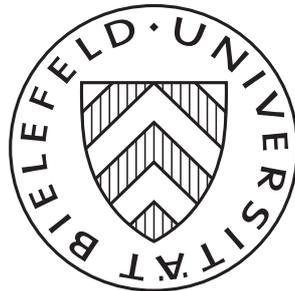


September 2015

Preferences under Ignorance

Olivier Gossner and Christoph Kuzmics



PREFERENCES UNDER IGNORANCE

OLIVIER GOSSNER AND CHRISTOPH KUZMICS

ABSTRACT. A decision maker (DM) makes choices from different sets of alternatives. The DM is initially fully ignorant of the payoff associated to each alternative, and learns these payoffs only after a large number of choices have been made. We show that, in the presence of an outside option once payoffs are learned, the optimal choice rule from sets of alternatives is one that is *as if* the DM had strict preferences over all alternatives. Under this model, the DM has preferences for preferences while being ignorant of what preferences are “right”.

Keywords: consistency, rationality, weak axiom of revealed preferences, strict preference

Journal of Economic Literature classification numbers: C73, D01, D11

1. INTRODUCTION

Recent empirical evidence shows that although individual agents’ decisions are at large consistent with a theory of preferences, these preferences vary wildly across agents. For instance, using scanner data of household purchases, Echenique, Lee, and Shum (2011) and Dean and Martin (2015) find that individual households make consistent choices.¹ On the other hand Dean and Martin (2011) and Crawford and Pendakur (2013), show that households exhibit significant heterogeneity in preferences over consumption bundles.²

From the point of view classical consumer theory (see e.g. Mas-Collel, Whinston, and Green (1995, Chapter 1) for a textbook treatment), evidence of consumer behavior “as if” they had rational preferences is reassuring.

We are very grateful to Tymon Tatur and Matthew White and conference participants at the conference on the Biological Basis of Economic Behavior, Vancouver, and the workshop on Strategic Information Acquisition and Transmission, Munich, as well as seminar participants at Northwestern University, London School of Economics, University of Warwick, Paris Game Theory Seminar, Roy Seminar, European University Institute, Center for Rationality at Jerusalem, Tel-Aviv University, Johns Hopkins University, and the University of British Columbia for helpful comments and suggestions.

¹To quote Echenique, Lee, and Shum (2011, p. 1205), “[i]t is fair to say that most of the empirical literature, using both field and experimental data, finds relatively few violations of GARP”

²Both findings of consistent and heterogenous behavior are confirmed by Choi, Fisman, Gale, and Kariv (2007) in the context of risk-preferences, see also Dean and Martin (2010, Section 5.2.3).

This theory also accounts for heterogeneity of preferences, as is evidenced by the literature on consumer preference aggregation as in e.g. Gorman (1953); Chiappori and Ekeland (1999).

From the point of view of the more recent literature trying to understand individual behavior from an evolutionary point of view (see, e.g. Robson, 2001b), the observed heterogeneity of consistent decision making is more of a puzzle. In fact, this literature explains the prevalence of behaviors by the evolutionary advantages provided. But then, one of the following two must be true. Either consumption affects fitness and better choices provide an evolutionary advantage, in which case agents using non fitness maximizing rules should gradually disappear given the evolutionary force is against them. Or consumption does not affect fitness, in which case all decision rules provide the same fitness. But in the second case, why would agents have preferences at all since no decision rules, rational or not, can provide any advantage over another?

In this paper we present a possible rationale for heterogeneous but always consistent choice behavior even if nature's objective function is the same for all individuals she is working on. In this we follow the literature on the principal agent approach to evolution as in Robson (2001a), Samuelson and Swinkels (2006), and Robson and Samuelson (2011), among others, by modeling nature as the principal who endows the individual agent with a rule (e.g. utility function) that the agent then uses to make choices.

More closely we follow Robson (2001a), Rayo and Becker (2007), Netzer (2009), and Steiner and Stewart (2014) in that we make nature choose a rule of behavior before she knows the exact environment the individual will find herself in.³

The model we provide can also be understood to provide normative recommendations to a decision maker who has to make repeated decisions without ex-ante knowing which alternatives are better than others. In this sense this paper is conceptually similar also to the literature on "rational inattention" as in e.g. Matejka and McKay (2015).

To fix ideas consider the following two scenarios. Two friends decide to go on a diet (with the purpose to lose weight, to feel better, to feel less tired, to combat an illness, or for some such goal). Both dieters are offered, each day, food from different menus, but they are ignorant as to what choices are good to achieve their objective. Independently of each other, they both choose a choice rule, i.e. a rule that specifies what choice to make depending on each

³Otherwise there is no point in nature giving us a utility function or making us think. Nature could just make us make the "right" choices. See e.g. Robson (2001b) for a discussion of this issue.

possible menu. After a while they meet and exchange their experience. The least successful dieter can then decide to adopt the most successful dieter's choice rule. How should they choose their diet to begin with?⁴

A contestant in a game show is informed that she will have to make repeated choices from subsets of a set of balls of different colors. Repeatedly, the contestant is presented with a different basket containing balls of different colors, out of which she always chooses a single one. After all choices are made, a lottery determines the dollar value of each color. The contestant can then either choose to earn a dollar amount equal to the sum of all the ball values, or opt out for a fixed prize. How should the contestant choose from each basket?

The model we provide to tackle our questions can be roughly sketched as follows. A decision maker (DM) will be asked to make repeated choices from subsets of a finite set of grand alternatives. The DM is asked to choose a choice rule that specifies what choice she would make for every possible subset of the set of all alternatives. A choice rule can be consistent (i.e. derived from a rational preference relation) but can also be non-consistent in the sense of exhibiting cycles or other non-transitivities. We allow all choice rules. At the time of this choice, the DM acts under a veil of ignorance and knows nothing about the value of the various alternatives to her. Nature (random nature, not mother nature) then randomly chooses a gain function that attaches material gains (or fitness) to each alternative. After some time the DM learns how well her choice rule is doing on average without learning how each alternative contributes to the overall material gain. The DM can then stick to her chosen rule and obtain the resulting average material payoff or adopt an outside option, the value of which is chosen ex-ante randomly. The DM evaluates material payoff with a fixed expected utility function.

We show that, provided the DM is not too risk averse, any optimal choice rule must be a strictly consistent choice rule. Moreover we identify conditions under which all strictly consistent rules are equally optimal.

The argument for this claim is roughly as follows. We show that in such an environment all choice rules produce the same expected material fitness. We then show, and this is the crucial result, that strictly consistent rules are in some sense the most risky rules. To be more precise we note that any choice rule induces (at the ex-ante level) a probability distribution over material gains. It can then be shown that for any choice rule there is a distribution over strictly consistent choice rules that induces a distribution over material

⁴Consider a witch doctor in some village choosing treatments from a set of available treatments (dictated for instance by the weather among other things) for his various patients. Or farmers choosing crops for their various fields. Or sea-farers stranded on an exotic island with strange fruits and animals.

gains that is a strict mean preserving spread over the distribution of material gains induced by the given choice rule. If the DM is not too risk averse she will then strictly prefer this distribution over strictly consistent rules over the given choice rule, because of the outside option.⁵ Thus, for any choice rule the DM will find a strictly consistent choice rule that she prefers strictly over the given choice rule.

If the DM is sufficiently risk averse, however, then the DM will find strictly consistent rules too risky and choose a choice rule that is not strictly consistent and may even exhibit non-transitivities.

Thus, under ignorance, consistent decision making is optimal when 1) the DM is not too risk averse, and 2) there is some form of outside option the DM can adopt at the end.

Our model provides a rationale for heterogeneous preferences for the two following reasons. First, we show there are conditions on the distribution over choice sets for the agents, under which even if all agents face the same such distribution over choice sets, all strictly consistent rules are equally good and all other rules suboptimal. Different agents may thus adopt different strictly consistent rules while each of them being optimal.⁶ Second, we show that while all agents have the same utility function, different distribution over choice sets or different distributions over outside options lead to different optimal strictly consistent rules. Hence, a strictly consistent rule that is optimal for one agent may be suboptimal for another even though they both share the same utility function and face the same complete uncertainty as to which alternatives are good for them.

The paper proceeds as follows. In Section 2 we provide the model. Section 3 we state the main theorem and sketch its proof, in the course of which we establish two additional results that are of interest in their own right. Section 4 provides a discussion of the exact role the assumptions play for the various results. In section 5 we use a simple example with two purposes. First, it should help the reader to understand both the model and its workings. Second, it demonstrates the boundaries of our results by highlighting what is not true in this model. Section 6 finally discusses a few possible extensions of the model.

2. MODEL

⁵This is the same logic as in finance, e.g. in Merton (1973) and Rasmusen (2007), where increasing risk increases option value.

⁶While each individual chooses one strictly consistent choice rule, the aggregate behavior will look like that of a random utility model as in Luce (1959) with implications as in Block and Marschak (1960), see also Gul and Pesendorfer (2013).

2.1. Choice. Let $K = \{1, \dots, |K|\}$, $|K| > 1$ be the set of all possible **alternatives**. Let $\mathcal{L} = \mathcal{P}(K) \setminus \{\emptyset\}$ denote the set of all non-empty subsets of K . We call an element in \mathcal{L} a **choice set**. A decision maker is repeatedly asked to make a choice from different choice sets.

Definition 1. A **choice rule** is a function $R : \mathcal{L} \rightarrow \mathcal{L}$ such that $R(L) \subseteq L$ for all $L \in \mathcal{L}$. Let \mathcal{R} denote the set of all such choice rules.

Following Uzawa (1956) and Arrow (1959) (see also Chapter 1.B in Mas-Collel, Whinston, and Green (1995)), let \succeq denote a binary (preference) relation over elements in K with the interpretation that when $i \succeq j$ an individual holding this preference relation weakly prefers i over j . The relation \succeq is complete if for any two $i, j \in K$, $i \succeq j$ or $j \succeq i$ (or both), it is transitive if $i \succeq j$ and $j \succeq l$ imply $i \succeq l$. A complete and transitive relation is called **consistent** (often also termed “rational”, see e.g. Definition 1.B.1 in Mas-Collel, Whinston, and Green (1995)). In this paper a special case of consistent preferences plays a prominent role, namely, strict preferences.

A relation \succeq is anti-symmetric if whenever $i \succeq j$ and $j \succeq i$ then $i = j$. We call a preference relation **strictly consistent** if it satisfies completeness, transitivity, and anti-symmetry.

These definitions extend from preference relations to the corresponding individual’s behavior.

Definition 2. A choice rule $R \in \mathcal{R}$ is **consistent** if there exists a complete and transitive preference relation \succeq such that, for every L , $R(L)$ is the set of maximal elements in L for \succeq . It is **strictly consistent** if it is consistent and $R(L)$ is a singleton for all $L \in \mathcal{L}$. Let \mathcal{R}^s denote the set of strictly consistent rules.

It is easily verified that a strictly consistent rule is one based on a strictly consistent preference relation.

A rule that plays a particular role in our set-up is one that accepts the whole offered set. It is defined by $R^I(L) = L$ for every L . Since this is the choice rule associated to an indifferent preference relation, we call the rule R^I the indifferent rule.

2.2. The environment. An environment consists of two components. First, nature chooses a gain function that associates gain levels to possible choices. It is useful to consider a fixed finite set of gain levels $G \subset \mathbb{R}_+$. A **gain function** $g : K \rightarrow G$ is then a function from the set of all possible choices to this set of possible gain levels, with the interpretation that $g(k) \in G$ is the gain an individual receives when choosing $k \in K$.

We extend any gain function to the set \mathcal{L} of choice sets by setting

$$g(L) = \frac{1}{|L|} \sum_{k \in L} g(k)$$

for $L \in \mathcal{L}$, with the natural interpretation that $g(L)$ is the expected gain for the decision maker when L is the set of accepted alternatives, thus assuming that each element in L is then chosen by the decision maker with equal probability.

Second a **distribution over choice sets** $p \in \Delta(\mathcal{L})$ describes the frequency with which choice sets are presented to the decision maker. We assume that enough choice sets are available with positive frequency, thus making the assumption that p has full support over the non-singleton subsets of \mathcal{L} .

In some cases, it is useful to consider **neutral** distributions, for which all alternatives play the same role.

Definition 3. *A distribution p over choice sets is **neutral** if, for every permutation π of K , and every choice set $L \subseteq K$, $p(L) = p(\pi(L))$.*

Obviously, the uniform distribution is neutral. Other examples of neutral distributions over choice sets are the uniform distributions over choice sets of fixed size l , for $1 \leq l \leq |K|$.

In order to define the average gain of a rule R given a distribution of choice sets p and a gain function g , we need to first specify the choices realized by the decision maker when facing the choice set L . If $R(L)$ is a singleton, the decision maker ends up getting the alternative $R(L)$, and obtains a payoff (for this period) of $g(R(L))$. If $R(L)$ is not a singleton we assume that the decision maker accepts all alternatives in $R(L)$ equally, and ends up with each of them with equal probabilities.⁷ Given our set extension of g , the average payoff received by the decision maker is also, in this case, $g(R(L))$.

Given gain function g and a distribution of choice sets p , the (average) material gain of any rule $R \in \mathcal{R}$ is then given by

$$g_p(R) = \mathbb{E}_p g(R(L)) = \sum_{L \in \mathcal{L}} p(L) g(R(L)).$$

Let \mathcal{G} be a finite set of gain functions, and let $q \in \Delta(\mathcal{G})$ be a distribution over gain functions. For a permutation $\pi : K \rightarrow K$ and a gain function $g : K \rightarrow G$ let $g^\pi : K \rightarrow G$ be the *permutation of g* defined by $g^\pi(l) = g(\pi(l))$ for all $l \in K$.

⁷We believe this to be an innocuous assumption, which, however, provides us with the property that the set of all decision rules is finite. We do not believe that any additional insight can be gained by relaxing this assumption.

Definition 4. A distribution over gain functions, $q \in \Delta(\mathcal{G})$, is **symmetric** if $g^\pi \in \mathcal{G}$ and $q(g) = q(g^\pi)$ for every gain function $g \in \mathcal{G}$ and for every permutation $\pi : K \rightarrow K$.

A decision maker holding a symmetric belief over gain functions is completely agnostic as to what choices carry better payoffs than others, and by how much. Furthermore, if the distribution over choice sets is neutral, nothing allows to distinguish different choices. Our objective in this paper is to show when such an agnostic decision maker is strictly better-off by using a strictly rational choice rule, thus behaving “as if” she had a strict preference over choices in K .

In what follows, we assume that the distribution q over gain functions is symmetric and that its support contains at least one non-constant gain function.

2.3. Risk aversion and utilities. We distinguish gains, which are interpreted as material or monetary payoffs, from utilities. The agent has a utility function $u : \mathbb{R} \rightarrow \mathbb{R}$, and $u(g)$ is the agent’s utility for a material gain of g . We assume that u is twice differentiable, increasing and concave, and that the coefficient of absolute risk aversion $\rho(g) = \frac{-u''(g)}{u'(g)}$ is bounded between a lower bound $\underline{\rho}$ and an upper bound $\bar{\rho}$. In what follows we say that the agent is *sufficiently risk averse* if $\underline{\rho}$ is large, and *not too risk averse* if $\bar{\rho}$ is close to 0.

2.4. Outside options. After observing the “average” material payoff corresponding to the rule R , the decision maker may either stick to the induced material payoff, or switch to an outside option with material gain g' . The value g' is random and statistically independent of q . The realized value of g' is observed by the decision maker after she learns the average material payoff induced by her chosen rule. We assume that g' has a positive density in the interval $[\min_k g(k), \max_k g(k)]$. This assumption excludes the trivial cases in which g' is either smaller than $\min_k g(k)$ with probability 1 and the outside option is never chosen, as well as the case in which it is larger than $\max_k g(k)$ with probability 1 and the outside option is always selected. Note however that it encompasses situations in which the outside option is available with positive probability only, as they are captured by distributions of g' that put positive probability on values less than $\min_k g(k)$.

2.5. The decision maker’s problem. The decision maker (DM) knows the set of alternatives K , the distribution of choice sets p the distribution of gain functions q as well as the distribution of the outside option g' . The timing of the decision problem is as follows. First, the DM chooses a rule in

\mathcal{R} . Then nature chooses a gain function according to q . This gain function is not known to the DM at this time. The DM makes choices according to her chosen rule in every choice set L which she faces with frequency $p(L)$. The DM then learns the average realized gain $g_p(R)$. The outside option value g' realized and is observed by the DM, who can then choose the maximum of this average realized gain and g' . In short, the DM chooses a rule $R \in \mathcal{R}$ in order to maximize her ex-ante expected utility

$$\mathbb{E}_{q,g'} [\max \{u(g_p(R)), u(g')\}].$$

The timing of events in the model is described in Table 1.

- 0 • DM chooses rule R
- 1 • gain function g and average gain $g_p(R)$ realizes
- 2 • outside option g' realizes
- 3 • DM receives material gain $\hat{g} = \max\{g_p(R), g'\}$ providing utility $u(\hat{g})$

TABLE 1. Timeline of events

3. RESULTS

In this section we first state the main result (Theorem 1) and then sketch its proof by providing two intermediate results that are of interest in their own right (Theorems 2 and 3).

3.1. Optimal choice: indifferent versus strictly consistent. The main result of this paper is the following theorem.

Theorem 1. *If the DM is sufficiently risk averse and p is neutral, then the indifferent choice rule R^I is optimal and all strictly consistent rules are suboptimal. If the DM is not too risk averse, then for every p every optimal rule is strictly consistent and for p neutral every strictly consistent rule is optimal.*

The proofs of all results are given in the Appendix. In what follows we provide a sketch of the proof and identify two intermediate results that are of interest in their own right.

First we note that, given the assumption of q symmetric, i.e. the assumption that the DM is fully ignorant as to what alternatives are good for her, all choice rules yield the exact same ex-ante expected gain. In other words, absent an outside option and for a risk neutral agent, all rules are equally good.

Lemma 1. *Let $R, R' \in \mathcal{R}$ be arbitrary decision rules. Then*

$$\mathbb{E}_q g_p(R) = \mathbb{E}_q g_p(R').$$

If all the rules give the same expected utility, they can still differ in the level of risk they provide.

Let R, R' be two rules. We say that R is strictly riskier than R' if the distribution of $g_p(R)$ under q is a strict mean-preserving spread of the distribution $g_p(R')$ under q . One distribution is a strict mean-preserving spread of another if it is a mean-preserving spread of and not identical to the other. If μ is a distribution over rules and R' is a rule, we say that μ is strictly riskier than R' if the distribution of $g_p(R)$ under q and μ is a strict mean-preserving spread of the distribution $g_p(R')$ under q .

The next theorem shows that the indifferent rule is risk-minimizing when p is neutral.

Theorem 2. *Assume that p is neutral. Then R^I is a least risky rule.*

The above theorem explains why R^I is an optimal decision rule for sufficiently risk averse agents. Note in passing that it also shows that, absent any outside option, R^I is optimal for any risk-averse agent.

Our next result shows that the strictly rational rules maximize risk in an unambiguous sense.

Theorem 3. *Let R be any non strictly consistent rule, then there exists a distribution μ over strictly consistent rules such that μ is strictly riskier than R . If p is neutral, then every strictly consistent rule is strictly riskier than any non strictly consistent rule.*

By force of this theorem a risk neutral DM, and in the presence of an outside option, when considering a non strictly consistent rule, will always find a distribution over strictly consistent rules (a mixed strategy putting weight only on strictly consistent rules) that she strictly prefers over the given rule. To then finish the argument we note that, as the DM strictly prefers this distribution over strictly consistent rules over the given rule, she must also strictly prefer one of these strictly consistent rules over the given rule.

We have thus explained how Theorems 2 and 3 can be used to proof the main result, Theorem 1. The proof of Theorems 2 and 3, identifying how rules can be partially ordered by the mean-preserving spread order, rests on a key lemma, which we establish in the next subsection.

3.2. Choice rules and choice distributions. A key to a better understanding a choice rule's performance in the decision maker's problem is to consider the probability distribution over choices in K induced by this choice rule and by the distribution over choice sets. Given the distribution p over

choice sets and a choice rule R , let $\lambda_p(R)(k)$ denote the overall probability with which an element $k \in K$ is selected under the rule R ; it is given by

$$\lambda_p(R)(k) = \sum_{L:k \in R(L)} \frac{p(L)}{|R(L)|}.$$

We call $\lambda_p(R)$ the **choice distribution** associated to R . This choice distribution summarizes the frequency with which each item in K is selected by R . This distribution is known to the agent.

Since the gain function $g: K \rightarrow G$ is not known to the agent, neither is the payoff $g_p(R)$ associated to R , but this payoff can be easily deduced from g and $\lambda_p(R)$, as

$$g_p(R) = \sum_k \lambda_p(R)(k)g(k).$$

For a fixed g , a rule's average payoff is entirely determined by its choice distribution. And for g unknown, the distribution of payoffs induced by R and g is entirely determined by $\lambda_p(R)$ and by the distribution of g . As we shall see, it is useful to think of the choice distribution induced by her rule as the object of choice for the agent.

For a given distribution p over choice sets, let Λ_p denote the set of all choice distributions available to the agent, i.e.

$$\Lambda_p = \{\lambda_p(R), R \in \mathcal{R}\}.$$

Similarly, denote by Λ_p^s the subset of Λ_p consisting of distributions induced by strictly consistent rules, i.e.

$$\Lambda_p^s = \{\lambda_p(R), R \in \mathcal{R}^s\}.$$

The following result locates the choice distributions induced by consistent rules as extreme points in the set of choice distributions. It shows that the extreme points of the convex hull of Λ_p consists of points in Λ_p^s only.

Lemma 2. *Every choice distribution in Λ_p is a convex combination of choice distributions in Λ_p^s .*

This lemma is proven in Appendix A. It is this lemma that provides the key insight needed to prove Theorem 3 by establishing that the strictly consistent rules are, in a certain sense made precise in the statement of the Theorem, most risky.

4. DISCUSSION OF THE ASSUMPTIONS

Here we briefly discuss the role played by the different assumptions in our main results. We first argue that the assumptions of full support for p and g' are not important and relaxing these changes the results only slightly. We

then discuss why some results require the assumption of p neutral and how the results change if q is not symmetric.

Section 2 assumes that p has full support over non-singleton choice sets. Now suppose that p does not have full support. Note first that the conclusions of Lemmas 1 and 2 still hold. The same applies to Theorem 2: the indifferent rule is still a least risky rule and to the first sentence in Theorem 1: the indifferent rule is an optimal rule for a sufficiently risk averse DM. The conclusions of Theorem 3 are slightly modified: it is true that for any non strictly consistent rule there is a distribution over strictly consistent rules that yields a mean-preserving spread in terms of distributions of gains, but this spread does not anymore have to be strict. The second sentence in Theorem 1 needs to be adapted to say that under sufficiently low risk aversion, there exists an optimal strictly consistent rule. This rule is not unique when p doesn't have full support, since choices outside the support of p do not affect payoffs, thus are irrelevant. In this case, it can be shown that all optimal rules must coincide with a strictly consistent rule on the support of p .

We also assumed that the outside option g' has full support over a sufficiently large interval. Note first that this assumption is only relevant for Theorem 1. Relaxing this assumption does not change the conclusion of the first sentence in Theorem 1 and only changes the conclusion of the second sentence in the same way as relaxing the assumption of p full support does: there exists a strictly consistent optimal rule, but not only strictly consistent rules may be optimal. To see this, observe for instance that if g' takes only values outside of the range of $g_p(R)$, all rules yield the same payoff hence are optimal.

The assumption of p neutral does not enter Lemmas 1 and 2, but does enter all three Theorems. Indeed it is important in Theorem 2 and the first sentence in Theorem 1. If p is non-neutral, then the indifferent rule is not necessarily least risky, and thus also not necessarily optimal for a highly risk averse DM. To see this consider $K = \{a, b\}$ with $p(\{a\}) = p(\{a, b\}) = \frac{1}{4}$ and $p(\{b\}) = \frac{1}{2}$. Then the indifferent rule yields a choice distribution $\frac{3}{8}a + \frac{5}{8}b$ and, under q symmetric, is more risky than the strictly consistent rule that corresponds to the preference $a \succ b$ which yields a choice distribution of $\frac{1}{2}a + \frac{1}{2}b$.

Note that even when p is neutral, under p full support, the indifferent rule R^I is never the unique risk-minimizing rule as long as K contains at least three elements. Consider all subsets L of K with $|K| - 1$ elements. By p neutral all these sets have the same weight under p . Consider any rule R that coincides with R^I on all sets except these sets L , such that $R(L)$ is a

singleton for every L and such that all $R(L)$ differ. Then, $\lambda_p(R) = \lambda_p(R^I)$. More generally a rule that has appropriate cycles, and is thus non-consistent, is also least risky.

The most interesting implication of p non-neutral is the role p plays in Theorem 3 and the second sentence in Theorem 1. The example in Section 5 below shows that, for p non-neutral, yet q symmetric, it is not the case that all strictly consistent rules are most risky and that all strictly consistent rules are equally good and all optimal for a DM with low risk aversion. There does not always exist a unique most risky strictly consistent rule, and different p 's imply different most risky rules (even keeping q the same).

We finally turn to the two assumptions made on q . Assuming that there is at least one non-constant gain function in the support of q only avoids that the model is trivial. The second assumption, q symmetric, makes the model interesting by assuming the decision maker has a veil of ignorance. We believe that it is under this condition that results showing the optimality of strictly consistent rules are the most striking. Nevertheless, it is still interesting to examine the implications of an asymmetric q . The first observation in this case is that the conclusion of Lemma 1 does generally not hold if q is not symmetric. In this case (for instance in the trivial case in which q is supported by one payoff function only), some rules can provide a higher expected gain than others. Interestingly, however, under the presence of an outside option and for a not too risk averse DM, the optimal rule is not generally the rule that maximizes the expected gain under the most likely gain function under q , as shown in Section 5 below.

It is still true, however, that even if q is non-symmetric, if p and g' have full support, the optimal rule for a DM with low risk aversion (second part of Theorem 1) is strictly consistent. The proof requires little adaptation. The key argument is the following. By Lemma 2, for every non strictly consistent rule, there exists a distribution over strictly consistent rules (as in Theorem 3) that produces a strict mean preserving spread in terms of choice distributions. This distribution also provides a strict mean preserving spread of payoffs for every q . Thus, a low risk averse DM will, for any q , prefer this distribution on strictly consistent rules over the given non strictly consistent rule. Hence, at least one of these strictly consistent rules provides a higher expected payoff than the non strictly consistent rule. Which of the strictly consistent rules is optimal can then depend on q and the distribution of the outside option g' .

5. AN EXAMPLE

We study an example in detail, showing in particular how the optimal choice rules can depend on the data of the problem when p is not neutral.

Let $K = \{a, b, c\}$, and p be given by: $p(\{a, b\}) = p(\{a, c\}) = \frac{1}{4}$, and $p(\{b, c\}) = p(\{b\}) = p(\{c\}) = \frac{1}{8}$, $p(\{a, b, c\}) = p(\{a\}) = \frac{1}{16}$. Note that b and c are symmetrically treated in p , but that p is not neutral.

Given the symmetries in the setup, there are, without loss of generality, only three strictly consistent rules with potentially different payoff distributions. The strict preferences corresponding to these rules are:

$$\begin{aligned} R^a & a \succ b \succ c \\ R^b & b \succ c \succ a \\ R^c & c \succ a \succ b \end{aligned}$$

Their corresponding choice distributions are $\lambda(R^a) = \frac{5}{8}a + \frac{1}{4}b + \frac{1}{8}c$, $\lambda(R^b) = \frac{1}{16}a + \frac{9}{16}b + \frac{3}{8}c$, and $\lambda(R^c) = \frac{5}{16}a + \frac{1}{8}b + \frac{9}{16}c$.

Let us consider gain functions that attach gain 1 to one element in K and 0 to then other two, and q the uniform distribution over these three gain functions. The payoff distributions of the strictly consistent rules under q are given in the following table (one \bullet represents a probability weight of $\frac{1}{3}$).

R	$\frac{1}{16}$	$\frac{2}{16}$	$\frac{3}{16}$	$\frac{4}{16}$	$\frac{5}{16}$	$\frac{6}{16}$	$\frac{7}{16}$	$\frac{8}{16}$	$\frac{9}{16}$	$\frac{10}{16}$
R^a		\bullet		\bullet						\bullet
R^b	\bullet					\bullet			\bullet	
R^c		\bullet			\bullet				\bullet	

It is seen that the payoff distributions of R^a and R^b are mean preserving spreads of the payoff distribution of R^c , but that neither the payoff distribution of R^a nor R^b is a mean preserving spread of the other. It follows that it is always the case that one of the two rules R^a or R^b is optimal.

We now show that which of R^a or R^b is optimal depends on the distribution of outside options. First consider a distribution of g' with full support that puts high probability on some value $x \in (\frac{9}{16}, \frac{10}{16})$, and for simplification think of the limit case in which the distribution puts probability 1 on x . Under R^b , the option is always chosen, hence the expected payoff is x , while under R^a the option is chosen with probability $\frac{2}{3}$ and the expected payoff is $\frac{1}{3}\frac{10}{16} + \frac{2}{3}x > x$. The option value is maximal under R^a which is then the only optimal rule. On the other hand, if the distribution of g' puts high probability (think of it as being 1) on some value $x \in (\frac{1}{16}, \frac{2}{16})$, the option is never chosen under R^a , which then yields an expected payoff of $\frac{1}{3}$, while it is chosen with probability $\frac{1}{3}$ under R^b which yields an expected payoff of

$\frac{1}{3}x + \frac{1}{3}\frac{6}{16} + \frac{1}{3}\frac{9}{16} > \frac{1}{3}$. Hence in this second case the option value is maximal under R^b which is now the only optimal rule.

Note that ex-ante all elements of K have the same chance of being the best choice. Nevertheless it is not true that all (strictly consistent) rules are equally good. Together with our result for p neutral and q symmetric, this implies that p has a subtle effect on which rules are good and which are bad. The optimal rule depends on p (just as much) as on q .

Now consider the same example but with a slightly different distribution over gain functions, denoted q' . Let q' be such that it is derived from q by taking a small $\epsilon > 0$ probability weight from all gain functions other than the $g^a \in \mathcal{G}$ with $g^a(a) = 1$ and $g^a(b) = g^a(c) = 0$ and move that total probability mass to that gain function g^a . Thus, g^a is the most likely gain function under q' . Let g' put high probability on some value $x \in (\frac{1}{16}, \frac{2}{16})$. Then for sufficiently small ϵ , rule R^b is strictly better than R^a , even though R^a is the unique optimal rule for gain function g^a . Thus, even if one gain function is more likely than all others, the strictly consistent rule associated with this gain function may not be optimal.

6. EXTENSIONS

We first point out one general robustness property of our results. Most of our main results in Theorems 1, 2, and 3 are of the following kind. We show that according to some partial order of “better than” certain decision rules are strictly better than other decision rules. This implies that any sufficiently small change to the model, i.e. any ϵ (positive yet small) addition of something cannot change the result (given the continuity of the DM’s objective function).

In what follows we show several directions in which the model can be extended.

6.1. Costly experimentation and impatience. The model studied so far considers that if the outside option is chosen, then the resulting utility is the one corresponding to the outside option’s gain. This means that experimentation of a rule R is costless in the sense that when the outside option is chosen, the payoff generated by R is irrelevant. We can instead consider that experimentation is costly in the following sense. The payoff from R materializes in a first stage, and the agent obtains the corresponding utility. Then in a second stage the agent may decide to switch to the outside option, or not. The agent has a discount factor of $0 < \delta < 1$, meaning that the objective is to maximize $(1 - \delta)$ times the utility in the first period plus δ times the utility in the second period. The agent’s problem then becomes

to maximize over all rules R the total expected utility:

$$\mathbb{E}_{p,g'} [\max\{u(g(R)), (1 - \delta)u(g(R)) + \delta u(g')\}].$$

Note that $\max\{u(g(R)), (1 - \delta)u(g(R)) + \delta u(g')\} = (1 - \delta)u(g(R)) + \delta \max\{u(g(R)), u(g')\}$. Thus, the DM's objective is to maximize

$$(1 - \delta)\mathbb{E}_{p,g'} [u(g(R))] + \delta\mathbb{E}_{p,g'} [\max\{u(g(R)), u(g')\}].$$

This new objective function differs from the one before only by the additional first term.

If the DM is risk averse the first term in the objective function is maximized at a least risky rule. Thus, the first sentence in Theorem 1 is still true: the indifferent rule, being least risky, is optimal for a sufficiently risk averse DM. The second sentence is also true in this setup. By Lemma 1 all rules yield the same expected gain. Thus, for low risk aversion, all rules produce almost the same expected utility. As risk aversion becomes smaller and smaller the first term in the objective function becomes irrelevant, and is dominated by the second term. This second term, however, coincides with the original objective function.

6.2. Finite sampling. In our main model, we consider that the agent observes the expected payoff $g_p(R) = \mathbb{E}_p g(R(L))$ before deciding whether to use the rule R or take the outside option. The payoff $g_p(R)$ can be understood as the average of $g(R(L))$ over an infinite sequence of realizations of the choice set L according to p . Now consider a variation of the model in which the agent gets to observe the average payoff $\frac{1}{n} \sum_t g(R(L_t))$ over a finite and iid. sequence with law p of choice sets L_1, \dots, L_n before deciding to take the outside option or not.

In this new decision problem, the DM is less informed than in the original model, as $g_p(R)$ is only observed with noise. Thus the value from choosing any rule is not higher in the new model than in the original. Note however that the payoff of the indifferent rule $g_p(R^I)$ equals $\frac{1}{|K|} \sum_k g(k)$ and is thus known in advance. This implies that whenever R^I is optimal in the original model, it is still optimal in the modified one for any value of n .

In the modified model, the choice as to whether to switch to an outside option or not depends on a subtle bayesian updating after observing $\frac{1}{n} \sum_t g(R(L_t))$. Still, the DM can use the following rule: switch to g' iff $g' > \frac{1}{n} \sum_t g(R(L_t))$. Since by the law of large numbers, $\frac{1}{n} \sum_t g(R(L_t))$ converges almost surely to $\mathbb{E}_p g(R(L))$ when n becomes large, this switching rule yields an expected payoff going to $\max\{g_q(R), g'\}$ when n becomes large. This implies that the choice of a rule in the modified problem gives

an expected payoff that becomes arbitrarily close to the payoff in the original problem. Therefore, whenever all optimal rules are strictly consistent in the original model, the same remains true with finite sampling, for n large enough.

Note finally that the results of this section extend to any model in which $g_p(R)$ is observed with noise. The result on the optimality of the indifferent rule is true for any noise structure, and the one on the optimality of strictly consistent rule holds as long as the noise is small enough.

6.3. Risk loving. The utility function u in our model could, of course, be also derived from evolutionary concerns as in e.g. Robson (1996a) and Robson (2001a). In many such models the endogenously determined utility function exhibits risk-neutrality or -aversion. In winner-take-all environments Robson (1996b) and Dekel and Scotchmer (1999) find that evolution favors risk-loving utility functions at least to some extent. Our results show that in this case, with or without the presence of an outside option, the decision maker's utility maximizing choice rule is a strictly consistent one.

6.4. Small and large stakes. Instead of considering high or low risk aversion levels, our results have a natural interpretation for a fixed risk aversion level, but considering different sizes of stakes.

Consider indeed the same decision problem as in our model, but in which all material payoffs, whether from decisions or from the outside option, are multiplied by a constant $\gamma > 0$. Then, the decision problem is the same as before, except that the decision maker's utility $u(g)$ is now replaced by $v(g) = u(\gamma g)$. Since the coefficient of risk aversion of v at g is γ times that of u at γg , it is bounded above by $\gamma\bar{\rho}$ and $\gamma\rho$.

We say that stakes are high when γ is sufficiently large, and that stakes are low when γ is sufficiently small. Our main theorem can be now rephrased in terms of stakes instead of risk aversion as follows:

Theorem 4. *Consider a risk averse decision maker. If the stakes are sufficiently high and p is neutral, then the indifferent choice rule R^I is optimal and all strictly consistent rules are suboptimal. If stakes are sufficiently low, then for every p every optimal rule is strictly consistent and for p neutral every strictly consistent rule is optimal.*

APPENDIX A. PROOFS

A.1. Proof of Lemma 1. Recall that, for a given rule $R \in \mathcal{R}$, the ex-ante expected payoff is given by

$$\begin{aligned}
\mathbb{E}_q g_p(R) &= \mathbb{E}_q \mathbb{E}_p g(R(L)) \\
&= \sum_{g \in \mathcal{G}} q(g) \sum_{L \in \mathcal{L}} p(L) g(R(L)) \\
&= \sum_{g \in \mathcal{G}} q(g) \sum_{L \in \mathcal{L}} p(L) \frac{1}{|R(L)|} \sum_{k \in R(L)} g(k) \\
&= \sum_{L \in \mathcal{L}} \frac{p(L)}{|R(L)|} \sum_{k \in R(L)} \sum_{g \in \mathcal{G}} q(g) g(k),
\end{aligned}$$

where the last equality follows from a simple change in the order of summation.

We complete the proof of Lemma 1 by showing that $\sum_{g \in \mathcal{G}} q(g) g(k)$ does not depend on k . Since q is symmetric, for every permutation π of K we have

$$\sum_{g \in \mathcal{G}} q(g) g(k) = \sum_{g \in \mathcal{G}} q(g) g(\pi(k)).$$

By averaging over all permutations π we obtain:

$$\begin{aligned}
\sum_{g \in \mathcal{G}} q(g) g(k) &= \frac{1}{|K|!} \sum_{\pi} \sum_{g \in \mathcal{G}} q(g) g(\pi(k)) \\
&= \frac{1}{|K|!} \sum_{g \in \mathcal{G}} q(g) \sum_{\pi} g(\pi(k)) \\
&= \frac{1}{|K|!} \sum_{g \in \mathcal{G}} q(g) \sum_{k'} \frac{|K|!}{|K|} g(k') \\
&= \frac{1}{|K|} \sum_{k'} g(k').
\end{aligned}$$

QED

A.2. Proof of Lemma 2. We prove that Λ_p^s contains the extreme points of the convex hull of Λ_p in $\mathbb{R}^{|K|}$. By the supporting hyperplane theorem, it suffices to prove that, for any vector $v = (v(k))_k \in \mathbb{R}^{|K|}$, $\max_{\lambda_p \in \Lambda_p} \sum_k \lambda_p(k) v(k)$ is attained at some $\lambda_p \in \Lambda_p^s$. Interpret $v(k)$ as a “fictitious utility” for the choice k . For $L \subseteq K$, let $v(L) = \frac{1}{|L|} \sum_{l \in L} v(l)$. Let π be a permutation of K that orders the coordinates of v such that $v(\pi(1)) \geq v(\pi(2)) \geq \dots \geq v(\pi(k))$. Maximizing $\sum_k \lambda_p(k) v(k)$ over $\lambda_p \in \Lambda_p$ is equivalent to maximizing the expected “fictitious utility” $\sum_{L \in \mathcal{L}} p(L) v(R(L))$ over all rules.

The rule R^π that selects the least element according to π in every choice set, $R(L) = \min\{l, \pi(l) \in L\}$, maximizes each term of the sum $\sum_{L \in \mathcal{L}} p(L)v(R(L))$, so it maximizes the sum. Also, R^π is strictly consistent, since it is the rule that corresponds to the preference relation $\pi(1) \succ \pi(2) \succ \dots \succ \pi(k)$. Hence, $\lambda_p(R^\pi)$ belongs to Λ_p^s , and achieves $\max_{\lambda_p \in \Lambda_p} \sum_k \lambda_p(k)v_k$.

QED

A.3. Proof of Theorem 3. In order to prove Theorem 3 the following two Lemmas are useful.

Lemma 3. *Let $R \in \mathcal{R}^s$ and $R' \in \mathcal{R}$. If $\lambda_p(R') = \lambda_p(R)$, then $R' = R$.*

Proof: Consider w.l.o.g. the strictly consistent rule R corresponding to the preference relation $1 \succ 2 \succ 3 \succ \dots \succ |K|$, and let R' be a rule such that $\lambda_p(R') = \lambda_p(R)$. Since $R(L) = \{1\}$ whenever $1 \in L$,

$$\lambda_p(R)(1) = \sum_{1 \in L} p(L) \leq \sum_{1 \in L, 1 \in R'(L)} \frac{p(L)}{|R'(L)|} = \lambda_p(R')(1).$$

Since p has full support, the inequality above is an equality iff. $R'(L) = \{1\}$ whenever $1 \in L$. Now we have

$$\lambda_p(R)(2) = \sum_{2 \in L, 1 \notin L} p(L) \leq \sum_{2 \in L, 1 \notin L, 2 \in R'(L)} \frac{p(L)}{|R'(L)|} = \lambda_p(R')(2).$$

Here again, equality holds only if $R'(L) = \{2\}$ whenever $2 \in L$ and $1 \notin L$.

By induction on k , we obtain that $R'(L) = \{k\}$ whenever $k \in L$ and $1, \dots, k-1 \notin L$, i.e. $R' = R$.

QED

Lemma 4. *For every non-constant vector $(a_k)_{k \in K} \in \mathbb{R}^{|K|}$ and every non-constant gain function g , there exists a permutation g^π of g such that $\sum_k a_k g^\pi(k) \neq 0$.*

Proof: Consider a vector $(a_k)_{k \in K} \in \mathbb{R}^{|K|}$ such that for all permutations g^π of a non-constant gain function g we have $\sum_k a_k g^\pi(k) = 0$. Consider the permutation π that only exchanges two indexes, $i, j \in K$. Then we have both

$$\sum_{k \neq i, j} a_k g(k) + a_i g(i) + a_j g(j) = 0$$

and

$$\sum_{k \neq i, j} a_k g(k) + a_i g(j) + a_j g(i) = 0.$$

The difference of these two expressions gives

$$a_i g(i) + a_j g(j) = a_i g(j) + a_j g(i),$$

or, equivalently,

$$(a_i - a_j)(g(i) - g(j)) = 0.$$

Thus, for every $i, j \in K$ we have $a_i = a_j$ or $g(i) = g(j)$. By assumption there exist $i, j \in K$ such that $g(i) \neq g(j)$, and thus for these we have $a_i = a_j$. Let $a = a_i = a_j$.

For every $k \neq i, j$, since we cannot have both $g(k) = g(i)$ and $g(k) = g(j)$ we have either $a_k = a_i = a$ or $a_k = a_j = a$. Therefore the vector $(a_k)_{k \in K}$ is constant.

QED

Proof of Theorem 3: Let $R \in \mathcal{R} \setminus \mathcal{R}^s$. By Lemma 2, $\lambda_p(R)$ is a convex combination of choice distributions in Λ_p^s . That is, there exists a distribution μ over \mathcal{R}^s such that

$$\lambda_p(R) = \sum_{R' \in \mathcal{R}^s} \mu(R') \lambda_p(R').$$

We now have for every g :

$$(A.1) \quad g_p(R) = \sum_{R' \in \mathcal{R}^s} \mu(R') g_p(R') = \mathbb{E}_\mu g_p(R').$$

Therefore, for every g , the distribution of $g_p(R')$ under μ is a mean preserving spread of the constant $g_p(R)$. This remains true when g is taken at random according to q : the distribution of $g_p(R')$ under q and μ is a mean preserving spread of the distribution of $g_p(R)$ under q .

We now show that this mean-preserving spread is strict. To show that it suffices to show that the mean preserving spread of equation (A.1) is strict for one g in the support of q . I.e. we need to prove that there exists g in the support of q and R' in the support of μ such that $g_p(R') \neq g_p(R)$.

By Lemma 3 there exists a rule $R' \in \mathcal{R}^s$ such that $\alpha(R') > 0$ and $\lambda_p(R') \neq \lambda_p(R)$. Let $a_k = \lambda_p(R')(k) - \lambda_p(R)(k)$. Since $\lambda_p(R') \neq \lambda_p(R)$, there exists k s.t. $a_k \neq 0$. But then since $\sum_k a_k = 0$, a is non-constant. Then, by q symmetric and Lemma 4, there exist g in the support of q s.t. $\sum_{k \in K} a_k g(k) \neq 0$. The results follows since $g_p(R) - g_p(R') = \sum_k a_k g(k)$.

To prove the final statement of Theorem 3 we note that under p neutral for any two strictly consistent rules $R', R'' \in \mathcal{R}^s$ their choice distributions $\lambda_p(R')$ and $\lambda_p(R'')$ are permutations of each other. But then under q symmetric the induced distribution over material gains g is the same for both rules. QED

A.4. Proof of Theorem 2. We first show that under p neutral, the choice distribution of the indifference rule R^I is uniform: For every k and every

permutation π of K we have

$$\begin{aligned}\lambda_p(R^I)(k) &= \sum_{L, k \in L} p(L) \frac{1}{|L|} \\ &= \sum_{L, k \in L} p(\pi(L)) \frac{1}{|\pi(L)|} \\ &= \sum_{L, k \in \pi^{-1}(L)} p(L) \frac{1}{|L|}\end{aligned}$$

Averaging over all permutations π we have:

$$\begin{aligned}\lambda_p(R^I)(k) &= \frac{1}{|K|!} \sum_L \sum_{\pi, k \in \pi^{-1}(L)} p(L) \frac{1}{|L|} \\ &= \sum_L \frac{|L|}{|K|} p(L) \frac{1}{|L|} \\ &= \frac{1}{|K|},\end{aligned}$$

hence the result.

Now consider an arbitrary rule $R \in \mathcal{R}$ with $R \neq R^I$. Assume $\lambda_p(R) = \lambda_p(R^I)$. Then R and R^I yield the same distribution over gains for any distribution q .

Thus, suppose $\lambda_p(R) \neq \lambda_p(R^I)$. Consider an arbitrary gain function g in the support of q . If g is constant then both rules R and R^I induce the same payoff distribution (equal to the constant value of g with probability 1). Now consider g non-constant. Then, by the assumed symmetry of q , all its permutations g^π satisfy $q(g^\pi) = q(g)$. Denote by the orbit of g the set of all permutations of g (all of which receive the same probability weight under q). Then rule R^I induces the same payoff (equal to $\frac{1}{|K|} \sum_k g(k)$) for all gain functions in the orbit of g .

This is not true for rule R . To see this let $a_k = \lambda_p(R)(k) - \lambda_p(R^I)(k)$. Since $\lambda_p(R) \neq \lambda_p(R^I)$, there exists k s.t. $a_k \neq 0$. But then since $\sum_k a_k = 0$, a is non-constant. Then, by Lemma 4, there exist a permutation g^π of g s.t. $\sum_{k \in K} a_k g^\pi(k) \neq 0$. Since $g_p^\pi(R) - g_p^\pi(R^I) = \sum_k a_k g^\pi(k)$, this implies that $g_p^\pi(R) \neq g_p^\pi(R^I)$.

Thus, while for any non-constant g in the support of q , the overall expected gain conditional on q realizing in the orbit of g is the same for both rules R and R^I , the gain from rule R^I is the same for all g in the orbit, while there is variation of this gain for rule R . Hence, conditional on this orbit, the payoff distribution of under R is a strict mean preserving spread of the payoff distribution under R^I .

To complete the proof, observe that the payoff distribution of a rule under q is the average of all its payoff distributions over the different orbits. Since the payoff distribution of R is a mean preserving spread of the payoff distribution of R^I on all orbits, and since there is one orbit with positive probability under q for which this mean preserving spread is strict, the payoff distribution of R under q is a strict mean preserving spread of the payoff distribution of R^I under q . QED

A.5. Proof of Theorem 1. We first prove the first statement of the theorem by showing that, under p neutral, if the DM is sufficiently risk averse, then for any rule $R \in \mathcal{R}$ and for any distribution of the outside option g' (given the full support assumption) we have

$$\mathbb{E}_{q,g'} [\max \{u(g_p(R^I)), u(g')\}] \geq \mathbb{E}_{q,g'} [\max \{u(g_p(R)), u(g')\}],$$

with a strict inequality for $R \in \mathcal{R}^s$.

The above inequality is satisfied with equality for any rule $R \in \mathcal{R}$ such that $\lambda_p(R) = \lambda_p(R^I)$. Thus, suppose $\lambda_p(R) \neq \lambda_p(R^I)$ (as is the case for any $R \in \mathcal{R}^s$ by Lemma 3). Then as q is symmetric and as it has at least one non-constant gain function in its support, note that the argument in the proof of Theorem 2 actually implied that rule R induces a distribution over gains under q that has a strictly wider support than the distribution over gains induced by rule R^I . (This is true for every orbit of any non-constant gain function g , and, thus, true for any q that satisfies our assumptions.)

As u is an increasing function we have that $\max \{u(g_p(R)), u(g')\} = u(\max \{g_p(R), g'\})$ (for any R).

For any $R \in \mathcal{R}$ let $X(R) = \max \{g_p(R), g'\}$. By the full support assumption for g' , we then have that $X(R)$, with R such that $\lambda_p(R) \neq \lambda_p(R^I)$, has strictly wider support than $X(R^I)$.

For any discrete random material gain distribution, X , the certainty equivalent, denoted $CE(X)$, tends to the lowest value in its support, as the lower bound of absolute risk aversion $\underline{\rho}$ tends to infinity. Thus, an infinitely risk averse DM strictly prefers rule R^I over R . By continuity, there then must be a minimal level of absolute risk aversion such that any DM with risk aversion uniformly higher than this level also strictly prefers rule R^I over R .

We now prove the second sentence of the theorem by proving that, for any p and for any $R' \in \mathcal{R} \setminus \mathcal{R}^s$ there is a $R^* \in \mathcal{R}^s$ such that a risk-neutral DM strictly prefers R^* over R' , i.e. such that $\mathbb{E}_{q,g'} [\max \{g_p(R'), g'\}] < \mathbb{E}_{q,g'} [\max \{g_p(R^*), g'\}]$. Then by continuity, there is a maximal level of absolute risk aversion such that any DM with risk aversion uniformly lower than this level also strictly prefers rule R^* .

By Theorem 3 for any $R' \in \mathcal{R} \setminus \mathcal{R}^s$ there is a distribution μ over strictly consistent rule that is strictly riskier than rule R' .

As the maximum is a convex function and as g' has full support (in particular it has support where the gains distributions induced by rule R' and distribution μ differ) we have

$$\mathbb{E}_{q,g'} [\max\{g_p(R'), g'\}] < \mathbb{E}_{q,g'} \left[\sum_{R \in \mathcal{R}^s} \mu(R) \max\{g_p(R), g'\} \right].$$

Interchanging the order of summation we have

$$\mathbb{E}_{q,g'} [\max\{g_p(R'), g'\}] < \sum_{R \in \mathcal{R}^s} \mu(R) \mathbb{E}_{q,g'} [\max\{g_p(R), g'\}].$$

Thus, there must be at least one $R^* \in \mathcal{R}^s$ such that

$$\mathbb{E}_{q,g'} [\max\{g_p(R'), g'\}] < \mathbb{E}_{q,g'} [\max\{g_p(R^*), g'\}].$$

To finish the proof note that under p neutral all strictly consistent rules induce the same distribution over gains. QED

REFERENCES

- ARROW, K. J. (1959): “Rational choice functions and orderings,” *Economica*, 26, 121–27.
- BLOCK, H. D., AND J. MARSCHAK (1960): “Random orderings and stochastic theories of responses,” in *Contributions to probability and statistics*, ed. by I. Olkin, S. Ghrye, W. Hoeffding, W. Madow, and H. Mann, pp. 97–132, Stanford, CA. Stanford University Press.
- CHIAPPORI, P.-A., AND I. EKELAND (1999): “Aggregation and Market Demand: An Exterior Differential Calculus Viewpoint,” *Econometrica*, 67, 1435–1457.
- CHOI, S., R. FISMAN, D. GALE, AND S. KARIV (2007): “Consistency and heterogeneity of individual behavior under uncertainty,” *The American Economic Review*, 97(5), 1921–1938.
- CRAWFORD, I., AND K. PENDAKUR (2013): “How many types are there? A simple, theory-consistent approach to unobserved heterogeneity,” *The Economic Journal*, 123(567), 77–95.
- DEAN, M., AND D. MARTIN (2010): “How rational are your choice data?,” mimeo.
- (2011): “Testing for rationality with consumption data: Demographics and heterogeneity,” *Available at SSRN 1894785*.
- (2015): “Measuring rationality with the minimum cost of revealed preference violations,” *Review of Economics and Statistics*, forthcoming.

- DEKEL, E., AND S. SCOTCHMER (1999): "On the evolution of attitudes towards risk in winner-take-all games," *Journal of Economic Theory*, 87(1), 125–143.
- ECHENIQUE, F., S. LEE, AND M. SHUM (2011): "The money pump as a measure of revealed preference violations," *Journal of Political Economy*, 119(6), 1201–1223.
- GORMAN, W. M. (1953): "Community preference fields," *Econometrica*, pp. 63–80.
- GUL, F., AND W. PESENDORFER (2013): "Random Utility Maximization with Indifference," mimeo, Princeton University.
- LUCE, R. D. (1959): *Individual Choice Behavior: A Theoretical Analysis*. Wiley, New York.
- MAS-COLLEL, A., M. D. WHINSTON, AND J. R. GREEN (1995): *Microeconomic Theory*. Oxford University Press, Oxford, UK.
- MATEJKA, F., AND A. MCKAY (2015): "Rational inattention to discrete choices: A new foundation for the multinomial logit model," *American Economic Review*, 105(1), 272–98.
- MERTON, R. (1973): "Theory of Rational Option Pricing," *The Bell Journal of Economics and Management Science*, 4, 141–183.
- NETZER, N. (2009): "Evolution of time preferences and attitudes toward risk," *The American Economic Review*, 99(3), 937–955.
- RASMUSEN, E. (2007): "When Does Extra Risk Strictly Increase the Value of Options?," *The Review of Financial Studies*, 20(5), 1647–1667.
- RAYO, L., AND G. S. BECKER (2007): "Evolutionary efficiency and happiness," *Journal of Political Economy*, 115(2), 302–337.
- ROBSON, A., AND L. SAMUELSON (2011): "The evolution of decision and experienced utilities," *Theoretical Economics*, 6(3), 311–339.
- ROBSON, A. J. (1996a): "A biological basis for expected and non-expected utility," *Journal of Economic Theory*, 68, 397–424.
- ROBSON, A. J. (1996b): "The evolution of attitudes to risk: Lottery tickets and relative wealth," *Games and Economic Behavior*, 14(2), 190–207.
- ROBSON, A. J. (2001a): "Why would nature give individuals utility functions?," *Journal of Political Economy*, 109, 900–14.
- (2001b): "The biological basis of economic behavior," *Journal of Economic Literature*, 38, 11–33.
- SAMUELSON, L., AND J. SWINKELS (2006): "Information, evolution and utility," *Theoretical Economics*, 1, 119–42.
- STEINER, J., AND C. STEWART (2014): "Perceiving Prospects Properly," mimeo.

UZAWA, H. (1956): “Note on preference and axioms of choice,” *Annals of the Institute of Statistical Mathematics*, 8, 35–40.

CNRS – COLE POLYTECHNIQUE AND LONDON SCHOOL OF ECONOMICS, LONDON,
GOSSNER@POLY.POLYTECHNIQUE.EDU

CENTER FOR MATHEMATICAL ECONOMICS, BIELEFELD UNIVERSITY, CHRISTOPH.KUZMICS@UNI-
BIELEFELD.DE