
Rationality and the speed of decision-making (extended abstract)

Michael Mandler

Department of Economics
Royal Holloway College, University of London
Egham, Surrey TW20 0EX
United Kingdom

Abstract

We consider agents who choose by proceeding through an ordered list of criteria and give the lower bound on the number of criteria that are needed for an agent to make decisions that obey a given set of preference rankings. Agents with rational preferences can always use lists with the lower-bound number of criteria while any agent with nonrational preferences must on some domains use strictly more criteria. We preview some of the results in Mandler (2009) and explain in more detail the order-theoretic link between rationality and rapid decision-making.

1 INTRODUCTION

One way to choose between alternatives is to divide the domain of possible alternatives into various categories and then decide on an ordering of those categories. If this step does not discriminate sufficiently one can then move on to a second categorization, and so on. For any pair of alternatives, an agent can proceed in this fashion through a sequence of orderings, choosing the alternative that is recommended by the first ordering that actually ranks the pair. While this choice procedure, which we call a *checklist of criteria*, does not look like classical economic behavior it does describe one of the rough-and-ready methods that people use to make decisions.

A checklist user's decisions implicitly define a preference relation: for any pair of alternatives that are fed into the choice procedure, we can label the chosen alternative to be strictly preferred, and if no criterion in a checklist discriminates between the alternatives we can label the pair unranked. Which types of checklist users can make decisions the most rapidly, the rational agents whose preferences are complete and transitive or the irrational agents? The answer to this question

turns out to be unambiguous: the rational agents can on some domain always use a strictly shorter checklist of criteria than the irrational agents can and on no domain do they have to use a longer list. Rational agents can in fact make preference discriminations with great efficiency; the number of discriminations they can make is an exponential function of the number of criteria at their disposal. And the speed advantage of rationality holds even when compared to agents whose sole source of irrationality is incompleteness.

Our conclusions depart from Herbert Simon's view that the assumption of utility maximization places 'a heavy (often unbearable) computational burden on the decision maker' (Simon (1990)). For checklist users, in contrast, it is the rational agents who have the easier computational task and therefore make decisions more quickly. The difficulty with Simon's and kindred views is that they take utility to be a set of external facts that an agent has to uncover. If instead we take an agent's decision-making procedure as primitive, then the computational burden of decision-making becomes a function of the procedure alone. We may then ask if the procedures that lead to rational decisions carry a heavier or lighter burden than the irrational procedures.

The greater decision-making efficiency of rational preferences derives from a fundamental fact about linear orders. On domains that consist of at least four elements, linear orders are the only binary relations such that any two subsets with the same number of elements are order-isomorphic. In our setting, a rational preference relation induces a linear ordering over its indifference classes and the order isomorphism fact implies that no matter how we constrain the criteria that an agent can use, the agent will be able to divide up his indifference classes into order-isomorphic subsets that a single criterion can then order simultaneously. As we will see, this is the source of the speed advantage of rational agents.

Checklists of criteria bear a close resemblance to Chip-

man’s (1960) lexicographic theory of utility, which was designed to provide representations for preferences that do not have classical utility functions (such as lexicographic preferences). In Chipman, the preference between two items is determined by the first utility in a sequence of real-valued utility functions that ranks the items. In our model a sequence of criteria replaces the sequence of utilities and, to gauge decision-making efficiency, we will use an accounting system that measures the discriminatory capacity of criteria. It turns out that even the least discriminating criteria, when strung together in a sequence, can generate a vast set of preference discriminations. The absence of a measure of discriminatory capacity in Chipman has hid from view the power of his representation concept.

More concretely, the present work follows in the footsteps of the checklists studied in Mandler, Manzini, and Mariotti (2008). The MMM checklists are constrained to be rational, however, which precludes the rational versus irrational horse race that we pursue here. Kalai, Rubinstein, and Spiegler (2002), Apesteguia and Ballester (2005, 2008), and Manzini and Mariotti (2007) also consider choice behavior that can be decomposed into ‘criteria.’ KRS and Apesteguia and Ballester (2005) have an agenda similar to the present paper but they seek out the most concise set of rationales for decisions, not the most efficient use of discriminatory capacity. The introduction of an accounting system for discriminatory capacity of criteria marks out a distinct agenda.

In addition to previewing some of the results in Mandler (2009), we will use this forum to explain in more detail the order-theoretic connection between rationality and the speed of decision-making.

2 CHECKLISTS OF CRITERIA

Fix some domain of alternatives X . We consider agents who make decisions between pairs of alternatives $x, y \in X$ by proceeding through a checklist of criteria C_1, \dots, C_L , where each C_i is an asymmetric binary relation on X , and choosing the alternative that is ranked superior by the first C_i that actually orders x and y . One should think of each criterion as dividing X into a relatively small number of categories or equivalence classes. Outside of asymmetry, criteria are unrestricted; they could be rational orderings, they could cycle, or they could leave some alternatives unranked.

An agent might construct a criterion C_i in two steps, first a division of X into categories that partition X and then a ranking (though not necessarily a complete ranking) of these categories. Formally, the categories are the equivalence classes of C_i , which we define in

section 3.

For example, X might be a large set of cars, criterion 1 could divide cars into cheap, mid-priced, and expensive categories, criterion 2 could divide cars into sedans, convertibles, SUV’s, and all others, and finally criterion 3 might divide cars by their country of origin. After the categorization step, the agent must figure out how to order each set of categories. Then, to decide between a pair of cars, the agent proceeds in sequence through the criteria: if one car has the better price ranking, then the agent buys that car; if not, the agent goes on to the next categorization, and so on until the agent comes to a categorization that does rank the pair. If the pair is ranked by none of the criteria, the agent declares either car to be acceptable.

We can define a preference relation \succ from a checklist by seeing for any pair in X which element of the pair is chosen. Formally, a preference \succ is another asymmetric binary on X . Given a preference \succ , we say that $C = (C_1, \dots, C_L)$ is a *checklist for \succ* if and only if, for all $x, y \in X$,

$$x \succ y \Leftrightarrow \exists i \text{ with } 1 \leq i \leq L \text{ such that } x C_i y \text{ and} \\ \text{not } y C_j x \text{ for all } j < i.$$

Notice that it is only the first criterion that ranks a pair that determines preference; if the first criterion C_i that ranks x and y has $x C_i y$ then some later criterion C_k can have $y C_k x$ and still $x \succ y$. We will also say that a preference \succ *has the checklist C* if C is a checklist for \succ . A checklist for a preference amounts to a mild extension of Chipman’s (1960) definition of a lexicographic utility.

Checklists of properties. Some seemingly crude sequential decision procedures may not at first appear to qualify as checklists, for example, the model in Mandler, Manzini, and Mariotti (2008). Suppose an agent chooses between two items x and y by proceeding through a list of ‘properties,’ P_1, P_2, P_3, \dots , where each P_i is a subset of X . The agent first checks if x and y have property 1 (i.e., if $x \in P_1$ and $y \in P_1$). If either or x or y has property 1 but the other does not then the agent chooses the item that does. If both x and y have property 1 or neither does then the agent proceeds to property 2, and so on. A list of properties qualifies as a checklist of criteria since the agent could make the same decisions by using a C_1 that ranks any item in P_1 as strictly superior to any item in $X \setminus P_1$, a C_2 that ranks item in P_2 as strictly superior to any item in $X \setminus P_2$, and so on.

Checklists can be seen as a decision procedure for two-element choice sets, but they can also be used to decide from a larger choice set A by letting the C_i sequentially eliminate C_i -dominated items from A . Our results

apply to this interpretation as well.

3 EQUIVALENCE CLASSES

We measure decision-making speed by the number of criteria that a checklist uses. Notice however that checklists that make only a handful of preference discriminations can make do with fewer criteria than checklists that make many discriminations. Similarly checklists whose criteria are themselves highly discriminating will enjoy a speed advantage. We therefore compare the efficiency of two checklists only when they lead to the same number of preference discriminations and the criteria of both checklists wield the same discriminatory power. To this end, we need a gauge of the number of discriminations that applies both to preferences and criteria.

Given an asymmetric relation $>$ on X , we define the binary relation \approx on X by

$$x \approx y \Leftrightarrow \{z \in X : z > x\} = \{z \in X : z > y\} \text{ and} \\ \{z \in X : x > z\} = \{z \in X : y > z\}.$$

We call \approx the *equivalence relation* of $>$. When $>$ is transitive, \approx is a textbook construction; see Fishburn (1970) and Mandler (2009) for applications in economics. For the equivalence relation of a preference \succ , we use the symbol \sim and call it the *indifference relation* of \succ . Given an asymmetric relation $>$ and its equivalence relation \approx , a *$>$ -equivalence class* is a nonempty $I \subset X$ such that (1) $x, y \in I \Rightarrow x \approx y$ and (2) $(x \in I \text{ and } x \approx y) \Rightarrow y \in I$.

Our measure of the discriminatory capacity of a checklist will be an upper bound imposed on the number of equivalence classes for the criteria in the checklist. Given a positive integer p , we say that $C = (C_1, \dots, C_L)$ is a *p -checklist* for \succ if and only if C is a checklist for \succ and each C_i has p or fewer C_i -equivalence classes. It is only for convenience that each C_i in a p -checklist has the same bound on its discriminatory capacity; in Mandler (2009) we show how to let p vary by criterion.

Our equivalence relations furnish a convenient definition of rationality. Given an asymmetric relation $>$ on X with equivalence relation \approx , define $>$ to be *complete* if and only if for every $x, y \in X$ either $x (> \cup \approx) y$ or $y (> \cup \approx) x$ or both. We define $>$ to be *rational* if and only if $>$ is complete and transitive. It is easy to confirm that an asymmetric relation $>$ is complete and transitive if and only if it is negatively transitive (i.e., where $x \not> y \not> z$ implies $x \not> z$), the traditional definition of rationality for asymmetric relations. We have followed the more circuitous path both because we in any event need the count of equivalence classes to define decision-making efficiency and because a ‘com-

plete and transitive’ definition of rationality gives a finer classification of possible failures of rationality.

4 THE AVAILABILITY OF CHECKLISTS

If each criterion in a checklist is rational, the preference defined by the checklist is rational.

Theorem 1 *If each criterion C_i in the p -checklist $C = (C_1, \dots, C_L)$ is rational and C is a checklist for \succ , then \succ is rational.*

It need not be difficult to choose rationally as Herbert Simon suggested. For a checklist user, rationality can just indicate that the agent uses rational criteria.

By way of proof, notice that for each $x \in X$ we can define an integer, written in base p , that has as its first digit the number of C_1 -equivalence classes that x is C_1 -superior to, has as its second digit the number of C_2 -equivalence classes that x is C_2 -superior to, and so on. Then $x \succ y$ if and only if there is a first digit for the integers assigned to x and y that differs and it is the integer assigned to x that is larger in that digit. The preference \succ thus inherits the natural ordering of some subset of the integers and hence must be complete and transitive.

Any criterion C_i with one or two equivalence classes must be complete and transitive: if C_i has two equivalence classes I and J then transitivity is vacuous and since either $I C_i J$ or $J C_i I$ must obtain C_i is complete, while if C_i has just one equivalence class then both completeness and transitivity are vacuous.¹ Thus all of the preferences generated by 2-checklists must be rational. Since the ‘checklist of properties’ model, discussed in section 2, implicitly imposes a two-equivalence-class speed limit on criteria, it therefore cannot stage an efficiency contest between rational and irrational agents: the irrational agents cannot enter the starting gate.

As long as $p \geq 3$, the choice behavior of essentially any agent could in principle arise from a checklist.

Theorem 2 *Any preference \succ with a finite number of indifference classes has a p -checklist if $p \geq 3$.*

For any two \succ -indifference classes I and J with $I \succ J$, define a criterion $C_{I,J}$ by

$$x C_{I,J} y \text{ if and only if } (x \in I \text{ and } y \in J).$$

The checklist that consists of all such criteria (arranged in any order) is evidently a checklist for \succ

¹For any binary relation R on X , any subsets $A, B \subset X$, and any $c \in X$, we use ARB to mean $a R b$ for all $a \in A$, $b \in B$ and use cRA to mean $c R a$ for all $a \in A$.

and it is a p -checklist for any $p \geq 3$ since each $C_{I,J}$ has three equivalence classes, namely I , J , and $X \setminus (I \cup J)$.

If the preference \succ has a finite number of indifference classes n , one may show there is no p -checklist for \succ with fewer than $\lceil \log_p n \rceil$ criteria.² A proof is in Mandler (2009). We will see in section 6 that there are checklists for rational preferences that achieve this theoretical minimum. If we fix p , then, as the number of \succ -indifference classes n increases, $\lceil \log_p n \rceil$ increases slowly (at a less-than-polynomial rate) and the ratio of the minimum number of criteria that a checklist requires to n will converge rapidly to 0 as n increases. Alternatively, if we take the number of criteria to be the exogenous fact, then we can conclude that the number of preference discriminations that a rational agent can make is an exponential function of the number of criteria the agent uses. For rational agents checklists can therefore serve as highly efficient tools for sifting through alternatives. I stress the verb ‘can’: any preference with a checklist also has checklists with arbitrarily many criteria (for example, a checklist could pointlessly repeat criteria or use criteria that simply reverse orderings made earlier in the checklist). But rational agents at least have the option of deciding as quickly as any other checklist user. In the next section, we lay out some examples that show that irrational agents can fall well short of the $\lceil \log_p n \rceil$ performance standard.

5 SLOW CHECKLISTS: EXAMPLES AND EXERCISES

We saw in the proof for Theorem 2 that we could build a 3-checklist for an arbitrary \succ by assigning one criterion for every pair of indifference classes that \succ strictly ranks. If \succ is complete the checklist that results will have $\binom{n}{2} = \frac{n(n-1)}{2}$ criteria, not an impressive performance.

We can do better, even when $p = 3$, by defining for each \succ -indifference class I a criterion C_I by $y C_I x \Leftrightarrow (x \in I \text{ and } y \succ x)$, and then forming a checklist by arranging the criteria defined in this fashion in arbitrary order. This construction defines a checklist for \succ since for any pair x, y with $y \succ x$ the only criterion in the checklist that ranks x and y is the C_I where I is the \succ -indifference class that contains x .

Thus, if we measure speed by the number of criteria, the worst conceivable case for an agent with n indifference classes is that the agent must use a checklist with n criteria. To see that the worst case can in fact occur, consider the n -cycles, the class of preference relations

²For any real number a , $\lceil a \rceil$ denotes the smallest integer b such that $b \geq a$.

\succ_n , each on the domain $\{1, 2, \dots, n\}$ and defined by

$$n \succ_n n-1 \succ_n \dots \succ_n 1 \succ_n n.$$

As long as $n > 4$, the shortest 3-checklist for \succ_n has n criteria. To sketch the argument for this conclusion, suppose there is a 3-checklist for \succ_n with fewer than n criteria. Then there must be a criterion C_i in the checklist that makes at least two of the rankings in \succ_n . For example suppose $x \succ_n y \succ_n z$ and that C_i has $x C_i y C_i z$. Then, since $p = 3$, $x + 1$ (where $n + 1$ denotes 1) must be in the same C_i -equivalence class as x, y , or z . But since $n > 4$, $x + 1$ is not \succ_n -ranked against y or z and thus cannot share an equivalence class with x, y , or z . We also cannot have a C_i that makes the rankings $w \succ_n x$ and $y \succ_n z$ for any four elements from distinct \succ -indifference classes since the four must all be in distinct C_i -equivalence classes (for instance w cannot share a C_i -equivalence class with y since if it is \succ_n -ranked against z it must be that $z \succ_n w$ rather than $w \succ_n z$), violating the requirement that $p = 3$.

The n -cycles illustrate that irrational agents may have to use criteria that sift alternatives slowly; unlike rational preferences, the ratio between the minimum number of criteria that a checklist requires to the number of indifference classes n need not be near 0 even when n is large.

We offer a couple of exercises, which involve arguments similar to those above, that the reader may find entertaining.

Exercise 1. Find a 3-checklist with 3 criteria for the 4-cycle (the preference \succ on $\{1, 2, 3, 4\}$ defined by $4 \succ 3 \succ 2 \succ 1 \succ 4$).

Exercise 2. Assuming \succ is complete and has n indifference classes, find a 3-checklist with $n - 1$ criteria.

The moral of these examples and exercises is that while the incompleteness of \succ may appear to bring a speed advantage (your criteria need to make fewer rankings in total), incompleteness forces an agent to make sure that any \succ -unranked pair is unranked by every criterion in any checklist for \succ . This burns up criterion equivalence classes and thus forces an agent to use more criteria.

6 RATIONAL VS IRRATIONAL AGENTS

For any rational \succ with n indifference classes and any discriminating capacity p for criteria there is always a p -checklist for \succ with the theoretical minimum number of criteria $\lceil \log_p n \rceil$. While there are cases when an irrational \succ will enjoy the same speed, such cases are

fragile; changes in the domain or in p will lead an irrational \succ to use a slower checklist that employs more criteria.

To illustrate this principle, consider the goal that a preference should always use the theoretical minimum number of criteria no matter how we restrict the domain of alternatives. Restricted domains are a pertinent test for decision-making speed since an agent might know that only certain elements in X could conceivably be available in an upcoming stretch of time; during that time the agent can make do with a checklist that can decide only between alternatives drawn from the restricted domain. Or the domain might be fixed once and for all but we are outside observers unsure of which subset of the largest imaginable set of alternatives is the true domain.

We say the binary relation \succ' is a *subrelation* of \succ if and only if there is a subdomain $X' \subset X$ such that

$$\succ' = \succ \cap (X' \times X').$$

A preference \succ has *quick checklists on all domains* if and only if for all subrelations \succ' of \succ and all $p \geq 3$ there is a p -checklist for \succ' with $\lceil \log_p n \rceil$ criteria.

Theorem 3. *If the preference \succ has n indifference classes, where $10 \leq n < \infty$, then \succ has quick checklists on all domains if and only if \succ is rational.*

Thus even if the only irrationality of a preference is incompleteness it will on some domain require more criteria than a fully rational preference. For instance, if we begin with a rational preference and omit some of its rankings of indifference classes – in such a way that the new preference remains transitive and retains the same number of indifference classes – then the new preference will on some domains require a checklist with more criteria.

The proof of the ‘only if’ half of Theorem 3 is a long and tedious enumeration of cases (Mandler (2009)). To see why a rational preference always has a checklist with the theoretical minimum number of criteria, consider a rational \succ with exactly p^2 indifference classes, in which case $\lceil \log_p n \rceil = 2$. If we label the indifference classes of \succ , going from worst to best, as $1, \dots, p^2$ then in the unique p -checklist with 2 criteria C_1 will divide the \succ -indifference classes into the p C_i -equivalence classes $\{1, \dots, p\}$, $\{p + 1, \dots, 2p\}$, ..., $\{(p - 1)p + 1, \dots, p^2\}$ which it then linearly orders in the sequence we wrote them, going from C_1 -worst to C_1 -best. Criterion C_2 also divides the \succ -indifference classes into p equivalence classes and linearly orders them, $\{1, p + 1, (p - 1)p + 1\}$, $\{2, p + 2, (p - 1)p + 2\}$, ..., $\{p, 2p, \dots, p^2\}$ going from C_2 -worst to C_2 -best. This (C_1, C_2) works as a checklist for \succ since C_1 orders its equivalence classes as \succ does, which implies that each

pair of \succ -indifference classes that C_1 does not rank lies *within* a C_1 -equivalence class, and C_2 orders the elements of each C_1 -equivalence class just as \succ does.

This example is special in that n is an exact power p . When this is not the case, then for a rational \succ we can still construct C_1 so that the number of \succ -indifference classes in each C_1 -equivalence class differs by at most one and let C_1 order these equivalence classes as \succ does; and then again choose C_2 so that it orders the elements of each C_1 -equivalence class as \succ does. Of course if $n > p^2$, then we will need more than two criteria.

What is it about a rational \succ that ensures that the above algorithm can be executed? The key is that a rational \succ satisfies a ‘nested order-isomorphism’ property: whenever we partition the indifference classes of a rational \succ into p cells, where each cell has the same number of \succ -indifference classes and is endowed with the \succ ordering, the cells will be order-isomorphic.³ This property in turn implies that any rational \succ satisfies a slightly stronger property: for any partition of the indifference classes of \succ into cells of possibly differing cardinalities then, for any pair of these cells \mathcal{I} and \mathcal{J} with $|\mathcal{I}| \leq |\mathcal{J}|$ (again endowed with the \succ ordering), there will be an order isomorphism between \mathcal{I} and any $\mathcal{J}' \subset \mathcal{J}$ with $|\mathcal{J}'| = |\mathcal{I}|$. So when, for example, we partition the indifference classes of a rational \succ into C_1 -equivalence classes, the C_1 -equivalence classes are guaranteed to satisfy the stronger property and hence we will be able to find a C_2 that can simultaneously partition and order each C_1 -equivalence class.

We can informally summarize this discussion with the conclusion that rational preferences have the richest possible supply of order-isomorphic subrelations. In fact the nested order-isomorphism property characterizes the complete and transitive binary relations.

Theorem 4. *Let R be an asymmetric binary relation with $n \geq 4$ equivalence classes. Then R is complete and transitive if and only if, for any two sets of R -equivalence classes \mathcal{I} and \mathcal{J} such that $|\mathcal{I}| = |\mathcal{J}|$, (\mathcal{I}, R) and (\mathcal{J}, R) are order-isomorphic.*

Constructing the isomorphism for the ‘only if’ half of Theorem 4 is straightforward. The ‘if’ half is almost as simple. Since $n \geq 4$, there must be two R -equivalence classes I and J such that $I R J$. So, if R has the order-isomorphism property given in the theorem then for any pair of R -equivalence classes K and L we must have either $K R L$ or $L R K$. Hence R is complete. It follows that R , seen as an ordering of

³Two sets, where each is endowed with a binary relation, (A, \geq_A) and (B, \geq_B) , are order-isomorphic if and only if there is a bijection $f : A \rightarrow B$ such that, for all $x, y \in A$, $x \geq_A y \Leftrightarrow f(x) \geq_B f(y)$.

R -equivalence classes, is weakly connected ($I \neq J \Rightarrow$ either IRJ or JRI). The transitivity of R then follows from the easily confirmed fact that for any weakly connected non-transitive binary relation T on a four element domain Y there is both a 3-element subset of Y on which T is complete and transitive and a 3-element subset of Y on which T cycles. So if R were not transitive, there would be \mathcal{I} and \mathcal{J} , each with three R -equivalence classes, that fail to be order-isomorphic.

Theorem 4 points out a sweeping feature of a rational preference \succ : any pair of equinumerous sets of \succ -indifference classes will be order-isomorphic. But an irrational \succ may of course have *some* sets of \succ -indifference classes that are order-isomorphic. This is the reason why in Theorem 3 we are forced to characterize rationality in terms of quick checklists on all domains. For a fixed domain and fixed p , an irrational \succ may well have a p -checklist with the minimum number of criteria. The easiest way to construct such a \succ is to pick some n for the cardinality of X that is an exact power of p , say $n = p^2$, then select a C_1 with p equivalence classes, each with p elements of X , and finally a C_2 with p equivalence classes, each of which contains exactly one element from each of the C_1 -equivalence classes. As long as either C_1 or C_2 fails to be complete and transitive, the \succ that has (C_1, C_2) as a checklist will not be rational. What Theorem 3 shows is that the \succ that emerges from this construction will not have enough order-isomorphic subsets; on some subdomain of X , \succ will fail to have a checklist with the minimum number of criteria.⁴

7 FURTHER DEVELOPMENTS

The analysis of the decision-making efficiency of checklists is developed in further directions in Mandler (2009).

1. We have seen that there are carefully constructed cases where for a specific p and subdomain, an irrational \succ has a p -checklist with the theoretical minimum number of criteria. How common are these exceptions? For any n , we can calculate (up to an order-isomorphism) both the total number of irrational preferences on a domain of n elements and the number of those preferences that have p -checklists with the minimum number of criteria. Fixing p , one may show that the ratio of the latter number to the former converges to 0 at a super-exponential rate as $n \rightarrow \infty$. The exceptions become rare very quickly.

⁴The same construction explains the domain restriction in Theorem 3: there are irrational preferences with 9 indifference classes that have a 3-checklist with two criteria, and any increase in p or restriction of the domain will not require a checklist with more than two criteria.

2. So far, we have discussed decision-making efficiency entirely in terms of the number of criteria in a checklist. But we can also calculate the expectation of the number of criteria that an agent must examine in order to choose from a pair of alternatives in X , assuming that the agent is equally likely to face each possible pair and that the agent's checklist correctly decides every pair. Subject to one proviso, one may show that a rational agent can always use a checklist such that the expected number of criteria examined is at least as small as that of any other agent with the same number of indifference classes. The proviso is that indifference classes must be singletons. Since a pair of indifferent elements cannot be ranked by any criterion in a checklist, an agent facing such a pair must proceed through his or her entire checklist; a rational agent with many indifferences might therefore decide more slowly than other agents.

3. Although we have taken the integer $p \geq 2$ to be a parameter, an agent who is inventing his own categorizations of alternatives could decide how many distinctions his criteria will make. What types of criteria is it optimal to invent? Suppose we measure the cost of creating a p -checklist by the total number of criterion rankings, which equals the number of rankings of equivalence classes in each criterion multiplied by the number of criteria. Since there is no gain in using more criteria than necessary, the cost of criteria is approximated by

$$c(p, n) = \binom{p}{2} \lceil \log_p n \rceil .$$

(Some minor adjustments to this expression are required when n is not an exact power of p to allow one of the criteria to use fewer than p equivalence classes.) For fixed n , $c(p, n)$ is minimized at $p = 2$. There is an irony here: the 'checklist of properties' model in section 2 seems like the crudest way to make decisions. Yet by the above standard it is the most efficient! And recall that any agent with a 2-checklist must be rational; hence we have another basis for the conclusion that rational agents are the most efficient decision-makers.

Acknowledgements

I am grateful to Jose Apestegua, Mordecai Kurz, Mark Machina, Paola Manzini, Marco Mariotti, Herbert Scarf, and particularly Claire Blackman for helpful discussions and advice. The suggestions of the TARK referees proved highly fruitful.

References

Apestegua J. and M. Ballester, 2005, 'Minimal books of rationales,' available at <ftp://ftp.econ.unavarra.es/pub/DocumentosTrab/DT0501.PDF>.

Apestegua J. and M. Ballester, 2008, 'A characterization of sequential rationality,' available at SSRN: <http://ssrn.com/abstract=1143272>.

Chipman, J., 1960, 'The foundations of utility' *Econometrica* 28: 193-224.

Fishburn, P., 1970, *Utility Theory for Decision Making*. Wiley: New York.

Kalai, G., A. Rubinstein, and R. Spiegler, 2002, 'Rationalizing choice functions by multiple rationales' *Econometrica* 70: 2481-88.

Mandler, M., 2009, 'Rational agents are the quickest' available at <http://personal.rhul.ac.uk/uhte/035/rational.agents.are.the.quickest.Mandler.pdf>

Mandler, M., 2009. Indifference and incompleteness distinguished by rational trade. *Games and Economic Behavior*, forthcoming.

Mandler, M., P. Manzini, and M. Mariotti, 2008, 'A million answers to twenty questions: choosing by checklist' available at <http://personal.rhul.ac.uk/uhte/035/million.answers.to.twenty.questions.pdf>

Manzini, P. and M. Mariotti, 2007, 'Sequentially rationalizable choice' *American Economic Review* 97: 1824-1839.

Simon, H., 1990, 'Bounded Rationality' in J. Eatwell, M. Milgate and P. Newman (eds.) *The New Palgrave: Utility and Probability*. Macmillan, London.