

How to Model Heterogeneity in Costly Punishment: Insights from Responders' Response Times

URS FISCHBACHER^{1,2}, RALPH HERTWIG^{3*} and ADRIAN BRUHIN⁴

¹Thurgau Institute of Economics, Kreuzlingen, Switzerland

²University of Konstanz, Germany

³Max Planck Institute for Human Development, Center for Adaptive Rationality, Berlin, Germany

⁴University of Lausanne, Lausanne, Switzerland

ABSTRACT

We investigate what processes may underlie heterogeneity in social preferences. We address this question by examining participants' decisions and associated response times across 12 mini-ultimatum games. Using a finite mixture model and cross-validating its classification with a response time analysis, we identified four groups of responders: one group takes little to no account of the proposed split or the foregone allocation and swiftly accepts any positive offer; two groups process primarily the objective properties of the allocations (fairness and kindness) and need more time the more properties need to be examined; and a fourth group, which takes more time than the others, appears to take into account what they would have proposed had they been put in the role of the proposer. We discuss implications of this joint decision–response time analysis. Copyright © 2013 John Wiley & Sons, Ltd.

KEY WORDS ultimatum game; response time; finite mixture model; heterogeneity; altruistic punishment; response time; heuristics

Our collective folklore is populated with characters, such as Zorro, Robin Hood, and the Lone Ranger, who fight to right what they perceive of as injustice and in the process punish norm violators. Dirty Harry, the screen cop made famous by Clint Eastwood, even tells the norm violators “Go ahead, make my day,” thus informing them that he will derive satisfaction from punishing them. Clearly, these characters are figments of our imagination, although some may have historical origins. The willingness to enforce norms even if doing so exacts costs to one, however, is not a prerogative of fictional characters but can also be observed in ordinary people. Take, for example, the legion of whistleblowers, such as Daniel Ellsberg, risking their career and retaliation to expose the misconduct on the part of an agency, organization, or a company.

Not everybody, however, finds a Robin Hood in himself or herself or finds satisfaction in punishing norm violators. Do economists and psychologists therefore have to assume fundamentally different character or personality traits to explain people's varying predilections to punish norm violators? Calling upon such distinct traits is indeed the predominant approach in moral philosophy (see Doris, 2002). Alternatively, however, these differences across people could be differences in degree, not kind. The same issue also arises concerning the cognitive processes underlying our choices, which are social in the sense that they also affect other people apart from the decision maker. When making these choices, do all people process the same chunks of information and motives similarly, except that they may weight them differently? Or, do some people recruit processes that are different in kind?

We will be concerned with heterogeneity in social choices. People are able to display an amazing range of behaviors

towards others. In particular, there is a rich body of empirical evidence showing that humans care about others. They are generous, and they reward the kind actions of others and punish unkind behavior. Their responses, however, can also betray envy and sometimes even spite (e.g., Berg, Dickhaut, & McCabe, 1995; Falk, Fehr, & Fischbacher, 2005; Fehr & Gächter, 2000; Güth, Schmittberger, & Schwarze, 1982). Observations of diverse other-regarding preferences challenge a frequent assumption in economic models, namely, that humans are assumed to be rational decision makers and to harbor purely self-regarding preferences (see Camerer & Fehr, 2006).

Several theories have been proposed to accommodate the panoply of other-regarding preferences. All retain the utility framework and account for other-regarding behaviors by either discarding or modifying the assumption of purely self-regarding preferences. *Social preference theories* modify it by stipulating that, in addition to the material payoff, people's choices may be guided by the outcomes and behaviors of others. Those other-regarding preferences (henceforth, social motives) are incorporated as additional terms into the utility function. The most prominent among these theories are those by Bolton and Ockenfels (2000), Charness and Rabin (2002), Dufwenberg and Kirchsteiger (2004), Falk and Fischbacher (2006), Fehr and Schmidt (1999), Kirchsteiger (1994), Levine (1998), and Rabin (1993). Common to all is the acknowledgement that humans are heterogeneous in the extent to which they are guided by selfish and social motives. The theories, however, differ in the motives (represented by parameters) that they postulate and in how they combine them. For example, the theories by Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) assume that people value equality and dislike inequity, whereas the theories by Charness and Rabin (2002), Dufwenberg and Kirchsteiger (2004), Falk and Fischbacher (2006), Levine

*Correspondence to: Ralph Hertwig, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. E-mail: hertwig@mpib-berlin.mpg.de

(1998), and Rabin (1993) assume that people like to reciprocate and hence reward kind and punish unkind behavior. Each theory aspires to describe all humans, capturing heterogeneity by parameterization of the assumed motives. In practice, each theory succeeds in explaining some data but fails in explaining all (see, e.g., Brandts & Sola, 2001; Engelmann & Strobel, 2004; Falk, Fehr, & Fischbacher, 2003, 2008; Falk et al., 2005). In this article, we take a different approach to shed light onto the phenomenon of heterogeneity. Specifically, we take advantage of response times (RTs) in the mini-ultimatum game, a variant of the classic and frequently studied ultimatum game (Güth et al., 1982). Next, we briefly introduce these games and then turn to the standard approach of modeling heterogeneity in social games. Note that a companion piece to the present article is that of Hertwig, Fischbacher, and Bruhin (2013). Their theoretical and empirical analyses focus on models of sequential decision trees. Here, we extend this work by reporting more detailed tests and analyses of people's behavior and RTs. Furthermore, the present focus is on establishing evidence for heterogeneity between people's decisions in the mini-ultimatum game and how it can be mapped onto underlying psychological processes.

THE ULTIMATUM GAME AND THE MINI-ULTIMATUM GAME

The ultimatum game involves two parties who play a single round in which one person, the proposer, suggests how to split a fixed monetary pie, typically provided by the experimenter. This proposed split represents a take-it-or-leave-it offer (an ultimatum) that the other person, the responder, chooses to accept or reject. The interaction between the players is anonymous. If the responder chooses to accept the offer, the division will be implemented. Should the responder decide to reject the proposed division, both players will go away empty-handed. In theory, a purely self-interested responder will accept any proposed positive payoff, no matter how small. Anticipating this choice, a self-interested proposer will offer nothing more than the smallest amount possible. The equilibrium offer (i.e., the division for which no player has anything to gain by doing something differently) thus allocates the smallest positive payoff to the responder and the lion's share to the proposer. This, however, is not what is typically observed in the laboratories. More than 30 years of research on the ultimatum game has consistently found that responders tend to reject low offers and thus behave at odds with the assumption that they simply maximize their self-interest (see Camerer, 2003; Güth & Tietz, 1990).

In our investigation, we employed a variant of this classic social game. The mini-ultimatum game is a sequential two-player game. The first mover can choose between two fixed allocations. Then, the second mover can accept or reject this choice. Acceptance means that the allocation will be implemented, whereas rejection means that both players will receive zero. We presented our participants with 12 different

mini-ultimatum games and classified them according to their *rejection behavior*. Subsequently, we used this classification in the analysis of participants' RTs.

HOW TO MODEL HETEROGENEITY?

Most theories of social preferences capture only one social motive, be it, for example, inequity aversion (Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999) or reciprocity (Levine, 1998; Rabin, 1993). Moreover, all the models fail to explain the behavior of *all* participants (e.g., Engelmann & Strobel, 2004). Within parameterized social preference theories, there are at least two ways to respond to this state of affairs. Accepting that the existing theories explain only a subset of respondents, one could model heterogeneity using distinct models for different people. Alternatively, one may aim for a unifying theory that encompasses multiple social motives and that explains heterogeneity by assuming that people differ in the strength and combinations of these motives, as has been done, for example, by Charness and Rabin (2002), Cox, Friedman, and Gjerstad (2007), Cox, Friedman, and Sadiraj (2008), or Falk and Fischbacher (2006).

The interpretation of degenerated parameter values illustrates how these two approaches differ. Degenerated values (mostly zero values), for instance, for parameters capturing social motives could be interpreted to mean that the person, in reality, operates on the basis of a selfish preference function. Consequently, the person's choices can be accommodated in terms of a simpler model that altogether omits parameters capturing social motives. The same seemingly selfish behavior, however, could also be consistent with very low values on these *social* parameter(s). Consequently, such parameters would not be zero but only close to zero; hence, such a person may indeed behave unselfishly in some other situations. The first interpretation of degenerated parameter values implies that heterogeneity is categorical in nature; the second assumes its nature to be gradual. Which is more appropriate?

To investigate the nature of heterogeneity in social choice, we use behavioral data and, in addition, a process measure, namely, RT. Our starting premise is as follows: A theory that aims to accommodate all heterogeneity and assumes the same set of parameters (although not parameter values) and the same process (e.g., calculations) for all respondents implies, *ceteris paribus*, relatively homogeneous response time patterns. Differences only arise to the extent that people differ in terms of processing speed but not in terms of different decision processes. If, however, observed RTs prove to be systematically different across people and games, then such heterogeneity suggests that across people, different processes are at work.

The issue of how to model heterogeneity is of importance far beyond research on social games. Take, for illustration, the modeling of individuals' risky choice between monetary gambles. Expected utility theory assumes heterogeneity to be a matter of degree, not kind, and recruits families of utility functions (such as constant absolute or relative risk aversion) to describe it. Similarly, cumulative prospect theory (Tversky & Kahneman, 1992) captures heterogeneity in terms of

different parameter values of its value and probability weighting functions (Glöckner & Pachur, 2012). Applying a finite mixture model (McLachlan & Peel, 2000) to choices in monetary gambles, Bruhin, Fehr-Duda, and Epper (2010), however, observed two distinct types of decision makers. One group's parameter values represented rational and risk-neutral choices; the other group's values revealed loss aversion and an S-shaped probability weighting. Even though there is heterogeneity within the second group (see, e.g., also Booij & van de Kuilen, 2009), their decision process is qualitatively different from that of members of the rational, risk-neutral group (e.g., overweighting of rare events and underweighting of common events versus linear treatment of probabilities).

Also, investigating risky choices, Brandstätter, Gigerenzer, and Hertwig (2006, 2008) turned to models of lexicographic heuristics. Within this framework, heterogeneity can also be described either parametrically (e.g., aspiration levels can be parameterized; see Rieskamp, 2008) or in terms of distinct heuristics (e.g., heuristics that completely ignore probability information such as minimax or those that sequentially process outcome and probability information such as the priority heuristic; see Brandstätter et al., 2006). Regarding the latter approach to modeling heterogeneity, it appears fair to conclude: Although the evidence for specific heuristics is often controversially debated (see, e.g., the extensive work on the priority heuristic; Ayal & Hochman, 2009; Birnbaum & Gutierrez, 2007; Birnbaum & LaCroix, 2008; Cokely & Kelley, 2009; Fiedler, 2010; Glöckner & Herbold, 2011; Johnson, Schulte-Mecklenbeck, & Willemsen, 2008; Rieger & Wang, 2008; Rieskamp, 2008 but also Brandstätter & Gussmack, 2012; Arieli, Ben-Ami, & Rubinstein, 2011), a widely accepted view in psychological research on behavioral decision making is that different people recruit different cognitive strategies to respond to the same task.

RESPONSE TIME: A WINDOW ON THE UNDERLYING PROCESS

Unlike in economics, in psychology, RT has often been employed as a measure of human behavior in its own right. Models of, for example, human categorization—how people assign a set of stimuli to a number of different groups or concepts—are often required to predict not only behavior but also the RTs (Lafond, Lacouture, & Cohen, 2009; Lamberts, 2000). Temporal dynamics have also played a crucial role in the long-lasting debate in psychology on the serial or parallel nature of information processing (Townsend, 1990). Moreover, in investigations of human decision making, RT has regularly been used to test a choice or inference model's predictions (e.g., Pachur & Hertwig, 2006) or to discriminate between the predictions of competing models (e.g., Brandstätter et al., 2006; Bröder & Gaissmaier, 2007; Payne, Bettman, & Johnson, 1993).

Bergert and Nosofsky (2007), for example, tested RT predictions of a lexicographic inference heuristic, the take-the-best heuristic (Gigerenzer & Goldstein, 1996), against the predictions of a weighted-additive model, embodying

rational decision making. In making a decision about which of two alternatives scores higher on some quantitative dimension (e.g., which of two companies has higher annual revenues), the take-the-best heuristic considers the properties (cues) of the alternatives (e.g., information about number of employees, the type of industry) in order of their diagnosticity (cue validity) and arrives at an inference based on the first one that distinguishes between the alternatives. The weighted-additive model, in contrast, assigns to each cue a weight, calculates the summed total evidence in favor of each alternative, and chooses the alternative with more total evidence. The models—one embodying restricted search, the other exhaustive search—imply distinct RT patterns, depending on the choices in question (only in the most extreme case, when merely the lowest-ranked cue distinguishes, do the models' RT predictions converge). As Bergert and Nosofsky showed, generalizations of both models yield formally identical predictions of choice probabilities but “embody dramatically different decision making processes” (p. 116). They, in turn, imply very different RTs, with take-the-best terminating search once a cue inspected discriminates and the weighted-additive model always inspecting all cues. The observed RT pattern pointed toward a search pattern consistent with the search-restricting take-the-best heuristic.

Economists have only recently begun to take advantage of RT as a window on cognitive processes. Rubinstein (2007), for example, proposed distinguishing between *instinctive* and *cognitive* choices, and used RT to investigate which decisions, for example, in the beauty contest game, require cognitive effort and which require no or little cognitive effort (thus counting as instinctive in Rubinstein's view). Recording RTs in the context of the ultimatum game, Brañas-Garza, León-Mejía, and Miller (2007) found that selfish responders decide faster. Relatedly, Piovesan and Wengström (2009) found that selfish dictators decide faster and that selfish decisions are made more swiftly.

The latter two investigations are related to ours insofar as they classify respondents on the basis of their overt decisions and then recorded RTs. We, however, go beyond the obvious distinction between selfish and non-selfish players. To preview, we collect choices and RTs in a series of mini-ultimatum games, and using a finite mixture model analysis, we classify participants according to their choices. The classification suggests four distinct types of participants, who differ in whether and how social motives matter for their decisions. Importantly, this heterogeneity maps onto systematic RT pattern: First, we find distinct behavioral classes to be reflected in distinct average RT. Second, RTs relate to the interaction of participant types and games. This becomes most obvious when comparing participants who appear to act on the basis of social motives to those without. Regardless of the details of a specific offer, the strictly selfish respondents can quickly accept it (given our games), with no RT difference between games. In contrast, a person who takes social motives into account will have to appraise a given offer in light of these motives. This takes time, and because different games evoke these motives to different degrees, *game-specific RTs* result. We indeed find distinct patterns of RTs between behavioral classes and games, suggesting that

qualitatively different processes are at work. Interestingly, however, some of these distinct patterns of RT could, in principle, also be predicted by social preference model. In what follows, we show how and contrast those predictions to those derived from models of heuristics.

RESPONSE TIME PREDICTIONS BASED ON SOCIAL PREFERENCE THEORIES AND LEXICOGRAPHIC HEURISTICS

We begin with a disclaimer: The aim of our investigation is not to test between different social preference theories. Rather, we will assess the relative importance of the postulated social motives on an individual level. Moreover, we will examine to what extent the social preference theories and a heuristic processing account can accommodate the observed RTs. We use the following terminology when analyzing rejection behavior. The payoffs for the proposer (P) and responder (R) in the chosen (c) and the forgone (f) allocation are as follows: P_c , R_c , P_f , and R_f . We assume that all payoffs are positive but refrain from further assumptions. In particular, we do not assume that the sum of the payoffs is constant (as in the standard ultimatum game). Let us now discern between three social motives, each one could cause rejection of a positive payoff.

Assuming rationality and selfish preferences, no offer in the mini-ultimatum games will ever be rejected as long as the responder's payoff (R_c) is positive. In the present games, R_c payoffs were always positive. Starting with Güth et al. (1982), however, rejections of low offers have been often demonstrated. Various social preference theories have been proposed that recruit different social motives causing such rejections. Arguably, the most prominent motive for rejection is inequity aversion as invoked in Bolton and Ockenfels (2000) and Fehr and Schmidt's (1999) models.

Inequity aversion

Applied to the mini-ultimatum games, inequity aversion means that if the responder fares worse than the proposer, one will have a reason to reject the allocation offered. Consequently, there is a trade-off between the responder's disutility resulting from inequity and the utility resulting from the absolute payoff. For people who behave according to inequity aversion, only the two payoffs of the chosen alternative are relevant, that is, R_c and P_c . We say that an allocation is *unfair* (i.e., inequitable) if the responder receives less than the proposer and that an offer will be rejected if it is unfair (i.e., $R_c - P_c < 0$).

Kindness

Lack of equity or fairness, however, appears to be not the only reason motivating rejection. Brandts and Sola (2001) and Falk et al. (2003) showed that the rejection rate of an offer also depends on the counterfactual allocation (i.e., the one that could have been chosen but was not). Taking it into account, a responder can evaluate an offer's kindness.

Kindness is a social motive that is featured prominently in Dufwenberg and Kirchsteiger's (2004) and Rabin's (1993) reciprocity theories. An offer is said to be *unkind* if it is smaller than the counterfactual allocation, and an offer will be rejected if it is unkind (i.e., $R_c - R_f < 0$).¹ By rejection, the responder can reduce the proposer's payoff, and this act can be a term in the responder's utility function.²

Mirror

As a third criterion for the acceptance or rejection of an offer, we investigate the possibility that the offer is evaluated by taking the role of the proposer and probing oneself as to whether one would have deemed the allocation acceptable and thus also proposed it. We do not suggest a normative measure. We take an empirical approach and assume that people differ in their norm of what is an (still) acceptable division of resources (see López-Pérez, 2008, for a formal theory of such an approach). They themselves abide by the norm and punish others who violate it. Stimulating the possible dilemma the proposer faces is tantamount to taking an honest look in the mirror and asking oneself: What would I do? Applied to the mini-ultimatum game, this social motive predicts that responders consider an offer as more acceptable if they also would have made it.

How can one find out what a person would do to others? Naturally, a person's choices as a proposer could reveal one's personal norm. Therefore, all our participants played both roles, that of the responder and proposer. The proposer, however, is likely to reckon with the fact that an offer could be rejected. So, in the example earlier, a person could consider the 800:200 offer as perfectly acceptable but nevertheless refrain from proposing it for fear of rejection. Hence, we expected choices in the dictator game might be better proxies of a person's norm of acceptability and therefore asked people also to respond in the same games as a dictator. To conclude, we expect that a person is more likely to reject an offer if one would not make that offer oneself (in an ultimatum or dictator game).

In sum, we distinguish between three non-exclusive social motives for rejections: unfairness, unkindness, and mirror (i.e., I would not have proposed this allocation). If a responder has any or all of these concerns, rejections will become more likely, relative to a responder who derives one's choice exclusively from one's narrow self-interest. On the basis of these social motives, we can now derive RT predictions for social preference theories and for lexicographic heuristics.

¹A *kind* offer refers to the kinder of the possible offers. The kind offer is thus a relative construct and not necessarily kind in an absolute sense. For example, choosing 800:800 over 200:500 is kinder but not in essence kind.

²Strictly, the theories do not compare the allocations but the expected allocations taking into account the responder's second-order beliefs about one's rejection. There are two problems with equilibria based on this assumption: First, there are many equilibria, thus making it nearly impossible to test the theories. Second, some of these equilibria are rather implausible. For example, an offer could be rejected because the responder thinks the proposer expects a rejection and, therefore, by choosing this offer, intends a payoff of zero for the responder. However, why would the proposer make an offer with the intention of having it rejected?

Response times and social preference models

Economic models have little to no pretension of making predictions about processes. Notwithstanding this fact, one can derive—based on the auxiliary assumption that increasing complexity requires more time (e.g., Payne et al., 1993)—the following two qualitative RT predictions. First, if responders invoke a utility calculation to evaluate the proposed split, and if they merely differ in how they weight their motives (self-interest vs. social motives) that are arguments in this function, then all people should display the same RT or at least the same pattern of RTs. Second, if for some people only a single motive matters (e.g., self-interest), their calculation will be simpler, relative to others who need to trade-off self-interest and social motives. Consequently, the former group's RTs should, on average, be shorter than the latter ones.

Let us illustrate what these predictions mean using Fehr and Schmidt's (1999) theory of inequity aversion. According to this theory, people experience, next to satisfaction with their material payoffs, disutility from advantageous and from disadvantageous social inequality. These two concepts are included in terms of two arguments in a more complicated (relative to a *neoclassical* utility function) social utility function, as described in Equation (1):

$$U_i = \pi_i - \frac{\alpha_i}{N-1} \sum_j \max[\pi_j - \pi_i, 0] - \frac{\beta_i}{N-1} \sum_j \max[\pi_i - \pi_j, 0] \quad (1)$$

Disutility is subtracted from the utility stemming from the material payoff π_i . Disadvantageous inequality has a weight of α_i , and advantageous inequality has a weight of β_i . Both parameters are assumed to be nonnegative. Finally, N denotes the number of players (two in our case).

Applied to our mini-ultimatum games, three behavioral predictions follow: First, if $\beta_i < 1$ (this assumption is made in the model), then no offer with advantageous inequality will be rejected. Thus, advantageous inequality is not relevant for the rejection decision. Second, an offer (P, R) involving disadvantageous inequality will be rejected if $R/(P+R) < \alpha_R/(1+2\alpha_R)$. Third, the response to an offer depends only on the offer made, regardless of the alternative. How do these behavioral predictions translate into RT predictions? A naïve initial assumption would be that complexity and, by extension, RTs are identical across responders. There is, however, heterogeneity in the parameters, and some heterogeneity could affect RTs. Specifically, for those responders with a parameter value of zero, the corresponding argument in the complex utility function can be canceled. On a process level, the calculations therefore become simpler and RTs, assuming serial processing, faster. That is, should one or both weighting parameters equal zero, one could argue that people behave according to a simpler variant of the complex utility function.

Assuming that in Fehr and Schmidt's (1999) theory, disadvantageous disutility needs to be considered only when disadvantageous inequality occurs, we can derive a simple RT prediction. To the extent that it can be swiftly determined whether an offer implies advantageous or disadvantageous inequality, then participants with $\alpha_i > 0$ can be expected to

have a longer RT for offers with disadvantageous inequality than participants with $\alpha_i = 0$.

Response times and lexicographic heuristics

There is an alternative to the assumption that individuals maximize a more complicated utility function that includes social motives. On this view, social motives are ordered according to, for example, their importance and players process them sequentially, making a decision as soon as a decision criterion is met. We briefly describe this modeling approach in terms of heuristics (Gigerenzer, Hertwig, & Pachur, 2011). There are several classes of heuristics, one obvious candidate for two-alternative choice problems (i.e., rejecting vs. accepting an allocation) being lexicographic rules. These rules order decision dimensions (e.g., aspects, reasons, cues, motives) according to some criterion, search through $m \geq 1$ dimensions, and render an immediate decision based on the first dimension that discriminates. Two prominent examples of lexicographic decision rules are elimination-by-aspects (Tversky, 1972) and the take-the-best heuristic (Gigerenzer & Goldstein, 1996).

Because very little is known about what kind of heuristics people may use—if they use them at all—in the context of mini-ultimatum games or ultimatum games in general, we do not estimate heuristic models in a behavioral econometric analysis. We take a different approach. Specifically, we use people's choices to classify them into different types of responders. These types, in turn, suggest possible lexicographic heuristics that responders could be using. Equally important, even the generic forms of the heuristics give rise to qualitative response time predictions. We then use empirical RT to test the predictions of the generic lexicographic heuristics and those suggested by the social preferences models. We assume that the heuristics process one or more of the three social motives identified before (unfairness, unkindness, and mirror), except that they, unlike the social preference models, process the different motives sequentially. For illustration, consider the following *generic* heuristic consisting of three motives:

- Step 1. If a responder receives more than the proposer, one will accept the offer without further ado.
- Step 2. If one receives less than the proposer, one will consider whether the offer was nevertheless kind and will accept it if it proves kind.
- Step 3. If the offer turned out to be unkind, the responder probes oneself: "Would I have made this offer?" If the answer is yes, one will accept the proposed allocation, otherwise one will reject it.

This generic lexicographic heuristic assumes that strong social motives (i.e., those that imply low behavioral variance, and in which, in the extreme case, all offers that pass or fail a test are accepted or rejected, respectively) have priority over weaker motives. Specifically, let us suggest that inequality aversion precedes kindness, and kindness precedes the kind of perspective taking as embodied in the mirror criterion. From this ensues, for example, the following key prediction: One and the same person will respond swifter when allocations meet the inequality aversion test relative to those allocations failing this test and the kindness test and that are,

Table 1. The 12 mini-ultimatum games investigated

Unkind offer		Kind offer		Choice of kind offer	
P proposer	R responder	P proposer	R responder	Dictator game	Ultimatum game
800	200	800	500	0.94	1.00
800	200	600	400	0.31	0.69
800	200	500	500	0.37	0.69
800	200	500	300	0.14	0.54
800	200	400	600	0.16	0.51
800	200	200	800	0.03	0.23
800	200	200	200	0.11	0.23
400	500	100	600	0.01	0.01
200	500	800	800	1.00	0.99
200	500	400	800	0.99	0.94
200	500	100	800	0.13	0.14
200	200	800	400	0.96	0.84

Note: The two rightmost columns report the proportion of participants (as proposers) who chose the kind offer in the dictator and ultimatum games.

consequently, decided on the basis of the mirror test. In what follows, we will examine whether there are any indications that people's choices and RT follow such a lexicographic processing style.

METHOD

Participants played 12 mini-ultimatum games. Table 1 lists the 12 mini-ultimatum games. Each line represents one game. For example, in the second game, the proposer could choose between an allocation of (800:200) or (600:400), where the first number refers to the proposer's payoff and the second number to the responder's payoff. The two columns to the left display the unkind offer; that is, the offer in which the responder would receive the lower payoff relative to the alternative.

We used the strategy method (Selten, 1967) to collect data for all potential outcomes from all participants. Responders had to decide for both allocations whether they would accept it, before they learnt which allocation was actually chosen by the proposer. To render possible the collection of RT data, we did not use the standard strategy method. Instead, we presented the two distributive variants of a game separately, amounting to 24 allocations, presented in random order. So in the example earlier, participants decided to accept or reject the 800:200 offer *given* the 600:400 alternative. Similarly, they decided to accept or reject the 600:400 offer given the 800:200 alternative. All players played the games in both roles. Half of the participants assumed first the role of the proposer and then switched to the responder role; for the other half, the order was reversed. In addition, all participants played the role of a proposer in dictator games based on the same allocations as in the mini-ultimatum games (i.e., they determined the allocation without the responder having a chance to reject it). Participants received written instructions explaining the scenarios (see Appendix). They answered two control questions testing their basic understanding of the procedure.

Material

The games in Table 1 were selected in such a way that previously suggested social motives for rejecting the proposed allocation were represented across the games. For 11 out of 24 situations, the responder receives less in the chosen allocation, relative to the forgone alternative. In two situations, these payoffs are tied. Moreover, relative to the proposer's payoff, the responder fares worse in 11 cases, and in four cases, they are tied.

We focus on the responder's decisions and describe the steps involved in these decisions in more detail. Players were presented with a distributive situation (Figure 1) consisting of four payoffs: the payoffs for the proposer and responder in the chosen allocation (P_c , R_c) and in the alternative, foregone, allocation (P_f , R_f). As shown in Figure 1, payoffs were presented in a table, with one dimension representing the allocations and the other dimension the recipients of the payoffs. The orientation of the two dimensions (rows vs. columns) varied across participants. Responders rejected or accepted an offer by clicking a button on the screen.

We measured RT from the moment the allocation appeared on the screen to the moment the responder pressed the button. RT was recorded when the server received the message that the button had been pressed, which meant that there was some delay in recording RT.³ Once all decisions were made, one mini-ultimatum game and one dictator game were randomly selected. Participants were paid according to their proposer and responder's decisions in these games (four in total). Each person was paired with different participants in these four games. We paid people for a random subset of games to encourage them to treat each decision as if it were a response to a one-shot game and to discourage participants from forming a meta-response policy to the set of games. Paying a few decisions ensures that the unfairness of a single decision becomes manifest in the final payoff. One hundred points were worth 2.50 Swiss francs (CHF, about 2.15\$) in the ultimatum

³This delay, however, was the same across all participants, and because they could continue with their task, it is not correlated with the decisions.

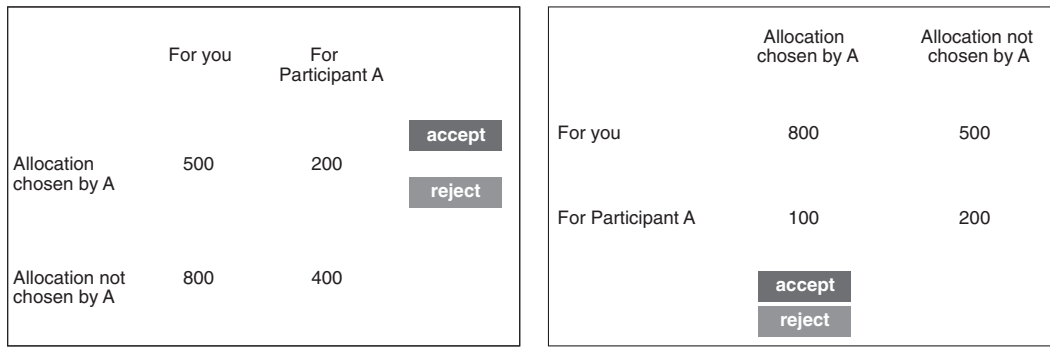


Figure 1. Two orientations of the screen for responder's (player B) decision used in the study. The empty border of the screen is cropped

games and 0.50 CHF (about 0.43\$) in the dictator games. Participants earned on average 39.65 CHF (about \$34 at the time of the experiment), including a showup fee of 10 CHF (about \$8.58).

Participants

Participants were 70 students from the Swiss Federal Institute of Technology, Zurich, and from the University of Zurich. They were recruited with an online recruitment system (Greiner, 2004). The sessions lasted between 60 and 75 minutes and were conducted in a computer lab. The experiment was programmed and conducted with z-Tree (Fischbacher, 2007).

RESULTS

Across responders, we observed heterogeneous patterns of acceptances and rejections in the 24 allocation decisions. To detect commonalities and differences in these patterns, we conducted a finite mixture model analysis. In contrast to standard regression techniques, estimating behavior only at the aggregate level, finite mixture models identify types of behavior and produce type-specific behavioral parameter estimates. On the basis of type-specific sets of parameter estimates, one can classify an individual into the type that best fits their choices (here: their rejection pattern).

Table 2 summarizes aggregate behavior by showing the results of a probit regression for responders' acceptance decision. In this analysis, we use the fairness and kindness criteria to analyze whether the responders chose this allocation in the ultimatum game and in the dictator game (i.e., the mirror criterion). The proportion of participants who chose the kind (and, by extension, the proportion who selected the unkind offer) in both games is reported in Table 1. The regression reveals that on the aggregate level, that is, across all participants, each criterion contributes to the decision, except the person's proposer behavior in the ultimatum game. The latter result confirmed our expectation that ultimatum game behavior is intertwined with strategic concerns. To find out which of the criteria matter most for whom, we analyzed type-specific rejection patterns in a finite mixture

model analysis. In doing so, we assume that there is latent heterogeneity, meaning the population consists of a finite mixture of *C* different behavioral types. However, ex ante individual-type membership is unobservable. Thus, the finite mixture model's log likelihood,

$$\ln L(\Theta; X) = \sum_{i=1}^N \ln \sum_{c=1}^C \pi_c f(\theta_c; x_i)$$

weights the individual type-specific likelihood contributions *f* ($\theta_c; x_i$)—here, the densities of a logit regression with type-specific parameters θ_c —by the probabilities π_c that individual *i* belongs to type *c*. Because we cannot observe type membership directly, these probabilities are the same for all individuals ex ante and correspond to the unknown share π_c of type *c* among the population. Maximizing $\ln L(\Theta; X)$ yields the maximum likelihood estimates for the type-specific behavioral parameters $\hat{\theta}_c$ and the corresponding relative type sizes $\hat{\pi}_c$.⁴

Once the types are characterized by their specific behavioral parameters and their relative sizes, we can classify each individual into the type that best corresponds to their rejection pattern. Formally, Bayes' rule defines the ex post probabilities,

$$\tau_{ic} = \frac{\hat{\pi}_c f(\hat{\theta}_c; x_i)}{\sum_{m=1}^C \hat{\pi}_m f(\hat{\theta}_m; x_i)}$$

that individual *i* is of type *c* given the type-specific characteristics $\hat{\theta}_c$ and $\hat{\pi}_c$. We classify each individual into the type associated with the highest individual ex post probability τ_{ic} .

Types based on the finite mixture analysis

Table 3 summarizes the results of the finite mixture estimation. It shows the distinct behavioral parameters and the number of

⁴The finite mixture model's log likelihood is highly nonlinear, even after taking logs. Furthermore, it is well known that the model's likelihood usually exhibits several local maxima and may even be unbounded. These properties render numerical maximization difficult (for more details see McLachlan & Peel, 2000). To overcome such problems, we employed a version of Demster, Laird, and Rubin's (1977) expectation-maximization algorithm for maximizing the model's likelihood.

Table 2. Probit regression of rejections

Dependent variable: reject	(1)	(2)	(3)	(4)	(5)
Unfair	0.832 (0.00)				0.874 (0.00)
Unkind		0.815 (0.00)			0.586 (0.00)
Not chosen in dictator game			0.638 (0.000)		0.616 (0.00)
Not chosen in ultimatum game				0.069 (0.65)	0.114 (0.54)
Constant	-1.707 (0.00)	-1.695 (0.00)	-1.600 (0.00)	-1.255 (0.00)	-2.497 (0.00)
Observations	1680	1680	1680	1680	1680
Pseudo r^2	0.0805	0.0777	0.0483	0.0006	0.1780
χ^2	27.63	53.31	32.00	0.209	115.5
Number of clusters	70	70	70	70	70

Note. *p*-values (in parentheses) are based on robust standard errors with clusters for participants in parenthesis.

Table 3. Result of finite mixture model (the model's dependent variable is rejection)

Classes	C ₀	C ₁	C ₂	C ₃
N	30	9	19	12
(Const)	—	-8.19 (0.44)	-18.59 (0.00)	-2.55 (0.61)
Unfair	—	6.44 (0.49)	15.58 (0.00)	0.15 (0.98)
Unkind	—	3.75 (0.46)	2.05 (0.08)	0.82 (0.61)
Mirror ^a	—	1.30 (0.76)	0.38 (0.49)	1.80 (0.01)

Note. *p*-values (in parentheses) are based on robust standard errors with clusters for participants.

^aBased on the not-chosen allocations in the dictator game.

individuals. Our analysis identified types C₁, C₂, and C₃.⁵ In this classification, in which each person is assigned to the class for which the person's posterior probability of type membership is highest, 90% of participants' membership probability exceeded 90%. The additional type C₀ is made up of 30 individuals, accepting all allocations. We refer to this group as the *always-accept type*. We suspect that the driving force behind their behavior is pure self-interest but, admittedly, other interpretations are possible (see Discussion section). This type had to be excluded from the finite mixture estimation.

What characterizes the other classes? In class C₁, the concerns of unfairness and unkindness proved most important. This result is not statistically significant, which is most likely

⁵We used the normalized entropy criterion (NEC) to determine the optimal number of behavioral types. Classical model selection criteria such as the Akaike information criterion or the Bayesian information criterion are designed to trade-off parsimony against goodness of fit, but they do not directly assess the ability of the finite mixture model to provide well-separated and non-overlapping behavioral types. Consequently, they risk to favor finite mixture specifications that incorporate too many behavioral types and to overfit the data. To remedy this problem, Celeux and Soromenho (1996) proposed using the NEC. It is based on the posterior probabilities of type membership and directly reflects the model's ability to provide a classification into clearly separated behavioral types. In our data, NEC favors a specification with three behavioral types (classes C₁ to C₃) for the participants with nonzero regression rates.

due to the few people in this class ($n=9$). In class C₂, again unfairness and unkindness were key. This group of people invariably accepts all *fair* allocations, and hence, the approximated coefficient is large. In class C₃, unfairness and unkindness lose power, and the mirror criterion appears to be the only one that captures people's behavior.

Table 4 lists the rejection rates obtained in classes C₀ to C₃, separately for each of the 24 distributive situations. Our statistical classification reflects quite distinct rejection patterns. We ordered the classes according to the average variance per game, that is, according to how well behavior in each class was defined, with rejection rates approximating 0% or 100%, respectively. There are large deviations from these endpoints, especially in classes C₂ and C₃, which imply that there is—not surprisingly—still unexplained heterogeneity.

Response time across types

On the basis of this analysis of behavioral heterogeneity, we can now turn to RT. We first report on the statistical properties of RT. Figure 2 shows the RT distribution and the logarithm (base 10) of RT. The RT distribution is strongly right skewed and deviates from the normal distribution. The distribution of the logarithm of RT time appears to follow a normal distribution.⁶ For this reason, we use the logarithm in the following analyses. More specifically, we report the logarithm with base 10 in Figure 2 (right panel), because it is easier to translate into the actual time, but use the natural logarithm in the regressions because a 1% change is approximated by a 0.01 difference in the natural logarithm.⁷ When we report averages, we report geometric means instead of

⁶A Shapiro–Wilk *W* Test rejects normality of $\ln(\text{RT})$ with a *p*-value of 2.2%. However, the data are not independent and individual heterogeneity could cause this result. For this reason, we applied a Shapiro–Wilk *W* Test to all 24 distributive situations, and there was no case with significance between 5% and 10%. Thus, it seems reasonable to accept normality—in particular when we control for the situation.

⁷Because $\ln(x) = \ln(10)\log_{10}(x)$, the two logarithms differ only by a factor and all statistical analyses are equivalent using either of them.

Table 4. Average rejection rate (%) in the four types of responders (C₀ to C₃) across the 24 distributive variants

Payoffs				Criterion ^a		Rejection rate (%)				Mean C ₁₋₄
P _c	R _c	P _f	R _f	Unfair	Unkind	C ₀	C ₁	C ₂	C ₃	
400	500	100	600	0	1	0	0	0	0	0
800	800	200	500	0	0	0	0	0	0	0
400	800	200	500	0	0	0	0	0	0	0
800	500	800	200	1	0	0	0	5	0	1
500	500	800	200	0	0	0	0	0	8	1
500	300	800	200	1	0	0	0	0	17	3
600	400	800	200	1	0	0	0	5	8	3
200	500	100	800	0	1	0	11	0	8	3
400	600	800	200	0	0	0	0	0	17	3
800	400	200	200	1	0	0	22	5	0	4
200	800	800	200	0	0	0	0	0	25	4
100	800	200	500	0	0	0	0	0	25	4
200	500	800	800	0	1	0	0	0	33	6
200	500	400	800	0	1	0	0	0	42	7
200	200	800	200	0	0	0	0	0	50	9
100	600	400	500	0	0	0	0	0	50	9
200	200	800	400	0	1	0	0	0	67	11
800	200	200	200	1	0	0	67	16	25	17
800	200	500	300	1	1	0	100	32	25	26
800	200	200	800	1	1	0	89	37	42	29
800	200	600	400	1	1	0	89	32	50	29
800	200	400	600	1	1	0	100	26	67	31
800	200	500	500	1	1	0	89	37	67	33
800	200	800	500	1	1	0	100	47	50	34

The columns *unfair* and *unkind* show how the games were classified (using a binary classification). P_c and R_c refer to the payoffs for the proposer and responder, respectively; P_f and R_f denote the foregone allocation.

^aA value of 1 equals an unfair (unkind) allocation, respectively; a value of 0 equals a fair (kind) allocation, respectively.




Figure 2. Histogram of response time (in seconds, left panel) and of base 10 logarithm of response time (right panel). Data in the left panel leave out 14 of 1680 observations because they exceeded 20 seconds (the maximum equals 56)

the arithmetic mean because the former are closer to the median in the current distribution.

We begin with a general observation. With experience, RT decreases considerably. For illustration, the average RTs for the first and last decisions were 8.7 versus 2.4 seconds, respectively. The decline in RT is rather monotonous and particularly pronounced across the first four decisions. Because the decisions appeared in random order, a control for experience is not necessary. However, we also tested all the results for stability with respect to experience. Of course, including experience considerably increases explanatory power of the regression reported later in Tables 5–7.

Do RTs differ across classes C₀–C₃? Let us first analyze the RT predictions suggested by Fehr and Schmidt’s (1999) theory. The first is that always-accept responders have shorter RTs because they can do without the inequity

Table 5. Ordinary least squares regressions with ln(RT) as dependent variable for the four classes

	C ₀	C ₁	C ₂	C ₃
Unfair	0.045 (0.49)	0.303 (0.00)	0.209 (0.00)	0.070 (0.3)
Unkind	0.047 (0.3)	0.017 (0.8)	0.130 (0.00)	0.071 (0.45)
Mirror ^a	0.168 (0.002)	0.021 (0.76)	0.134 (0.001)	0.147 (0.02)
Constant	0.813 (0.00)	1.102 (0.00)	1.195 (0.00)	1.506 (0.00)
Observations	720	216	456	288
r ²	0.014	0.072	0.058	0.024
Number of clusters	30	9	19	12

Note. p-values (in parentheses) are based on robust standard errors with clusters for participants.

^aBased on the not-chosen allocations in the dictator game.

Table 6. Ordinary least squares regressions with ln(RT) as dependent variable

Selection	(1)	(2)	(3)	(4)	(5)	(6)
	C ₁ unfair	C ₁ fair	C ₂ unfair	C ₂ fair	C ₁ , C ₂ unfair	C ₁ , C ₂ fair
Unkind	0.075 (0.45)	-0.023 (0.84)	0.223 (0.00)	0.110 (0.03)	0.175 (0.00)	0.068 (0.17)
Constant	1.385 (0.00)	1.128 (0.00)	1.412 (0.00)	1.277 (0.00)	1.403 (0.00)	1.229 (0.00)
Observations	99	117	209	247	308	364
r ²	0.00445	0.000415	0.0330	0.00800	0.0213	0.00311
Number of clusters	9	9	19	19	28	28

Note. *p*-values (in parentheses) are based on robust standard errors with clusters for participants.

Table 7. Ordinary least squares regression of ln(RT)

	(1)	(2)	(3)	(4)
Rejected	0.471 (0.00)	0.241 (0.00)	0.410 (0.00)	0.131 (0.03)
C ₁		0.253 (0.05)		0.284 (0.03)
C ₂		0.454 (0.00)		0.465 (0.00)
C ₃		0.638 (0.00)		0.669 (0.00)
Unfair			0.066 (0.06)	0.106 (0.00)
Unkind			0.024 (0.43)	0.052 (0.07)
Mirror ^a			0.098 (0.00)	0.129 (0.00)
Constant	1.179 (0.00)	0.939 (0.00)	1.095 (0.00)	0.802 (0.00)
Observations	1680	1680	1680	1680
r ²	0.041	0.151	0.047	0.164
Adjusted r ²	0.0406	0.149	0.0449	0.160
F	41.94	16.51	21.74	20.68
Number of cluster	70	70	70	70

Note. *p*-values (in parentheses) based on robust standard errors with clusters for participants.

^aBased on the not-chosen allocations in the dictator game.

disutility terms (i.e., both their α and β parameters equal 0; see Equation (1)). Consistent with this prediction, we found that the average RT of 2.6 seconds in C₀ was shorter, relative to 3.5 (C₁), 4.1 (C₂), and 5.2 (C₃), respectively. All differences, with the exception of the difference between C₁ and C₂, are significant at the 5% level (pairwise comparison using regressions with clustering on participant level).

The average RTs suggest a difference between the classes (which were defined exclusively on the basis of behavior and not RT), and the RT differences echo the behavioral differences. In Table 5, we report regression of ln(RT) with our three decision criteria as regressors. First, according to the *r*² value, the game parameters contribute least to the RTs in class C₀. Because these players appear only to consider whether their payoff is positive (Table 4), they are indeed exempt from examining any other properties of the games. Consequently, they can decide without further ado. Interestingly, though, even this strictly self-interested group of players experiences a moment of soul-searching: For offers that these participants would not have made themselves (in the dictator game), it takes them slightly longer to accept the offer.

For classes C₁ and C₂, we find that unfairness plays an important role in the RT and, as predicted by Fehr and Schmidt's (1999) theory, RT is longer for the unfair situations than for the fair ones. The same prediction would also be made if fairness were the first criterion in a sequential decision process based on a lexicographic heuristic. Indeed, the unfairness criterion has the highest coefficient in both classes, and its negative value means that choices require less time if fairness is met. This, in turn, is consistent with a sequential process that (i) instantaneously accepts a fair allocation and (ii) defers to the next criterion (*kindness*) when an offer is unfair. In this sequential processing, kindness will only be called upon when an outcome is unfair. In Table 6, we report regressions in which we analyze whether RT depends on kindness, separately for fair and unfair offers. When offers were unfair, then further criteria would need to be evaluated, with kindness being the next obvious criterion. Consistent with sequential processing, the regressions show that when unfair allocations are *kind*, then RT is significantly shorter, relative to unfair and unkind allocations. Moreover, when offers are fair, the impact of kindness on RT is lower relative to unfair offers, consistent with sequential processing. In sum, behavior in classes C₁ and C₂ and the RT pattern are consistent with a step-by-step consideration of the fairness and kindness criteria.

What about people in C₃? For them, probing themselves in terms of the mirror criterion is the most important behavioral motive, and as the regression in Table 5 reveals, it is also the most important determinant for the RT. Rather than processing first the more basic concerns of unfairness and unkindness, these players appear to probe themselves and use their inferred proposer behavior as a guideline for their responder decision.⁸

As mentioned previously, the always-accept responders reached their decisions faster. Moreover, acceptance decisions in general were also reached quicker; that is, reaching the decision to reject took more time. In Table 7, we report regression of ln(RT) using a dummy for whether the offer was rejected or not and, additionally, we used our main explanatory variables and dummies for the different classes. This allows us to analyze whether it is rejection per se that takes time and/or whether

⁸Alternatively, as one reviewer suggested, participants in this group may act upon what they perceive to be the reasonable choice—for instance, in terms of *descriptive norms* (Cialdini & Goldstein, 2004)—and punish unreasonable choices (e.g., the choice of 200:800 over 400:800). We analyzed the acceptance decisions and found that this proxy of an objective descriptive norm does not offer much explanatory power.

the situations that commonly result in rejection require time. The regressions show that the different types of responders are the main reason why accepted decisions are quicker. Although the rejection dummy has a coefficient of 0.47 when no further explanatory variables are present in the regression, it is reduced to 0.24 when type dummies are introduced. Moreover, the parameters of the games also explain part of this difference. Introducing them further reduces the dummy to 0.13. This means that these two constructs—types of responders and kind of game—account for more than 70% of the explanatory power of the rejection dummy. This, in turn, suggests that it is not simply the decision to reject that explains why rejections take more time.

To conclude, we made three major observations based on RT. First, average RT differs across the four types of responders, identified by our statistical classification. Second, within two types, the C_1 and C_2 responders, RT differs as a function of properties of the game (e.g., unfair vs. fair games). Third, these game-specific differences in RTs suggest a sequential and lexicographic processing of decision criteria, with fairness being the most important concern, followed by kindness (see also Hertwig et al., 2013). Next, we discuss the implications of these findings.

DISCUSSION

We used a variant of the classic ultimatum game, the mini-ultimatum game, to examine the degree and kind of heterogeneity in responder's behavior across 24 allocation scenarios. The heterogeneity was substantial, with some responders categorically accepting any allocation and others rejecting up to 46%. To detect common types among responders, we conducted a finite mixture model analysis and cross-validated its classification with an RT analysis. We observed four groups of responders: one group (C_0) takes little to no account of the proposed split or the foregone allocation and swiftly accepts any positive offer. Two groups (C_1 and C_2) process in a sequential fashion primarily the overt and objective properties of the allocations (fairness and kindness) and take more time the more properties need to be scrutinized. A fourth (C_3), which takes more time, appears to take into account, what they would have proposed had they been put in the role of the proposer. In what follows, we discuss the types of responders, the limitations of our investigation, and the nature of heterogeneity in the games studied.

Types of responders: behavior and response time

About two-fifths (43%) of our participants consistently accepted all allocations. Their RT did not vary as a function of other properties of the allocations (such as whether they were fair or kind). One interpretation of this type of responder is that they act out of a purely selfish concern and swiftly accept any positive offer, no matter how fair or kind it is. Paying attention to this sole concern enables a very simple decision policy: In this domain, economic rationality can be implemented in terms of a one-reason decision-making rule (Gigerenzer et al., 2011). Admittedly, the motivation

to accept any positive offer, however, could also stem from other considerations, and in theory, even from a motive opposite to pure self-interest, namely, extreme generosity: "If the other player chooses to propose an allocation that clearly favors him, relative to me, I give him the benefit of the doubt and assume he needs it."

For another two-fifths of participants (40%), unfairness and unkindness mattered on the behavioral level and were reflected in distinct and systematic patterns of RT. A social preference model incorporating both social motives into the utility function could describe these behavioral results. The RT pattern, however, suggests that people use different processes in their decisions. Specifically, the RT patterns for these responders suggest that the very same responder takes more or less time as a function of the specific game, that is, the number of *tests* performed before accepting (or eventually rejecting) an allocation. Even if an allocation is ostensibly unfair, it may still prove to be kind. Examining its kindness after the fairness test has been failed, however, takes time (for details, see Hertwig et al., 2013). This kind of lexicographic and noncompensatory processing is difficult to reconcile with the existing social preference models.

Finally, a small portion of participants (17%) could not be modeled in terms of the established social and nonsocial motives (self-interest, fairness, and kindness). They appear to be most strongly influenced by whether they would have proposed the allocation had it been their turn to make the proposal and had they no fear of rejection. We found that the best proxy of this kind of projecting oneself into the role of the proposer is people's behavior (in C_3) in the dictator game.

Limitations of our investigation

Our implementation of the mini-ultimatum game is different from the classic implementation of the ultimatum game. We required a larger number of decisions per person to be able to compare RTs within the same person across games (and to have enough statistical power for the finite mixture analysis). The ultimatum game, in contrast, is often played only once. We cannot exclude the possibility that requiring people to make more than one decision prompted them to adopt sequential heuristics that they otherwise would not have recruited. Yet, let us point out that heuristics are commonly used to simplify and, consequently, speed up a decision process, and, perhaps, these benefits are indeed most welcome in a context involving repeated decisions. In our case, the sequential and noncompensatory heuristics we observed (in particular in C_1 and C_2) require more time and encompass more computation than the one-reason decision-making policy of the responders behaving economically rational (C_0). That is, to the extent that the repeated decisions encourage the use of simple heuristics, one would expect the simplest policy to become widely used. In our study, this policy coincided with rational and selfish behavior. Although a substantial portion of responders behaved in this way, more than half did not.

As we have demonstrated, social preference models—although often interpreted in terms of as-if models—can be used, supplemented by auxiliary assumptions, to derive

RT predictions. It is conceivable that the ability of these models to describe (fit) qualitative RTs would be further increased if one designed probabilistic versions of, say, the inequity aversion model. A probabilistic version could, for example, posit that the disutility terms are processed with probabilities proportional to the values of α (disadvantageous inequality) or β (advantageous inequality). Or a disutility term might only be considered if it exceeds a certain threshold, and this threshold may vary between participants in a parametric fashion. Clearly, at this point, our data cannot rule out that probabilistic alternatives can successfully fit systematic difference in RT across participants. However, such an increased mimicking (fitting) power (by further increasing the number of adjustable parameters) would be likely to come at the expense of loss of predictive power (an issue already of importance to the discrete version of the inequity aversion model; Blanco, Engelmann & Normann, 2011).

CONCLUSION

Taken together, the combination of choice and RT data suggests heterogeneity not only on the level of choice but also on the level of the underlying cognitive processes. A comprehensive theory of punishment behavior in the mini-ultimatum game and the ultimatum game needs to accommodate the differences in choice and process and needs to provide an explanation for both. Which kind of theory it will be is open to debate, and the three authors of this article (two economists and one psychologist) do not fully agree. One approach is to further refine social preferences models, by adding ever more terms into the utility function and by designing probabilistic versions thus accommodating systematic RT differences (see Discussion section). One major challenge to this modeling approach is robustness or, in other words, how predictive will the parameter values prove in predicting individual behavior in other games (see, e.g., Blanco et al., 2011). Another approach is to design sequential and noncompensatory models of heuristics, one model for each type of responder. Again the challenge is robustness. Will the same person use the same heuristic to make decisions in other related games, and how does a person decide what heuristic to select in case they recruit different ones. Although the authors of this article may not fully agree on which modeling approach to take, we are fully convinced that entertaining both approaches and having them compete promise to best advance our theoretical understanding.

APPENDIX A: INSTRUCTIONS FOR PARTICIPANT B

You are now participating in an economic experiment that has been funded by various research support foundations. The instructions that you have received serve exclusively as your private information. *During the experiment, communication is strictly prohibited.* If you have questions, please direct them to us. Failure to observe this rule will result in exclusion from the experiment and forfeit of all earnings.

You will participate in several experiments. We will now describe the procedure for the first experiment. In this experiment, there are participants in two roles; there is *participant A* and *participant B*. Each half of the participants is either in the role of participant A or in the role of participant B. *You are a participant B.*

In this experiment, you will be required to make decisions in *different situations*. In each situation, there will be one participant A and one participant B forming a group. At the end, one situation will be drawn for which the payoff is relevant. You and the participant A who is in the group with you will receive your payment according to your decisions in the relevant situation. You will not learn with whom you are in a group. Likewise, the other person will not learn that you are in a group with her.

In the instructions, we will speak of points. The following conversion is applied:

$$100 \text{ points} = 2.50 \text{ Swiss francs}$$

In the relevant situation, you will receive 2.5 cents per point. Additionally, you will receive a fixed sum of 10 Swiss francs.

All situations have the same structure, and the structure will be described in the following.

Structure of a situation

First, participant A will choose one of two point allocations, X and Y. This allocation determines how many points you and participant A receive.

For example: in allocation X, participant A receives 100 points and you receive 500 points; in allocation Y, participant A receives 200 points and you receive 300 points. As you see in this example, the sum of points in each allocation is not necessarily the same.

You can choose whether to accept or reject the decision of participant A. If you accept the decision, you and participant A receive the chosen point values; if not, you both receive nothing.

At the end of the experiments, one relevant situation will be drawn, for which you will be paid.

Detailed procedure

1. Decision of participant A

In this experiment, there is a series of situations. In all situations, participant A must choose one of the two point allocations, X or Y.

2. Your decision as participant B

As participant B, you will specify in all situations for both possible decisions of participant A whether you accept or reject the decision. You will make your decision before you know what decision A made. You will decide for each of these cases one after another.

The following decision situation is displayed on the screen that follows:

- Participant A had the choice between the allocations *400 for you and 100 for A* or *200 for both participants*.
- Participant A chose the allocation *400 for B and 100 for A*. The allocation chosen by A is always on the left, and

the one not chosen is always on the right. (It could also be that A chooses the other allocation. This also happens and in this case 200 for both would be on the left.)

	A's chosen allocation	A's non-chosen allocation
For you	400	200
For Participant A	100	200

You must now indicate whether you accept or reject this decision by pressing the corresponding button. After each entry, a continue button will appear which will bring you to the next entry. This will continue until you have decided for all situations.

3. The draw

At the end of the experiments, the relevant situation will be drawn. On the computer screen, all situations will be listed and shown with a number. The participant at seat 1 will roll a die to determine the *relevant situation*. The same situation will be drawn for all participants. After the draw, you and the participant A in your group will learn which allocation participant A chose in the drawn situation, whether you accepted or rejected the decision, and your income and the income of the participant A in your group.

Practice questions

- Participant A can choose between 100 points for herself herself and 300 points for B and 200 points for herself herself and 100 points for B. A chooses the first variant. Participant B accepts the decision.

What is A's income?
What is B's income?

- Participant A can choose between 100 points for herself herself and 300 points for B and 200 points for herself herself and 100 points for B. A chooses the second variant. Participant B rejects participant A's decision.

What is A's income?

What is B's income?

INSTRUCTIONS FOR PART 2

We will now go through a second experiment. It also consists of several situations like the first experiment and again there are participants A and participants B. *You are, however, in this*

experiment a participant A. The situations are structured the same as those in the first experiment. One situation will again be drawn and likewise will be paid with the conversion rate

1 point = 2.5 cents.

The draw for both experiments will take place at the end. At the end of the experiments, you will receive the 10 Swiss francs as an entry fee and the income from both experiments.

You will now be with a participant B in a group. In the relevant situation, you will be in a group with a person different from the one you were with in the relevant situation in the previous experiment.

As mentioned previously, the procedure in this experiment is as in the first experiment, only in this experiment you are a participant A. For this reason, a detailed explanation of how to make your entry as participant A will now follow.

YOUR DECISION AS PARTICIPANT A

In this experiment, there is again a series of situations. In all situations, you must choose one of two point allocations, X and Y. The two allocations will be displayed in a table on the screen as in the hypothetical example that follows. In the example represented, you would have to decide between the following allocations: In allocation X, you would receive 100 points and participant B 400 points; in allocation Y, you would both receive 200 points. You decide on an allocation by clicking on one of the two buttons under the allocations. For example, if you wanted to choose allocation X, you would click on the left button; if you wanted to choose allocation Y, check the right button. After each entry, a continue button appears that will bring you to the next entry. This will continue until you have decided for all situations.

	Allocation X	Allocation Y
For you	100	200
For Participant B	400	200

ADDITIONAL EXPERIMENT FOR YOU AS PARTICIPANT A

Because participant B has more decisions than participant A, you will take part in an additional experiment in which only you, as participant A, have a decision to make. The additional experiment is also concerned with allocation decisions between you and a participant B. This participant B is a different person from any you have been in a group with so far. In this additional experiment, participant B will have no possibility to reject. One situation from these decisions will also be drawn

and paid out. The conversion rate amounts to 50 cents per 100 points. You will be informed on the screen when these decisions begin.

When you have understood the instructions press *continue* on the screen.

ACKNOWLEDGEMENTS

We are grateful for the support from the Research Priority Program of the University of Zürich on “The Foundations of Human Social Behavior: Altruism Versus Egoism”. We also thank Gerald Eisenkopf, Simeon Schudy, and Verena Utikal and the participants of seminars in Maastricht, Jerusalem, Nürnberg, and Zürich for their helpful comments, and Laura Wiles for editing the manuscript.

REFERENCES

- Arieli, A., Ben-Ami, Y., & Rubinstein, A. (2011). Tracking decision makers under uncertainty. *American Economic Journal: Microeconomics*, 3, 68–76.
- Ayal, S., & Hochman, G. (2009). Ignorance or integration: The cognitive processes underlying choice behavior. *Journal of Behavioral Decision Making*, 22, 455–474.
- Berg, J., Dickhaut, J. W., & McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10, 122–142. doi: 10.1006/game.1995.1027
- Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 107–129. doi: 10.1037/0278-7393.33.1.107
- Birnbaum, M. H., & Gutierrez, R. J. (2007). Testing for intransitivity of preferences predicted by a lexicographic semi-order. *Organizational Behavior and Human Decision Processes*, 104, 97–112.
- Birnbaum, M. H., & LaCroix, A. R. (2008). Dimension integration: Testing models without trade-offs. *Organizational Behavior and Human Decision Processes*, 105, 122–133.
- Blanco, M., Engelmann, D., & Normann, H. T. (2011). A within-subject analysis of other-regarding preferences. *Games and Economic Behavior*, 72, 321–338.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90, 166–193. doi: 10.2307/117286
- Booij, A. S., & van de Kuilen, G. (2009). A parameter-free analysis of the utility of money for the general population under prospect theory. *Journal of Economic Psychology*, 30, 651–666.
- Brañas-Garza, P., León-Mejía, A., & Miller, L. M. (2007). Response time under monetary incentives: The ultimatum game. Jena Economic Research Papers in Economics, No. 2007-70. Friedrich-Schiller-University Jena, Max-Planck-Institute of Economics.
- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, 113, 409–432. doi: 10.1037/0033-295X.113.2.409
- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2008). Risky choice with heuristics: Reply to Birnbaum (2008), Johnson, Schulte-Mecklenbeck, & Willemsen (2008) and Rieger and Wang. *Psychological Review*, 115, 281–289. doi: 10.1037/0033-295X.115.1.281
- Brandstätter, E., & Gussmack, M. (2012). The cognitive processes underlying risky choice. *Journal of Behavioral Decision Making*. doi: 10.1002/bdm.1752
- Brandts, J., & Solà, C. (2001). Reference points and negative reciprocity in simple sequential games. *Games and Economic Behavior*, 36, 138–157. doi: 10.1006/game.2000.0818
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin and Review*, 1, 895–900.
- Bruhin, A., Fehr-Duda, H., & Epper, T. (2010). Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica*, 78(4), 1375–1412. doi: 10.3982/ECTA7139
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- Camerer, C. F., & Fehr, E. (2006). Measuring social norms using experimental games: A guide for social scientists. In J. Henrich, R. Boyd, S. Bowles, C. F. Camerer, E. Fehr, & H. Gintis (Eds.), *Foundations of human sociality: Experimental and ethnographic evidence from 15 small-scale societies* (pp. 55–95). New York: Oxford University Press.
- Celeux, G., & Soromenho, G. (1996). An entropy criterion for assessing the number of clusters in a mixture model. *Journal of Classification*, 13, 195–212. doi: 10.1007/BF01246098
- Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics*, 117, 817–869. doi: 10.1162/003355302760193904
- Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and Conformity. *Annual Review of Psychology*, 55, 591–621.
- Cokely, E. T., & Kelley, C. M. (2009). Cognitive abilities and superior decision making under risk: A protocol analysis and process model evaluation. *Judgment and Decision Making*, 4, 20–33.
- Cox, J. C., Friedman, D., & Gjerstad, S. (2007). A tractable model of reciprocity and fairness. *Games and Economic Behavior*, 59, 17–45. doi: 10.1016/j.geb.2006.05.001
- Cox, J. C., Friedman, D., & Sadiraj, V. (2008). Revealed altruism. *Econometrica*, 76, 31–69. doi: 10.1111/j.0012-9682.2008.00817.x
- Demster, A., Laird, N., & Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Ser. B*, 39, 1–38.
- Doris, J. M. (2002). *Lack of character: Personality and moral behavior*. New York: Cambridge University Press.
- Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, 47, 268–298. doi: 10.1016/j.geb.2003.06.003
- Engelmann, D., & Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, 94, 857–869. doi: 10.1257/0002828042002741
- Falk, A., Fehr, E., & Fischbacher, U. (2003). On the nature of fair behavior. *Economic Inquiry*, 41, 20–26. doi: 10.1093/ei/41.1.20
- Falk, A., Fehr, E., & Fischbacher, U. (2005). Driving forces behind informal sanctions. *Econometrica*, 73, 2017–2030. doi: 10.1111/j.1468-0262.2005.00644.x
- Falk, A., Fehr, E., & Fischbacher, U. (2008). Testing theories of fairness—Intentions matter. *Games and Economic Behavior*, 62, 287–303. doi: 10.1016/j.geb.2007.06.001
- Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, 54, 293–315. doi: 10.1016/j.geb.2005.03.001
- Fehr, E., & Gächter, S. (2000). Fairness and retaliation: The economics of reciprocity. *Journal of Economic Perspectives*, 14, 159–181. doi: 10.1257/jep.14.3.159
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, 114, 817–868. doi: 10.1162/003355399556151
- Fiedler, K. (2010). How to study cognitive decision algorithms: The case of the priority heuristic. *Judgment and Decision Making*, 5, 21–32.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10, 171–178. doi: 10.1007/s10683-006-9159-4
- Gigerenzer, G., & Goldstein, D. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650–669. doi: 10.1037/0033-295X.103.4.650

- Gigerenzer, G., Hertwig, R., & Pachur, T. (Eds.). (2011) Heuristics: The foundations of adaptive behavior. New York: Oxford University Press.
- Glöckner A., & Herbold A.-K. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes, *Journal of Behavioral Decision Making*, 24, 71–98.
- Glöckner, A., & Pachur, T. (2012). Cognitive models of risky choice: Parameter stability and predictive accuracy of prospect theory. *Cognition*, 123, 21–32.
- Greiner, B. (2004). An online recruitment system for economic experiments. In K. Kremer, & V. Macho (Eds.), *Forschung und wissenschaftliches Rechnen 2003* (pp. 79–93). GWDG Bericht 63. Göttingen, Germany: Gesellschaft für wissenschaftliche Datenverarbeitung.
- Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization*, 3, 367–388. doi: 10.1016/0167-2681(82)90011-7
- Güth, W., & Tietz, R. (1990). Ultimatum Bargaining Behavior – A Survey and Comparison of Experimental Results. *Journal of Economic Psychology*, 11, 417–449.
- Hertwig, R., Fischbacher, U., & Bruhin, A. (2013). Simple heuristics in a social game. In R. Hertwig, U. Hoffrage, & the ABC Research Group (Eds.), *Simple heuristics in a social world*. (pp. 39–65) New York: Oxford University Press.
- Johnson, E. J., Schulte-Mecklenbeck, M., & Willemsen, M. C. (2008). Process models deserve process data: Comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review*, 115, 263–272.
- Kirchsteiger, G. (1994). The role of envy in ultimatum games. *Journal of Economic Behavior and Organization*, 25, 373–389. doi: 10.1016/0167-2681(94)90106-6
- Lafond, D., Lacouture, Y., & Cohen, A. L. (2009). Decision-tree models of categorization response times, choice proportions, and typicality judgments. *Psychological Review*, 116, 833–855. doi: 10.1037/a0017188
- Lamberts, K. (2000). Information-accumulation theory of speeded categorization. *Psychological Review*, 107, 227–260. doi: 10.1037/0033-295X.107.2.227
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiments. *Review of Economic Dynamics*, 1, 593–622. doi: 10.1006/redo.1998.0023
- López-Pérez, R. (2008). Aversion to norm-breaking: A model. *Games and Economic Behavior*, 64, 237–267. doi: 10.1016/j.geb.2007.10.009
- McLachlan, G. J., & Peel, D. (2000). *Finite mixture models*. New York: John Wiley and Sons.
- Pachur, T., & Hertwig, R. (2006). On the psychology of the recognition heuristic: Retrieval primacy as a key determinant of its use. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 983–1002. doi: 10.1037/0278-7393.32.5.983
- Payne, J. W., Bettman J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.
- Piovesan, M., & Wengström, E. (2009). Fast or fair? A study of response times. *Economics Letters*, 105, 193–196. doi: 10.1016/j.econlet.2009.07.017
- Rabin, M. (1993). Incorporating fairness into game-theory and economics. *American Economic Review*, 83, 1281–1302. doi: 10.1126/science.151.3712.867-a
- Rieger, M. O., & Wang, M. (2008). What is behind the priority heuristic? A mathematical analysis and comment on Brandstätter, Gigerenzer, and Hertwig (2006). *Psychological Review*, 115, 274–280.
- Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1446–1465. doi: 10.1037/a0013646
- Rubinstein, A. (2007). Instinctive and cognitive reasoning. A study of response times. *Economic Journal*, 117, 1243–1259. doi: 10.1111/j.1468-0297.2007.02081.x
- Selten, R. (1967). Die Strategiemethode zur Erforschung des eingeschränkt rationalen Verhaltens im Rahmen eines Oligopolexperimentes [The strategy method to investigate bounded rational behavior in the context of an oligopoly experiment]. In H. Sauermann (Ed.), *Beiträge zur experimentellen Wirtschaftsforschung* (pp. 136–168). Tübingen, Germany: J. C. B. Mohr (Paul Siebeck).
- Townsend, J. T. (1990). Serial vs. parallel processing: Sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science*, 1, 46–54. doi: 10.1111/j.1467-9280.1990.tb00067.x
- Tversky, R. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79, 281–299. doi: 10.1037/h0032955
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297–323. doi: 10.1007/BF00122574

Authors' biographies:

Urs Fischbacher is a professor at the University of Konstanz and the head of the Thurgau Institute of Economics in Kreuzlingen, Switzerland. His research is focused on behavioral and experimental economics, in particular the investigation of non-selfish preferences.

Ralph Hertwig is the director of the Center for Adaptive Rationality at the Max Planck Institute for Human Development, Berlin. He studies bounded and social rationality, experienced-based decision making, and the methodology of the social sciences.

Adrian Bruhin is a tenure-track assistant professor of Economics at the University of Lausanne, Switzerland. He is mainly interested in risk, time, and social preferences as determinants of microeconomic behavior.

Authors' addresses:

Urs Fischbacher, Thurgau Institute of Economics, Kreuzlingen, Switzerland and University of Konstanz, Germany.

Ralph Hertwig, Max Planck Institute for Human Development, Center for Adaptive Rationality, Berlin, Germany.

Adrian Bruhin, University of Lausanne, Lausanne, Switzerland.