

A Note on Mixed-Nash Implementation*

Navin Kartik[†]

Olivier Tercieux[‡]

January 30, 2012

Abstract

This note considers (complete information) Nash-implementation when mixed strategies are properly accounted for and the outcome space is infinite. We first construct an example in which preferences over lotteries fail the Archimedean axiom and show that, even under the classical sufficient conditions for implementation, the canonical mechanism for implementation fails: there exists a mixed-Nash equilibrium in which the outcome is undesirable with probability one. We then prove that the canonical mechanism works if Bernoulli utilities are bounded.

1 Setup

Our setup and notation are standard in implementation theory; the reader is referred to Maskin (1999) or Maskin and Sjoström (2002, MS hereafter) for details. The environment is a tuple (A, N, Θ) , where A is a set of outcomes, $N := \{1, \dots, n\}$ is a finite set of agents/players, and Θ is a set of possible states of the world. A and Θ are arbitrary topological spaces. Throughout, we treat any topological space, X , as a measurable space with its Borel sigma-field and restrict our attention to the set of Borel probability measures, denoted $\Delta(X)$. Product spaces are endowed with the product topology while the space of Reals is endowed with its usual topology. All functions are assumed to be measurable.

Each agent i has a von Neumann-Morgenstern (vNM) utility function, $u_i : A \times \Theta \rightarrow \mathbb{R}$. Agent i 's lower contour set at $(a, \theta) \in A \times \Theta$ is $L_i(a, \theta) := \{b \in A \mid u_i(b, \theta) \leq u_i(a, \theta)\}$. A social choice rule (SCR) is a function $F : \Theta \rightarrow 2^A \setminus \{\emptyset\}$, i.e. a non-empty-valued correspondence. A social choice function (SCF) is a single-valued SCR, i.e. a function $f : \Theta \rightarrow A$. A SCR F is (*Maskin*) *monotonic* if for all $(a, \theta, \theta') \in A \times \Theta \times \Theta$ the following is true: if $a \in F(\theta)$ and $L_i(a, \theta) \subseteq L_i(a, \theta')$ for all $i \in N$, then $a \in F(\theta')$. A SCR F satisfies *no veto power* (NVP) if for all $(a, j, \theta) \in A \times N \times \Theta$: if $u_i(a, \theta) \geq u_i(b, \theta)$ for all $b \in A$ and all $i \neq j$, then $a \in F(\theta)$.

A mechanism (or game form) is a pair $\Gamma = (M, g)$ where $M = \times_{i \in N} M_i$ and each M_i is the message space for player i and $g : \times_{i \in N} M_i \rightarrow A$ is an outcome function. We denote a message of player i by m_i and a profile of messages by m . We assume that each player i 's preferences over lotteries are represented by a utility function $U_i : \Delta(A) \rightarrow \overline{\mathbb{R}}^1$ such that for any probability measure, $P \in \Delta(A)$, $U_i(P) = \int_A u_i(a) dP(a)$. Pure- and mixed-strategy Nash equilibria are defined in the usual way.

*This paper grew out of our research on implementation with evidence (Kartik and Tercieux, 2011). In a 2008 version of this paper, we developed a mechanism that we have since replaced with Maskin and Sjoström's (2002) mechanism. We thank Eric Maskin for pointing us to Maskin and Sjoström (2002).

[†]Columbia University, Department of Economics. Email: nkartik@gmail.com.

[‡]Paris School of Economics and CNRS. Email: tercieux@pse.ens.fr.

¹ $\overline{\mathbb{R}} \equiv \mathbb{R} \cup \{-\infty, +\infty\}$ is the affinely-extended Reals.

A mechanism Γ implements a SCR F in mixed-strategy Nash equilibrium, or simply MSNE-implements F , if at each state θ the game induced by Γ satisfies: (1) for all $a \in F(\theta)$, there exists a pure Nash equilibrium m^* such that $g(m^*) = a$, and (2) in any mixed-Nash equilibrium $\mu \equiv (\mu_i)_{i \in N} \in \times_{i \in N} \Delta(M_i)$, $\mu(m) > 0 \implies g(m) \in F(\theta)$. A SCR F is MSNE-implementable if there exists a mechanism which MSNE-implements F . Note that we use a slightly less demanding notion than that of Maskin 1999 (p. 25): instead of condition (2) in Maskin, we only require that $g(m) \in F(\theta)$ for all m such that $\mu(m) > 0$. In other words, we ignore outcomes that are undesirable but occur with 0 probability. This will make no difference if A is countable.

2 A counter-example to MSNE implementation

The classic result (Maskin, 1999) on Nash implementation says that a monotonic SCR is implementable if it satisfies NVP and there are three or more players. The mechanism in Maskin 1999 works insofar as one considers pure-strategy implementation. MS provide a canonical mechanism to obtain MSNE-implementation.

In this section, we construct an example in which MSNE-implementation cannot be achieved by the mechanism proposed by MS. The reason we focus on the MS mechanism is because we are not aware of any other canonical mechanism in the literature that achieves MSNE-implementation when the MS mechanism cannot.² The key feature of our example is a failure of continuity of preferences over lotteries; more precisely, the Archimedean axiom is violated. We suspect, but have not proved, that in our example MSNE-implementation is impossible.

The MS mechanism is as follows (readers should consult MS for details). For each player i ,

$$M_i = A \times \Theta \times \{\alpha_i \mid \alpha_i : A \times \Theta \rightarrow A \text{ and } \alpha_i(a, \theta) \in L_i(a, \theta) \text{ for all } (a, \theta)\} \times \mathbb{N}.$$

There are three Rules governing outcomes:

Rule (i): If there is $j \in N$ s.t. $m_i = (a, \theta, \cdot, 1)$ for all $i \neq j$, $a \in F(\theta)$ and $n^j = 1$, take $g(m) = a$.³

Rule (ii): If there is $j \in N$ s.t. $m_i = (a, \theta, \cdot, 1)$ for all $i \neq j$, $a \in F(\theta)$ and $n^j > 1$ take $g(m) = \alpha_j(a, \theta)$.

Rule (iii): If neither (i) nor (ii) applies, then $g(m) = a_{i_*}$, where $i_* = \max \{i \mid n^i = \max_j n^j\}$ and $m_{i_*} = (a_{i_*}, \cdot, \cdot, \cdot)$

MS assert that, with three or more players, the above mechanism MSNE-implements any monotonic SCR satisfying NVP. MS only provide a formal argument for mixed-Nash equilibria with finite support. We will show through a counter-example that implementation fails when one considers all mixed-Nash equilibria, including those with infinite support. In fact, in our counter-example the MSNE is “bad” in a strong sense because *all* points in the support yield undesirable outcomes.

²Maskin (1999) offers a different mechanism whose goal was to achieve MSNE-implementation under the same conditions as MS. This mechanism also fails to work in our example.

³To be precise, we have modified Rule (i) from MS. MS’s Rule (i) is written as follows: If there is $j \in N$ s.t. $m_i = (a, \theta, \cdot, 1)$ for all $i \neq j$, and $n^j = 1$, take $g(m) = a$. However, the mechanism defined in this way does not work. Here is a counter-example: $\Theta = \{\theta, \theta'\}$, $A = \{a, b\}$ and $n = 3$. At each state, each agent strictly prefers a to b . The SCF f is a constant, $f(\cdot) = a$. Plainly, f is monotonic and satisfies NVP. But at each state, there is an undesirable (pure-strategy) equilibrium where $m_i = (b, \theta, \alpha, 1)$ for all i (i.e. where we fall into Rule (i) as defined above). This is an equilibrium because if some player i deviates, we would fall into Rule (ii) where the outcome can only be in $L_i(b, \theta) = \{b\}$, and hence no deviation would be profitable.

The counter-example is as follows. Assume three players $i = 1, 2, 3$. The set of outcomes is $A = \mathbb{N}^* \cup \{a\} \cup \{b\}$. There are two states of the world, $\Theta = \{\theta, \theta'\}$. The SCF is defined by $f(\theta) = a$ and $f(\theta') = b$. Fix any probability distribution $\mu \in \Delta(\mathbb{N}^*)$ such that $\mu(k) > 0$ for all $k \in \mathbb{N}^*$. We now stipulate some conditions on preferences over lotteries. We do this in terms of the Bernoulli utility functions $u_i : A \rightarrow \mathbb{R}$; recall that preferences over lotteries are represented by $U_i(\cdot)$ such that for any $P \in \Delta(A)$, $U_i(P) = \sum_{x \in A} P(x)u_i(x)$. Our first requirement is that at state θ :

$$(a1) \quad \forall k \in \mathbb{N}^* : u_1(k, \theta) = u_2(k, \theta) = -\frac{k}{\mu(k)}.$$

$$(a2) \quad \forall k \in \mathbb{N}^* : u_3(k, \theta) = u_3(a, \theta) = u_3(b, \theta).$$

$$(a3) \quad u_1(b, \theta) < u_1(1, \theta).$$

The second requirement is that at state θ' :

$$(b1) \quad \forall k \in \mathbb{N}^* : u_3(k, \theta') > u_3(b, \theta') = u_3(a, \theta')$$

$$(b2) \quad u_1(1, \theta') \leq u_1(b, \theta').$$

The third requirement is that

$$\arg \max_{x \in A} u_1(x, \theta) = \arg \max_{x \in A} u_2(x, \theta) = \{a\} \text{ and } \arg \max_{x \in A} u_1(x, \theta') = \arg \max_{x \in A} u_2(x, \theta') = \{b\}.$$

This condition ensures that f satisfies NVP. Note that the f is Maskin-monotonic because of the following two observations:

- $u_3(k, \theta) = u_3(a, \theta)$ and $u_3(k, \theta') > u_3(a, \theta')$;
- $u_1(1, \theta') \leq u_1(b, \theta')$ and $u_1(b, \theta) < u_1(1, \theta)$.

Now consider MS's mechanism and assume the true state is θ . Consider the following profile of mixed strategies: for $i = 1, 2$: μ_i is such that $\mu_i((\theta', b, \alpha_i(\cdot, \cdot), 1)) = 1$, while $\mu_3((\theta', b_k, \alpha_3(\cdot, \cdot), k)) = \mu(k)$ where $\alpha_i(\cdot, \cdot) \in L_i(\cdot, \cdot)$ satisfies $\alpha_i(b, \theta') \equiv b$ for all $i = 1, 2, 3$. In addition, $b_1 = b$ and for $k > 1$: $b_k = k$. So in this profile, players 1 and 2 are playing pure strategies and claiming that the state is θ' , although the true state is θ , and requesting outcome $b \neq f(\theta)$; player 3 is playing a mixed strategy where he always claims the state is θ' , but only puts probability $\mu(1)$ on the pure strategy that 1 and 2 are playing. We will argue that this profile is an MSNE at θ .

Obviously, player 3 has no incentive to deviate (he is indifferent between any outcome). Consider player $i = 1, 2$. Note that if player i sticks to the equilibrium strategy, then with probability $\mu(1)$, we fall into Rule (i) while with probability $1 - \mu(1)$ we fall into Rule (ii). When we fall into Rule (ii), the outcome is b because of the specification of player 3's proposed function $\alpha_3(\cdot, \cdot)$. Hence, i 's expected payoff is

$$\mu(1)u_i(b, \theta) + \sum_{k=2}^{\infty} \mu(k) \times u_i(b, \theta) = u_i(b, \theta).$$

If player i deviates to $m'_i = (\hat{\theta}, \hat{b}, \hat{\alpha}, \hat{k}) \neq (\theta', b, \alpha, 1)$, with probability $\sum_{k=\hat{k}+1}^{\infty} \mu(k)$ we fall into Rule (iii) and player 3 is the "dictator" who chooses outcome $b_k = k$. Hence i 's expected payoff from the deviation

is bounded from above by

$$\sum_{k=1}^{\hat{k}} \mu(k) \max_{c \in A} u_i(c, \theta) + \sum_{k=\hat{k}+1}^{\infty} \mu(k) \times u_i(k, \theta) = \sum_{k=1}^{\hat{k}} \mu(k) u_i(a, \theta) + \sum_{k=\hat{k}+1}^{\infty} \mu(k) \times \left(-\frac{k}{\mu(k)} \right) = -\infty.$$

So player i has no incentive to deviate, proving that the mixed strategy profile is indeed an equilibrium.

Note in addition that $m \in \text{Supp}(\mu) \Rightarrow g(m) = b \neq a = f(\theta)$, hence any point in the support of μ yields an undesirable outcome at θ .

Remark 1. The example uses quite crucially a representation of preferences where the Bernoulli utility function u_i for $i = 1, 2$ is unbounded below (at state θ). Thus, these agents' preferences over lotteries do not satisfy the Archimedean Axiom. Note that this is generally the case in implementation theory when the outcome space is augmented with transfers that are unbounded and agents have quasi-linear utilities in the transfers. As we will show below, unboundedness of the Bernoulli utilities is necessary for failure of MSNE-implementation.⁴

3 Bounded utilities are sufficient for MSNE implementation

In this section, we prove that if the Bernoulli utilities are bounded, then the MS mechanisms achieves MSNE implementation.

Assumption 1. *For any i and any θ , there exists $M \geq 0$ such that $M \geq u_i(\cdot, \theta) \geq -M$.*

Note that this Assumption was violated by the counter-example above.

Remark 2. The assumption is trivially satisfied if the outcome space is finite. The assumption is also satisfied if one requires that preferences over lotteries satisfy the Archimedean Axiom (and don't restrict the space of lotteries), given the vNM representation.

Theorem 1. *Assume $n \geq 3$. Under Assumption 1, any monotonic SCR satisfying NVP is MSNE-implementable.*

The rest of the section is devoted to the proof of this Theorem. It is easy to see, as usual, that "truthful messages" form a desirable pure Nash equilibrium, so we will only prove that there are no undesirable mixed equilibria. Let $\mu \equiv (\mu_i)_{i \in N}$ be a (mixed) Nash equilibrium when the true state is θ' , and fix any realization $m^* \equiv (m_i^*)_{i \in N}$ such that $\mu(m^*) > 0$. We must show that $g(m^*) \in F(\theta')$.

Suppose first that m^* is a realization such that there is j for which

$$\text{for all } i \neq j : m_i^* = (a, \theta, \cdot, 1), a \in F(\theta) \text{ while } n^j = 1$$

but that the profile θ differs from the true profile that is θ' . From Rule (i), the equilibrium outcome is $a \in F(\theta)$. For any player i , consider a sequence of $\{b_i^k\}_{k=1}^{\infty}$ such that $\lim_{k \rightarrow \infty} u_i(b_i^k, \theta') = \sup_{x \in A} u_i(x, \theta')$.⁵ Now

⁴Neither MS nor Maskin (1999) state any assumptions on preferences over lotteries. The only statement we could find about lottery preferences is fn. 7 in Maskin (1999), which states that the mixed-Nash implementation result holds for "any risk preferences."

⁵If there is an outcome that maximizes i 's preferences under θ' — as is the case if A is finite — then we could just set b_i^k to be that outcome. This more cumbersome definition accounts for the possibility that is no maximal element when A is infinite.

imagine that player i plays $m_i^k = (b_i^k, \theta, \alpha_i^k(\cdot, \cdot), k)$ where $\lim_{k \rightarrow \infty} u_i(\alpha_i^k(\tilde{a}, \tilde{\theta}), \theta') = \sup_{x \in L_i(\tilde{a}, \tilde{\theta})} u_i(x, \theta')$ for all $(\tilde{a}, \tilde{\theta}) \in A \times \Theta$. In the sequel, we show that if for some $b \in L_i(a, \theta)$, we have $u_i(b, \theta') > u_i(a, \theta')$, then, for k large enough (and under Assumption 1 made above) player i is better off using m_i^k rather than m_i^* against μ_{-i} , which contradicts the assumption that μ is a Nash equilibrium. We conclude that

$$\text{for all } i, L_i(a, \theta) \subseteq L_i(a, \theta').$$

Hence, since F is monotonic and $a \in F(\theta)$, we have $a \in F(\theta')$ as required.

So let us assume that for some $b \in L_i(a, \theta)$, we have $u_i(b, \theta') > u_i(a, \theta')$, and let us show that for k large enough, player i is better off using m_i^k rather than m_i^* against μ_{-i} .

Claim 1. $\lim_{k \rightarrow \infty} u_i(g(m_i^k, m_{-i}^*), \theta') > u_i(g(m_i^*, m_{-i}^*), \theta')$.

Proof. Either (m_i^k, m_{-i}^*) falls into Rule (ii) for all $k > 1$, or it falls into Rule (iii) for all $k > 1$. If (m_i^k, m_{-i}^*) falls into Rule (ii) for all $k > 1$, it follows that

$$\begin{aligned} \lim_{k \rightarrow \infty} u_i(g(m_i^k, m_{-i}^*), \theta') &= \lim_{k \rightarrow \infty} u_i(\alpha_i^k(a, \theta), \theta') \\ &= \sup_{x \in L_i(a, \theta)} u_i(x, \theta') \geq u_i(b, \theta') > u_i(a, \theta') = u_i(g(m_i^*, m_{-i}^*), \theta') = a \end{aligned}$$

where the first inequality comes from the fact that $b \in L_i(a, \theta)$. On the other hand, if (m_i^k, m_{-i}^*) falls into Rule (iii) for all $k > 1$, it follows that

$$\lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{s}_{-i}), \theta') = \sup_{x \in A} u_i(x, \theta') \geq u_i(b, \theta') > u_i(a, \theta') = u_i(g(m_i^*, m_{-i}^*), \theta'). \quad \blacksquare$$

Claim 2. For any realization \hat{m}_{-i} of μ_{-i} , $\{u_i(g(m_i^k, \hat{m}_{-i}), \theta')\}_{k > 1}$ converges and $\lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{m}_{-i}), \theta') \geq u_i(g(m_i^*, \hat{m}_{-i}), \theta')$.

Proof. Fix any \hat{m}_{-i} . Consider the case where (m_i^*, \hat{m}_{-i}) falls into Rule (i), and let $g(m_i^*, \hat{m}_{-i}) = \hat{a}$. In such a case, for $k > 1$, (m_i^k, \hat{m}_{-i}) falls either into Rule (ii) or into Rule (iii). If for some $k > 1$, (m_i^k, \hat{m}_{-i}) falls into Rule (ii), we must have that for all $k > 1$, (m_i^k, \hat{m}_{-i}) falls into Rule (ii) and hence, $g(m_i^k, \hat{m}_{-i}) = \alpha_i^k(\hat{a}, \hat{\theta})$ for some $\hat{\theta}$ s.t. $\hat{a} \in F(\hat{\theta})$. By definition of $\alpha_i^k(\cdot, \cdot)$,

$$\lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{m}_{-i}), \theta') = \sup_{x \in L_i(\hat{a}, \hat{\theta})} u_i(x, \theta') \geq u_i(\hat{a}, \theta') = u_i(g(m_i^*, \hat{m}_{-i}), \theta').$$

Next, suppose that for some $k > 1$, (m_i^k, \hat{m}_{-i}) falls into Rule (iii). Then we must have that for all $k > 1$, (m_i^k, \hat{m}_{-i}) falls into Rule (iii) and hence, $g(m_i^k, \hat{m}_{-i}) = b_i^k$. Thus,

$$\lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{m}_{-i}), \theta') = \sup_{x \in A} u_i(x, \theta') \geq u_i(g(m_i^*, \hat{m}_{-i}), \theta').$$

Finally, whenever (m_i^*, \hat{m}_{-i}) falls into Rule (ii) or (iii), for k large enough, (m_i^k, \hat{m}_{-i}) falls into Rule (iii) and $g(m_i^k, \hat{m}_{-i}) = b_i^k$. Hence, a similar reasoning as above applies. \blacksquare

The following Lemma is useful, which is where we first use Assumption 1:

Lemma 1. $\lim_{k \rightarrow \infty} \int_{\hat{m}_{-i}} u_i(g(m_i^k, \hat{m}_{-i}), \theta') d\mu_{-i} = \int_{\hat{m}_{-i}} \lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{m}_{-i}), \theta') d\mu_{-i}$.

Proof. Given Assumption 1 and the pointwise convergence from Claim 2, this is an application of the Lebesgue dominated convergence theorem. ■

We complete the argument by now observing that

$$\liminf_{k \rightarrow \infty} \int_{\hat{m}_{-i}} u_i(g(m_i^k, \hat{m}_{-i}), \theta') d\mu_{-i} = \int_{\hat{m}_{-i}} \lim_{k \rightarrow \infty} u_i(g(m_i^k, \hat{m}_{-i}), \theta') d\mu_{-i} > \int_{\hat{m}_{-i}} u_i(g(m_i^*, \hat{m}_{-i}), \theta') d\mu_{-i},$$

where the equality is by Lemma 1 and the strict inequality is by Claims 1 and 2 combined with the fact that $\mu_{-i}(m_{-i}^*) > 0$ and that, by Assumption 1, $\int_{\hat{m}_{-i}} u_i(g(m_i^*, \hat{m}_{-i}), \theta') d\mu_{-i}$ is finite. This implies that for k large enough,

$$\int_{\hat{m}_{-i}} u_i(g(m_i^k, \hat{m}_{-i}), \theta') d\mu_{-i} > \int_{\hat{m}_{-i}} u_i(g(m_i^*, \hat{m}_{-i}), \theta') d\mu_{-i},$$

and hence player i is better off using m_i^k rather than m_i^* against μ_{-i} , a contradiction.

Finally, in case m^* falls into Rule (ii) or Rule (iii) and $g(m^*) = a$, a similar reasoning as above can be used to show that if a is not the most preferred outcome at θ' for at least $n - 1$ players, then there is some player i for whom m_i^k is a profitable deviation. Hence, by NVP, we must also have $a \in F(\theta')$.

4 Concluding remarks

Remark 3. Based on the proof of Theorem 1, it is clear that we can extend the argument to all *objective* correlated equilibria (CE). In other words, if μ is a CE at θ' and $\mu(m^*) > 0$, then $g(m^*) \in F(\theta')$, using only the assumptions for Nash implementation.⁶

On the other hand, since the set of all action profiles played with positive probability by some subjective correlated equilibrium is a superset of the set of correlated rationalizable action profiles (Brandenburger and Dekel (1987)), the results of Bergemann, Morris, and Tercieux (2011) imply that the assumptions in Theorem 1 are not enough to ensure implementation in *subjective* correlated equilibria.

Remark 4. Although we have assumed here that preferences over lotteries have a vNM representation, note that the proof of Theorem 1 proceeds by showing that for an *any* opponents' pure strategy profile, we can find a strategy that dominates the equilibrium strategy for the relevant player (as was also noted by MS). Therefore, it appears that one needs only minimal assumptions on preferences over lotteries for MSNE-implementation. We are currently developing this point.

References

- [1] Bergemann, D., Morris, S. and O. Tercieux, (2011) "Rationalizable Implementation," *Journal of Economic Theory*, 146, 1253-1274.

⁶As a logical matter, one cannot generally conclude that implementation in CE implies Nash-implementation (nor the converse, of course). The reason is that, in general, the desirable equilibria in a mechanism that implements in CE may not even be mixed-Nash equilibria.

- [2] Brandenburger, A., and E. Dekel, (1987) "Rationalizability and Correlated Equilibria," *Econometrica*, 55, 1391-1402.
- [3] Kartik, N. and O. Tercieux, (2011) "Implementation with Evidence," forthcoming in *Theoretical Economics*.
- [4] Maskin, E., (1999) "Nash Equilibrium and Welfare Optimality," *Review of Economic Studies*, 66, 23-38.
- [5] Maskin, E. and T. Sjostrom, (2002) "Implementation Theory," in *Handbook of Social Choice and Welfare*, ed. by K. J. Arrow, A. K. Sen, and K. Suzumura, Amsterdam: North Holland, vol. 1, 237-288.