

When Fair Isn't Fair: Understanding Choice Reversals Involving Social Preferences

James Andreoni

University of California, San Diego

Deniz Aydin

Washington University in St. Louis

Blake Barton

Stanford University

B. Douglas Bernheim

Stanford University

Jeffrey Naecker

Wesleyan University

In settings with uncertainty, tension exists between ex ante and ex post notions of fairness. Subjects in an experiment most commonly select the ex ante fair alternative ex ante and switch to the ex post fair alternative ex post. One potential explanation embraces consequentialism and construes reversals as time inconsistent. Another abandons consequentialism in

We thank participants at the 2015 Stanford Institute for Theoretical Economics Psychology and Economics Workshop, the 2016 American Economic Association meetings, the 2016 New England Experimental Economics Workshop, the 2016 Early Career Behavioral Economics Conference, the 2017 Economic Science Association meetings, the 10th Maastricht Behavioral and Experimental Economics Symposium, and seminars at Columbia, Texas

Electronically published March 23, 2020

[*Journal of Political Economy*, 2020, vol. 128, no. 5]

© 2020 by The University of Chicago. All rights reserved. 0022-3808/2020/12805-0002\$10.00

favor of deontological (rule-based) ethics and thereby avoids the implication that revisions imply inconsistency. We test these explanations by examining contingent planning and the demand for commitment. Our findings suggest that the most common attitude toward fairness involves a time-consistent preference for applying a naive deontological heuristic.

I. Introduction

Suppose that 20 lottery tickets will be divided between two equally deserving households, A and B. Ten of the tickets are red and 10 are blue. One of the 20 will be chosen at random, and the household holding it will win a cash prize. Household A already holds all the red tickets but has done nothing to earn them. Your task is to allocate the blue tickets. How would you divide them up? Most people express a strict preference for giving all 10 blue tickets to household B to even out the chances of winning, presumably in the interest of fairness.

Now suppose that, after assigning all 10 blue tickets to B, you learn that the winning ticket is blue. You are then given a chance to reallocate the blue tickets. What would you do? As we show, most people express a strict preference for reallocating the blue tickets equally between A and B, again to even out the chances of winning.

In our experience, the choice pattern just described strikes most people as eminently reasonable, at least initially. However, if one adopts a consequentialist perspective on decision-making as is standard throughout economics, these choices violate the principle of time consistency. From this perspective, the objective of initially allocating all blue tickets to B is to ensure that each household has a 50% chance of winning the prize before the resolution of pertinent uncertainty (*ex ante*). However, in light of the subsequent revision, A's *ex ante* chances of winning are 75%: there is a 50% chance that A wins because the winning ticket is red and a 25% chance that A wins because the decision maker reallocates tickets after learning that the winning ticket is blue. Thus, the revision is inconsistent with a consequentialist interpretation of the *ex ante* objective.

The current paper has two main objectives. First, we document the types of choice reversals described above in a laboratory experiment. Indeed, we show that the single most common behavioral pattern is for subjects to select the *ex ante* fair alternative *ex ante* (i.e., before learning the color of

A&M, Cornell, the University of Southern California, Emory, Claremont Graduate University, the University of Connecticut, Yale, Bocconi, Wissenschaftszentrum Berlin für Sozialforschung, the University of Lyon, the University of Zurich, Chapman University, the University of San Diego, the RAND Institute, Santa Clara University, and California Institute of Technology for helpful comments. Data are provided as supplementary material online.

the winning ticket) and switch to the ex post fair alternative ex post (i.e., after learning the color of the winning ticket). This pattern does not diminish with experience, and the preferences of most subjects are strict. Second, we attempt to distinguish between two classes of potential explanations. One embraces consequentialism and construes the reversals as manifestations of time inconsistency. This inconsistency arises naturally from the tension between ex ante and ex post perspectives on fairness: if the ex ante perspective is compelling ex ante and the ex post perspective is compelling ex post, then a decision maker may shift from the first to the second as events evolve. The other class of explanations rejects consequentialism along with the notion that revisions imply inconsistency. Instead, it depicts decision makers as having time-consistent preferences for applying simple (and arguably naive) deontological principles (rule-based ethics), even though a consequentialist would deem their implications time inconsistent. A rule that prescribes egalitarian treatment of indistinguishable parties will likewise produce ex ante fair allocations for actions executed ex ante and ex post fair allocations for actions executed ex post.

We employ two strategies to distinguish between these explanations. First, we assess the demand for commitment among decision makers who have observed their proclivity to switch. In the preceding example, if the decision maker is a time-consistent consequentialist, she will recognize that her ex post choices will shift A's ex ante odds of winning to 75% and will therefore seek to remove opportunities for revision. In contrast, if she has a time-consistent preference for applying a simple deontological rule requiring egalitarianism, she will exhibit an aversion to any commitment that would preclude her from responding ethically. For example, she will preserve the flexibility to reallocate the blue tickets equally upon learning that the winning ticket is blue. Our second strategy is to examine contingent planning. Instead of allowing decision makers to revise their choices ex post, we require them to specify contingent plans for their revisions ex ante. Under the hypothesis of time-inconsistent consequentialism, the decision maker will choose a plan that is ex ante fair accounting for the revision. In contrast, under our hypothesis concerning naive deontological ethics, she will specify that "live tickets" should be reallocated equally ex post.

Our data on the demand for commitment require careful interpretation. Roughly 40% of our subjects strictly prefer commitment to flexibility, while roughly 30% prefer flexibility to commitment. Taken at face value, this finding suggests that time-inconsistent consequentialists are a bit more numerous than those who practice simple deontological ethics. However, the observed preference for commitment likely overstates the prevalence of time-inconsistent consequentialists. Subjects who are prone to exhibit the characteristic choice pattern (switching from ex ante to ex post equalizing allocations) also disproportionately manifest a preference for retaining

flexibility over making commitments. Conversely, a preference for commitment is most prevalent among those who are least likely to switch, which suggests that many of those who choose commitment do so to avoid the annoyance of having to submit their preferences twice, rather than to preempt revisions. Our analysis of contingent planning corroborates these inferences: many subjects choose an initial allocation that is *ex ante* fair but instruct us to reallocate their tickets evenly if it turns out that the winning ticket is one of theirs. These results support our main conclusion: the most common attitude toward fairness involves a time-consistent preference for applying naive deontological principles, even though a consequentialist would deem their implications time inconsistent.

Our findings have important practical implications. Even when people agree about the importance of fairness, they may disagree as to what constitutes a fair decision. An important dimension of disagreement concerns the question of whether a fair society should pursue equality of opportunity or equality of outcomes. Those who favor standards based on equality of original opportunities tend to view fairness from an *ex ante* perspective. They tolerate even highly unequal outcomes provided all parties had comparable shots at success. In contrast, those who favor standards based on equality of updated opportunities or outcomes tend to think about fairness from an *ex post* perspective. Differences of opinion concerning the relative importance of these principles can produce conflict over policy issues. However, our analysis suggests that those differences may not be stable. In particular, we have shown that as information is revealed, people readily shift from the *ex ante* to the *ex post* perspective. A society populated by such individuals would design policies *ex ante* to promote equality of opportunity, only to undermine the objectives of those policies by consensus (at least from the perspective of a consequentialist planner) once winners and losers emerge.¹ Examples of potential applications include rules governing the treatment of preexisting conditions by health insurers, social insurance, policies impacting access to education, and the assignment of priority for organ recipients; see the concluding section VII for further discussion.

Our paper contributes to a large and growing body of empirical research on attitudes toward fairness. The importance of fairness as a behavioral motivation is by now well established.² Several previous experimental studies have examined whether people care about *ex ante* fairness, *ex post* fairness, or both; see Bolton, Brandts, and Ockenfels (2005); Karni, Salmon, and Sopher (2008); Krawczyk and Le Lec (2010); Brock, Lange, and Ozbay

¹ Coate (1995) makes a similar point in a setting where the inconsistency arises from a different source (the Samaritan's dilemma).

² Classic experimental results include the tendency to divide a prize equally in the dictator game and reject lopsided offers in the ultimatum game. See, e.g., Forsythe et al. (1994); Hoffman, McCabe, and Smith (1996); Camerer (1997); Bohnet and Frey (1999); Andreoni,

(2013); Kircher, Ludwig, and Sandroni (2013); and Trautmann and van de Kuilen (2016). Our contribution involves exploring the existence and causes of choice reversals arising from the tension between ex ante and ex post fairness.³

The rest of this paper is organized as follows. We provide conceptual background in section II. We then describe the basic framework for our experiment in section III. Section IV studies the effect of ex ante versus ex post framing of allocation problems and investigates whether concerns for fairness generate choice reversals (revisions). Sections V and VI test competing explanations by examining the demand for commitment and contingent planning. Section VII concludes.

II. Conceptual Issues

Distinguishing between consequentialist and nonconsequentialist explanations for the apparent choice reversals described at the outset of section I requires a clear understanding of the pertinent theories. This section explains these theories and describes their testable implications.

A. *Consequentialism and Concern for Fairness as a Source of Time Inconsistency*

According to the doctrine of consequentialism, “whether an act is morally right depends only on consequences.”⁴ Moral philosophers have devised many variants of consequentialism, and there is disagreement as to what the doctrine entails. However, for the allocation problem described at the outset of section I, a “plain vanilla” interpretation of consequentialism

Brown, and Vesterlund (2002); Andreoni and Miller (2002); Andreoni, Castillo, and Petrie (2003); and Andreoni and Bernheim (2009). Early attempts to model concerns about fairness include Rabin (1993), Fehr and Schmidt (1999), Bolton and Ockenfels (2000), and Charness and Rabin (2002). Related behavioral patterns are commonly observed in the field. For instance, equal sharing is common in the context of joint ventures among corporations (Veugelers and Kesteloot 1996; Dasgupta and Tao 1998; Hauswald and Hege 2003), share tenancy in agriculture (Agrawal 2002; DeWeaver and Roumasset 2002), bequests to children (Menchik 1980, 1988; Wilhelm 1996; Bernheim and Severinov 2003), and arbitration (Bloom 1986).

³ To our knowledge, only one previous study (Trautmann and van de Kuilen 2016) offers evidence potentially related to the issue of fairness and time inconsistency. It examines a two-player allocation game in which nature randomly tilts the outcomes in favor of one player or the other and shows that players revise their ex ante choices after the resolution of uncertainty with modest frequency. While it is obviously related to our work, the experimental design implicates considerations other than fairness, in that each player has a selfish interest in the outcome. Ex post and ex ante behavior may differ for three confounding reasons: opportunities for reciprocity exist only ex ante, subjects may succumb to self-serving narratives ex post, and subjects may revise their beliefs about other players' choices. In addition, the study investigates neither contingent planning nor the demand for commitment vs. flexibility and consequently sheds no light on potential explanations for revisions.

⁴ See Sinnott-Armstrong (2015).

would imply that one should judge an allocation of lottery tickets based solely on its implications concerning the probabilities with which each household would win the prize.

Within a consequentialist framework, concerns for fair divisions of probabilistic claims on a prize go hand in hand with time inconsistency. The objective of this subsection is to explain this connection.

To build intuition, we begin with a simple model. A natural hypothesis concerning fairness is that, at any given point in time, the decision maker is concerned about the distribution of expected utility (EU). Focusing on the allocation task described in section I, we can write household i 's EU as $EU_i = p_i U_i^W + (1 - p_i) U_i^L$, where p_i is the probability that i wins the prize, U_i^W is i 's utility if i is the winner, and U_i^L is i 's utility if i is the loser. Allocating lottery tickets amounts to selecting the probabilities $p_A, p_B \in [0, 1]$ such that $p_B = 1 - p_A$. Assume that the decision maker's preferences are governed by a strictly quasi-concave objective function of the form $W(EU_A, EU_B)$.⁵ If the problem is sufficiently symmetric so that $W(U_A^W, U_B^L)$ is close to $W(U_A^L, U_B^W)$, the optimal choice—call it p_A^* —is interior (and the optimal choice of p_B is $p_B^* = 1 - p_A^*$). Indeed, with perfect symmetry, the decision maker's ideal choice is to set $p_A^* = p_B^* = .5$.

What happens when the decision maker is allowed to reallocate tickets after learning that some are definitely losers, so that the ex post probabilities of winning the prize (conditional on the initial ticket allocation) depart from p_A^* and p_B^* ? Reoptimizing W over probabilities yields the same solution as before. Consequently, the decision maker revises her initial allocation to achieve a division of the remaining "live" tickets that reinstates the probabilities p_A^* and p_B^* .

The implied ex post revision is time inconsistent; it reflects a failure to follow through on a contingent plan that already specifies a desired outcome for every possible state of nature.⁶ As a result, it induces ex ante odds that the decision maker finds unattractive from the ex ante perspective. If she is sophisticated (in the sense that she anticipates her choice reversal), she will seek to remove opportunities for revision by undertaking commitments before the resolution of uncertainty. Similarly, regardless of whether she is sophisticated or naive, if she were asked to specify a contingent plan for her revision before learning anything about the realization, she would simply reaffirm her preference for her initial allocation.

One should not infer from the preceding example that fair consequentialists are inevitably time inconsistent. To illustrate, suppose that the

⁵ Note that any departure from linearity renders W nonlinear in probabilities and hence inconsistent with the independence axiom.

⁶ The phenomenon of time inconsistency is commonly associated with the notion of present focus and in particular with quasi-hyperbolic discounting (the β - δ model). It is important to bear in mind that present focus is merely an example of time inconsistency. The type of time inconsistency studied in this paper does not involve present focus.

decision maker maximizes $E(W(U_A, U_B))$ instead of $W(EU_A, EU_B)$. Notice that we can rewrite this objective function as

$$p_A W(U_A^W, U_B^L) + (1 - p_A) W(U_A^L, U_B^W).$$

Accordingly, the decision maker allocates all tickets to household A when $W(U_A^W, U_B^L)$ exceeds $W(U_A^L, U_B^W)$ and allocates all tickets to household B when this inequality is reversed. While she could also choose an interior allocation in the knife-edge case where $W(U_A^W, U_B^L) = W(U_A^L, U_B^W)$, her preference would not be strict—indeed, she would be indifferent among all possible allocations. The same decision rule is optimal regardless of how tickets outside the decision maker's control are distributed and applies with the same force both ex ante and ex post. Therefore, the decision maker is always content to stick with her preferred ex ante allocation upon reaching the ex post position.⁷

In our first example, the decision maker has a strict preference for interior probabilities and is also time inconsistent. In our second example, the decision maker has a weak preference for boundary allocations and is time consistent. These examples are representative, in that strict preferences for interior values of p_A^* and p_B^* always imply time inconsistency. To understand why this is the case, notice that $p_A^* \in (0, 1)$ cannot be a strict optimum unless utility is nonlinear in probabilities, which means that preferences violate the independence axiom.⁸ In other words, someone who allocates probabilities in this way cannot be an EU maximizer.

The next step is to recall that within a consequentialist framework, EU preferences are time consistent, while non-EU preferences are not.⁹ To be sure, a time-consistent individual may wish to revise choices when new information becomes available. However, she will never do so if the original choice specifies contingent actions tailored to each possible realization of that information.¹⁰ The connection between time consistency and the independence axiom is intuitive: in effect, time consistency requires

⁷ When the decision maker sees the two households as equally meritorious ($W(U_A^W, U_B^L) = W(U_A^L, U_B^W)$), she is indifferent about the division of lottery tickets both ex ante and ex post. Consequently, she is also indifferent about making revisions and commitments. Under this hypothesis, behavioral patterns would likely be haphazard, but a fortuitous resolution of indifference could nevertheless produce almost any choice pattern. To falsify this hypothesis, one must therefore demonstrate that preferences for initial allocations, revisions, and/or commitments are strict. We examine the strictness of preferences in secs. IV.E and V.D.

⁸ Classical discussions of the inconsistency between a preference for ex ante fairness and the independence axiom include Harsanyi (1955) and Diamond (1967). For more recent perspectives, see Fudenberg and Levine (2012) and Saito (2013).

⁹ Classic references include Markowitz (1968) and Raiffa (1968). See also the excellent discussion in Machina (1989), who emphasizes the role of consequentialism.

¹⁰ Consider, e.g., the decision problem described in the introduction. There are 20 states of nature, each corresponding to the selection of a particular lottery ticket. Any allocation

that the preferences governing choices at a given node in a decision tree are independent of the probability with which the node is reached, as well as the consequences of following any other positive-probability path. Putting these two classical observations together, one naturally arrives at the conclusion that strict consequentialist preferences for interior probabilities generally imply time inconsistency.¹¹

B. Nonconsequentialism and Time Consistency

One should not conclude from section II.A that strict preferences for interior probabilities necessarily imply time inconsistency in all instances. Stepping outside the consequentialist framework, other possibilities arise.

1. Naive Deontological Ethics

In moral philosophy, an important class of alternatives to consequentialism involve deontological ethics. According to the deontological perspective, whether an action is morally right or wrong depends on its conformance to specific rules and notions of duty. Kant's categorical imperative is perhaps the best-known theory of this type. Irrespective of their normative validity, deontological perspectives such as Kant's are likely too complicated to provide good positive theories of ethical reasoning. Even so, people routinely appear to employ simple deontological rules, such as "do not lie" and "do not steal."

When performing allocation tasks such as those described in section I, people may similarly apply simple deontological principles. While we will not be able to pin down a specific rule, we take the following to be a normatively appealing exemplar from a larger class:

When differences in claims on resources are the result of chance or other past events outside the control of the affected individuals, the ethical course of action is to impose equality.

The mechanical application of this deontological rule replicates some of the main empirical implications of fair consequentialism (under symmetry). In the ex ante setting, the decision maker allocates the blue tickets

of the tickets between the households is a complete state-contingent plan specifying an assignment of the prize for every state of nature. Thus, a time-consistent decision maker would not want to change the allocation upon learning that certain states did not materialize. As we have already emphasized, any such revision alters the ex ante probability of winning.

¹¹ A version of this point appears in Machina (1989). See also Trautmann and Wakker (2010).

to achieve an equal split overall, because the division of the red tickets was outside the control of both potential recipients. Yet in the *ex post* setting, she divides the blue tickets equally, because the original ticket assignment and the selection of a blue ticket were both outside the recipients' control. Thus, her decision criterion generates an apparent choice reversal.

Significantly, this theory departs from fair consequentialism with respect to empirical implications concerning the demand for commitment and for contingent planning. It implies an aversion to any commitment that would preclude the decision maker from responding ethically—for example, by reallocating the blue tickets to achieve equal division in accordance with her deontological rule upon discovering that the winning ticket is blue. In such instances, the decision maker would have a strict preference for flexibility. It can also rationalize instances in which the decision maker instructs an agent *ex ante* to equalize holdings of the blue tickets if it turns out that the winner is blue, because this instruction ensures an ethical action *ex post*.

In short, the naive deontological heuristic discussed above leads the decision maker to agree with the fair consequentialist that the *ex ante* perspective is compelling *ex ante*, and that the *ex post* perspective is compelling *ex post*, but in contrast to consequentialism does not require her to acknowledge any inconsistency between these judgments. On the contrary, the apparent choice reversal reflects a time-consistent ethical preference for conforming to a fixed deontological rule.

An individual who adopts the simple deontological ethic outlined in the preceding paragraphs and applies it in the manner we have hypothesized may be naive in either or both of the following senses. First, she may be philosophically naive, in that her application of the rule may involve arbitrary narrow framing for which there is no ready normative justification. In particular, when making her *ex ante* choice, her evaluation frame encompasses both the blue tickets and the red tickets. However, when engaging in contingent planning *ex ante*, her evaluation frame encompasses only the blue tickets, because she interprets the problem from the perspective of the stated *ex post* contingency (instead of reformulating it from the perspective of her *ex ante* position before applying her rule). Second, she may be cognitively naive, in that her lack of concern for the consequential inconsistency of her *ex ante* and *ex post* choices may result from a logical lapse rather than a deliberate judgment. To be clear, the simple deontological ethic does not require cognitive naiveté: the decision maker might acknowledge the consequential equivalence of allocating all 10 lottery tickets to household B when A also holds 10 possible winners, and doing so after learning that A's tickets are no longer "live," but nevertheless deny the ethical equivalence of these alternatives. In this paper, we do not attempt to distinguish between competing theories such as these that might account for the adoption of a naive deontological

ethic (although we do offer some suggestive observations in the concluding section). Rather, we simply ask whether people behave in a way that is consistent with the hypothesized ethic.

2. Resolute Nonconsequentialism

Another hypothesis worth investigating is that a fair decision maker violates the precepts of consequentialism because she cares about process. In particular, she may behave consistently as time passes because she takes past uncertainty (risks already borne) into account at each moment in a manner consistent with her earliest choices. People who behave in this manner are called “resolute.”¹² A decision maker with resolute preferences for ex ante fairness engages in the same reasoning as a consequentialist ex ante but continues to evaluate actions from this ex ante perspective as events unfold. A decision maker with resolute preferences for ex post fairness mimics the ex post reasoning of a consequentialist, regardless of whether she contemplates the problem from an ex ante or an ex post perspective.

Machina (1989) offers the following appealing illustration of resolute non-EU preferences. Mom has two children, Ben and Abby, as well as a single treat. She cares about outcome fairness and would ideally split the treat between them, but regrettably it is indivisible, so she must give it to one or the other. Imagine that she strictly prefers a coin flip over either sure outcome. Mom flips the coin, and Abby wins. After pouting briefly, Ben has a sudden inspiration: he points out to Mom that, in light of her stated and revealed preferences, she would be better off flipping the coin again. Mom’s response: “sorry kid, you had your chance.” In this example, Mom strictly prefers egalitarian allocations of chances to win a prize, but her preferences are resolute, so she is time consistent.

The empirical implications of resoluteness are straightforward. A resolute decision maker with strict fairness preferences is time consistent and exhibits no choice reversals. She either selects the ex ante fair allocation both ex ante and ex post or selects the ex post fair allocation both ex ante and ex post. She is indifferent between committing to an ex ante choice and retaining flexibility, and she specifies no changes in ticket allocations when making contingent plans.

III. Experimental Framework

Our study consists of a collection of related experiments. In this section, we summarize shared aspects of the experimental design, data analysis,

¹² The phrase “resolute preferences” appears to originate with McClennen (1989), but there are conceptual antecedents. See the discussion in Machina (1989).

and implementation. In subsequent sections, we provide additional detail concerning the individual experiments and summarize our results.

A. Basic Allocation Tasks

Each decision task involves the allocation of 20 lottery tickets between two impoverished Kenyan families (A and B). The division of 10 tickets (numbered 11–20) is fixed in advance and varies from task to task (the “computer’s” ticket allocation). The subject allocates the remaining 10 tickets (numbered 1–10). We then select one ticket at random, and the family holding that ticket receives a \$10 donation.¹³

We implement donations with the cooperation of a well-established charity, GiveDirectly.¹⁴ The organization operates a platform for donating money directly to needy households in poor African nations. We selected the households viewed by our subjects from lists of GiveDirectly’s potential recipients.

We examine multiple variants of the allocation task, which differ with respect to the subject’s knowledge and the timing of her decision. In all cases, the subject learns the computer’s allocation before assigning her own tickets. The main variants are as follows:¹⁵

Ex ante decisions.—The subject makes her decision immediately after learning the computer’s allocation, without receiving any other information. We display her ticket allocation on the screen and ask her to review it; if desired, she can submit an updated allocation. She repeats this step until she confirms her choice. We then select a ticket at random, which determines the winner.

Ex post decisions.—We tell the subject that we have selected the winning lottery ticket at random. She also learns whether it is one of the computer’s tickets or one of hers. In the latter case, she then allocates her own tickets without knowing which is the winner. We display her ticket allocation on the screen and ask her to confirm or revise it. She repeats this step until she confirms her choice. We then reveal the winning ticket.

Ex ante decisions with surprise ex post revisions.—After making one or more decisions in the ex ante frame, we return to these decisions and one at a time reveal to the decision maker whether the winning ticket is one of the

¹³ The randomness of the outcome likely heightens fairness considerations. Cappelen et al. (2013) show that people are particularly disinclined to accept ex post differences that result from luck rather than choice.

¹⁴ See <http://www.givedirectly.org>. GiveDirectly is recognized as one of the most efficient charities serving this sector. It was cofounded by a University of California, San Diego (UCSD) faculty member, a fact that may have enhanced its credibility with our UCSD undergraduate subjects.

¹⁵ We explore additional variants in later sections.

computer's tickets or one of hers. In the latter case, she does not learn the number of the winning ticket. We then display her ticket allocation on the screen again and, as in an ex post decision, ask her to confirm or revise it. (We do not advise her in advance that she will have another opportunity to revise her choices after learning whether the winning ticket is one of hers.) She repeats this step until she confirms her choice. The randomly selected ticket then determines the winner.

We structure the presentation of each task to ensure that subjects view the two Kenyan households as equally deserving. At the outset of each task, subjects view photos of 16 potential recipients, including their households A and B. We obtained the photographs from GiveDirectly, and they are of the actual recipients. The composition of families within each group is uniform. In particular, the recipients were shown in one of the following groups: single younger women, single older women, couples with one child, or single men. To discourage subjects from searching for and inflating the significance of minor differences between families, we do not indicate which household within a group is A and which is B. Subjects are informed that the households and their respective roles are assigned before they allocate their tickets. After viewing the photos, subjects make their allocation choice by selecting whether each of their 10 tickets should go to household A or household B. Subjects see the computer's allocation of tickets on this same interface, and they are warned that they cannot change this allocation. The few subjects who attempted to do so received an error message and were returned to the allocation interface to try again. For more details on the instructions and interface, see the screenshots in appendix C.2 (apps. A–C are available online).

B. Categorization of Choices

To streamline our analysis of the data, we group allocations into five categories. Table 1 illustrates this categorization for the case in which the computer allocates eight tickets to household A.

Ex ante equalizing.—The subject allocates tickets so that each potential recipient ends up with the same number in total. For example, if the computer allocates eight tickets to recipient A and two to B, the subject allocates two tickets to A and eight to B.

Ex post equalizing.—The subject allocates five tickets to both potential recipients.

Overcompensating.—The subject allocates enough tickets to the potential recipient who received fewer from the computer to skew the overall distribution in that recipient's favor, overcompensating for the disparity. For example, if the computer allocates eight tickets to A and two to B, the subject allocates one ticket to A and nine to B.

TABLE 1
 ILLUSTRATION OF ALLOCATION CATEGORIES
 COMPUTER'S CHOICE = (8, 2)

	Overcompensating (1)	Ex Ante Equalizing (2)	Mixed (3)	Ex Post Equalizing (4)	Reinforcing (5)
Subject's choice	(0, 10)	(1, 9)	(2, 8)	(3, 7)	(4, 6)
				(5, 5)	(6, 4)
				(7, 3)	(8, 2)
				(9, 1)	(10, 0)

NOTE.—For this example, we assume that the computer assigned eight tickets to household A. We order and categorize choices according to the number of tickets that the subject allocates to household A. To equalize probabilities ex ante, the subject would allocate two tickets to household A and eight to B. To equalize probabilities ex post, the subject would allocate five tickets to each household. Other possible allocations fall into one of three ranges: overcompensating, mixed, and reinforcing.

Mixed.—The subject allocates more tickets to the potential recipient who received fewer from the computer but does not completely compensate for the disparity. For example, if the computer allocates eight tickets to A and two to B, the subject allocates four tickets to A and six to B.

Reinforcing.—The subject allocates more tickets to the potential recipient who receives more from the computer. For example, if the computer allocates eight tickets to A and two to B, the subject allocates six tickets to A and four to B.

C. *Details concerning Implementation*

We conducted the experiment at the UCSD Economics Laboratory within the guidelines of an Institutional Review Board–approved human subjects protocol. Subjects viewed these instructions on computer screens and followed along as the study leader read them aloud. Participants made all responses using a computer interface programmed with Qualtrics survey software. We separated subjects with partitions to ensure that they felt that their allocations were private. At the end of the experiment, subjects completed a short questionnaire in lieu of individual debriefing. A total of 702 subjects participated in the experiment across all treatments. Each subject received \$15 for participating. Typically, the experiment lasted 45 minutes.¹⁶

After completing all survey tasks, subjects filled out a short survey on demographics, including questions designed to elicit political inclinations. We did not find any robust relationships between behavior and political views, but it is worth noting that our sample includes relatively few subjects who self-identified as strongly conservative.

IV. Framing Effects and Choice Reversals

In this section, we demonstrate that subjects tend to choose ex ante equalizing allocations when initially confronting tasks with ex ante framing and tend to choose ex post equalizing allocations when initially confronting tasks with ex post framing. Furthermore, the initial framing does not lock them into a perspective on fairness either across or within tasks. As a result, in tasks with ex ante decisions and surprise revisions, the single most common behavioral pattern is for subjects to select the ex ante fair alternative ex ante and then switch to the ex post fair alternative ex post. This pattern does not diminish with experience, and the preferences of most subjects are strict.

¹⁶ For more details on treatment balance, as well as screenshots of all instructions and decision tasks, see app. C.

TABLE 2
MAIN TREATMENTS

Treatment	Rounds 1 and 2	Rounds 3 and 4	Rounds 5–8	Number of Subjects
4A_4A ^R	Ex ante	Ex ante	Ex ante with surprise revision	71
4P_4A ^R	Ex post	Ex post	Ex ante with surprise revision	72
2A2P_4A ^R	Ex ante	Ex post	Ex ante with surprise revision	48
2P2A_4A ^R	Ex post	Ex ante	Ex ante with surprise revision	48

A. *Experimental Design*

In our four main treatments, each subject performs eight allocation tasks. We divide these tasks into four sets of two, with sets separated by 1-minute breaks. Subjects understand that they will perform at most one task involving any given household, and they view 16 new potential recipients in every round. We also advise them in advance that we will implement only one of the eight allocations, chosen at random at the end of the experiment.

Table 2 summarizes the structure of the four main treatments and indicates the number of subjects who participated in each. The first column lists treatment labels, which describe each treatment's composition using simple shorthand notation: "A" denotes an ex ante task, "P" denotes an ex post task, and "A^R" denotes an ex ante task with surprise ex post revision. Thus, the label 2A2P_4A^R indicates that the treatment starts with two ex ante tasks ("2A") followed by two ex post tasks ("2P"), followed by four rounds involving ex ante tasks with surprise revisions ("4A^R"). Importantly, all revisions take place after the subject makes initial allocations in rounds 5–8. The main treatments have a common structure: in the first four rounds, subjects perform either ex ante tasks, ex post tasks, or a mixture of the two, while the last four rounds (listed after the underscore) always consist of ex ante decisions with surprise ex post revisions.¹⁷

We vary the computer's ticket allocation by round, as shown in table 3. In light of this variation, ex ante fair choices exhibit a distinctive fingerprint. Subjects do not see this table in advance; rather, they learn the computer's allocation at the start of each round.

B. *Basic Framing Effects*

This section documents two findings concerning the initial allocation chosen in each task (i.e., before any revisions). First, subjects tend to choose ex ante equalizing allocations when initially confronting tasks with ex ante framing and tend to choose ex post equalizing allocations when initially

¹⁷ Explanations of other treatments, which we used to examine the strictness of preferences and to test between competing theories of choice reversals, appear in subsequent sections.

TABLE 3
FIXED ALLOCATION OF COMPUTER'S TICKETS, BY ROUND

	Round							
	1	2	3	4	5	6	7	8
Tickets to household A	7	2	10	1	8	3	9	0
Tickets to household B	3	8	0	9	2	7	1	10

confronting tasks with ex post framing. Second, perspectives on fairness exhibit no persistence: subjects readily switch between ex ante and ex post perspectives across tasks, and responses to the initial framing of a task do not depend on the framing of previously encountered tasks.

To establish the first of these two findings, we focus on the first four rounds of treatments $4A_4A^R$, in which subjects start with four ex ante allocation tasks, and $4P_4A^R$, in which subjects start with four ex post tasks. Figure 1 shows the distributions of choices across the five categories defined in section III.B. Panels *A* and *B* pertain to subjects performing tasks with ex ante and ex post framing, respectively, during the first four rounds. The height of each bar indicates the fraction of choices that fell within a given category. The shading reflects the consistency of subjects' choices—it indicates the extent to which the choices in a given category were made by subjects whose decisions fell into that category in every

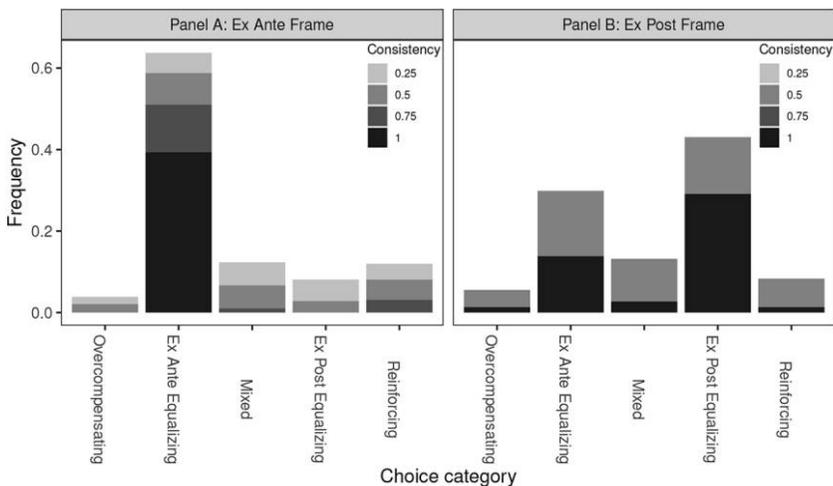


FIG. 1.—Distributions of initial choices conditional on initial framing. Panel *A* is based on the first four rounds of treatment $4A_4A^R$ (284 observations). Panel *B* is based on the first four rounds of treatment $4P_4A^R$ (144 observations). Shading indicates the extent to which the choices in a given category were made by subjects whose decisions fell into that category in every round (darkest shading), three-quarters of the rounds, half of the rounds, or one-quarter of the rounds (lightest shading).

round (darkest shading), three-quarters of the rounds, half of the rounds, or one-quarter of the rounds (lightest shading). Note that subjects made four ex ante choices but only two ex post choices, because their tickets were selected only half of the time. We highlight consistency across rounds because it could be an indication of the seriousness and deliberateness with which subjects approached the tasks and acted on coherent decision principles.

The differences between the distributions depicted in panels *A* and *B* of figure 1 are striking. For panel *A*, which pertains to initial tasks with ex ante framing, most choices are ex ante equalizing (i.e., fully offsetting). Furthermore, all fully consistent choosers were ex ante fair. In contrast, for panel *B*, which pertains to initial tasks with ex post framing, the modal choice is ex post fair (i.e., it involves no offsetting). Indeed, moving from panel *A* to panel *B*, the primary change is that the frequency of ex ante fair choices declines by 34 percentage points, while the frequency of ex post fair choices rises by 35 percentage points. Notably, ex ante fairness remains reasonably common in the ex post frame (consistent with findings in Cappelen et al. 2013), even among consistent choosers, while ex post fairness is relatively rare in the ex ante frame.¹⁸

Standard tests for the equality of distributions, such as Pearson's χ^2 test, are inapplicable here because they do not account for within-subject correlation across the four rounds. More specifically, any test that treats multiple observations of choices by the same subject as independent will tend to exaggerate the statistical significance of the differences across treatments. A resolution of this issue requires assumptions about the structure of the underlying statistical process. Accordingly, we pool the data from the two treatments, estimate a multinomial logit model with category-specific constants and category-treatment interactions, and perform a χ^2 test of the hypothesis that all the coefficients for the interaction terms are zero, clustering standard errors at the subject level. For the distributions depicted in figure 1, we reject equality decisively ($p < .001$).

So far, we have seen that the framing of the four initial decisions strongly influences the initial perspective on fairness. That finding does not necessarily imply that our subjects will exhibit choice reversals. After all, our experiment involves decision tasks that few, if any, subjects have previously encountered. Perhaps someone who initially performs such a task with one type of framing thinks through the class of tasks from that perspective and then adheres to the resulting decision principles through subsequent

¹⁸ The differences between initial decisions made with ex ante and ex post framing do not dissipate over the course of the first four rounds of treatments 4A_4A^R and 4P_4A^R, as subjects have more time to think through their attitudes toward these types of decision tasks. See fig. C.3 (figs. C.1–C.28 are available online). Also, in fig. C.14, we show that the differences between the distributions illustrated in panels *A* and *B* of fig. 1 are primarily attributable to consistent choosers.

tasks, even if the framing changes. In that case, subjects would exhibit no reversals. Thus, we ask whether, subject by subject, choices change as the decision frame changes.

To investigate these issues, we focus on treatments with changing decision frames, beginning with 2A2P_4A^R—in which subjects performed two tasks with ex ante framing, then two with ex post framing, and then four with ex ante framing (followed by surprise revisions)—as well as 2P2A_4A^R—in which subjects performed two tasks with ex post framing, then two with ex ante framing, and then an additional four with ex ante framing (followed by surprise revisions). For now, when examining rounds 5–8, we will focus on the original choices, leaving the analysis of revisions to section IV.C.

Figure 2 displays distributions of choices over the same five categories as figure 1, except that here we report results separately for rounds 1 and 2, 3 and 4, and 5–8. The first row pertains to treatment 2A2P_4A^R, while the second pertains to 2P2A_4A^R. For comparison, we also include treatments 4A_4A^R and 4P_4A^R in the third and fourth rows, respectively. We have highlighted the shifting frames both with text (labeled “EA” or “EP” in the corner) and with shading (darker background for the ex post frame).

Looking at this figure, one sees a striking similarity between the distributions of choices made within a given frame, regardless of the preceding choices. All of the choice distributions for ex ante frames closely resemble the distribution in panel *A* of figure 1, in that ex ante fair choices are predominant. All of the choice distributions for ex post frames resemble the distribution in panel *B* of figure 1, in that ex post fairness is the most common decision type. Thus, framing effects exhibit little, if any, persistence: choices depend on the framing of the current task, but not to any significant degree on the framing of initial or previous tasks. Subjects readily shift their perspectives on fairness back and forth along with the decision frame. Formal statistical tests confirm these visual impressions.¹⁹

C. Choice Reversals

The previous section documented a pronounced and stable tendency for subjects to adopt an ex ante perspective on fairness when making decisions with ex ante framing and an ex post perspective when making decisions

¹⁹ We reject the hypothesis that frame has no effect on allocations in rounds 1–4 of treatments 4P_4A^R, 4A_4A^R, 2A2P_4A^R, and 2P2A_4A^R ($p < .0001$), while we fail to reject the hypothesis that treatment dummies jointly have no effect on allocations in those same treatments ($p = .19$). Last, we fail to reject the hypothesis that the distributions of allocations in the last four rounds of these treatments are indistinguishable from each other ($p = .36$). From the figures, it is apparent that the failure to reject stems from the similarity of the distributions rather than from low power.

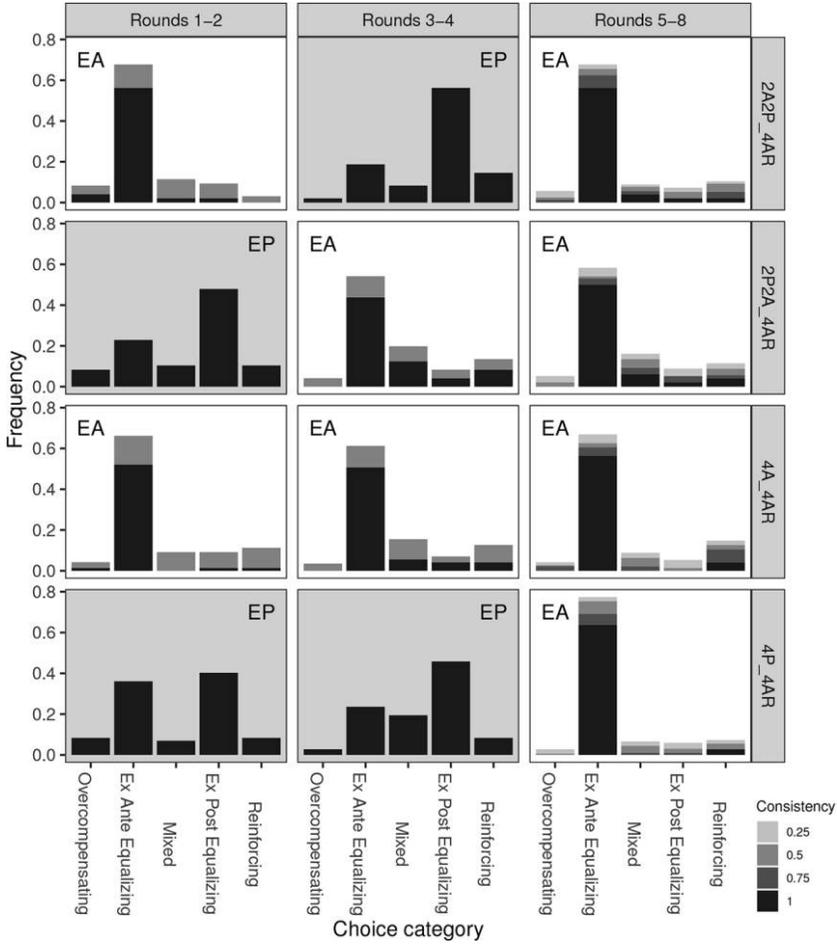


FIG. 2.—Distributions of choices in treatments with changing decision frames. Rows 1 and 2 are based on rounds 1–8 of treatment 2A2P_4A^R (48 subjects) and 2P2A_4A^R (48 subjects), respectively. Rows 3 and 4 are based on rounds 1–8 of treatment 4A_4A^R (71 subjects) and 4P_4A^R (72 subjects), respectively. Results for rounds 5–8 reflect original choices, not revisions.

with ex post framing. Those findings point to a potential source of choice reversals but do not actually establish that such reversals occur. It is one thing to invoke different decision criteria in completely separate tasks and potentially quite another to revise the choice made in a given task after arriving at a set of applicable principles for that task. Conceivably, people could apply their principles resolutely within each task while failing to do so across tasks.

To determine whether choice reversals actually occur, we examine the decisions that subjects make when they are unexpectedly allowed to revise

decisions ex post after allocating tickets ex ante. (Recall that we frame these opportunities neutrally as a second round of “confirming” their choices to mitigate possible experimenter demand effects.) We first focus on revisions made for rounds 5–8 of treatment 4A_4A^R. The subjects in this treatment encounter only tasks with ex ante framing before learning that they can revise the last four choices ex post. Revisions were the rule rather than the exception. Subjects revised 68.3% of the original round 5–8 choices, and 78.9% of subjects revised at least one choice.²⁰ Consistent with the notion that the ex post perspective on fairness becomes compelling once the ex post position is reached, switches to 50–50 were by far the most common type of revision (71.1%).

Figure 3 displays the distributions for original and final choices (left and center panels, respectively). A comparison of the two panels reveals the effect of unexpected revision opportunities on the distribution of allocations. The overall distribution shifts dramatically from one in which ex ante fair choices predominate to one in which ex post fair choices predominate. Indeed, there is a striking resemblance between figure 3 and figure 1. The right panel of figure 3 shows that revisions generally implemented equal division ex post.

In principle, the choice reversals by subjects in treatment 4A_4A^R could be the result of subjects not considering the ex post perspective until they find themselves with ex post opportunities to revise allocations. Conceivably, those who consider both the ex ante and the ex post perspectives might reconcile the conflict internally and display greater consistency as a result. Figure 2 suggests not: subjects continue to adopt ex ante perspectives on fairness in tasks with ex ante framing and ex post perspectives in separate tasks with ex post framing, even after exposure to both frames. However, that evidence stops short of demonstrating that subjects continue to reverse ex ante decisions when provided with opportunities to make ex post revisions.

To address this set of issues, we examine patterns of revisions in the three treatments that expose subjects to the ex post perspective in rounds 1–4: 4P_4A^R, 2A2P_4A^R, and 2P2A_4A^R. The frequency of revisions in these treatments is 69.4%, 69.8%, and 53.1%, respectively. As in treatment 4A_4A^R, subjects who made revisions primarily switched to ex post equalizing allocations. Moreover, differences in the distributions of revision types (whether the subject moved away from, toward, to, or past ex post fairness) between treatment 4A_4A^R on the one hand and treatments

²⁰ The revision frequency started out at 75.8% in round 5, dropped to 64.7% in round 6, and then rebounded a bit in rounds 7 and 8 (65.8% and 67.6%, respectively). Overall, there is no indication that the tendency to revise dissipates once subjects become aware of their behavior.

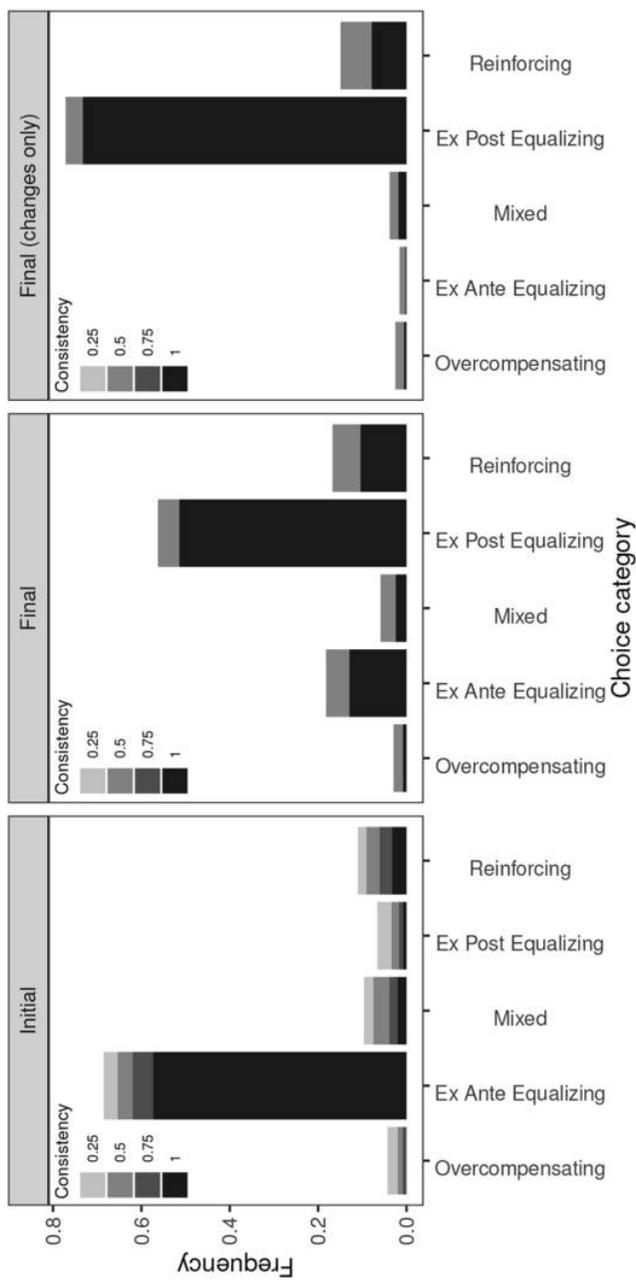


FIG. 3.—Marginal distributions of original and final choices. This figure is based on the final four rounds of treatment 4A_4A^R (284 observations).

2A2P_4A^R, 2P2A_4A^R, and 4P_4A^R on the other were not statistically significant.²¹

Figure 4 shows the joint distribution of the original and final choices for rounds 5–8, pooling over all four treatments. The figure consists of five panels with five bars each. There is one panel for each possible type of the original choice, which is indicated along the top of the figure. Within each group, there is one bar for each possible type of the final choice, as indicated by the label. Types of choices are displayed in the same order as in figure 1, both for original and for final choices. Frequencies are expressed as percentages of the total number of round 5–8 original-final choice pairs, so it is easier to see which patterns are most prevalent. The figure reveals that the most common original-final choice pair, by a wide margin, is an ex ante equalizing allocation followed by a revision to an ex post equalizing allocation (44.3% of observations). The second-most common choice pair, also by a wide margin, involves resolute ex ante fairness: the subject chooses the ex ante equalizing allocation at the outset and declines to revise it (17.2%).

These patterns are essentially the same for all four treatments, regardless of the framing experienced in the first four rounds. Thus, the predominance of the main pattern—initial ex ante fair choices followed by ex post fair revisions—is undiminished when subjects gain experience with the tension between the ex ante and ex post perspectives.

D. Subjects' Explanations for Choice Reversals

An unincentivized survey at the end of our experiment included two open-response questions designed to shed some light on motives and reasoning: “What factors did you consider when making your initial allocation decisions?” and “What factors did you consider when deciding whether to revise your initial allocations?” Virtually all subjects who switched from ex ante fairness to ex post fairness explained that, in both settings, they were trying to equalize the chances of winning. For example, one subject wrote, “When I knew my ticket was chosen, I reallocated the tickets to ensure both households had an equal opportunity.” Many articulated the same simple rule or criterion for both settings, in some cases using precisely the same words, such as “equal chance between the two householders,” “fairness; equal probability,” and “I consider equality.” None of our subjects commented on the tension between their ex ante and ex post choices. On the contrary, some of them explicitly stated that they saw their revision as following the same principles as their initial allocation. For

²¹ We fail to reject the hypothesis that revision behavior does not depend on treatment ($p = .88$). Once again, the failure to reject reflects the similarity of the distributions rather than low power.

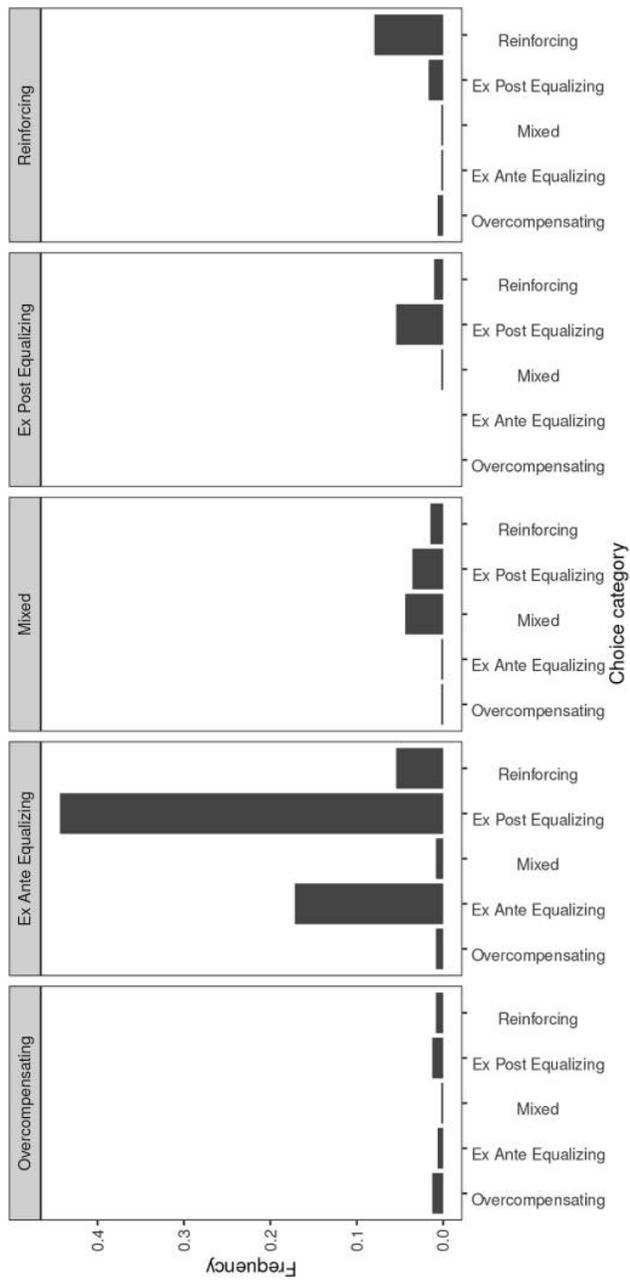


FIG. 4.—Joint distribution of original and final (revised) choices. This figure is based on the last four rounds of treatments 4A_4A^R, 4P_4A^R, 2A2P_4A^R, and 2P2A_4A^R (478 observations). The panel labels indicate the classification of the original allocations, while the labels of the bars indicate the classification of the final allocation.

example, one subject explained their revision as follows: "I used the same logic as I did before, giving the indistinguishable households equal chances of winning." Another wrote, "Same thing as the first time. I wanted to make the lottery fair so both households had an equal chance of winning the lottery." A third responded, "Same as above. To even out the odds for both families if my ticket was chosen." Shorter answers concerning the logic of a revision included "Same as before" and "Same as above."

These subjective responses point toward the mechanical application of a simple deontological heuristic of the type hypothesized in section II.B.1. We readily acknowledge the limitations of this type of survey evidence and concede that it is weak when evaluated in isolation. However, when considered alongside the behavioral patterns documented in sections V and VI, the survey responses help to complete the picture.

E. Strictness of Allocation Preferences

It is important to verify that the patterns documented in the previous subsections reflect strict preferences rather than the arbitrary resolution of indifference. To this end, we added treatments in which subjects performed one of the following three modified decision tasks:

Ex ante allocations with incentivized redistributions.—After a subject chooses an allocation in the ex ante frame, we present her with an unanticipated opportunity to enlarge the prize by reallocating all of her tickets to the household she treated less favorably. For example, if the subject chose to give eight tickets to household A and two to household B, the alternative would allocate all 10 tickets to household B while increasing the prize from \$10 to $\$(10 + x)$, where $x \in \{0.10, 0.50, 1, 2, 5\}$.²² Subjects make decisions for all five values of x . This modified task, A^S , allows us to evaluate the strictness of preferences for the initial allocation. We incentivize the reallocation by adjusting the size of the prize rather than through payments to the subject to avoid introducing a confounding factor (variation in the degree of altruism across subjects).

Ex ante allocations with surprise ex post revisions plus incentivized redistributions.—After a subject chooses an allocation in the ex ante frame, she learns whether the winning ticket is one of hers and then receives an unanticipated chance to revise her allocation, as in the A^R task. We then present her with an unanticipated opportunity to enlarge the prize (as in the A^S task) by reallocating all her tickets to the household she treated less favorably. This modified task, A^{RS} , allows us to evaluate the strictness of preferences for the final allocation.

²² If the subject initially divides the tickets equally, the alternative allocates all tickets to a randomly selected household in return for enlarging the prize.

TABLE 4
STRICTNESS OF ALLOCATION PREFERENCE TREATMENTS

Treatment	Rounds 1–4	Number of Subjects
4A ^S	Ex ante with incentivized redistributions	72
4A ^{RS}	Ex ante with surprise revision and incentivized redistributions	70
4A ^{RI}	Ex ante with surprise revision and incentive to abandon	55

Ex ante allocations with surprise ex post revisions plus incentives to abandon the revisions.—After a subject chooses an allocation in the ex ante frame, she learns whether the winning ticket is one of hers and then receives an unanticipated chance to revise her allocation, as in the A^R task. We then ask her to choose between her revised ticket allocation with a prize of \$10 and her initial ticket allocation with a prize of $\$(10 - x)$, where $x \in \{0.10, 0.50, 1, 2, 5\}$. This modified task, A^{RI}, allows us to evaluate the strictness of preference for the revised allocation over the initial allocation.

As detailed in table 4, we fielded one treatment for each type of task. Because these are relatively time-consuming tasks, we limited these treatments to four rounds. We informed subjects at the outset that we would implement their decision for one randomly chosen task and value of x .

When facing a small (\$0.10) incentive, subjects were unwilling to abandon their chosen allocations in 75% of A^S tasks, 67% of A^{RS} tasks, and 56% of A^{RI} tasks. Increasing x to \$0.50 produced only modest declines in these percentages. For the A^S and A^{RS} tasks, the reluctance to switch remained high even with much larger incentives. For example, in the 4A^S treatment, subjects declined a \$5 bonus nearly half (47%) of the time. Subjects may have exhibited weaker preferences in the 4A^{RI} treatment because the alternative—their initial allocation (usually 50–50)—was generally less draconian or because they were more averse to prize reductions than attracted to increases. For complete results, see figure 5.

V. Commitment Opportunities

In section II, we saw that consequential fairness preferences give rise to time inconsistency and, if decision makers are sophisticated, to a demand for commitment. We also explained that a time-consistent preference for implementing a simple and arguably appealing deontological rule can produce the same pattern of seemingly inconsistent ex ante and ex post choices but nevertheless yields a strict preference for flexibility, so that the decision maker can respond to information in accordance with her ethical principles. In this section, we distinguish between these classes of theories by investigating whether a demand for commitment arises in the current context.

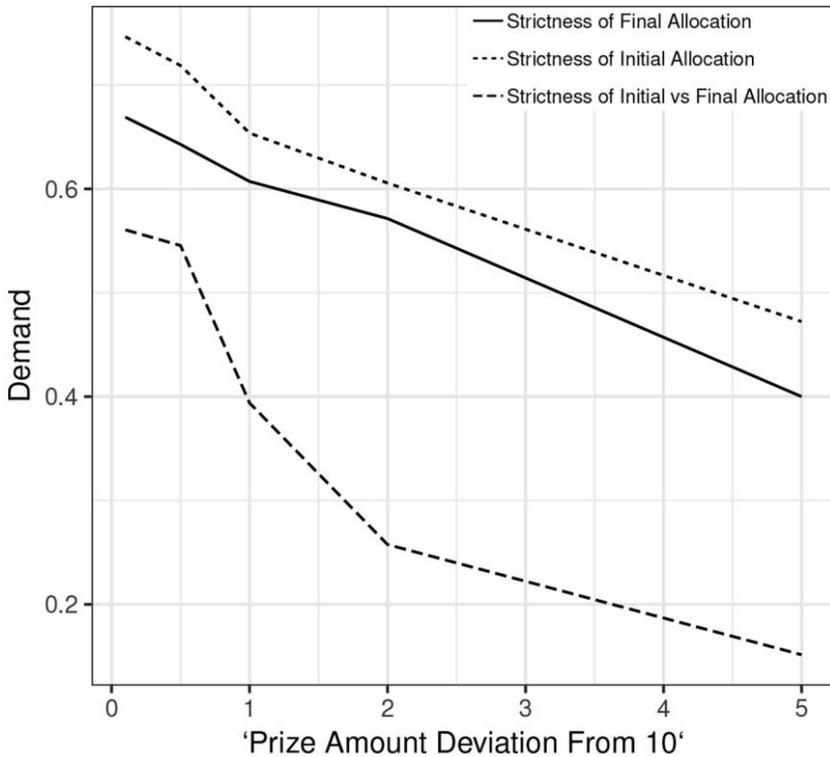


FIG. 5.—Strictness of allocation preferences. This figure is based on treatments $4A^S$, $4A^{RS}$, and $4A^{Rf}$. For the first two treatments, the vertical axis indicates the percentage of subjects not willing to redistribute their tickets in return for enlarging the prize by the amount indicated on the horizontal axis. For treatment $4A^{Rf}$, the vertical axis indicates the percentage of subjects not willing to restore their initial allocation in return for preserving the prize instead of reducing it by the amount indicated on the horizontal axis.

A. Experimental Design

To evaluate the demand for commitment, we introduce another variation of the allocation task:

Ex ante decisions with commitment.—After a subject chooses an allocation in the ex ante frame and confirms her choice (as in an ex ante task), we inform her that she will have an opportunity to reallocate her tickets after learning whether the winner is one of hers (but before learning who holds the winning ticket), unless she wishes to forgo that opportunity. At that point, she must express a preference for flexibility (“I definitely want the opportunity to revise”), a preference for commitment (“I definitely do not want the opportunity to revise”), or indifference (“I do not

care about having an opportunity to revise"). If a subject expresses a preference for flexibility, she learns whether the winning ticket was one of the computer's tickets or one of hers. In the latter case, she does not learn the number of the winning ticket but receives an opportunity to reallocate her tickets (as in an ex post task). If a subject expresses a preference for commitment, she makes no other decisions. If a subject expresses indifference, we implement a 50–50 randomization between these two alternatives. We then select a ticket at random, which determines the winner. Subjects learn all these rules in advance.

We implemented this variation of the allocation task in a treatment (4A^R_4A^C) with 72 subjects. During the first four rounds, subjects have opportunities to experience decision-making in both the ex ante and the ex post frames, as well as to notice their own tendencies to make revisions.²³ During the final four rounds, they start by making ex ante decisions but are given options to forgo subsequent revision opportunities.

As shown in section V.B, many subjects choose to make commitments, which mitigate the tendency to shift from ex ante fair to ex post equalizing allocations. However, additional findings presented in section V.C lead us to conclude that the apparent demand for commitment exaggerates the prevalence of time-inconsistent consequentialism. As a group, those who are inclined to switch from ex ante to ex post fairness actually avoid making commitments to a greater extent than other subjects. Apparently, many of them prefer to have and to exercise the flexibility to switch. That preference is consistent with the theory of naive deontological ethics discussed in section II.B.1.

B. Overall Effects of Commitment Opportunities

In this subsection, we address three questions. First, do subjects choose to forgo future flexibility when given the opportunity? Second, does the availability of these commitment opportunities reduce the frequency of revisions? Third, does it change the distribution of final choices?

Our first finding is that subjects choose commitment and flexibility with reasonably high frequency. They expressed a strict preference for commitment 40.6% of the time, a strict preference for flexibility 30.2%

²³ In rounds 1–4 of treatment 4A^R_4A^C, subjects generally exhibited the same patterns observed in rounds 5–8 of treatment 4A_4A^R, documented in sec. IV. For instance, 60.4% of the original choices were ex ante fair, while only 10.4% were ex post fair, and subjects revised 65.3% of choices ex post when given the opportunity. Of the revised choices, 69.1% were ex post equalizing, while only 1.1% were ex ante equalizing. We do not reject the equivalence of initial ($p = .37$) or final ($p = .78$) behavior in rounds 1–4 of treatment 4A^R_4A^C and rounds 5–8 of treatment 4A_4A^R (due to the similarity of the distributions rather than the absence of statistical power).

of the time, and indifference 29.2% of the time. These frequencies do not vary systematically across rounds.

Making a commitment does not necessarily change the outcome. For example, those with no inclination to revise may opt for commitments to avoid the inconvenience of reiterating their choices. Despite that possibility, our second finding is that commitment opportunities significantly reduce the frequency of revisions. Subjects revised only 36.8% of decisions in the last four rounds of $4A^R_4A^C$,²⁴ which is a little more than half of the comparable frequencies from the first four rounds of the same treatment (65.3%) and the last four rounds of treatment $4A^R_4A^R$ (68.3%); moreover, these differences are statistically significant ($p < .001$ in both cases).²⁵

Our third finding is that commitment opportunities significantly change the distribution of final choices. Comparing the distributions of the original allocations, we see very little difference between the first four rounds and the last four rounds of treatment $4A^R_4A^C$.²⁶ In contrast, there are striking and statistically significant differences between the distributions of final outcomes ($p = .03$).

Under the hypothesis that our subjects are fair consequentialists who seek commitments to mitigate time inconsistency, we would expect the frequency of ex ante fair allocations to be higher and that of ex post equalizing allocations to be lower, with commitments. That is indeed what we find: the frequency of ex ante equalizing allocations is 11 percentage points higher (49.0% vs. 37.8%) in the last four rounds (with commitment) than in the first four (without commitment), and the frequency of ex post fair allocations is about 8 percentage points lower (24.0% vs. 32.3%).

A closer look at the joint distribution of initial and final choices confirms that commitment opportunities mostly suppress migration from ex ante to ex post fair choices. Resolute ex ante behavior increases from 16.7% to 35.4%, while revisions from ex ante to ex post fairness decrease from 36.8% to 17.4%.²⁷

²⁴ For much of the analysis in this section, including the calculation of this figure, we focus on the tasks that the subject would have been allowed to revise if she had chosen flexibility.

²⁵ Similarly, 51.4% of subjects revised at least one decision in the last four rounds of $4A^R_4A^C$, compared with 80.6% in the first four rounds of the same treatment and 78.9% in the last four rounds of treatment $4A^R_4A^R$; these differences are also statistically significant ($p < .001$ in both cases).

²⁶ In fact, we do not reject the hypothesis that these two distributions are identical ($p = .43$). This finding reflects the similarity between the distributions rather than the lack of statistical power.

²⁷ The fractions of individuals choosing and sticking with three of the other four options also decline, but the changes are modest by comparison. As we discuss in app. B, offering commitment also suppresses migration from ex ante equalizing allocations to ex post equalizing allocations among subjects whose choices were consistent across rounds.

C. *Understanding the Demand for Flexibility and Commitment*

We have seen that subjects make commitments with high frequency and that these commitments reduce the frequency of revisions, primarily from ex ante to ex post equalizing allocations. Moreover, it is also the case that many subjects opt for flexibility and then revise their allocations. How can we account for both findings?

One possibility is that the theories discussed in section II are correct but the population is heterogeneous. Under this view, one attributes the preference for and effects of commitment to time inconsistency among sophisticated subjects with consequential non-EU preferences and the preference for flexibility and switching to subjects who embrace a naive deontological ethic. However, there are other possibilities. In principle, naiveté (lack of self-awareness) among time-inconsistent subjects could explain why some subjects maintain flexibility and then revise their allocations, and experimenter demand effects could account for all of these observations.²⁸

In this subsection, we present a series of findings that cast additional light on subjects' reasons for making or not making commitments. These findings speak to two questions. First, which subgroups exhibit the greatest demand for commitment? Second, what do subjects do with flexibility when they intentionally retain it?

1. Which Subgroups Exhibit the Greatest Demand for Commitment?

If the primary purpose of commitments is to impede undesired revisions from ex ante fair to ex post equalizing allocations, then the demand for commitment should be greater among subjects who choose initial allocations that entail a degree of ex ante fairness and especially among those who then tend to switch to ex post equalizing allocations when no commitments are allowed. In contrast, if migration from ex ante fair to ex post fair allocations reflects a naive deontological ethic, those same groups should exhibit a greater demand for flexibility. As we explain next, the evidence points to the latter hypothesis.

First, we find that the demand for commitment is lower and the demand for flexibility is higher when subjects choose allocations that they are more likely to revise (specifically, ones that entail a degree of ex ante fairness). When subjects started out by selecting the ex post fair allocation, the frequency with which they chose commitment was roughly three times as high as that with which they chose flexibility (52.6% vs. 15.8%).

²⁸ See app. B.3 for more discussion of these points.

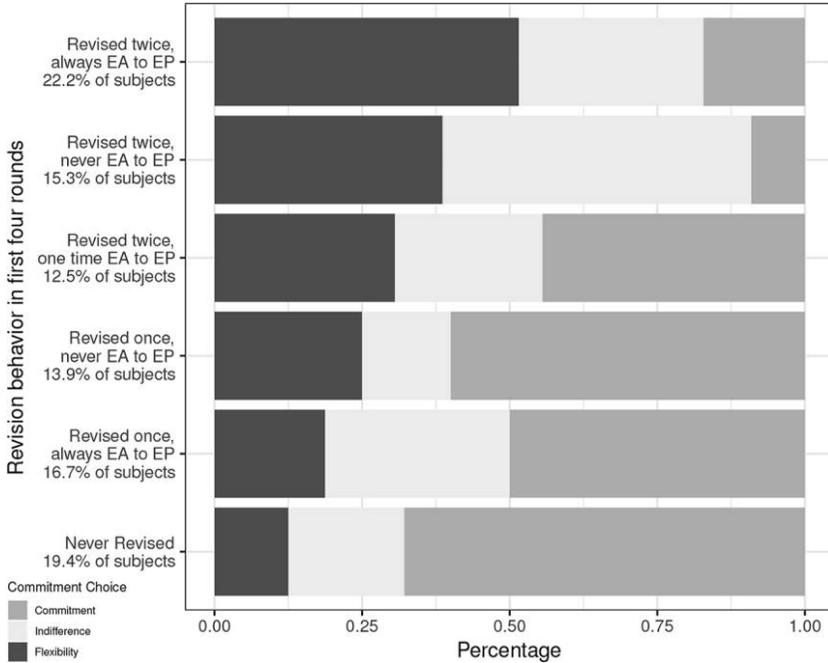


FIG. 6.—Commitment choices by migration patterns during the first four rounds. Data are from treatment $4A^R_4A^C$ (72 subjects). Revision categories are based on behavior in the first four rounds, while commitment choices are from the last four rounds. An allocation is classified as revised if the participant changed the numerical allocation of tickets, even if this revision did not move them to a different choice category.

In contrast, when subjects started by selecting the ex ante fair option, the frequency with which they chose commitment was only slightly larger than that with which they chose flexibility (42.4% vs. 33.7%).²⁹

Second, we find that the demand for commitment is lower and the demand for flexibility is higher among subjects who exhibit a greater tendency to migrate from ex ante fair to ex post equalizing allocations when no commitments are allowed. Recall that every subject had two opportunities to revise initial allocations during the first four rounds and no opportunities to make commitments. In figure 6, we have divided the subjects into six groups according to the patterns of their initial choices and revisions during those rounds. For each group, we display the frequencies

²⁹ When they started by selecting reinforcement (the only other nonoffsetting category), the relative prevalence of commitment choices (41.8% vs. 18.6%) was nearly as large as when they selected the ex post fair allocation. When they started by choosing either an overcompensating or a mixed allocation, the relative frequency of a preference for commitment (20.0% vs. 40.0% and 15.0% vs. 45.0%, respectively) was even lower than when they chose the ex ante fair allocation.

with which those subjects expressed a preference for flexibility, a preference for commitment, and indifference during the last four rounds. Those who revised twice in the first four rounds, always from the ex ante equalizing allocation to the ex post equalizing allocation, opted for flexibility more than 50% of the time and for commitment only 17% of the time. In sharp contrast, those who never revised in the first four rounds opted for commitment more than 65% of the time and for flexibility only 12% of the time. More generally, the figure establishes that the demand for flexibility was concentrated among those who revised more frequently in the first four rounds, while the demand for commitment was concentrated among those who revised less frequently. The differences between these frequencies are statistically significant ($p = .027$).

Relatedly, we show in appendix B.3 that those who chose flexibility and then made revisions likely understood their propensity to revise, because they had frequently revised allocations in the first four rounds. Consequently, the tendency to retain and then use flexibility does not appear to flow from naive or uninformed decision-making.

2. How Do Subjects Exercise Flexibility When They Intentionally Retain It?

If the primary purpose of commitments is to impede undesired revisions from ex ante fair to ex post equalizing allocations (as implied by fair consequentialism), then we would expect to find that the subjects who opt for flexibility exhibit relatively low rates of revision and migration from ex ante fair to ex post fair allocations. In contrast, if migration from ex ante fair to ex post fair allocations reflects the consistent application of a naive deontological ethic, that pattern should be particularly prevalent among those who affirmatively choose flexibility. As we explain next, the evidence again points to the latter hypothesis.³⁰

First, we find that in tasks with commitment options, the revision rate is exceptionally high among those who opt for flexibility. Overall, subjects revised 85.4% of decisions in tasks where they chose flexibility over commitment. Significantly, that figure is higher, not lower, than the comparable figures for the first four rounds (65.3%) and the last four rounds (68.3%) of treatment 4A_4A^R.

Second, we find that those who opt for flexibility are disproportionately inclined to migrate from ex ante fair to ex post equalizing allocations. Focusing on the migration patterns for those who affirmatively retained the

³⁰ Here we acknowledge that experimenter demand effects may establish a baseline frequency for revisions. However, that possibility does not explain the specific observation that revisions by those who choose flexibility tend to yield ex post fair outcomes. Closer examination of revisions allows us to differentiate between the hypotheses of interest.

flexibility to revise, 66.7% of the original choices were ex ante fair, and of those, 80.8% were revised to ex post fair choices. Thus, migration from ex ante to ex post fairness predominates among uncommitted choices—it accounts for 51.2% of the choice pairs. This pattern suggests that many of those who migrate from ex ante fair to ex post fair choices actually prefer the flexibility to migrate. Focusing on those who said they were indifferent between commitment and flexibility, only 4.8% of the choice pairs exhibited migration from ex ante to ex post fairness, and the most common pattern was to select the ex ante fair allocation and stick with it. This contrast again suggests that those who intentionally avoid commitments affirmatively value the ability to switch from an ex ante fair choice to an ex post fair one and have no desire to preclude this migration.

D. Strictness of Preferences for Commitment and Flexibility

It is once again important to verify that the patterns documented in the previous subsections reflect strict preferences rather than the arbitrary resolution of indifference. To this end, we added another type of decision task:

Ex ante allocations with incentivized commitment.—This task adds an additional stage to task A^c. For those choosing flexibility, we ask whether they would be willing instead to commit to their original choice if we increased the total prize from \$10 to \$(10 + x), where $x \in (0.25, 0.50, 1, 2.50, 5)$. Similarly, for those choosing commitment, we ask whether they would be willing instead to retain flexibility if we increased the total prize by the same amounts. This modified task, A^{cs}, allows us to evaluate the strictness of preference for commitment and flexibility.

We implemented this variation of the allocation task in a treatment (4A^R_4A^{CS}) with 79 subjects. Similar to other treatments, 69% of initial choices in rounds 1–8 were ex ante equalizing, and 75% of revisions in rounds 1–4 were ex post equalizing. In rounds 5–8, subjects chose flexibility 34% of the time, commitment 36% of the time, and indifference 30% of the time. We find that preferences for commitment and flexibility are typically strict, in that roughly 80% of subjects are unwilling to switch for the smallest prize bonus (\$0.25). Also, the demand for flexibility is more robust among those who revised from the ex ante fair to the ex post equalizing allocation at least once in the first four rounds, while the demand for commitment is more robust among those who did not make this revision. See figure 7 for complete results.

VI. Contingent Planning

In section II, we explained that someone who follows the hypothesized naive deontological ethic will reaffirm her desire to switch to the ex post

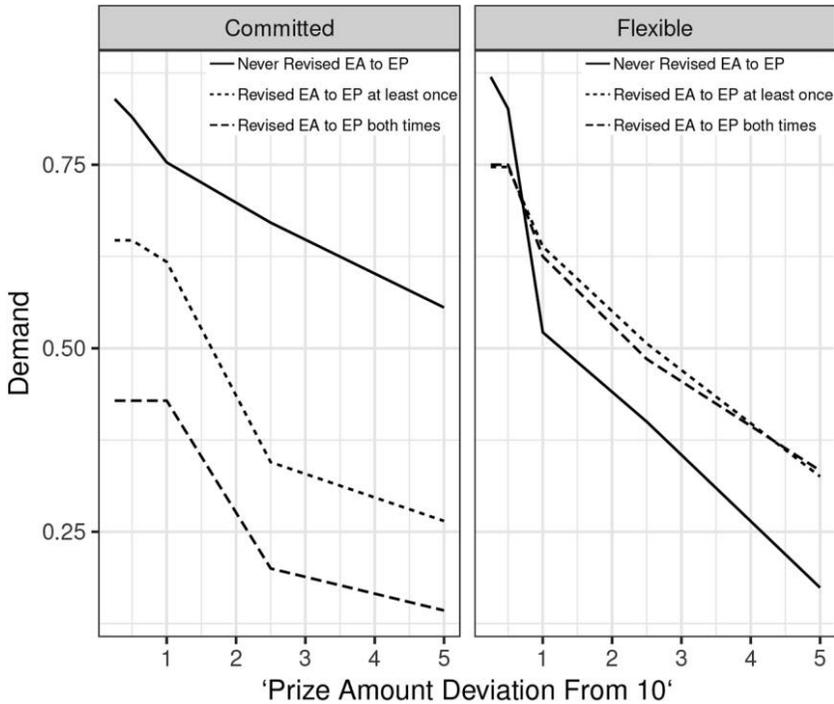


FIG. 7.—Strength of preference for commitment and flexibility. Data are from treatment $4A^R_4A^{CS}$. Revision categories are based on behavior in the first four rounds, while commitment choices are from the last four rounds.

fair allocation when asked to specify a contingent plan for her revision before learning any information about the realization. In contrast, a time-inconsistent consequentialist with strict fairness preferences will choose a contingent plan that reiterates her ex ante choice. Thus, to test these theories, we examine another variation of the allocation task:

Ex ante allocations with planned ex post revisions.—The subject allocates her tickets immediately after learning the computer's allocation, without receiving any other information. As in an ex ante decision, we ask her to confirm or revise it. Later, once all initial allocations have been entered, we revisit each allocation problem again. We explain that the participant will soon learn whether the winning ticket is one of hers, and we ask her to provide us with instructions for that contingency. She may reenter her initial allocation, or she may provide a revised allocation. Aside from committing to revisions before rather than after the receipt of information, this task, labeled A^P , is identical to task A^R . It allows us to evaluate whether subjects lock in their ex ante or ex post objectives when explicitly adopting contingent plans.

In this task, a time-inconsistent consequentialist will always choose a plan that delivers the same outcome as the ex ante task (task A)—typically ex ante fairness. Indeed, from a consequentialist perspective, tasks A and A^p are equivalent, because the initial ticket allocation already specifies a fully contingent plan (“If this ticket is the winner, then this household will receive the prize”). In contrast, if a subject implements the naive deontological heuristic as we have hypothesized, she will implement a plan that delivers the same outcome as the ex post task—typically ex post fairness. From her perspective, the difference in framing makes the tasks A and A^p nonequivalent.

We implemented this variation of the allocation task in a treatment (4A_4A^p) with 46 subjects. Subjects first made four decisions with ex ante framing, followed by four tasks with planned revisions.

The results are shown in figure 8, which displays the marginal distributions of original and planned revisions during the final four rounds of treatment 4A_4A^p. Similar to other treatments, a majority of ex ante choices were ex ante equalizing.³¹ In contrast, a majority of the contingent plans were ex post equalizing.³² Overall, there is a striking similarity between figure 8 and the first two panels of figure 3, which show the original and final choices in the last four rounds of treatment 4A_4A^r. Many subjects choose an initial allocation that is ex ante fair but instruct us to reallocate their tickets evenly if it turns out that the winning ticket is one of theirs. The similarity between the distributions of planned revisions and revised choices indicates that subjects plan ex ante to make the same selections that they would prefer ex post. This pattern is at odds with the hypothesis of time-inconsistent consequentialism but confirms the hypothesis of naive deontological ethics.

VII. Conclusions

In this paper, we have experimentally explored how people think about fairness in settings where there is a tension between ex ante and ex post perspectives. We demonstrated that most people robustly pursue ex ante fairness in the ex ante position and ex post fairness in the ex post position. Most importantly, however, when we reveal information that converts an ex ante frame into an ex post frame, subjects deliberately switch from ex

³¹ In the last four ex ante decisions that came before contingent planning decisions, 66% (122/184) of decisions were ex ante equalizing, compared to 67% (190/284) of initial decisions in the last four rounds of 4A_4A^r ($p = .93$).

³² Specifically, 57% of contingent allocations were ex post equalizing, while just 16% were ex ante equalizing. The comparable frequencies in rounds 5–8 of treatment 4A_4A^r were 53% and 15%, respectively. Altogether, 62% of revisions were from ex ante equalizing to ex post equalizing in the contingent planning tasks, compared to 65% in the final four rounds of 4A_4A^r ($p = .99$).

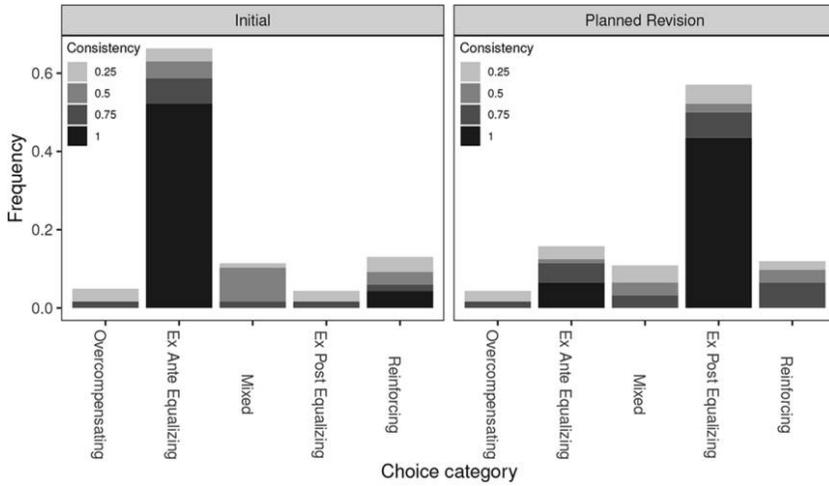


FIG. 8.—Marginal distributions of original and planned revisions. This figure is based on the final four rounds of the contingent planning treatment (184 observations).

ante fair choices to ex post fair choices, despite the fact that these revisions make the final allocation clearly unfair from an ex ante perspective.

We have considered two classes of explanations for this pattern. The first holds that our subjects are fair consequentialists and that switching reflects time inconsistency that emerges naturally from the conflict between the ex ante and ex post perspectives. The second depicts decision makers as having time-consistent preferences for applying naive deontological heuristics, even though a consequentialist would deem their implications time-inconsistent. In particular, a rule that prescribes egalitarian treatment of indistinguishable parties will likewise produce ex ante fair allocations for actions executed ex ante and ex post fair allocations for actions executed ex post.

How can we tell whether we are observing undesired time inconsistency? We give people the chance to learn that they are time-inconsistent and offer them ex ante commitment opportunities. How can we tell whether we are observing deliberate and therefore time-consistent desires to follow a simple deontological rule? After giving them some experience, we ask them to specify fully contingent plans. While a modest demand for commitment indicates some time-inconsistent consequentialism, the weight of the evidence points to naive deontological ethics.

These findings have potentially important implications for public policy. The tension between concerns for equality of opportunity (a notion of ex ante fairness) and equality of outcomes (a notion of ex post fairness) is evident in many public policy debates. Our analysis implies that notions

of fairness will tend to evolve systematically as information concerning outcomes progressively emerges. As a result, citizens who care about fairness may end up supporting policies that undermine the consequential implications of their original positions.

As an example, consider the question of whether the government should permit health insurance companies to account for preexisting conditions when setting premiums. Ex ante fairness requires only that everyone has the opportunity to purchase fairly priced insurance prior to developing any medical condition. Under that policy, some who initially decide against insurance will subsequently develop conditions that imply elevated health risks, at which point insurance companies may be willing to provide coverage only at prohibitive rates. Ex post fairness may then argue for rules that make insurance affordable—for example, by prohibiting underwriting based on preexisting conditions. Unfortunately, that policy discourages people from purchasing insurance ex ante and thereby exposes insurance companies to adverse selection, which renders these markets less efficient. Similar considerations arise in the context of social insurance. For instance, concerns over ex post fairness toward victims of natural disasters may account for the US government's heavy reliance on ex post disaster relief (i.e., through the Federal Emergency Management Agency), which undermines ex ante insurance and risk minimization.

On a more conceptual level, one can think of this paper as a positive investigation of normative ethics. We do not attempt to derive criteria for judging whether a choice is ethical. Instead, our research sheds light on the criteria that people actually use. It points toward a deontological perspective, wherein people judge the morality of an action nonconsequentially, according to its consistency with ethical rules. Whether the judgments we identify resonate with a particular flavor of deontology (such as Kantianism) is an interesting question but one that ventures beyond the more pragmatic objectives of the current study.

Our findings raise other important questions that are worth consideration in future research. For example, while we have attempted to distinguish between two broad classes of explanations for the tendency to switch between ex ante and ex post perspectives on fairness, much remains unclear about the particular structure of preferences. As we have noted, an objective function of the form $W(EU_A, EU_B)$, which captures concern for the distribution of EU, can give rise to a preference for equal division of lottery tickets, but so can other specifications. Consider, for example, the possibility that decision makers employ probability weighting, an assumption for which there is substantial precedent in the literature on risk and uncertainty (e.g., Kahneman and Tversky 1979). In the context of our split-the-tickets task, we would write the objective function as

$$\pi(p_A)E(W(U_A^W, U_B^L)) + \pi(1 - p_A)(W(U_A^L, U_B^W)).$$

A preference for equal division of tickets emerges under the assumptions of symmetry and concavity of π . While the implications of these two preference specifications are essentially indistinguishable for split-the-tickets tasks, they diverge sharply in related contexts. We refer the interested reader to appendix A, where we demonstrate that it is possible to differentiate these models by examining a related class of decision tasks.³³

Another important question is whether those practicing the simple deontological ethic are philosophically naive or cognitively naive. The hypothesis of cognitive naiveté encompasses the possibility that people may misapply their underlying ethical principles in ex ante and/or ex post settings because they have difficulty reasoning out the full implications of their choices. For example, to understand the consequences of choices in our experiment, subjects must engage in contingent reasoning and also reduce a compound lottery. On the one hand, prior research has shown that people sometimes find these tasks challenging.³⁴ On the other hand, the close correspondence between contingent plans and ex post choices suggests that subjects had a strong grasp of the relevant contingency in this instance. Moreover, in appendix A, we show that qualitatively similar patterns arise with a prize-splitting (as opposed to a ticket-splitting) task, for which the structure of contingencies is even simpler and compound lotteries are avoided. While this finding suggests that our results are at least partly attributable to philosophical rather than cognitive naiveté, further investigation of this question is clearly warranted.

It is also important to know whether the robustness with which people switch between ex ante and ex post fairness reflects the habitual application of a familiar ethical rule or the thoughtful application of a coherent value system. Would they continue to migrate freely among these perspectives if they had a direct stake in the outcome, or would they rationalize a self-serving ethical perspective? Would a particular perspective become more compelling if one of the recipient households were arguably more deserving? Do political beliefs and other socioeconomic factors predict the mix of preference types? Investigating these and other important questions raised by this study will, we hope, contribute to a deeper and more complete understanding of social preferences.

³³ In app. A, we describe a split-the-prize task, in which we specify an arbitrary allocation of a fixed dollar prize between the two parties and the decision maker selects an alternate allocation. A coin flip determines whether we implement the fixed or chosen allocation. Each subject chooses their allocation ex ante but can revise it ex post on learning that the coin flip has selected it. In this setting, the implications of the two preference formulations, $W(EU_A, EU_B)$ and $\pi(p_A)E(W(U_A^W, U_B^L)) + \pi(1 - p_A)(W(U_A^L, U_B^W))$ differ sharply. Additional treatments discussed in the appendix suggest that the population may include individuals with both types of preferences.

³⁴ Esponda and Vespa (2014) document failures involving contingent reasoning. Harrison, Martínez-Correa, and Swarthout (2015) find that people have difficulty reducing compound lotteries in some contexts but not in others.

References

- Agrawal, Arun. 2002. "Indigenous Knowledge and the Politics of Classification." *Internat. Soc. Sci. J.* 54 (173): 287–97.
- Andreoni, James, and B. Douglas Bernheim. 2009. "Social Image and the 50–50 Norm: A Theoretical and Experimental Analysis of Audience Effects." *Econometrica* 77 (5): 1607–36.
- Andreoni, James, Paul M. Brown, and Lise Vesterlund. 2002. "What Makes an Allocation Fair? Some Experimental Evidence." *Games and Econ. Behavior* 40 (1): 1–24.
- Andreoni, James, Marco Castillo, and Ragan Petrie. 2003. "What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game." *A.E.R.* 93 (3): 672–85.
- Andreoni, James, and John Miller. 2002. "Giving according to GARP: An Experimental Test of the Consistency of Preferences for Altruism." *Econometrica* 70 (2): 737–53.
- Bernheim, B. Douglas, and Sergei Severinov. 2003. "Bequests as Signals: An Explanation for the Equal Division Puzzle." *J.P.E.* 111 (4): 733–64.
- Bloom, David E. 1986. "Empirical Models of Arbitrator Behavior under Conventional Arbitration." *Rev. Econ. and Statis.* 68 (4): 578–85.
- Bohnet, Iris, and Bruno S. Frey. 1999. "Social Distance and Other-Regarding Behavior in Dictator Games: Comment." *A.E.R.* 89 (1): 335–39.
- Bolton, Gary E., Jordi Brandts, and Axel Ockenfels. 2005. "Fair Procedures: Evidence from Games Involving Lotteries." *Econ. J.* 115 (506): 1054–76.
- Bolton, Gary E., and Axel Ockenfels. 2000. "ERC: A Theory of Equity, Reciprocity, and Competition." *A.E.R.* 90 (1): 166–93.
- Brock, J. Michelle, Andreas Lange, and Erkut Y. Ozbay. 2013. "Dictating the Risk: Experimental Evidence on Giving in Risky Environments." *A.E.R.* 103 (1): 415–37.
- Camerer, Colin F. 1997. "Progress in Behavioral Game Theory." *J. Econ. Perspectives* 11 (4): 167–88.
- Cappelen, Alexander W., James Konow, Erik O. Sorensen, and Bertil Tungodden. 2013. "Just Luck: An Experimental Study of Risk-Taking and Fairness." *A.E.R.* 103 (4): 1398–413.
- Charness, Gary, and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Q.J.E.* 117 (3): 817–69.
- Coate, Stephen. 1995. "Altruism, the Samaritan's Dilemma, and Government Transfer Policy." *A.E.R.* 85 (1): 46–57.
- Dasgupta, Sudipto, and Zhigang Tao. 1998. "Incomplete Contracts, Ownership Rights, and the Optimality of Equity Joint Ventures." *J. Econ. Behavior and Org.* 37 (4): 391–413.
- DeWeaver, Mark, and James Roumasset. 2002. "Risk Aversion as Effort Incentive: A Correction and Prima Facie Test of the Moral Hazard Theory of Share Tenancy." *Econ. Bull.* 15:1–16.
- Diamond, Peter A. 1967. "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty." *A.E.R.* 57 (4): 759–76.
- Esponda, Ignacio, and Emanuel Vespa. 2014. "Hypothetical Thinking and Information Extraction in the Laboratory." *American Econ. J.: Microeconomics* 6 (4): 180–202.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. "A Theory of Fairness, Competition, and Cooperation." *Q.J.E.* 114 (3): 817–68.
- Forsythe, Robert, Joel L. Horowitz, N. E. Savin, and Martin Sefton. 1994. "Fairness in Simple Bargaining Experiments." *Games and Econ. Behavior* 6 (3): 347–69.

- Fudenberg, Drew, and David K. Levine. 2012. "Fairness, Risk Preferences and Independence: Impossibility Theorems." *J. Econ. Behavior and Org.* 81 (2): 606–12.
- Harrison, Glenn W., Jimmy Martínez-Correa, and J. Todd Swarthout. 2015. "Reduction of Compound Lotteries with Objective Probabilities: Theory and Evidence." *J. Econ. Behavior and Org.* 119:32–55. doi:10.1016/j.jebo.2015.07.012.
- Harsanyi, John C. 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Welfare." *J.P.E.* 63 (4): 309–21.
- Hauswald, Robert, and Ulrich Hege. 2003. "Ownership and Control in Joint Ventures: Theory and Evidence." CEPR Working Paper no. 4056, Centre Econ. Policy Res., London.
- Hoffman, Elizabeth, Kevin A. McCabe, and Vernon L. Smith. 1996. "On Expectations and the Monetary Stakes in Ultimatum Games." *Internat. J. Game Theory* 25 (3): 289–301.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–91.
- Karni, Edi, Tim Salmon, and Barry Sopher. 2008. "Individual Sense of Fairness: An Experimental Study." *Experimental Econ.* 11 (2): 174–89.
- Kircher, Philipp, Sandra Ludwig, and Alvaro Sandroni. 2013. "On the Difference between Social and Private Goods." *B.E. J. Theoretical Econ.* 13 (1): 151–77.
- Krawczyk, Michal, and Fabrice Le Lec. 2010. "'Give Me a Chance!' An Experiment in Social Decision under Risk." *Experimental Econ.* 13 (4): 500–511.
- Machina, Mark J. 1989. "Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty." *J. Econ. Literature* 27 (4): 1622–68.
- Markowitz, Harry M. 1968. *Portfolio Selection: Efficient Diversification of Investments*, vol. 16. New Haven, CT: Yale Univ. Press.
- McClennen, Edward F. 1989. "Justice and the Problem of Stability." *Philosophy and Public Affairs* 18:3–30.
- Menchik, Paul L. 1980. "Primogeniture, Equal Sharing, and the U.S. Distribution of Wealth." *Q.J.E.* 94 (2): 299–316.
- . 1988. "Unequal Estate Division: Is It Altruism, Reverse Bequests, or Simply Noise?" In *Modelling the Accumulation and Distribution of Wealth*, edited by Denis Kessler and Andre Masson, 105–16. Oxford: Oxford Univ. Press.
- Rabin, Matthew. 1993. "Incorporating Fairness into Game Theory and Economics." *A.E.R.* 83 (5): 1281–302.
- Raiffa, Howard. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, MA: Addison-Wesley.
- Saito, Kota. 2013. "Social Preferences under Risk: Equality of Opportunity versus Equality of Outcome." *A.E.R.* 103 (7): 3084–101.
- Sinnott-Armstrong, Walter. 2015. "Consequentialism." <https://plato.stanford.edu/entries/consequentialism>.
- Trautmann, Stefan T., and Gijs van de Kuilen. 2016. "Process Fairness, Outcome Fairness, and Dynamic Consistency: Experimental Evidence for Risk and Ambiguity." *J. Risk and Uncertainty* 53 (2): 75–88.
- Trautmann, Stefan T., and Peter P. Wakker. 2010. "Process Fairness and Dynamic Consistency." *Econ. Letters* 109 (3): 187–89.
- Veugelers, Reinhilde, and Katrien Kesteloot. 1996. "Bargained Shares in Joint Ventures among Asymmetric Partners: Is the Matthew Effect Catalyzing?" *J. Econ.* 64 (1): 23–51.
- Wilhelm, Mark O. 1996. "Bequest Behavior and the Effect of Heirs' Earnings: Testing the Altruistic Model of Bequests." *A.E.R.* 86 (4): 874–92.