Supplemental Material

Choices Without Preferences: Principles of Rational Arbitrariness

1. Endowment Effect Study

This section provides a full report of the study summarized in Table 4 of the article. We adapted a task in which Bonini et al. (1999) demonstrated vague assertion gaps to probe for endowment-like effects both with and without endowments. In one of their studies, Bonini et al. asked some of their Italian participants the following question:

> When is a man tall? Of course, very big men are tall and very small men are not tall. We're interested in your view of the matter. Please indicate the smallest height that in your opinion makes a man tall.
> A man is tall if his height is greater than or equal to ____ centimetres.

Other participants were instead asked, "When is a man not tall?", responding that "A man is not tall if his height is less than or equal to ____ centimetres." The smallest "tall" man (M = 181.5 cm) was considerably taller than the tallest "not-tall" man (M = 160.5 cm), a gap which suggests avoidance of assertion when the application of a vague term is indeterminate. We used a similar procedure to obtain monetary cutoff points for buying and selling a mug, for choosing a mug over money, and for choosing the money over the mug.

Method

In a paper-and-pencil task, 470 undergraduates at the Rady School of Management (250 males, 220 females, M_{age} = 21.4) evaluated a pictured UCSD mug. Materials and data for this study are available at https://osf.io/zbmvk/?view_only=5a4b74d4de7146cc8d0582fdf34910cf. Two of four between-subjects conditions provided a test of the standard endowment effect. In a "seller" ["buyer"] condition, participants were told to imagine that they owned [could buy] the mug, and then stated the smallest [largest] amount for which they would sell [buy] the mug. The tasks were introduced with wording modeled after Bonini et al. (1999), introducing the idea of a cutoff between selling and non-selling [buying and non-buying] amounts. For example, sellers read the following text:

> Whether or not you would sell the UCSD mug probably depends on the amount of money you are offered. If the amount of money is very very large, you'd probably sell the mug. If the amount is very very small, you probably wouldn't sell the mug.
>
> We would like to know your cutoff for selling the mug. That is, <u>what is the smallest amount of money you would be willing to accept for the mug</u>? Think about whether you would sell the mug for different amounts of money, and fill in the blank to complete the sentence below.
>
> For any amount of money above $__________, I would sell the UCSD mug.

Buyers read corresponding text explaining the cutoff for buying prices, and then completed the following sentence: "For any amount of money below $__________, I would buy the UCSD mug."

In two additional "chooser" conditions, were asked to imagine that they had to choose between the UCSD mug and some amount of money. Those in the "choose-money" condition were asked to state the smallest amount for which they would choose the money over the mug, while those in the "choose-mug" condition stated the largest amount for which they would choose the mug over the money. For example, "choose-money" participants read the following text:

> Whether you would take the UCSD mug or the money probably depends on the amount of money you are offered. If the amount of money is very very large, you'd probably take the money. If the amount is very very small, you'd probably take the mug.

> We would like to know your cutoff for taking the money. That is, <u>what is the smallest amount of money for which you would still take the money over the mug</u>? Think about what you would choose for different amounts of money, and fill in the blank to complete the sentence below.
>
> For any amount of money above $__________, I would take the money (rather than the UCSD mug).

"Choose-mug" participants read corresponding text for the choice of the mug over the money, and then completed the following sentence: "For any amount of money below $__________, I would take the UCSD mug (rather than the money)."

Participants then answered demographic questions on the next page. Because we had recently run a similar study at a different subject pool at the same university, in the Psychology Department, we also asked participants if they had taken part in previous experiments both at the Rady School of Management and in the UCSD Psychology Department. (We asked about both pools to make sure participants would distinguish between them in their responses.) For the main analysis reported here, participants are included regardless of their response to these items. But to ensure that potential cross-contamination couldn't have contributed to the effect, we also analyzed the data excluding participants who had taken part in a Psychology Department study (a conservative approach, since they could have taken part in any study run there). Excluding these participants had no effect on the direction or significance of any effects reported below. This additional analysis is available on the article's OSF site.

Results and Discussion

We restricted the analysis to precise monetary cutoffs below $50, resulting in a final sample of 464 (1.28% of responses excluded). Mean monetary cutoffs were analyzed in a 2 x 2 ANOVA with task (trade vs. choice) and focal act (take money vs. take mug) as independent variables. Thus selling is a trade in which the focal act is taking the money (in exchange for the mug), while buying is a trade in which the focal act is taking the mug (in exchange for money). Table S1 reproduces the results from Table 4 in the article, broken down according these two independent variables.

Table S1
Monetary Cutoffs by Task and Focal Act

Task	Focal Act	Question	*N*	*M* (*SD*)	*p*
Trade	Take $	Sell above...	115	$11.23 (7.62)	< .001
	Take mug	Buy below...	118	$7.02 (4.89)	
Choice	Take $	Take $ above...	116	$9.50 (6.94)	.01
	Take mug	Take mug below...	115	$7.36 (5.55)	

There was a significant main effect of focal act, $F(1, 460) = 29.14$, $p < .001$, no main effect of task, $F(1, 460) = 1.40$, $p = .237$, and the interaction fell short of significance, $F(1, 460) = 3.11$, $p = .079$.[1] Of greatest interest are the simple effects of focal act in the two tasks. Both were significant, with higher monetary cutoffs when the focal act was taking the money, both in the trade conditions, $t(193.5) = 5.01$, $p < .001$, and in the choice conditions, $t(219.2) = 2.59$, $p = .01$. As discussed in the article, the gaps in both conditions are predicted by conservatism. A reluctance to make an arbitrary assertion ("I would choose $*x* over a mug," in the range of uncertain valuation), like a reluctance to engage in arbitrary transactions (selling the mug for $*x* in the same range), appears to drive up the stated cutoff for taking the money.

[1] A Brown-Forsythe test revealed significant heterogeneity of variance ($p < .01$), driven by inflated variance in the seller condition. Because the design is balanced with a large sample size, this is unlikely to strongly affect the analysis, and we report separate t-tests in the analysis of simple effects.

2. Computer Simulations

Computer simulations are useful in evaluating implementation policies when specific assumptions are made about (i) the DM's underlying preference structure, (ii) the structure of the environment, and (iii) computational constraints to which the DM may be subject. Because there is no simple utility metric for outcomes when preferences are incomplete, these assessments center around pairwise comparisons: We can ask, for two DMs with the same preference structure under matched conditions, how often it happens that one DM's outcome is definitely ranked above the other's outcome according to their shared preference structure. Whenever a DM's outcome is dominated, we can conclude that they would have been better off in this scenario, according to their own preferences, had they adopted the alternative implementation policy.[2]

We report one set of illustrative simulations. We consider a space of goods defined on two attributes, where the DM's incomplete preferences are defined by dominance, as in Figure 1B in the article. That is, each good is represented by an ordered pair of real numbers between 0 and 1, with $(a_1, a_2) \succcurlyeq (b_1, b_2)$ if and only if $a_1 \geq b_1$ and $a_2 \geq b_2$ (and the preference is strict in case at least one of the a_i is not equal to the corresponding b_i). In the Monte Carlo simulations reported here, each DM holds an ordered bundle of two such goods, allowing us to examine the subtle incoherence traps noted by Mandler (2005) and described in footnote 7 of the article. On each round, the DM is offered the chance to trade one of the goods in their bundle (with a randomly selected index) for an alternative good randomly drawn from the two-attribute space. Each simulation run features a pair of yoked DMs who receive matched trading offers (i.e.,

[2] While these dominance rates clarify the (dis)advantages of different implementation policies, their comparative nature requires that they be interpreted with caution. For example, it is possible for one implementation policy to often dominate a second, for the second to often dominate a third, but for the first to never dominate the third.

identical random sequences of bundle-index selections and alternative goods) but use different implementation policies when their decision is not dictated by the common preference structure. We then ask whether the resulting bundle of DM1 dominates, is dominated by, or is not comparable to DM2's bundle (i.e., whether it is possible to line up the goods in the two bundles so that both goods in DM1's bundle are weakly preferred to the corresponding good in DM2's bundle, with at least one good strongly preferred).

The simulations compare three types of implementation policy – a Random policy, in which choices not dictated by the preference order are always made at random; a Status Quo Maintenance (SQM) policy in which trades not required by the preference order are always declined; and a Choices-as-Precedents (CP) policy, in which the DM trades when required to do so either by (i) the preference structure, or (ii) precedents established by past choices as well as their transitive implications. The full specification of a CP policy requires, in addition, a rule for arbitrary choices not determined by preferences and precedents. Here, we assume that the DM chooses at random in such cases. Alternative variants of the CP policy are considered below.

Unlike the Random and SQM policies, the CP policy requires memory for and inference from past precedents. This raises questions about CP's robustness both to computational costs and to occasional errors in memory and/or inference – questions that do not arise for Random and SQM policies. In the simulations reported here, we examine a simple constraint on memory, in which each precedent has some probability of being permanently deleted from memory in each round. Figures S1-S3 depict results from simulations in which this forgetting rate was set to 0% (no memory error), 1%, and 10%, respectively.[3]

[3] High-resolution versions of Figures S1-S3 are available at the article's OSF site.

Each panel shows results from 1 million runs of a simulation in which, across 50 trading periods, two DMs using different implementation policies were exposed to matched trading opportunities. The curves indicate, for each period, how often one DM's bundle dominated the other's at the end of that period. Figure S1 shows the pattern of dominance relations in the absence of memory error, which agree with the theoretical analysis in the article: The left panel shows that CP often dominates Random, but is never dominated by Random. The middle panel shows that SQM often dominates Random, and is dominated by Random only very rarely (0.24% of the time), presumably corresponding to runs where the Random policy's choices happened to coincide with those that would result from CP. Finally, the right panel shows that dominance is very rare in comparisons between CP and SQM: SQM never dominates CP, while CP dominates SQM 0.34% of the time. The rare instances of dominance presumably arise when CP exploits but SQM misses the subtle opportunities for upward trade discussed by Mandler (2005).

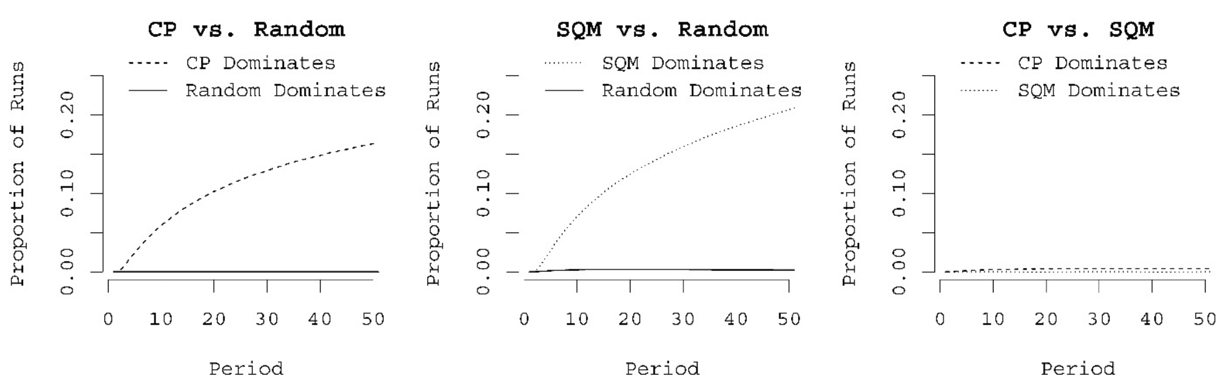

Figure S1. Dominance rates in pairwise comparisons between Random, CP, and SQM policies when precedents are never forgotten.

As Figures S2-S3 show, the picture changes when CP is subject to occasional forgetting of precedents. When the forgetting rate is 1% (Figure S2), SQM starts to sometimes dominate CP (right panel), and while CP retains a strong advantage over Random, there are very rare instances where Random dominates CP (i.e., the solid line in the left panel is not perfectly flat). When the forgetting rate is raised to 10% (Figure S3), SQM's advantage over CP dramatically grows (right panel), while CP's advantage over Random noticeably shrinks (left panel).

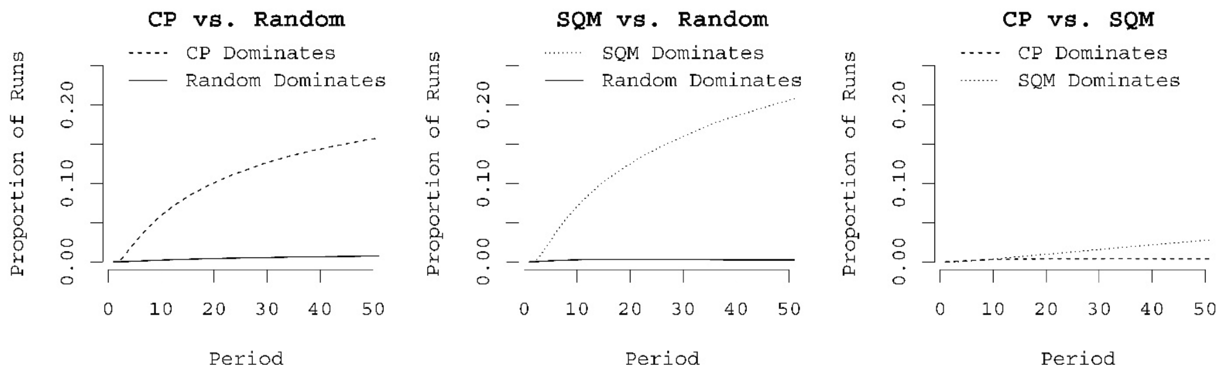

Figure S2. Dominance rates in pairwise comparisons between Random, CP, and SQM policies when each precedent has a 1% probability of deletion on each round.

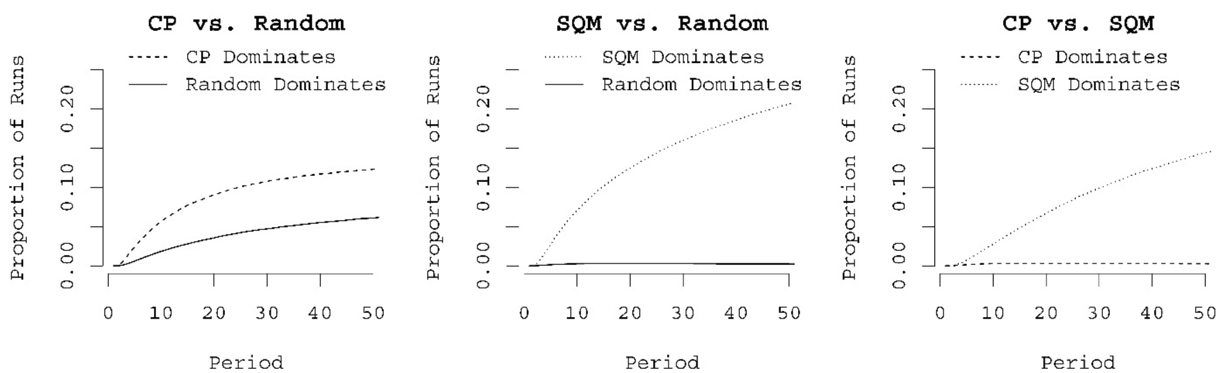

Figure S3. Dominance rates in pairwise comparisons between Random, CP, and SQM policies when each precedent has a 10% probability of deletion on each round.

These simulations illustrate how the relative performance of different implementation policies can depend on performance limitations. When a simple type of memory error is added to the simulation, the coherently arbitrary policy we examine here – which initially makes arbitrary decisions at random, and then follows remembered precedents – sees its subtle advantage over SQM reverse and its pronounced advantage over Random weaken.

The R code for these simulations is available on the article's OSF site. Parameters in the code can be adjusted to simulate other implementation policies, trading conditions, and rates of precedent deletion. The five parameters are: the number of simulated time periods (tHorizon), the number of items per bundle (nItems), the probability with which each precedent is forgotten in each time period (memoryLapse), whether the DMs are yoked in the sense that the new items they are offered in each period are the same (bYoked), and the type of choices-as-precendents policy (cpType: 1 for random, 2 for conservative, 3 for restless).

The last parameter requires some explanation. In the simulations shown here, CP's arbitrary initial choices (i.e., those not constrained by preferences or past precedents) are random (cpType = 1). However, a CP policy may combine any rule for arbitrary initial choice with subsequent obedience to precedents. For example, the rule may be conservative (cpType = 2), where the DM declines arbitrary offers unless past precedents require their acceptance, or restless (cpType = 3), where the DM accepts such offers. Intuitively, one would expect a CP policy of Type 2 to be more robust, and a CP policy of Type 3 to be less robust, to memory errors. Other simulations, not shown here, confirm these intuitions. It may be possible to join the strengths of conservatism and coherent arbitrariness in a single implementation policy.

Future simulations could explore other computational limits, along with other variants of a coherently arbitrary policy. For example, DMs might be saddled with costs corresponding to computing time, or subjected to time pressure, with a decision forced after some number of computational steps. One could also simulate variants of the CP policy in which precedents are simplified to ease the memory demands and computational complexity of precedent-based inference.[4] Explorations of this kind could clarify the bounded rationality of implementation when computation is costly and/or error-prone.

[4] For example, precedents could take the form of lower and upper bounds for the relative weight of Attribute 2, with the constraint that the DM's choices always be consistent with the use of a weighted-additive choice rule for *some* weight in that range. As new arbitrary choices are made, the range would progressively narrow. This policy only requires two numbers to be stored in memory, and implications of precedents are relatively easy to compute.