Ordinal dominance and risk aversion*

Bulat Gafarov
Pennsylvania State University
and Higher School of Economics
bzg134@psu.edu

Bruno Salcedo
Pennsylvania State University
brunosalcedo.com
bruno@psu.edu

January 2013
This version: 7 · 5 · 2014

Abstract We find that, for sufficiently risk-averse agents, strict dominance by *pure or mixed* actions coincides with dominance by *pure* actions in the sense of Börgers (1993), which, in turn, coincides with the classical notion of strict dominance by pure actions when preferences are asymmetric. Since risk-aversion is a *cardinal* feature, all finite single-agent choice problems with *ordinal* preferences admit compatible utility functions which are sufficiently risk-averse as to achieve equivalence between pure and mixed dominance. This result extends to some infinite environments.

Keywords Rationalizability · dominance · ordinal preferences · risk aversion

JEL classification D81 · C72

Suppose that a rational agent must choose between three actions: betting that an event E occurs, betting that E does not occur, or not betting at all. The agent's preferences are represented by the von Neumann-Morgenstern (vNM) utility function summarized in Figure 1. Notice that the ordinal ranking of action-state pairs remains unchanged as long as $0 < \gamma < 2$. Also, not betting is not strictly dominated by any *pure* action, and, if $\gamma \geq 1$, it is also not strictly dominated by any *mixed* action. However, if $\gamma < 1$, then it becomes strictly dominated by the mixed action which mixes uniformly between betting on E and betting on not E.

*Submitted to *Economic Theory*. This paper originated form a conjecture by Edward Green. We are thankful for his guidance and support, as well as the useful comments from Lisa Posey, Lidia Kosenkova, two anonymous referees, and the attendants of the 2014 Spring Midwest Trade and Theory Conference at IUPUI, and the 25$^{\text{th}}$ International Game Theory Conference at Stony Brook University. All remaining errors are our own. We gratefully acknowledge the Human Capital Foundation, (http://www.hcfoundation.ru/en/) and particularly Andrey P. Vavilov, for research support through the Center for the Study of Auctions, Procurements, and Competition Policy (http://capcp.psu.edu/) at the Pennsylvania State University.

	E	not E
bet on E	2	0
bet on not E	0	2
do not bet	γ	γ

Figure 1 Payoff matrix for introductory example, where $\gamma \in (0, 2)$.

Here, γ can be thought of as measuring the degree of concavity or risk-aversion of the agent's vNM utility function. Hence, we see that dominance by pure strategies coincides with dominance by mixed strategies if the agent is sufficiently risk-averse, and there exist a sufficiently risk-averse utility function which is compatible with the ordinal preferences of the agent. In the rest of the paper, we will show that these two observations essentially continue to hold for a large class of decision problems under uncertainty with *ordinal* preferences.

We compare strict dominance by *pure or mixed* actions (M_u-dominance) with the notion of dominance by *pure* actions (P-dominance) introduced by Börgers (1993).[1] An action is P-dominated if and only if it is weakly dominated by a pure action, conditional on any given set of states.[2] P-dominated actions are always M_u-dominated, but the converse need not be true.

A mixed action could dominate an action that is not P-dominated, because mixing enables the agent to average good and bad outcomes corresponding to different action-state pairs. However, mixing also exacerbates the agent's uncertainty about the outcome of the environment, by adding uncertainty about the result of using her own randomization device. Hence, the more risk-averse the agent is, the less appealing mixing will be.[3] We find that M_u-dominance reduces to P-dominance for sufficiently risk-averse agents, according to a specific measure which we call *timidity* (propositions 3 and 5).[4] In particular, the set of sufficiently timid utility functions includes all

[1] We index M_u-dominance by u to highlight the fact that it depends on the *cardinal* information embedded in von Neumann-Morgestern (vNM) utility functions. In contrast, P-dominance only depends on the agent's *ordinal* state-contingent preferences over actions.

[2] In general, when indifference is allowed, for an action to be strictly dominated by a pure action implies that it is P-dominated, which implies in turn that it is weakly dominated by a pure action. In the generic case in which all state-contingent preferences are strict, all three notions of pure dominance coincide.

[3] Our research is in a Bayesian framework, so we use "risk" and "uncertainty" as synonyms. Nevertheless, our intuition is closely related to the work of Klibanoff (2001). He asks under which conditions would an uncertainty adverse agent be willing to choose mixed actions. As it turns out, the trade-off between averaging outcomes (uncertainty) and increasing variance (risk) plays a prominent role.

[4] These results are similar in spirit to the main lemma in Chen and Luo (2012), which implies that, in "concave-like" games, an action is M_u-dominated if and only if it is strictly dominated by a pure action. However, their result is interesting only for uncountable environments (including mixed

CARA functions that are sufficiently risk-averse in the familiar sense.

A vNM utility function is said to be *strongly compatible* with the environment if it represents the ordinal preferences of the agent *over action-state pairs*. Any strictly concave and strictly monotone transformation of utility preserves strong compatibility while increasing timidity. In this manner, we find that if either the action space or the state space is finite, then there exists a strongly compatible vNM utility function which generates equivalence between P-dominance and M_u-dominance (Corollary 4). However, the degree of timidity required grows linearly with the size of the environment, and there are countable environments in which strong compatibility precludes dominance equivalence.

By relaxing the definition of compatibility, it is still possible to obtain dominance equivalence in a large class of infinite environments. If preferences are interpreted as revealed choices, then it is meaningless to compare rankings across states. We say that a vNM utility function is *compatible* with the environment if it represents the given *state-contingent* ordinal preferences over actions. If only compatibility is required, dominance equivalence is possible in all countable environments satisfying a discreteness assumption (Corollary 6).

Our work is closely related to Börgers (1993). Using our language, Börgers' main result can be expressed as follows. For finite environments, if an action is not P-dominated, then there exists a strongly compatible vNM utility function *–which may depend on the action–* according to which the action is also not M_u-dominated. Also, while Ledyard (1986) works in a very different context, some of his results have important implications for our environment. In particular, his Corollary 5.1 implies that every finite choice environment without P-dominated actions admits a compatible vNM utility function *–which may not be strongly compatible–* according to which there are no M_u-dominated actions.

We extend Börger's result by showing that a single vNM utility function can be used for all actions. While his result has the logical form: "for every act, there is a utility function such that…"; our result has the logical form: "there is a utility function such that, for every action…". We extend Ledyard's result by showing that this is possible even if strong compatibility is imposed. Also, we establish equivalence of the entire dominance relations and not just the undominated sets, we provide tight, intuitive, and sufficient conditions on utility, and we show that dominance equivalence is attainable in some infinite environments.

Dominance relations are important for rationalizability as a solution concept for games (Bernheim, 1984, Pearce, 1984). Under standard assumptions, rationalizability is equivalent to iterated M_u-dominance. Börgers' result thus implies that, when only ordinal preferences are common knowledge, then rationalizability is equivalent to iterated P-dominance (Epstein, 1997, Bonanno, 2008).[5] Our analysis implies that the

extensions of finite environments). In finite or countable environments –like the ones we consider– if an agent has concave-like preferences, then there exists a pure action which P-dominates every other action.

[5] Lo (2000) extends this result all models of preference satisfying Savage's P3 axiom

equivalence extends to situations in which utility functions are common knowledge among the players, but only ordinal preferences are known to an outside observer. Furthermore, it also allows to relate observations arising from different situations, as in generalized revealed preference theory (Chambers et al., 2010).

1. Single-agent choice problems

We consider a single-agent environment characterized by $(A, X, \succcurlyeq)$. $X = \{x, y, \ldots\}$ is a nonempty set of states of Nature, $A = \{a, b, \ldots\}$ is a set of (pure) actions, and $\succcurlyeq$ is a transitive and complete preference relation on $A \times X$. $\succcurlyeq_x$ denotes state-contingent preferences over actions conditional on state x, i.e., $a \succcurlyeq_x b$ if and only if $(a, x) \succcurlyeq (b, x)$.

Let $[a]_x = \{b \in A \mid a \sim_x b\}$ denote the set of actions that are indifferent to a conditional on x. Throughout the paper we impose the following assumption, which essentially requires the quotient set $A/\sim_x$ to be isomorphic to a subset of $\mathbb{Z}$, for every state x. While the assumption does limit the applicability of the results, it is satisfied by all finite environments, and it leaves sufficient space to accommodate many interesting infinite environments.

Assumption 1 The collection of equivalence classes $\{[c]_x \mid a \succ_x c \succ_x b\}$ is finite for every pair of actions a and b and every state x.

A vNM utility function $u \in \mathbb{R}^{A \times X}$ is *compatible* with the environment if it preserves state-contingent preferences, i.e., if $u(a, x) \geq u(b, x)$ if and only if $a \succcurlyeq_x b$. It is *strongly compatible* if it also preserves preferences across states, i.e., if $u(a, x) \geq u(b, y)$ if and only if $(a, x) \succcurlyeq (b, y)$. We extend the domain of utility functions to mixed actions $\alpha \in \Delta(A)$ and beliefs $\mu \in \Delta(X)$ in the usual way, and we denote payoff vectors associated to pure or mixed actions by $\vec{u}(\alpha) = (u(\alpha, x))_{x \in X}$.

Example 1 Going back to the motivating example from the introduction, let $X = \{1, 2\}$ and $E = \{1\}$, and let a_1 correspond to betting on E, a_2 to betting on $X \setminus E$, and a_0 to not betting. A vNM utility function u is strongly compatible if and only if it can be written as in Figure 1 after a positive affine transformation. Notice that this implies that that $u(a_0, 1) = u(a_0, 2)$. In contrast, u is compatible as long as $u(a_x, x) < u(a_0, x) < u(a_y, x)$ for all $x, y \in X$ with $x \neq y$. Notice that this does not impose any restrictions on the differences $u(a, x) - u(b, y)$ when $x \neq y$. Figure 2 shows a strongly compatible vNM utility function (left panel), and a vNM utility function which is compatible but not strongly compatible (right panel).

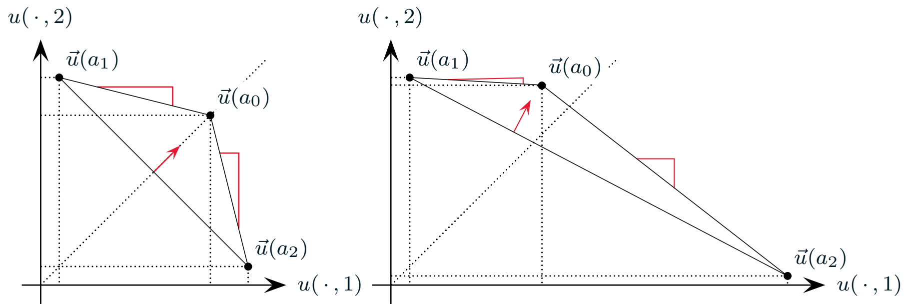

Figure 2 vNM utility functions for Example 1 with $X = \{1, 2\}$.

2. Pure and mixed dominance

Loosely speaking, an action is dominated if there exist different actions yielding preferred outcomes regardless of the state. By dominance *by pure actions*, we mean the notion introduced by Börgers (1993), according to which an action is dominated if and only if it is weakly dominated conditional on each subset of states.

Definition 1 An action a is *P-dominated* in $B \subseteq A$, if for every nonempty set $Y \subseteq X$ there exists some $b \in B$ such that $b \succsim_y a$ for all $y \in Y$, with strict preference for at least one state $y \in Y$. $P(B)$ denotes the set of P-dominated actions in B.

P-dominance extends the classical notion of strict dominance by pure actions, and both notions coincide when preferences are asymmetric. An agent who maximizes expected utility would never choose P-dominated actions, even if it they were not strictly dominated by pure actions. This is because, for an action to be a best response to some belief, it cannot be weakly dominated over the support of such beliefs. However, not being P-dominated is also not sufficient for being potentially optimal. It is well known that an action is potentially optimal if and only if it is not strictly dominated by *a pure or mixed action* according to the following definition.[6]

Definition 2 An action a is *M_u-dominated* in $B \subseteq A$ given a compatible vNM utility function u, if there exists a mixed action α such that $\text{supp}(\alpha) \subseteq B$ and $u(\alpha, x) > u(a, x)$, for every state $x \in X$. $M_u(B)$ denotes the set of M_u-dominated actions in B.

All P-dominated actions are also M_u-dominated relative to any compatible vNM utility function, but actions that are not P-dominated could still be M_u-dominated. We are interested in utility functions which guarantee that, if an action is M_u-dominated by a mixture α, then it is also P-dominated in the support of α. This requirement is equivalent to the following definition.

[6] This result can be traced back to Wald (1947).

Definition 3 [Dominance equivalence] A compatible vNM utility function u generates *dominance equivalence* if $P(B) = M_u(B)$ for all $B \subseteq A$.

What conditions over vNM utility functions imply dominance equivalence? When does there exist a compatible or strongly compatible vNM utility function satisfying such conditions? The answer to these questions is closely related to risk aversion, as measured by the *timidity* coefficient introduced in the next section. Before proceeding, it is instructive to revisit our motivating example to illustrate the role of risk aversion.

Example 2 Consider once again the example from the introduction. Clearly, there are no P-dominated actions, and every action other than a_0 is optimal conditional on some state. Hence, dominance equivalence holds if and only if $a_0 \notin M_u(A)$, which holds if and only if $\vec{u}(a_0)$ is above the line containing $\vec{u}(a_1)$ and $\vec{u}(a_2)$, see Figure 2. This simply means that the upper boundary of the set of feasible payoffs is concave, or, equivalently, that $u(\,\cdot\,, x)$ exhibits *decreasing differences* (on average for compatibility, and always for strong compatibility). In environments with more states, dominance equivalence requires that finite differences should decrease sufficiently fast.

3. Timidity

Given a compatible vNM utility function u and a state x, we use the following notation. The set of possible payoffs given x is denoted by $U(x) = \{u(a,x) | a \in A\}$, and its supremum and infimum are denoted by $\bar{u}(x) = \sup U(x)$, and $\underline{u}(x) = \inf U(x)$. Also, let $u^-(a,x) = \sup\{u_0 \in U(x) \,|\, u_0 < u(a,x)\}$ denote the best possible payoff conditional on x which is worse than $u(a,x)$. Similarly, let $u^+(a,x) = \inf\{u_0 \in U(x) \,|\, u_0 > u(a,x)\}$. Assumption 1 implies that $u^-(a,x) < u(a,x)$ whenever $u(a,x) > \underline{u}(x)$, and $u^+(a,x) > u(a,x)$ whenever $u(a,x) < \bar{u}(x)$.

Definition 4 Given a compatible vNM utility function u and an action state pair (a,x), the *timidity* coefficient of u at (a,x) is the number $\tau_u(a,x)$ given by $\tau_u(n,x) = +\infty$ if $u(a,x) \in \{\underline{u}(x), \bar{u}(x)\}$, and otherwise given by:

$$\tau_u(a,x) = \frac{u(a,x) - u^-(a,x)}{\bar{u}(x) - u(a,x)}. \tag{1}$$

In order to understand what timidity entails, it is useful to compare it with familiar measures of risk-aversion. Using finite differences instead of derivatives, the analogue of the Arrow-Pratt coefficient of absolute risk aversion for our discrete setting could be expressed as follows (see for instance Bohner and Gelles (2012)):

$$\rho_u(a,x) = 1 - \frac{u^+(a,x) - u(a,x)}{u(a,x) - u^-(a,x)}. \tag{2}$$

This coefficient is large when the *local* gain $(u^+(a,x) - u(a,x))$ is small compared with the local loss $(u(a,x) - u^-(a,x))$. In contrast, timidity compares the *global* gain $(\bar{u}(x) - u(a,x))$ with the local loss $(u(a,x) - u^-(a,x))$. Timidity requires that the potential loss of getting a slightly worse outcome should be more important than the potential gain of switching to the best possible outcome. A timid agent would refuse to spend a single dollar on a lottery ticket that promises to pay (with sufficiently low probability) more money that she could spend during a hundred lifetimes.

Example 3 Suppose each action-state results in a monetary prize given by a function $z \in \mathbb{Z}_{++}^{A \times X}$ such that $\{z(a,x) \mid a \in A\} = \mathbb{N}$ for all $x \in X$. Further suppose that the agent's preferences only depend on her preferences over money, represented by $v \in \mathbb{R}^{\mathbb{R}_{++}}$. In this case, $u = v \circ z$ is a strongly compatible vNM utility function. For a CARA agent with $v(m) = -\exp(-rm)$, $r > 0$, the agent also exhibits constant timidity:

$$\tau_u(a,x) = \frac{-\exp(-r\,z(a,x)) + \exp(-r\,(z(a,x) - 1))}{\exp(-r\,z(a,x))} = \exp(r) - 1. \tag{3}$$

Before proceeding to the main results, we conclude our analysis of timidity by noting that it satisfies one of Pratt's classic criteria. The following proposition implies that, if an agent becomes uniformly more risk-averse as measured by ρ_u, then she also becomes uniformly more timid.

Proposition 1 *Fix an action a, a state x and two compatible vNM utility functions u and v. If the set of mixed actions that are preferred to a given u and x is contained in the set of mixed actions that are preferred to a given v and x, then u is more timid than v at (a,x).*

4. Dominance equivalence and risk aversion

Let $W_x(a) = \{b \in A \mid a \succ_x b\}$ denote the set of actions that are worse than a conditional on x, and consider any compatible vNM utility function u. The following lemma states that if u is sufficiently timid at (a,x), then a is not M_u-dominated by any mixed action α that assigns sufficient probability to $W_x(a)$. The rest of our results looks for conditions on u that guarantee that this can be done whenever a is not P-dominated. The conditions essentially require u to be sufficiently timid relative to the size of the environment.

Lemma 2 *Given a compatible vNM utility function u, a pure action a, and a mixed action α, if there exists a state x such that $(\tau_u(a,x) + 1) \cdot \alpha(W_x(a)) \geq 1$, then a is not dominated by α given u.*

4.1. Finite environments

Let $K = \min\{\|A\|, \|X\|\}$. When K is finite, Caratheodory's theorem (Rockafellar, 1996, Theorem 17.1) implies that an action is M_u-dominated if and only if it is dominated by a mixed action which mixes at most K distinct actions. For every such α, there exists some action a such that $\alpha(a) \geq 1/K$. Therefore, the condition of Lemma 2 holds for all P-undominated actions, whenever the timidity coefficient is weakly greater than $K - 1$.

Proposition 3 *Given a compatible utility function u, if $\tau_u(a, x) \geq K - 1$ for all x and a, then u generates dominance equivalence.*

Suppose that K is finite and $A \times X$ is countable. Then there exist strongly compatible vNM utility functions $n^* \in \mathbb{Z}^{A\times X}$ which only take integer values. For example, if $A \times X$ were finite, n^* could be the rank function defined by:

$$\operatorname{rank}(a, x) = \left\| \left\{ (b, y) \;\middle|\; (a, x) \succcurlyeq (b, y) \right\} \right\|. \tag{4}$$

In words, the rank of an action-state pair is the number of action-state pairs that are weakly worse than it. Let u^* be the utility function defined by:

$$u^*(a, x) = -\exp\Big(-\log(K)\, n^*(a, x)\Big). \tag{5}$$

If we thought of $n^*(a, x)$ as a monetary prize, then u^* would represent the preferences of a CARA agent with with coefficient of risk aversion equal to $\log(K)$. Since u^* is a strictly monotone transformation of n^*, it is strongly compatible. Furthermore, we have that $u^*(a, x) \leq 0$ and $u^{*-}(a, x) \geq K u^*(a, x)$ for all a and x, which implies that $\tau_{u^*} \geq K - 1$. Therefore, Proposition 3 implies the following corollary:

Corollary 4 *If either X or A is finite and $A \times X$ is countable, then u^* is a strongly compatible vNM function, and yields dominance equivalence.*

When both X and A are infinite, u^* is not well defined. The following example provides a countable environment which does not admit any strongly compatible vNM utility function generating dominance equivalence.

Example 4 Let $X = \mathbb{N}$, and suppose the agent must choose a lottery from $A = \{a_0\} \cup \{a_x \,|\, x \in X\}$. Lottery a_0 represents an outside option corresponding to keeping her initial wealth. Lottery a_x represents a fair bet of one dollar *against* state x, i.e., it pays 1 if the true state is different from x and -1 otherwise. Further suppose that the agent has state-independent strictly monotone preferences over monetary holdings. After a positive affine normalization, any strongly compatible vNM utility function u

can be written as:

$$u(a,x) = \begin{cases} \gamma & \text{if} \quad a = a_0 \\ 0 & \text{if} \quad a = a_x \\ 1 & \text{otherwise} \end{cases} . \tag{6}$$

for some $\gamma \in (0,1)$. Take any such u, and consider any belief $\mu \in \Delta(\mathbb{N})$. For all $x \in \mathbb{N}$, we have that $u(a_0, \mu) = \gamma$ and $u(a_x, \mu) = (1 - \mu(x))$. If a_0 were a best response to μ, then it would be the case that $u(a_0, \mu) \geq u(a_x, \mu)$ and, consequently, $\mu(x) \geq \gamma > 0$ for all x. This would contradict the fact that μ is a probability measure. Hence, it follows that a_0 is M_u-dominated, despite the fact that it is not P-dominated.

4.2. Countable environments

When both X and A are infinite, guaranteeing dominance equivalence requires unbounded degrees of timidity. This is possible for countable environments if we do not require strong compatibility, because we may choose utility functions whose degree of timidity is always finite, but diverges to infinity along a sequence of states.

Proposition 5 *Given a compatible vNM utility function u, if X is countable and:*

$$\sum_{x \in X} \frac{1}{1 + \tau_u(a,x)} \leq 1, \tag{7}$$

for every action a, then u generates dominance equivalence.

The following example shows that the proposition is tight, in that, given any finite or countable X and a sufficiently large action space A, there always exist preferences such that: a compatible vNM utility function generates dominance equivalence *if and only if* it satisfies (7) for every action.

Example 5 Let X be any finite set with at least two elements, and let A and $\succcurlyeq$ be as in Example 4. Let u be *any* compatible vNM utility function such that:

$$\frac{1}{T} \equiv \sum_{x \in X} \frac{1}{1 + \tau_u(a_0, x)} > 1. \tag{8}$$

Simple algebra shows that a_0 is strictly dominated by the mixed action α given by $\alpha(a_0) = 0$ and $\alpha(a_x) = T/(1 + \tau_u(a_0, x))$ for $x \in X$.

If X is countable, then there exists an injective function $h \in \mathbb{N}^X$. Also, by Assumption 1, there exists a compatible vNM utility function $n^{**} \in \mathbb{Z}^{A \times X}$ which only takes integer values. Fix any such functions, and let u^{**} be the vNM utility function given by:

$$u^{**}(a,x) = -\exp\Big(-h(x)\, n^{**}(a,x)\Big). \tag{9}$$

Clearly, u^{**} is also compatible with the environment. Furthermore, we have that $u^{-**}(a,x) \geq e^{h(x)}u^{**}(a,x)$ and $u^{**} < 0$ for all a and x. Therefore:

$$\sum_{x \in X} \frac{1}{1+\tau_{u^{**}}(a,x)} = \sum_{x \in X} \frac{\bar{u}^{**}(x) - u^{**}(a,x)}{\bar{u}^{**}(x) - u^{-**}(a,x)} < \sum_{k \in \mathbb{N}} \frac{1}{e^k} < 1. \tag{10}$$

Proposition 5 thus implies that:

Corollary 6 *If X is countable, then u^{**} is a compatible vNM function and yields dominance equivalence.*

5. Summary and discussion

A vNM utility function guarantees that P-dominance coincides with M_u-dominance if it is sufficiently timid. For countable environments with discrete action spaces, it is always possible to find a sufficiently timid vNM utility function that is compatible with ordinal preferences over actions conditional on states. For finite environments with ordinal preferences over action-state pairs, it is always possible to find a sufficiently timid vNM utility function that is *strongly* compatible. In what follows, we discuss the application of the results to multi-agent environments, as well as some lines for further inquiry.

Rationalizability.– A strategic form game can be thought of as a collection of simultaneous single-agent decision problems. Rationalizability is then equivalent to the iterated removal of strategies that are not M_u-dominated. Our results then imply that, given any finite collection of finite games with ordinal payoffs, there exists a profile of compatible vNM utility functions such that, in each game, the set of rationalizable strategies corresponds to the set of strategies surviving the iterated removal of P-dominated strategies. In this sense, rationalizability and iterated P-dominance are equivalent in the absence of cardinal information.

Worst case vs. average bounds.– The degree of timidity assumed in our main results guarantees that dominance equivalence holds even in pathological scenarios with intricate preferences. In particular, it guarantees that a P-undominated action that yields the second worst outcome conditional on every state is potentially optimal, even if there are other actions yielding very good outcomes in all but one state.

An interesting problem not solved in this paper is to look for *expected* (rather worst-case scenario) bounds on timidity. For instance, one could ask for the probability that a uniformly generated utility function u will generate equivalence. One step further, having fixed only the size of the environment, one could ask for the expectation of this probability given uniformly generated preferences.

Uncountable environments.– Proposition 5 can be easily extended to accommodate environments with uncountable state spaces satisfying some technical assumptions.[7] On the other hand, our proofs depend crucially on Assumption 1, which, for most practical purposes, requires the action space to be countable. This is because the definition of timidity heavily relies on the fact that there exists some $\delta > 0$ such that $|u(a, x) - u(a', x)| \geq \delta$ whenever $u(a, x) \neq u(a', x)$.

References

Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028.

Bohner, M. and Gelles, G. M. (2012). Risk aversion and risk vulnerability in the continuous and discrete case. *Decisions in Economics and Finance*, 35(1):1–28.

Bonanno, G. (2008). A syntactic approach to rationality in games with ordinal payoffs. In Bonanno, G., Van Der Hoek, W., and Wooldridge, M., editors, *Logic and the Foundatons of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 59–85. Amsterdam University Press.

Börgers, T. (1993). Pure strategy dominance. *Econometrica*, 61(2):423–430.

Chambers, C. P., Echenique, F., and Shmaya, E. (2010). General revealed preference theory. Social Science Working Paper 1332, California Institute of Technology.

Chen, Y. and Luo, X. (2012). An indistinguishability result on rationalizability under general preferences. *Economic Theory*, 51(1):1–12.

Epstein, L. G. (1997). Preference, rationalizability and equilibrium. *Journal of Economic Theory*, 73(1):1–29.

Klibanoff, P. (2001). Characterizing uncertainty aversion through preference for mixtures. *Social Choice and Welfare*, 18(2):289–301.

Ledyard, J. O. (1986). The scope of the hypothesis of Bayesian equilibrium. *Journal of Economic Theory*, 39(1):59–82.

Lo, K. C. (2000). Rationalizability and the Savage axioms. *Economic Theory*, 15(3):727–733.

Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050.

Rockafellar, R. T. (1996). *Convex analysis*, volume 28 of *Princeton Landmarks in Mathematics and Physics*. Princeton University Press.

Wald, A. (1947). An essentially complete class of admissible decision functions. *The Annals of Mathematical Statistics*, 18(4):549–555.

A. Proofs

Proof of Proposition 1. Fix an action $a \in A$ and a state $x \in X$, and let u and v be compatible vNM utility functions such that

$$\{\alpha \in \Delta(A) \mid u(\alpha, x) \geq u(a, x)\} \subseteq \{\alpha \in \Delta(A) \mid v(\alpha, x) \geq v(a, x)\}. \tag{11}$$

[7] Details are available upon request.

We want to show that $\tau_u(a,x) \geq \tau_v(a,x)$. If a is either $\succcurlyeq_x$-maximum or $\succcurlyeq_x$-minimum, then $\tau_u(a,x) = +\infty$ and $\tau_v(a,x) = +\infty$ by definition, and the result is trivial. Hence, we assume for the rest of the proof that $\underline{u}(x) < u(a,x) < \bar{u}(x)$.

By Assumption 1, there exists an action $b \in A$ such that $u(b,x) = u^-(a,x)$, and, consequently, $v(b,x) = v^-(a,x)$. Let (a_m) be a sequence of actions such that $a_m \succ_x a$ for all m, $\lim_{m\to\infty} u(a_m,x) = \bar{u}(x)$, and $\lim_{m\to\infty} v(a_m,x) = \bar{v}(x)$. Also, for each $\theta \in [0,1]$ and each $m \in \mathbb{N}$, let $\alpha_{m,\theta}$ be the mixed action that plays a_m with probability θ, and b with probability $1-\theta$. For all such m we have that $u(\alpha_{m,1},x) > u(a,x) > u(\alpha_{m,0},x)$. Hence, since expected utility is continuous in the mixing probabilities, there exists $\theta(m) \in (0,1)$ such that $u(\alpha_{m,\theta(m)},x) = u(a,x)$. After some simple algebra this implies that:

$$\frac{u(a,x) - u^-(a,x)}{u(a_m,x) - u(a,x)} = \frac{\theta(m)}{1-\theta(m)}. \tag{12}$$

By (11), we have that $v(\alpha_{m,\theta(m)},x) \geq v(a,x)$, which implies that:

$$\frac{v(a,x) - v^-(a,x)}{v(a_m,x) - v(a,x)} \leq \frac{\theta(m)}{1-\theta(m)}. \tag{13}$$

Using (12) and (13) and taking limits as m goes to infinity thus yields the desired result

$$\tau_v(a,x) = \lim_{m\to\infty} \frac{v(a,x) - v^-(a,x)}{v(a_m,x) - v(a,x)} \leq \lim_{m\to\infty} \frac{u(a,x) - u^-(a,x)}{u(a_m,x) - u(a,x)} = \tau_u(a,x). \tag{14}$$

■

Proof of Lemma 2. Let $\beta = \alpha(W_x(a))$. Being that $u(b,x) \leq u^-(a,x)$ for $b \in W_x(a)$, and $u(b,x) \leq \bar{u}(x)$ for $b \in B$, it follows that:

$$\begin{aligned} u(\alpha,x) - u(a,x) &\leq \beta\big(u^-(a,x) - u(a,x)\big) + (1-\beta)\big(\bar{u}(x) - u(a,x)\big) \\ &= -\beta\left(\frac{u^-(a,x) - u(a,x)}{\bar{u}(x) - u(a,x)}\right)\big(\bar{u}(x) - u(a,x)\big) + (1-\beta)\big(\bar{u}(x) - u(a,x)\big) \\ &= \big(1 - \beta\cdot(\tau_u(a,x)+1)\big)\big(\bar{u}(x) - u(a,x)\big) \leq 0. \end{aligned} \tag{15}$$

■

Proof of Proposition 3. Fix a set $B \subseteq A$, an action $a \in A \setminus P(B)$, and a mixture α with $\alpha(B \setminus \{a\}) = 1$. There exists some $Y \subseteq X$ conditional on which a is not weakly dominated in B. Assume without loss of generality that for all $b \in B \setminus \{a\}$ there exists some $x \in Y$ such that $b \not\succ_x a$. This implies that for all $b \in B \setminus \{a\}$ there also exists some $x \in Y$ such that $a \succ_x b$, i.e., $B \setminus \{a\} \subseteq \cup_{x\in Y} W_x(a)$. Since $K = \min\{\|A\|, \|X\|\} < +\infty$, there exist a finite subset $Z = \{x_1, \ldots, x_k\} \subseteq Y$ with cardinality $k \leq K$, and such that $B \setminus \{a\} \subseteq \cup_{x\in Z} W_x(a)$. Therefore:

$$\sum_{x\in Z} \alpha(W_x(a)) \geq \alpha(B \setminus \{a\}) = 1 \geq \frac{k}{K} = \sum_{x\in Z} \frac{1}{K} \geq \sum_{x\in Z} \frac{1}{\tau_u(a,x)+1}. \tag{16}$$

This implies that there exists a state x such that $(\tau_u(a,x)+1)\alpha(W_x(a)) \geq 1$, and the result thus follows from Lemma 2. ■

Proposition 5. Let a, B, α and Y, be as in the proof of Proposition 3. As before, we know that $B \setminus \{a\} \subseteq \cup_{x \in Y} C_x(a, B)$, and thus:

$$\sum_{x \in Y} \alpha(C_x(a, B)) \geq 1 \geq \sum_{x \in X} \frac{1}{1 + \tau_u(a, x)} \geq \sum_{x \in Y} \frac{1}{1 + \tau_u(a, x)}. \tag{17}$$

Hence, there exists $x \in Y$ such that $(\tau_u(a, x) + 1)\alpha(C_x(a, B)) \geq 1$, and the result follows from Lemma 2. ■

B. Uncountable state space

Dominance equivalence is still possible in some environments with countable action spaces, and uncountable state spaces. Suppose that $(X, \mathscr{X}, \lambda)$ is a measure space, and let $Z(a, b) = \{x \in X \mid a \succ_x b\} \in \mathscr{X}$. We only require the two following assumptions.

Assumption 2 $(\exists \delta > 0)(\forall a, b \in A)(Z(a, b) \neq \emptyset \Rightarrow Z(a, b) \in \mathscr{X} \wedge \lambda(Z(a, b)) \geq \delta)$.

Assumption 3 There exists a measurable function $f : X \to (0, 1)$ such that $\int_X f \, d\lambda \leq 1$.

Assumption 2 requires that if an action a is preferred to an action b in at least one state, then it has to be preferred to it in a sufficiently large set. It plays a similar role as our requirement that A should be discrete, but it is significantly weaker. Assumption 3 makes it possible to have timidity grow "sufficiently fast" as to guarantee that the condition of lemma 2 is satisfied at some point. It is satisfied, for instance, when $X \subseteq \mathbb{R}$ and λ is the Lebesgue measure.

Since A is assumed to be countable, there exists some compatible vNM utility function $n \in \mathbb{R}^{A \times x}$. Let $h^\star(x) = -\log(\delta f(x))$, and define the vNM utility function $u^\star$ by:

$$u^\star(a, x) = -\exp(-h^\star(x)\, n(a, x)) \tag{18}$$

Proposition 7 *Under assumptions 2 and 3, $u^\star$ is a compatible vNM utility function which generates dominance equivalence.*

Proof. Since $u(\,\cdot\,, x)$ is a monotone transformation of $n(\,\cdot\,, x)$, $u^\star$ is compatible. For dominance equivalence, fix a set $B \subseteq A$, a P-undominated action $a \in A \setminus P(B)$, and a mixed action α with $\alpha(B \setminus \{a\}) = 1$. As, in the proof of proposition 3, the fact that $a \notin P(B)$ implies that $B \setminus \{a\} \subseteq \cup_{x \in X} W_x(a)$. Therefore:

$$\begin{aligned} \int_X \alpha(W_x(a))\, d\lambda &= \int_X \sum_{a \in W_x(a)} \alpha(a)\, d\lambda = \int_X \sum_{a \in A} \mathbb{1}(a \in W_x(a))\alpha(a)\, d\lambda \\ &= \sum_{a \in A} \alpha(a) \int_{Z(a,b)} d\lambda = \sum_{a \in A} \alpha(a)\lambda(Z(b, a)) \geq \sum_{a \in A} \alpha(a)\delta = \alpha(A)\delta = \delta, \end{aligned} \tag{19}$$

where we used the fact that $x \in Z(a,b)$ if and only if $b \in W_x(a)$. Now, let $g_x^\star(m) = -\exp(-h^\star(x)\, m)$ so that $u(a,x) = g_x^\star(n(a,x))$. Since $g_x^\star(m-1) = e^{h^\star(x)} g^\star(m)$ and $g^\star(m) < 0$ for all x and $m \in N(x)$, it follows that:

$$\int\limits_X \frac{1}{1+\pi_u(a,x)}\, d\lambda < \int\limits_X \exp(-h^\star(x))\, d\lambda = \delta \int\limits_X f(x)\, d\lambda \leq \delta. \tag{20}$$

From (19) and (20), that there exists x such that $(\tau_u(a,x)+1)\alpha(W_x(a)) \geq 1$, and the result thus follows from Lemma 2. ■