in

[11]Note that we do not make any assumption regarding the correlation of the different coordinates of type v_i for any $i \in \mathcal{I}$.

alternative k is $u_i^k(v_i, v_{-i})$.

For simplicity of exposition, we denote

$$Q_i^k(v_i) = \int_{V_{-i}} q^k(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i})$$

for the interim expected allocation probability (from agent i's perspective) that alternative k is implemented. Also write

$$T_i(v_i) = \int_{V_{-i}} t_i(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i})$$

for the interim expected transfer from agent i to the mechanism designer. Agent i's interim expected utility is

$$\begin{aligned} U_i(v_i) &= \int_{V_{-i}} \left[\sum_{1 \le k \le K} u_i^k(v_i, v_{-i})q^k(v_i, v_{-i}) - t_i(v_i, v_{-i}) \right] \lambda_{-i}(\mathrm{d}v_{-i}) \\ &= \int_{V_{-i}} \left[\sum_{1 \le k \le K} u_i^k(v_i, v_{-i})q^k(v_i, v_{-i}) \right] \lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i). \end{aligned}$$

A mechanism is Bayesian incentive compatible (BIC) if for each agent $i \in \mathcal{I}$ and each type $v_i \in V_i$,

$$\begin{aligned} & U_i(v_i) \ge 0, \text{ and} \\ & U_i(v_i) \ge \int_{V_{-i}} \left[\sum_{1 \le k \le K} u_i^k(v_i, v_{-i})q^k(v_i', v_{-i}) - t_i(v_i', v_{-i}) \right] \lambda_{-i}(\mathrm{d}v_{-i}) \end{aligned}$$

for any alternative type $v_i' \in V_i$.

A mechanism (q, t) is said to be "deterministic" if for almost all type profiles, the mechanism implements some alternative k for sure. That is, for λ-almost all $v \in V, q^k(v) = 1$ for some $1 \le k \le K$.

2.3 Mechanism equivalence

We shall employ the following notion of mechanism equivalence.

Definition 1. *Two mechanisms (q, t) and $(\tilde{q}, \tilde{t})$ are equivalent if and only if they deliver the same interim expected allocation probabilities and the same interim expected utilities for all agents, and the same ex ante expected social surplus.*

Remark 1. *Our equivalence is stronger than the prevailing mechanism equivalence notion. For example, Manelli and Vincent (2010) and Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) define two mechanisms to be equivalent if they deliver the same interim expected utilities for all agents and the same ex ante expected social surplus.*[12]

Remark 2. *The equivalent deterministic mechanism also guarantees the same ex post monetary transfers, and hence the same expected revenue; see Theorem 1.*

3 Examples

3.1 An illustration of equivalent deterministic mechanism

In the first example, we illustrate our mechanism equivalence notion in a single-unit auction environment.[13] The example is kept deliberately simple and its only purpose is to illustrate what we mean by equivalent deterministic mechanism. Our main result is far more general and the proof is much more complex.

Example 1. *There are two bidders, whose valuations are uniformly distributed in $[0, 1]$. Consider the following mechanism. Types are divided into intervals of equal probability and types in the same interval are treated equally. If agents' types belong to the same interval, each agent receives the object with probability $\frac{1}{2}$ and if agents' types belong to different intervals, the agent whose type belongs to $[\frac{1}{2}, 1]$ gets the object. In each cell, the first number is the probability that agent 1 gets the object and the second number is the probability that agent 2 gets the object.*

	$[\frac{1}{2}, 1]$	$[0, \frac{1}{2})$
$[\frac{1}{2}, 1]$	$\frac{1}{2}, \frac{1}{2}$	$1, 0$
$[0, \frac{1}{2})$	$0, 1$	$\frac{1}{2}, \frac{1}{2}$

[12] Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013, Section 4.1) show that the BIC-DIC equivalence breaks down when requiring the same interim expected allocation probability. They also note that "this notion (of interim expected allocation probabilities) becomes relevant when, for instance, the designer is not utilitarian or when preferences of agents outside the mechanism play a role".

[13] With slight adjustments, this example applies to the irregular case in Myerson's setting where the agents' ironed virtual values are the same in some interval.

It is immediate that, the following deterministic mechanism is equivalent in terms of interim expected allocation probabilities. Keeping the transfers unchanged, it is also easy to see the deterministic mechanism is equivalent in terms of interim expected utilities and ex ante social welfare.

	$[\frac{3}{4}, 1]$	$[\frac{2}{4}, \frac{3}{4})$	$[\frac{1}{4}, \frac{2}{4})$	$[0, \frac{1}{4})$
$[\frac{3}{4}, 1]$	$1, 0$	$0, 1$	$1, 0$	$1, 0$
$[\frac{2}{4}, \frac{3}{4})$	$0, 1$	$1, 0$	$1, 0$	$1, 0$
$[\frac{1}{4}, \frac{2}{4})$	$0, 1$	$0, 1$	$1, 0$	$0, 1$
$[0, \frac{1}{4})$	$0, 1$	$0, 1$	$0, 1$	$1, 0$

In Section 4, we show that for whatever randomized mechanism that the mechanism designer may choose to use, however complicated, there exists an equivalent mechanism that is deterministic. In other words, going from mechanisms that are deterministic to randomized mechanisms in general does not enlarge the set of obtainable outcomes.

3.2 Self purification and mutual purification

In this section, we provide two examples to demonstrate the conceptual difference between the existing approach of "self purification" and our approach of "mutual purification".

The first example is motivated by the game of matching pennies, while the second example is a single unit auction. Both games have two agents, and share the same information structure as follows.

1. Agent 1's type is uniformly distributed on $(0, 1]$ with the total probability $1 - \lambda_1(0)$, and has an atom at the point 0 with $\lambda_1(0) > 0$.
2. Agent 2's type is uniformly distributed on $[0, 1]$.
3. Agents' types are independently distributed.

Example 2 below illustrates the idea of "self purification". The behavioral strategy of agent 2 can be purified since the distribution of agent 2's type is atomless, while the behavioral strategy of agent 1 cannot be purified since agent 1's type has an atom.

Example 2. *Consider an $m \times m$ zero-sum generalized "matching pennies" game with incomplete information, where the positive integer m is sufficiently large such that $\frac{1}{m} < \lambda_1(0)$. The information structure is described in the beginning of this subsection. The action space for both agents is $A_1 = A_2 = \{a_1, a_2, \ldots, a_m\}$. The payoff matrix for agent 1 is given below. Notice that the payoffs of both agents do not depend on the type profile.*

Agent 2

		a_1	a_2	a_3	$\cdots$	a_m
	a_1	1	-1	0	$\cdots$	0
	a_2	0	1	-1	$\cdots$	0
Agent 1	a_3	0	0	1	$\cdots$	0
	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
	a_m	-1	0	$\cdots$	0	1

Suppose that both agents adopt the behavioral strategy $f_1(v) = f_2(v) = \frac{1}{m}\sum_{1 \le s \le m} \delta_{a_s}$, where δ_{a_s} is the Dirac measure at the point a_s. It is easy to see that (f_1, f_2) is a Bayesian Nash equilibrium and the expected payoffs of both agents are 0.

Claim 1. *Agent 2 has a pure strategy f_2' such that (f_1, f_2') is still a behavioral-strategy equilibrium and provides both agents the same expected payoffs, while agent 1 does not have such a pure strategy.*

Example 3 below shows how a purification for an agent relies on the dispersed information of the other agent, which partially illustrates the idea of "mutual purification". In particular, for some given randomized mechanism in the 2-agent setting with independent information as specified above, agent 1 who has an atom in her type space can achieve the same interim expected payoff by some deterministic mechanism,[14] while there does not exist such a deterministic mechanism for agent 2 who has dispersed information.

Example 3. *Consider a single unit auction with two agents. The information structure is described as above. The payoff function of agent i is $\epsilon v_i + (1 - v_j)^m$ for $i, j = 1, 2$ and $i \neq j$, where m is sufficiently large and ϵ is sufficiently small such that*

$$\frac{\lambda_1(0)}{2} > \epsilon + \frac{1}{m+1}.$$

[14]For simplicity, we only consider such an equivalence in terms of interim expected payoffs.

The allocation rule q is defined as follows. Let $q^i(v)$ be the probability that agent i gets the object, and $q^1(v_1, v_2) = q^2(v_1, v_2) = \frac{1}{2}$ for any (v_1, v_2). The interim expected payoff of agent 1 with value v_1 is

$$\int_{V_2} (\epsilon v_1 + (1 - v_2)^m) q^1(v_1, v_2) \lambda_2(\mathrm{d}v_2) = \frac{\epsilon v_1}{2} + \frac{1}{2(m+1)}.$$

The interim expected payoff of agent 2 with value v_2 is

$$\int_{V_1} (\epsilon v_2 + (1 - v_1)^m) q^2(v_1, v_2) \lambda_1(\mathrm{d}v_1) = \frac{\epsilon v_2}{2} + \frac{\lambda_1(0)}{2} + (1 - \lambda_1(0)) \frac{1}{2(m+1)}.$$

Claim 2. *There exists a deterministic mechanism which gives agent 1 the same interim expected payoff; but there does not exist such a deterministic mechanism for agent 2.*

We hasten to emphasize the key difference between our approach and the purification method used in the literature. With the classical purification method, each agent uses her own dispersed information to purify her behavioral strategy, which we call "self purification". In contrast, the purification approach we adopt to achieve our main result is to purify the randomized mechanism via other agents' dispersed information while keeping each agent's interim expected allocation probability and interim expected payoff unchanged *simultaneously*, which we call "mutual purification".

4 Results

This section establishes the main result of this paper. We consider a general environment in which agents could have nonlinear and interdependent payoffs. In particular, we assume that all agents have "separable payoffs" in the following sense.

Definition 1. *For each $i \in \mathcal{I}$, agent i is said to have separable payoff if for any outcome $k \in \mathcal{K}$ and type profile $v = (v_1, v_2, \ldots, v_I) \in V$, her payoff function can be written as follows:*

$$u_i^k(v_1, \ldots, v_I) = \sum_{1 \le m \le M} w_{im}^k(v_i) r_{im}^k(v_{-i}),$$

where M is a positive integer, and w_{im}^k (resp. r_{im}^k) is λ_i-integrable (resp. λ_{-i}-integrable) on V_i (resp. on V_{-i}) for $1 \le m \le M$.

That is, the payoff of each agent i is a summation of finite terms, where each term is a product of two components: the first component only depends on agent i's own type, while the second component depends on other agents' types. This setup is sufficiently general to cover most applications. In particular, it includes the interdependent payoff function as in Jehiel and Moldovanu (2001), and obviously covers the widely adopted private value payoffs as a special case.

Theorem 1. *Suppose that for each agent $i \in \mathcal{I}$, his payoff function is separable. Then for any mechanism (q, t), there exists a deterministic allocation rule $\tilde{q}$ such that*

1. *q and $\tilde{q}$ induce the same interim expected allocation probability;*

2. *$(\tilde{q}, t)$ delivers the same interim expected utility with (q, t) for each agent $i \in \mathcal{I}$.*

Thus, if (q, t) is BIC, then $(\tilde{q}, t)$ is also BIC.

Remark 3. *We prove a stronger result. First, it is clear from the proof of Theorem 1 that the equivalent deterministic mechanism $(\tilde{q}, t)$ also guarantees the same ex post monetary transfers. Therefore, our deterministic mechanism equivalence result does not require transferable utility. Second, the equivalence result is immune against coalitions; that is, when there is sharing of information between the coalition members (except for the grand coalition).[15] The second point is proved explicitly.*

5 Discussions

5.1 Benefit of randomness revisited

Chawla, Malec, and Sivan (2015) consider *multi-agent* setting and focus on the case where the agents' values are independent both across different agents' types and different coordinates of an agent's type. In particular, Chawla, Malec, and Sivan (2015, Theorem13) establish a constant factor upper bound for the benefit of randomness when the agents' values are independent. In the special case of multi-unit multi-item auctions, they show that the revenue of any Bayesian incentive compatible, individually rational randomized mechanism is at most 33.75 times the revenue of the optimal deterministic mechanism. In this paper, we push this

[15] Jackson and Sonnenschein (2007) also consider the issue of coalitional incentive compatibility. They show that the "linking mechanisms" are immune to manipulations by coalitions.

result to the extreme and show that the revenue maximizing auction can be deterministically implemented.[16]

5.2 An implementation perspective

We have motivated our result broadly, in terms of revenue, social surplus, interim expected allocation probabilities, interim expected utilities and even ex post payments. Alternatively, we may take an implementation perspective to formulate our result. Beyond the equivalence notion discussed throughout the paper, the deterministic allocation rule can also be required to pick some allocation in the support of the randomized allocation in the stochastic mechanism for each type profile v. Therefore, when a stochastic mechanism implements some social goal (i.e., at every type profile v, every realized allocation is consistent with the social goal), our equivalent deterministic mechanism also has the same property. We shall explain this point in the following paragraph.

Suppose that q is a random allocation rule. Given the K alternatives, the set of all nonempty subsets of $\{1, \dots, K\}$ can have at most $2^K - 1$ elements $\{C_j\}_{1 \le j \le 2^K - 1}$. As a result, the set of type profiles V can be divided into $2^K - 1$ disjoint subsets $\{D_j\}_{1 \le j \le 2^K - 1}$ such that

1. the support of $q(v)$ is C_j for all $v \in D_j$;
2. $\lambda(\cup_{1 \le j \le 2^K - 1} D_j) = 1$.

We define $2^K - 1$ functions $\{\beta_j\}_{1 \le j \le 2^K - 1}$ such that $\beta_j = 1 + \mathbf{1}_{D_j}$ for each j; that is, β_j is the summation of 1 and the indicator function of the set D_j. Instead of working with the function h, we can work with the new function $h' = (h, \beta_1, \dots, \beta_{2^K - 1})$. Lemma 1 and Proposition 1 (in the Appendix) still hold, and we can obtain a deterministic mechanism $\tilde{q}$ such that

$$\int_V q \beta_j \, \mathrm{d}\lambda = \int_V \tilde{q} \beta_j \, \mathrm{d}\lambda$$

for each j, and

$$\int_V q \, \mathrm{d}\lambda = \int_V \tilde{q} \, \mathrm{d}\lambda.$$

That is, $\int_{D_j} q \, \mathrm{d}\lambda = \int_{D_j} \tilde{q} \, \mathrm{d}\lambda$ for each j. Since $\sum_{k \in C_j} q^k(v) = 1$ for λ-almost all $v \in D_j$, $\int_{D_j} \sum_{k \in C_j} q^k(v) \lambda(\mathrm{d}v) = \lambda(D_j)$, which implies that $\int_{D_j} \sum_{k \in C_j} \tilde{q}^k(v) \lambda(\mathrm{d}v) = \lambda(D_j)$. As a result,

[16] Chawla, Malec, and Sivan (2015, p. 316) remarked that "our bounds on the benefit of randomness are in some cases quite large and we believe they can be improved".

for λ-almost all $v \in D_j$, $\tilde{q}^k = 1$ for some $k \in C_j$. This proves our claim that the deterministic allocation rule lies in the support of the random allocation rule.

5.3 Assumptions

This subsection discusses the assumptions behind our equivalence result. The requirement of multiple agents needs no further explanation. Atomless distribution is an indispensable requirement for almost all purification results. See Example 3 where we cannot purify the allocation for agent 2 while keeping her interim expected utility unchanged because agent 1's type distribution has an atom, let alone the stronger requirement that the deterministic mechanism requires such purification for all agents simultaneously. While our result requires independence, it is worth mentioning that we only require independence across agents and we do not make any assumption regarding the correlation of the different coordinates of type v_i for any agent $i \in \mathcal{I}$. Though separable payoff is a restriction, this setup is sufficiently general to cover most applications; see Section 4 for details.

6 Conclusion

We prove the following mechanism equivalence result: in a general social choice environment with multiple agents, for any stochastic mechanism, there exists an equivalent deterministic mechanism. On the one hand, our result implies that it is without loss of generality to work with stochastic mechanisms, even if the mechanism designer does not have access to a randomization device, or cannot fully commit to the outcomes induced by a randomization device. On the other hand, our result implies that the requirement of deterministic mechanisms is not restrictive in itself. Even if one is constrained to employ only deterministic mechanisms, there is no loss of revenue or social welfare. Therefore, our result provides a foundation for the use of deterministic mechanisms in mechanism design settings, such as auctions, bilateral trades, etc.

References

Aliprantis, C. D., and K. C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer.

Arkin, V. I., and V. L. Levin (1972): "Convexity of Values of Vector Integrals, Theorems on Measurable Choice and Variational Problems," *Russian Mathematical Surveys*, 27(3), 21–85.

Bester, H., and R. Strausz (2001): "Contracting with Imperfect Commitment and the Revelation Principle: the Single Agent Case," *Econometrica*, 69(4), 1077–1098.

Chawla, S., D. L. Malec, and B. Sivan (2015): "The Power of Randomness in Bayesian Optimal Mechanism Design," *Games and Economic Behavior*, 91, 297–317.

Chen, Y.-C., W. He, J. Li, and Y. Sun (2016): "Supplement to 'Equivalence of Stochastic and Deterministic Mechanisms'," mimeo, National University of Singapore and Chinese University of Hong Kong.

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

Deb, R., and M. Pai (2015): "Discrimination via Symmetric Auctions," mimeo, University of Toronto and University of Pennsylvania.

Diestel, J., and J. J. Uhl (1977): *Vector Measures*. Mathematical Surveys, Vol. 15, American Mathematical Society.

Dvoretzky, A., A. Wald, and J. Wolfowitz (1950): "Elimination of Randomization in Certain Problems of Statistics and of the Theory of Games," *Proceedings of the National Academy of Sciences of the United States of America*, 36(4), 256–260.

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.

Govindan, S., P. J. Reny, and A. J. Robson (2003): "A Short Proof of Harsanyi's Purification Theorem," *Games and Economic Behavior*, 45(2), 369–374.

Harsanyi, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-strategy Equilibrium Points," *International Journal of Game Theory*, 2(1), 1–23.

Hart, S., and P. J. Reny (2015): "Maximal Revenue with Multiple Goods: Nonmonotonicity and Other Observations," *Theoretical Economics*, 10(3), 893–922.

Jackson, M. O., and H. F. Sonnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," *Econometrica*, 75(1), 241–257.

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

Khan, M. A., K. P. Rath, and Y. Sun (2006): "The Dvoretsky–Wald–Wolfowitz Theorem and Purification in Atomless Finite-action Games," *International Journal of Game Theory*, 34(1), 91–104.

Laffont, J.-J., and D. Martimort (2002): *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.

Manelli, A., and D. Vincent (2006): "Bundling as an Optimal Selling Mechanism for a Multiple-good Monopolist," *Journal of Economic Theory*, 127(1), 1–35.

——— (2007): "Multidimensional Mechanism Design: Revenue Maximization and the Multiple-good Monopoly," *Journal of Economic Theory*, 137(1), 153–185.

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

McAfee, R. P., and J. McMillan (1988): "Multidimensional Incentive Compatibility and Mechanism Design," *Journal of Economic Theory*, 46(2), 335–354.

Milgrom, P. R., and R. J. Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10(4), 619–632.

Morris, S. (2008): "Purification," in *The New Palgrave Dictionary of Economics*, ed. by S. N. Durlauf, and L. E. Blume. Palgrave Macmillan.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–71.

Pavlov, G. (2011): "Optimal Mechanism for Selling Two Goods," *The BE Journal of Theoretical Economics*, 11(1), Article 3.

Pycia, M. (2006): "Stochastic vs Deterministic Mechanisms in Multidimensional Screening," mimeo, University of California Los Angeles.

Radner, R., and R. W. Rosenthal (1982): "Private Information and Pure-Strategy Equilibria," *Mathematics of Operations Research*, 7(3), 401–409.

Riley, J., and R. Zeckhauser (1983): "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, 98(2), 267–289.

Rochet, J.-C., and J. Thanassoulis (2015): "Stochastic Bundling," mimeo, University of Zurich and University of Warwick.

Royden, H. L., and P. M. Fitzpatrick (2010): *Real Analysis, Fourth Edition.* Prentice Hall.

Strausz, R. (2003): "Deterministic Mechanisms and the Revelation Principle," *Economics Letters*, 79(3), 333–337.

——— (2006): "Deterministic versus Stochastic Mechanisms in Principal-Agent Models," *Journal of Economic Theory*, 128(1), 306–314.

Thanassoulis, J. (2004): "Haggling over Substitutes," *Journal of Economic Theory*, 117(2), 217–245.

A Appendix

A.1 Proof of Claims 1 and 2

Proof of Claim 1. It is easy to see that the following pure strategy f_2' gives agent 2 the same expected payoff and (f_1, f_2') is still a Bayesian Nash equilibrium, where

$$f_2'(v) = \begin{cases} a_s, & v \in [\frac{s-1}{m}, \frac{s}{m}), 1 \leq s \leq m-1; \\ a_m, & v \in [\frac{m-1}{m}, 1]. \end{cases}$$

We next show that there does not exist a pure strategy g_1 of agent 1 such that g_1 is a component of a Bayesian Nash equilibrium with each agent's expected payoff being 0. Suppose that (g_1, g_2) is a Bayesian Nash equilibrium such that g_1 is a pure strategy of agent 1. Let $D_s = \{v_1 \in V_1 \colon g_1(v_1) = a_s\}$ for $1 \leq s \leq m$. Without loss of generality, we assume that $0 \in D_1$. Let $S = \arg\max_{1 \leq s \leq m} \lambda_1(D_s)$. Since $\lambda_1(D_s) \geq \lambda_1(D_1) \geq \lambda_1(0) > \frac{1}{m}$ for each $s \in S$,

S must be a strict subset of $\{1,\dots,m\}$. Without loss of generality, we assume that $s^* \in S$ and $s^*+1 \notin S$. Given agent 1's strategy g_1, agent 2 can adopt the pure strategy $g_2'(v_2) = a_{s^*+1}$ for any $v_2 \in V_2$. Then the expected payoff of agent 2 is $\lambda_1(D_{s^*}) - \lambda_1(D_{s^*+1}) > 0$ with the strategy profile (g_1, g_2'). Since (g_1, g_2) is a Bayesian Nash equilibrium, the expected payoff of agent 2 must be at least $\lambda_1(D_{s^*}) - \lambda_1(D_{s^*+1})$ with the strategy profile (g_1, g_2), which is strictly positive. This is a contradiction. □

Proof of Claim 2. We first construct a deterministic mechanism which gives agent 1 the same interim expected payoff. Define a function G on $V_1 \times V_2 = [0,1]^2$ by letting

$$G(v_1, v_2) = \int_0^{v_2} \left[\epsilon v_1 + (1 - v_2')^m\right] \lambda_2(\mathrm{d}v_2') - \left[\frac{\epsilon v_1}{2} + \frac{1}{2(m+1)}\right],$$

for any $(v_1, v_2) \in V_1 \times V_2$. It is clear that for any $v_1 \in [0,1]$, $G(v_1, 0) < 0 < G(v_1, 1) = \frac{\epsilon v_1}{2} + \frac{1}{2(m+1)}$. One can also check that $\frac{\partial G}{\partial v_2} = \epsilon v_1 + (1 - v_2)^m > 0$ for any $v_1 \in [0,1]$ and $v_2 \in [0,1)$. Hence, for each $v_1 \in [0,1]$, there exists a unique number $g(v_1) \in (0,1)$ such that $G(v_1, g(v_1)) = 0$. By the usual implicit function theorem, g must be differentiable, and hence measurable. Let $\hat{q}^1(v_1, v_2) = 1$ if $0 \le v_2 \le g(v_1)$ and 0 otherwise, and $\hat{q}^2(v_1, v_2) = 1 - \hat{q}^1(v_1, v_2)$. Then the mechanism $\hat{q}$ gives agent 1 the same interim expected payoff.

We next show that there does not exist any deterministic mechanism that gives agent 2 the same interim expected payoff. Suppose that there exists a deterministic mechanism $\tilde{q}$ that gives agent 2 the same interim expected payoff. Fix value $v_2 \in V_2 = [0,1]$.

Suppose that $\tilde{q}^2(0, v_2) = 1$. Then the interim expected payoff of agent 2 with value v_2 is

$$\int_{V_1} (\epsilon v_2 + (1 - v_1)^m) \tilde{q}^2(v_1, v_2) \lambda_1(\mathrm{d}v_1) \ge (\epsilon v_2 + 1)\lambda_1(0).$$

Recall that $\frac{\lambda_1(0)}{2} > \epsilon + \frac{1}{m+1}$. Hence we have

$$(\epsilon v_2 + 1)\lambda_1(0) \ge \lambda_1(0) > \frac{\lambda_1(0)}{2} + \epsilon + \frac{1}{m+1} > \frac{\epsilon v_2}{2} + \frac{\lambda_1(0)}{2} + (1 - \lambda_1(0))\frac{1}{2(m+1)}.$$

Thus, the interim expected payoff of agent 2 under the mechanism $\tilde{q}$ is strictly greater than the interim expected payoff of agent 2 under the mechanism q. This is a contradiction. Therefore, we must have $\tilde{q}^2(0, v_2) = 0$ since $\tilde{q}$ is a deterministic mechanism.

Next, since $\tilde{q}^2(0, v_2) = 0$, the interim expected payoff of agent 2 is

$$\int_{V_1} (\epsilon v_2 + (1 - v_1)^m) \tilde{q}^2(v_1, v_2) \lambda_1(\mathrm{d}v_1) = \int_{(0,1]} (\epsilon v_2 + (1 - v_1)^m) \tilde{q}^2(v_1, v_2) \lambda_1(\mathrm{d}v_1)$$

$$\leq (1-\lambda_1(0))\int_0^1 (\epsilon v_2 + (1-v_1)^m)\,\mathrm{d}v_1 = (1-\lambda_1(0))\epsilon v_2 + \frac{1-\lambda_1(0)}{m+1}$$
$$< \epsilon + \frac{1}{m+1} < \frac{\lambda_1(0)}{2}$$
$$< \frac{\epsilon v_2}{2} + \frac{\lambda_1(0)}{2} + (1-\lambda_1(0))\frac{1}{2(m+1)}.$$

That is, the interim expected payoff of agent 2 under the mechanism $\tilde{q}$ is strictly less than the interim expected payoff of agent 2 under the mechanism q. This is also a contradiction. Therefore, there does not exist any deterministic mechanism that gives agent 2 the same interim expected payoff. □

A.2 Proof of Theorem 1

Let h be a function from V to $\mathbb{R}_{++}^{IKM+1}$ such that $h_0(v) \equiv 1$, and $h_{ikm}(v) = r_{im}^k(v_{-i})$[17] for each $i \in \mathcal{I}$, $1 \leq k \leq K$ and $1 \leq m \leq M$.[18] Let $\mathcal{J}$ be the set of all nonempty proper subsets of $\mathcal{I}$, and Υ be the set of all allocation rules. That is, given any $\tilde{q} \in \Upsilon$, $\tilde{q}$ is a measurable function and $\sum_{k\in\mathcal{K}} \tilde{q}^k(v) = 1$ for λ-almost all $v \in V$. For any coalition $J \subseteq \mathcal{I}$, denote $\lambda_J = \bigotimes_{j\in J} \lambda_j$.

Fix a Bayesian incentive compatible mechanism (q,t). We consider the allocation rule $\tilde{q} \in \Upsilon$ such that for any $J \in \mathcal{J}$ and λ_J-almost all $v_J \in V_J$,

$$\mathbb{E}\left(\tilde{q}h_j | v_J\right) = \mathbb{E}\left(q h_j | v_J\right) \tag{1}$$

for $j = 0$ or $j = ikm$, $i \in \mathcal{I}$, $1 \leq k \leq K$ and $1 \leq m \leq M$.

Definition 2. *We define the following set* Υ_q*:*

$$\Upsilon_q = \{\tilde{q} \in \Upsilon \colon \tilde{q} \text{ satisfies Equation (1)}\}.$$

In what follows, we first provide the following characterization result for the set Υ_q: Υ_q is a nonempty, convex and weakly compact set in some Banach space. Therefore, the classical Krein-Milman Theorem (see Royden and Fitzpatrick (2010, p. 296)) implies that Υ_q admits

[17] Throughout this paper, IKM is the product of the integers I, K and M. However, the subscript ikm is not the product of the numbers i, k and m, but refers to the vector (i,k,m) identifying the function r_{im}^k.

[18] Denote $\mathbb{R}_{++}$ as the strictly positive real line. We assume that h is strictly positive without loss of generality. Indeed, we can work with the function h' from V to $\mathbb{R}_{+}^{2IKM+1}$ such that $h'_0(v) \equiv 1$, $h'^1_{ikm}(v) = |r_{im}^k(v_{-i})| + 1$, and $h'^2_{ikm}(v) = r_{im}^k(v_{-i}) + |r_{im}^k(v_{-i})| + 1$ for each $i \in \mathcal{I}$, $1 \leq k \leq K$ and $1 \leq m \leq M$. The function h' is strictly positive and suffices for our purpose.

extreme points. We proceed by showing that all extreme points of the set Υ_q are deterministic mechanisms.[19] The existence of a deterministic mechanism that is equivalent in terms of interim expected allocation probabilities immediately follows. The equivalence in terms of interim expected utilities and ex ante expected social surplus follows from Equation (4) and the separable payoff assumption. The incentive compatibility of the deterministic mechanism follows from Equation (4) and the assumption that types are independent.

The following lemma characterizes the set Υ_q.

Lemma 1. Υ_q *is a nonempty, convex and weakly compact subset.*

Proof of Lemma 1. Clearly, the set Υ_q is nonempty and convex. We first show that Υ_q is norm closed in $L_1^\lambda(V, \mathbb{R}^K)$, where $L_1^\lambda(V, \mathbb{R}^K)$ is the L_1 space of all measurable mappings from V to $\mathbb{R}^K$ under the probability measure λ.

Suppose that the sequence $\{q_m\} \subseteq \Upsilon_q$ and $q_m \to q_0$ in $L_1^\lambda(V, \mathbb{R}^K)$. Then by the Riesz-Fischer Theorem (see Royden and Fitzpatrick (2010, p. 398)), there exists a subsequence $\{q_{m_s}\}$ of $\{q_m\}$, which converges to q_0 λ-almost everywhere. Since $\sum_{k\in\mathcal{K}} q^k_{m_s}(v) = 1$ for λ-almost all v, $\sum_{k\in\mathcal{K}} q^k_0(v) = 1$ for λ-almost all v. As a result, $q_0 \in \Upsilon$.

For any $k \in \mathcal{K}$, $J \in \mathcal{J}$, and $\mathcal{B}(V_J) \otimes \left(\bigotimes_{1\le j\le I, j\notin J}\{V_j, \emptyset\}\right)$-measurable bounded mapping $p \colon V \to \mathbb{R}^K$,

$$\int_V (q^k_0 h_j) p\lambda(\mathrm{d}v) = \lim_{s\to\infty} \int_V (q^k_{m_s} h_j) p\lambda(\mathrm{d}v) = \int_V (q^k h_j) p\lambda(\mathrm{d}v)$$

for $j = 0$ or $j = ikm$. The first equality is due to the dominated convergence theorem, and the second equality holds since $\{q_{m_s}\} \subseteq \Upsilon_q$. Thus, $q_0 \in \Upsilon_q$, which implies that Υ_q is norm closed in $L_1^\lambda(V, \mathbb{R}^K)$.

Since Υ_q is convex, Υ_q is also weakly closed in $L_1^\lambda(V, \mathbb{R}^K)$ by Mazur's Theorem (see Royden and Fitzpatrick (2010, p. 292)). As Υ is weakly compact in $L_1^\lambda(V, \mathbb{R}^K)$, we have that Υ_q is weakly compact in $L_1^\lambda(V, \mathbb{R}^K)$, and hence has extreme points. □

Since Υ_q is a nonempty, convex and weakly compact set, Υ_q has extreme points. The following result shows that all extreme points of Υ_q are deterministic allocations.

[19] Manelli and Vincent (2007) use a related technique in the screening literature. Manelli and Vincent (2007) consider revenue maximizing multi-product monopolist and study the extreme points of the set of feasible mechanisms. They show that, with multiple goods, extreme points could be stochastic mechanisms. In contrast, we work with the mechanism design setting, study a particular set of interest Υ_q and show that all extreme points are deterministic. Apart from this general approach, the technical parts of the proofs are dramatically different.

Proposition 1. *All extreme points of Υ_q are deterministic allocations.*

Proof of Proposition 1. Pick an allocation rule $\tilde{q} \in \Upsilon_q$ which is not deterministic, we shall show that $\tilde{q}$ is not an extreme point of Υ_q.

Since $\tilde{q}$ is not deterministic, there is a positive number $0 < \delta < 1$, a Borel measurable set $D \subseteq V$ such that $\lambda(D) > 0$, and indices j_1, j_2 such that $\delta \leq \tilde{q}^{j_1}(v), \tilde{q}^{j_2}(v) \leq 1 - \delta$ for any $v \in D$. For any $J \in \mathcal{J}$, let D_J be the projection of D on $\prod_{j \in J} V_j$. For any $v_J \in D_J$, let $D_{-J}(v_J) = \{v_{-J} \colon (v_J, v_{-J}) \in D\}$ (abbreviated as D_{v_J}).

Consider the following problem on $\alpha \in L_\infty^\lambda(D, \mathbb{R})$: for any $J \in \mathcal{J}$ and $v_J \in D_J$,

$$\int_{D_{-J}(v_J)} \alpha(v_J, v_{-J}) h(v_J, v_{-J}) \lambda_{-J}(\mathrm{d}v_{-J}) = 0. \tag{2}$$

Recall that h is a function taking values in $\mathbb{R}^{IKM+1}$. For simplicity, denote $l_0 = IKM+1$. Define the set $\mathscr{E}$ as

$$\mathscr{E} = \{h(v) \cdot \sum_{J \in \mathcal{J}} \psi_J(v_J) \colon \psi_J \in L_\infty^\lambda(D_J, \mathbb{R}^{l_0}), \forall J \in \mathcal{J}\}.$$

Then a bounded measurable function α in $L_\infty^\lambda(D, \mathbb{R})$ is a solution to Problem (2) if and only if $\int_D \alpha \varphi d\lambda = 0$ for any $\varphi \in \mathscr{E}$. We show in the Supplemental Material (Chen, He, Li, and Sun (2016, Lemma A.3)) that $\mathscr{E}$ is not dense in $L_1^\lambda(D, \mathbb{R})$. By Corollary 5.108 in Aliprantis and Border (2006), Problem (2) has a nontrivial bounded solution α.

Without loss of generality, we assume that $|\alpha| \leq \delta$. We extend the domain of α to V by letting $\alpha(v) = 0$ when $v \notin D$. For every $v \in V$, define

$$\hat{q}(v) = \tilde{q}(v) + \alpha(v) \left(e_{j_1} - e_{j_2}\right);$$

$$\overline{q}(v) = \tilde{q}(v) + \alpha(v) \left(e_{j_2} - e_{j_1}\right).$$

Then $\sum_{k \in \mathcal{K}} \hat{q}^k(v) = \sum_{k \in \mathcal{K}} \overline{q}^k(v) = \sum_{k \in \mathcal{K}} \tilde{q}^k(v) = 1$. If $v \in D$, then $0 \leq \hat{q}^{j_1}(v), \overline{q}^{j_2}(v) \leq 1$ as $\delta \leq \tilde{q}^{j_1}(v), \tilde{q}^{j_2}(v) \leq 1 - \delta$, and $\hat{q}^j(v) = \overline{q}^j(v) = \tilde{q}^j(v)$ for $j \neq j_1, j_2$. If $v \notin D$, then $\hat{q}(v) = \overline{q}(v) = \tilde{q}(v)$ as $\alpha(v) = 0$. Thus, $\hat{q}$, $\overline{q} \in \Upsilon$.

For any $J \in \mathcal{J}$ and $\mathcal{B}(V_J) \otimes \left(\bigotimes_{1 \leq j \leq I, j \notin J} \{V_j, \emptyset\}\right)$-bounded measurable mapping $p \in L_\infty^\lambda(V, \mathbb{R}^K)$,

$$\int_V (\hat{q} \cdot p) h \lambda(\mathrm{d}v) = \int_V (\tilde{q} \cdot p) h \lambda(\mathrm{d}v) + \int_V \alpha(v) \left((e_{j_1} - e_{j_2}) \cdot p(v)\right) h(v) \lambda(\mathrm{d}v).$$

Since

$$\int_V \alpha(v) \left((e_{j_1} - e_{j_2}) \cdot p(v)\right) h(v) \lambda(\mathrm{d}v)$$

$$
\begin{aligned}
&= \int_{V_J} \int_{V_{-J}} \alpha(v) \left((e_{j_1} - e_{j_2}) \cdot p(v) \right) h(v) \lambda_{-J}(\mathrm{d}v_{-J}) \lambda_J(\mathrm{d}v_J) \\
&= \int_{V_J} (e_{j_1} - e_{j_2}) \cdot p(v) \int_{V_{-J}} \alpha(v) h(v) \lambda_{-J}(\mathrm{d}v_{-J}) \lambda_J(\mathrm{d}v_J) \\
&= 0,
\end{aligned}
$$

we have that

$$\int_V (\hat{q} \cdot p) h \lambda(\mathrm{d}v) = \int_V (\tilde{q} \cdot p) h \lambda(\mathrm{d}v),$$

which implies that $\hat{q} \in \Upsilon_q$. Similarly, one can show that $\overline{q} \in \Upsilon_q$. Since $\hat{q}$ and $\overline{q}$ are distinct and $\tilde{q} = \frac{1}{2}(\hat{q} + \overline{q})$, $\tilde{q}$ is not an extreme point of Υ_q. □

Now we are ready to prove our main result.

Proof of Theorem 1. Fix a mechanism (q, t). The proof is then divided into two steps. In the first step, we obtain a deterministic allocation rule $\tilde{q}$ which has the same interim expected allocation probability with q. In the second step, we verify that $(\tilde{q}, t)$ and (q, t) deliver the same interim expected utility for each agent.

By Proposition 1, every extreme point of Υ_q is a deterministic allocation rule. Therefore, we can fix a measurable allocation rule $\tilde{q}$ such that

1. $\tilde{q}^k(v) = 0$ or 1 for λ-almost all $v \in V$ and $1 \le k \le K$;

2. for any agent i and λ_i-almost all $v_i \in V_i$,

$$\int_{V_{-i}} \tilde{q}(v_i, v_{-i}) \lambda_{-i}(\mathrm{d}v_{-i}) = \int_{V_{-i}} q(v_i, v_{-i}) \lambda_{-i}(\mathrm{d}v_{-i}), \tag{3}$$

   and

$$\int_{V_{-i}} \tilde{q}(v_i, v_{-i}) h_{jkm}(v_i, v_{-i}) \lambda_{-i}(\mathrm{d}v_{-i}) = \int_{V_{-i}} q(v_i, v_{-i}) h_{jkm}(v_i, v_{-i}) \lambda_{-i}(\mathrm{d}v_{-i}) \tag{4}$$

   for any $j \in \mathcal{I}$, $1 \le k \le K$ and $1 \le m \le M$.

Let D_i be the subset of V_i such that Equation (3) or (4) does not hold. Then $\lambda_i(D_i) = 0$. Define a new allocation rule $\hat{q}$ such that

$$\hat{q}(v) = \begin{cases} q(v), & \text{if } v_i \in D_i \text{ for some } i \in \mathcal{I}; \\ \tilde{q}(v), & \text{otherwise.} \end{cases}$$

Then $\hat{q}^k(v) = 0$ or 1 for λ-almost all $v \in V$ and $1 \le k \le K$.

Fix agent i and $v_i \in V_i$. If $v_i \in D_i$, then $\hat{q}(v_i, v_{-i}) = q(v_i, v_{-i})$ and $\int_{V_{-i}} \hat{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) = \int_{V_{-i}} q(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i})$. If $v_i \notin D_i$, then

$$\begin{aligned}
\int_{V_{-i}} \hat{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) &= \int_{D_{-i}} \hat{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) + \int_{V_{-i}\setminus D_{-i}} \hat{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) \\
&= \int_{D_{-i}} q(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) + \int_{V_{-i}\setminus D_{-i}} \tilde{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) \\
&= 0 + \int_{V_{-i}} \tilde{q}(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) \\
&= \int_{V_{-i}} q(v_i, v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}),
\end{aligned}$$

where $D_{-i} = \cup_{j\in\mathcal{I}, j\neq i}\left(D_j \times \prod_{s\in\mathcal{I}, s\neq i,j} V_s\right)$. The first equality holds by dividing V_{-i} as D_{-i} and $V_{-i} \setminus D_{-i}$. The second equality is due to the definition of $\hat{q}$. The third equality holds since $\lambda_{-i}(D_{-i}) = 0$. The last equality is due to the condition that $v_i \notin D_i$. As a result, Equation (3) holds for $\hat{q}$ and every $v_i \in V_i$. Similarly, one can check that Equation (4) also holds for $\hat{q}$ and every $v_i \in V_i$.

Suppose that the mechanism $(\hat{q}, t)$ is adopted. By Equation (3), the allocation rules q and $\hat{q}$ induce the same interim expected allocation. We need to check that they induce the same interim expected utility. If agent i observes the state v_i but reports v_i', then his payoff is

$$\begin{aligned}
&\int_{V_{-i}} \left[\sum_{1\leq k\leq K} u_i^k(v_i, v_{-i})\hat{q}^k(v_i', v_{-i}) - t_i(v_i', v_{-i})\right]\lambda_{-i}(\mathrm{d}v_{-i}) \\
=& \sum_{1\leq k\leq K}\sum_{1\leq m\leq M} \int_{V_{-i}} w_{im}^k(v_i) r_{im}^k(v_{-i})\hat{q}^k(v_i', v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i') \\
=& \sum_{1\leq k\leq K}\sum_{1\leq m\leq M} w_{im}^k(v_i) \int_{V_{-i}} r_{im}^k(v_{-i})\hat{q}^k(v_i', v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i') \\
=& \sum_{1\leq k\leq K}\sum_{1\leq m\leq M} w_{im}^k(v_i) \int_{V_{-i}} h_{ikm}(v)\hat{q}^k(v_i', v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i') \\
=& \sum_{1\leq k\leq K}\sum_{1\leq m\leq M} w_{im}^k(v_i) \int_{V_{-i}} h_{ikm}(v)q^k(v_i', v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i') \\
=& \sum_{1\leq k\leq K}\sum_{1\leq m\leq M} w_{im}^k(v_i) \int_{V_{-i}} r_{im}^k(v_{-i})q^k(v_i', v_{-i})\lambda_{-i}(\mathrm{d}v_{-i}) - T_i(v_i') \\
=& \int_{V_{-i}} \left[\sum_{1\leq k\leq K} u_i^k(v_i, v_{-i})q^k(v_i', v_{-i}) - t_i(v_i', v_{-i})\right]\lambda_{-i}(\mathrm{d}v_{-i}).
\end{aligned}$$

The first and second equalities follow from the separable payoff assumption. The fourth equality follows from Equation (4) and also the assumption that types are independent. All other equalities are simple algebras. Thus, these two mechanisms (q, t) and $(\hat{q}, t)$ deliver the same interim expected utility for every agent. If (q, t) is Bayesian incentive compatible, then $(\hat{q}, t)$ is clearly Bayesian incentive compatible. This completes the proof. □