games

MDPI

Article

Investigating the Social Boundaries of Fairness by Modeling Ultimatum Game Responders' Decisions with Multinomial Processing Tree Models

Marco Biella [1,2,*], Max Hennig [3] and Laura Oswald [4]

[1] Faculty of Business and Economics, University of Basel, 4052 Basel, Switzerland
[2] Faculty of Psychology, Eberhard Karls Universität Tuebingen, 72076 Tübingen, Germany
[3] Department of Psychology, Julius-Maximilians-Universität Wuerzburg, 97070 Würzburg, Germany; maximilian.hennig@uni-wuerzburg.de
[4] Department of Psychology, Albert-Ludwigs-Universität Freiburg, 79104 Freiburg im Breisgau, Germany; laura.oswald@psychologie.uni-freiburg.de
* Correspondence: marco.biella@unibas.ch or marcobiella@live.com

Abstract: Fairness in competitive games such as the Ultimatum Game is often defined theoretically. According to some of the literature, in which fairness is determined only based on resource allocation, a proposal splitting resources evenly (i.e., 5:5) is generally assumed as fair, and minimal deviation (i.e., 4:6) is considered enough to classify the proposal as unfair. Relying on multinomial processing tree models (MPTs), we investigated where the boundaries of fairness are located in the eye of responders, and pit fairness against relative and absolute gain maximization principles. The MPT models we developed and validated allowed us to separate three individual processes driving responses in the standard and Third-Party Ultimatum Game. The results show that, from the responder's perspective, the boundaries of fairness encompass proposals splitting resources in a perfectly even way and include uneven proposals with minimal deviance (4:6 and 6:4). Moreover, the results show that, in the context of Third-Party Ultimatum Games, the responder must not be indifferent between favoring the proposer and the receiver, demonstrating a boundary condition of the developed model. If the responder is perfectly indifferent, absolute and relative gain maximization are theoretically unidentifiable. This theoretical and practical constraint limits the scope of our theory, which does not apply in the case of a perfectly indifferent decision-maker.

Keywords: fairness; competitive games; Ultimatum Game; multinomial processing tree; relative gain maximization; utility theory

check for updates

Academic Editors: Ulrich Berger, Ramzi Suleiman, Amnon Rapoport, Vernon Smith and Guillermina Jasso

Received: 20 August 2024
Revised: 13 December 2024
Accepted: 27 December 2024
Published: 3 January 2025

Citation: Biella, M., Hennig, M., & Oswald, L. (2025). Investigating the Social Boundaries of Fairness by Modeling Ultimatum Game Responders' Decisions with Multinomial Processing Tree Models. *Games*, *16*(1), 2. https://doi.org/10.3390/g16010002

1. Introduction

When engaging in economic transactions, people often face a trade-off between maximizing their own profit and following the social norm of fairness (Messick & Schell, 1992). This reasoning can be extended to a broader set of social interactions, such as helping behaviors or favor exchanges that do not involve economic aspects in strict terms.

Social interactions of this kind have been investigated by researchers in an experimental paradigm called the Ultimatum Game (Güth et al., 1982). This simple game, and its later developments (Biella & Sacchi, 2018; Civai et al., 2013), offer an opportunity to observe how people engage with strategic resources splitting and how they deal with (potentially) unfair situations. Specifically, the ultimatum has the potential to show people's behavior when confronted with the decision between maintaining an equitable outcome at the cost of losing personal monetary utility on one hand, or gaining some monetary utility at the cost of violating the norm of fairness on the other hand. Apart from the debate about when one of the two options is preferred over the other, the investigation of the cognitive processes underlying observable decisions is far from over. In this paper, we propose an innovative way of framing the processes driving decision-making in the Ultimatum Game, which lends itself to experimental testing and overcomes some of the open controversies on the topic.

1.1. Ultimatum Game and the Fairness/Utility Trade-Off

The Ultimatum Game is a simple sequential game in which two players strategically interact to divide some resources (Güth et al., 1982). These two players are the first-mover, generally referred to as the "proposer", and the second-mover, generally referred to as the "receiver" or "responder". In our research, we will focus on the responder's perspective. In addition to facilitating the discussion of our reasoning, this focus is necessary, as multinomial processing tree models (see next paragraph) require categorical responses. In the standard Ultimatum Game, the proposer is informed about the total amount of resources available (i.e., 10 Euro) and is asked to come up with a split proposal (i.e., 6 Euro/4 Euro). This proposal will be evaluated by the second-mover, the responder/receiver. The second-mover has two options to determine the final payoff of both players. If the second-mover accepts the offer, each player receives the designated amount (i.e., 6 Euro to the proposer and 4 Euro to the receiver), but if they reject the offer, both players receive nothing. If monetary gain is the only driver, the proposer should craft an offer that assigns the greater amount to him/herself, leaving only a minimal but non-zero quantity of resources for the receiver. The responder should accept such an offer, as rejecting it would imply receiving nothing, which is less that the non-zero monetary utility received if they accept. For the same reason, the responder should accept any offer that grants him/her any non-zero monetary gain, regardless of how much the proposer is earning. However, these predictions hold only if one assumes that monetary resources are the only component of the utility functions of both the responder and the proposer. These predictions are rarely met in the literature and violations of monetary utility maximization are often reported (Güth et al., 1982; Biella & Sacchi, 2018; Aina et al., 2020; Civai, 2013; Ruessmann & Topolinski, 2020). Indeed, a reasonable proposer might maximize their own payoff by crafting fair offers, which are more likely to be accepted. Indeed, any offer that is likely to be rejected will produce no gain for the proposer. Therefore, advancing such offers is not a reasonable behavior. The violation of predictions assuming that monetary gain maximization is the only component of both players' utility functions should not be seen as a failure of economic theory but rather as a signal that utility does not entail only monetary gain maximization. For example, fairness can be embedded into the utility functions of both the proposer and the receiver.

Several accounts have been proposed to explain deviations from the predictions derived using monetary gain maximization as the only component of the utility function. Some authors endorse an emotional explanation. Specifically, they postulate that receiving an unfair, disadvantageous offer, an offer in which the amount destined to the receiver is lower than the amount destined to the proposer, triggers negative emotions responsible for the rejection and the violation of economic principles (Aina et al., 2020; Civai et al., 2010; Pillutla & Murnighan, 1996). Indeed, the "wounded pride/spite model" endorses such an explanation (Pillutla & Murnighan, 1996), which is supported by neuroscientific evidence reporting greater activation of the anterior insula and dorsolateral prefrontal cortex, two areas related to anger and disgust, when the receiver is exposed to unfair, disadvantageous offers (Sanfey et al., 2003). Similarly, rejections may be explained in terms of reciprocity (Fehr & Schmidt, 1999). Thus, the negative behavior of rejecting an offer can be understood as reciprocating the negative behavior of proposing an unfair offer in the first place. Other accounts suggest that the rejection of Ultimatum Game offers, and the consequent violation of economic assumptions, is based on cognitive heuristics and the norm of fairness (Messick & Schell, 1992; Messick, 1995). Such a norm postulates that, when there is no reason to do otherwise, both players should receive the same share of resources. This definition of fairness is very stringent, and some adjustments are required for games that are inherently asymmetric, such as the Ultimatum Game (Kamas & Preston, 2012; Kravitz & Gunto, 1992; Suleiman, 2017; Suleiman, 2022). However, this decision rule has the benefit of leveling the relative monetary utility received by both players (Messick & Thorngate, 1967) and is in line with other economic accounts suggesting that a receiver rejecting the proposer's offer is engaging in negative reciprocity toward the proposer (Rabin, 1993). The latter is even in line with evolutionary theories and the literature on justice within social economic games, which suggest that if the norm of fairness is broken, a punishment should be imposed (Hardin, 1968; Schroeder et al., 2003) and that such a social control device promotes survival in the long run (Alexander, 1987; Boehm, 1999; Darwin, 1871). Indeed, there is an extensive literature on the evolution of fairness in resource-sharing games (Kahneman et al., 1986) in general, and related to the Ultimatum Game, specifically (Camerer & Thaler, 1995; Debove et al., 2016; Thaler, 1988). However, the standard paradigm confounds fairness and monetary utility, as any disadvantageous offer is both unfair and of lower monetary utility for the receiver. Therefore, an evolution of the Ultimatum Game has been proposed. In the Third-Party Ultimatum Game (Civai et al., 2013; Haruvy & Roth, 2022), an additional player called the decision-maker is introduced. Introducing the decision-maker effectively separates the role of the receiver and of the responder. The decision-maker has the role of evaluating the proposer's offer on the behalf of the receiver. If the decision-maker accepts the offer, both the proposer and the receiver receive part of the resources, and if they reject, proposer and receiver receive nothing. In both cases, the decision-maker does not receive any money and, therefore, should not be affected or guided by monetary utility. This version removes the utility/fairness confound. Moreover, this paradigm allowed researchers to demonstrate that people show inequity aversion (Bolton, 1991; Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999) when monetary utility is not a decision driver, but tolerate it when such a utility is at stake and in their favor, as in the case of an unfair advantageous offer in the standard Ultimatum Game, in which the decision-maker/receiver receives part of the resources (Civai et al., 2013). In this paradigm, researchers claim that inequity aversion is tolerated, as they document participants accepting uneven offers. Additionally, it has been shown that, if the decision-maker and the receiver are socially close (Biella et al., 2023) or belong to the same minimal social group (Biella & Sacchi, 2018), the monetary utility directed toward the receiver affects the decision made by the decision-maker such that inequity aversion is more tolerated if the monetary utility is in favor of the receiver. Again, researchers used decision-makers' acceptance of uneven offers as key evidence for inequity aversion tolerance.

After reviewing the accounts aimed at explaining decision-making in the Ultimatum Game, we can identify three main drivers that play a crucial role. These drivers are the innate sense of fairness that is captured by inequity aversion (Bolton, 1991; Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999), the economic principle of absolute gain maximization, and a situationally grounded effort to maximize relative gain (Rabin, 1993). Beyond the debate on whether such processes are cognitive (i.e., heuristics) or emotional (i.e., anger/disgust, wounded pride/spite model) in nature (Civai et al., 2013), the fact that these tree processes are considered the main drivers of decision-making in the Ultimatum Game stands, and it is the subject of our present analysis.

It is worth noting that our investigation does not attempt to refute any of the theoretical accounts advanced to explain decisions in the Ultimatum Game. On the contrary, we build upon such theories, and we aim at developing a model that complies with such theories' predictions. Our contribution mostly comes from the theoretical work required to validate the multinomial processing tree model we are proposing (see next section). Such a validation requires formalizing a sufficient set of processes responsible for the decisions made by the responders, deriving theoretically sound predictions (coming from the accounts we just reviewed), and comparing the model's performance against such predictions.

1.2. Multinomial Processing Tree Models

Multinomial processing tree models (MPT) are powerful analytical tools that offer the capability of disentangling co-occurrent processes guiding responses within a specific task (Hütter & Klauer, 2016). Relying on categorical responses, these models quantify the probability of each process driving the outcome, conditioned upon the contribution of all other processes. Adopting the responder/receiver's perspective allows us to leverage the categorical nature of responses to model the underlying processes behind decisions. In this framework, researchers can obtain process-pure quantification of the processes in the context of the paradigm of interest. Moreover, MPT models allow for a straightforward test of differences in parameter values across experimental conditions (Klauer et al., 2011; Riefer & Batchelder, 1988). Crucially, MPT models require the researcher to formulate precise hypotheses, as the formalized model cannot be developed if the number, nature, and composition of the processes are not clearly specified. Additionally, the most important processes playing a role in the paradigm of interest must be included for the model to work properly. Therefore, model specification must be driven by careful theoretical analysis. Once the model has been specified, it can be fitted using the maximum likelihood method, determining the parameter values that make the response data most likely.

The parameters can be interpreted as the probability of each process to drive responses. More specifically, they are the conditional probabilities depending on the previous process (Hütter & Klauer, 2016). Each branch in the tree represents the operation of a single process or a succession of processes leading to an observable response. High parameter values denote that the process has a strong influence on response production. Crucially, the estimation of the parameters is dependent on the order of the processes. Such dependency has two main consequences. First, parameter comparison across models must be carried out on models enforcing the same parameter order. Second, parameters must be interpreted as probabilities conditioned on previous processes. Once the parameters have been estimated, the model fit can be evaluated using a chi-square test (Hu & Batchelder, 1994) and more fine-grained metrics such as Cohen's w (Cohen, 1988). Regarding the chi-square test, a canonical threshold for significance testing can be assumed, $\alpha = 0.05$, but additional care is warranted, as the chi-square test is known to be oversensitive to large sample sizes (Foldnes & Henning Olsson, 2015; Powell & William, 2001). Regarding Cohen's w, we assume $w = 0.10$ as a reasonable threshold for a small magnitude. Any deviation from the ideal fit larger than that will be considered too much, leading to the conclusion that satisfactory fit has not been reached.

Several research programs, such as the investigation of moral dilemmas (Conway & Gawronski, 2013; Hennig & Hütter, 2020) and the automaticity of attitude acquisition (Hütter & Sweldens, 2018; Hütter et al., 2012), have already benefited from the application of these models, and our goal is to apply such a framework in the context of the Ultimatum Game.

1.3. MPT Model for the Ultimatum Game

To the best of our knowledge, modeling the Ultimatum Game's responders' decisions using MPT models has never been attempted. The existing literature already investigates the drivers of proposals generation and has highlighted the central role of fairness (Forsythe et al., 1994; Harrison & McCabe, 1996; Hoffman et al., 1994). Although successful, these prior attempts took the proposer's perspective or relied on experimental designs to investigate the decision's drivers in isolation. We think that the literature on the Ultimatum Game can greatly benefit from the application of MPT models, as proper quantification of the processes driving responses in this paradigm will shed new light on the debate around the processes themselves, and on the role of fairness in particular. Moreover, many economic games, such as the Ultimatum Game, exhibit similar features in terms of structure, response type, and driving processes that closely resemble other tasks (i.e., moral dilemmas) that have already benefited from the use of MPT models (Conway & Gawronski, 2013; Hennig & Hütter, 2020).

1.3.1. Underlying Processes

To make our approach fruitful, however, our formal model must be guided by careful theoretical analysis, starting from the identification of the driving processes. Based on the literature above, we consider the norm of fairness, maximization of relative gain, and maximization of absolute gain as the most important processes playing a role in the Ultimatum Game. The norm of fairness is defined as the implicit and socially established principle that unbalanced resources splits should not be proposed neither accepted. Relative gain is defined as the ratio between the resources destined to the receiver over the resources destined to the proposer. If an offer allocates more resources to the receiver than to the proposer, it can be considered an unfair offer, with relative gain favoring the receiver. Finally, absolute gain is defined as the absolute payoff earned by the receiver if the offer is accepted. Our theoretical analysis does not divide processes into cognitive (Messick & Schell, 1992; Rabin, 1993; Bolton & Ockenfels, 2000; Fehr & Schmidt, 1999) versus emotive (Pillutla & Murnighan, 1996; Sanfey et al., 2003) but rather merges these two sides of the same coin. For example, negative emotions arising from receiving an unfair, disadvantageous offer are equally embedded in the fairness-based and relative-gain-based process. Indeed, negative emotions might arise from both perceiving the unfair, disadvantageous offer as offensive, due to the violation of the norm of fairness, and as an attempt to reduce the receiver's relative gain. On the other hand, an unfair, advantageous offer might trigger fewer negative emotions, as only the latter, the attempt to reduce the receiver's relative gain, is missing, while the former, the violation of the norm of fairness, is still aversive and potentially triggering. Similarly, expectations can overlay with such processes, as long as they provide predictions that are coherent with the process itself. The existing literature, for example (Harrison & McCabe, 1996; Suleiman, 1996), postulates that decision-makers and receivers might expect proposers to advance fair offers. In. It is not a theoretical problem if emotions, expectations, or any other phenomena on a different level of analysis are conflated with the process of interest, as long as these processes lead to clear, non-contradictory theoretical predictions and show no logical inconsistencies, which is the case, as detailed in the upcoming paragraph.

To estimate the model, the order of the processes must be specified. However, the order of the model does not have any theoretical implications, as the processes are assumed to act in parallel (Hütter & Klauer, 2016). Analytically speaking, the order specified imposes only one constraint, namely, that the same processes order must be specified to compare two models. Similarly, the parameters should not be interpreted in isolation but in the larger context of the whole model. Therefore, we ordered the processes based on the theoretical level in which they reside. We started with fairness, which is the most related to social norms. This norm exists in the presence of an agent, at least one interaction partner, and a larger social group enforcing the norm. Second, we introduced relative gain maximization, which implies some level of interaction with a social other. This process resides at the interpersonal level, in between the decision-makers and the players of the game. Finally, we added absolute gain maximization, which can be placed at the individual level. This last process requires only the decision-maker and an individual preference for each possible outcome. This ordering, from the most social to the most individual, makes sense from the theoretical point of view and, as long as model comparison happens between models enforcing the same ordering, it is not analytically problematic (Hütter & Klauer, 2016).

To summarize, we developed a model in which the response in the Ultimatum Game is driven by three processes. These processes are embedded in an MPT model with the following parameters: f for fairness, $r1$, $r2$, $r3$, and $r4$ for the maximization of relative gain, and a for the maximization of absolute gain. Relative gain maximization requires four parameters to encode the different intensities of relative gain maximization at varying offer levels. Theoretically, r parameters represent the same process, but the varying intensity must be represented by separate parameters.

1.3.2. The Model and Its Theoretical Predictions

After the identification of the processes of interest, clear and testable theory-driven predictions must be specified to formalize the model. The expected decision outcome for each offer level must be specified under the assumption that each process guides the behavior (Hütter & Klauer, 2016). Moreover, such predictions must avoid logical inconsistencies and must make the model identifiable (Singmann & Kellen, 2013).

The fairness-based process has fairly straightforward predictions. Every time an offer splits resources unevenly, this process leads to rejection, while acceptance is expected if the offer distributes resources evenly. However, this formalization assumes that the boundaries of fairness are clear-cut. Especially in repeated-game versions of the Ultimatum Game (versions in which the proposer and the receiver go through several bargaining rounds); however, the boundaries of fairness might be more lenient (Forsythe et al., 1994; Harrison & McCabe, 1996). A receiver might consider a slightly uneven disadvantageous offer as fair with the expectation that the proposer will consider, and provide, the advantageous version of the same slightly uneven offer. Crucially, under both formalizations the predictions are the same.

Regarding the relative-gain-based process, predictions can be easily derived along the whole continuum of possible offers except for a perfectly even offer. Specifically, if the amount destined to the receiver is lower than the amount destined to the proposer, this process imposes rejection, while acceptance is expected if the relative proportion of resources favors the receiver. In the singular case of a perfectly even offer, relative gain does not favor any of the two players. Here, this process is silent, as it is logically impossible to use relative gain to make a decision in the case of an offer that does not show any relative gain advantage for either side. To reflect this, our model does not have a relative-gain-based prediction for perfectly even offers. Such an instance of a process lacking predictions is not an issue, as long as the combination of the remaining processes and their predictions make the model identifiable. Additionally, some evidence suggests that the more uneven the offer, the stronger the reaction (Sanfey et al., 2003; Messick & Thorngate, 1967; Van't Wout et al., 2006). To properly capture this increased intensity, relative gain maximization will be modeled as several processes with identical predictions (i.e., one parameter for each level of intensity).

Finally, the absolute gain process is the one with the most straightforward predictions. In line with the classical economics account, this process predicts that each offer that allocates at least some resources to the receiver will be accepted. Any rejections suggest that this process has been overturned.

Collecting all the predictions for each process under each offer level, we can derive the formal MPT model and its parameters. Specifically, we have an f parameter modeling the fairness-based process, multiple r parameters modeling the different stages of the relative-gain-based process, and an a parameter modeling the absolute-gain-based process (Figure 1A).

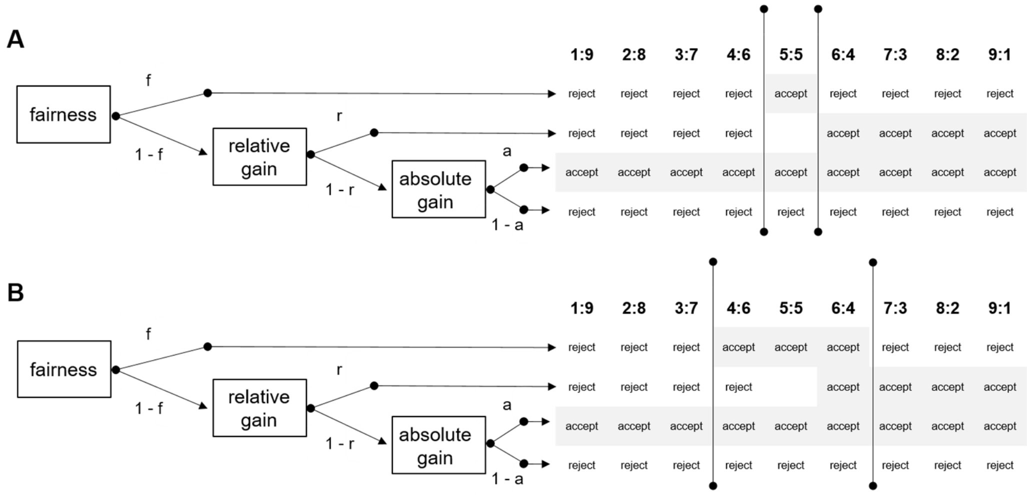

Figure 1. Multinomial processing tree models embedding (**A**) "Strict" and (**B**) "Lenient" fairness conceptualization and their respective predictions for all offer levels. Vertical lines represent fairness boundaries. Shaded text represents offer acceptance, while non-shaded text represents offer rejection.

The only theoretical unknown is at what offer level fairness boundaries can be placed. To allow for a more lenient formulation of fairness (Forsythe et al., 1994; Harrison & McCabe, 1996), which allows for less clear-cut boundaries, a second model can be derived. In such a model, the fairness-based process predicts acceptance in a larger range of offer levels centered on the perfectly even one. Here, fairness and absolute gain maximization converge with relative gain maximization only on the right-hand side, where offers are advantageous for the receiver, while it contradicts relative gain maximization but still converges with absolute gain maximization (Figure 1B).

It is worth noting that, in addition to having a relatively high number of parameters, both models yield theoretically meaningful and consistent predictions. Moreover, both models take all three processes into account while requiring the lowest number of parameters. Indeed, alternative models that consider all three processes are possible at the cost of an increased number of parameters. Therefore, the present models are the ones that capitalize the least on chance due to the model complexity in terms of the number of parameters.

Additionally, the model flexibility in terms of fairness boundaries and multiple r parameters is theoretically justified and never yields logical inconsistencies. Less flexible models (i.e., a model with a single r parameter) are possible at the cost of violating theoretically meaningful predictions (i.e., assuming that relative gain maximization is the same for slightly and heavily uneven offers).

1.4. The Present Research

In the present research, we aim to (a) use the MPT framework to predict Ultimatum Game responses, obtaining satisfactory model fit; (b) validate the model by investigating how it performances under different theoretically meaningful conditions (i.e., the standard and Third-Party Ultimatum Game); and (c) locate the boundaries of perceived fairness by testing predictions of parameters implementing strict and lenient conceptualizations of fairness. We base our investigation on three studies.

The first goal (a) of the present research is achieved mainly by our first experiment, and marginally by the remaining two. In Experiment 1, we aim at testing the proof of concept of our model. We simply run a standard Ultimatum Game to test if the MPT framework can model the data properly. Specifically, we test if the models (with "strict" and "lenient" fairness boundaries) fit the observed data.

The second goal (b) is achieved by the comparison of data from the first experiment with the data from the Third-Party Ultimatum Game of Experiment 2. Here, our expectation is that the model fits properly on both experiments separately, and that the comparisons of the models' parameters fit our theoretical predictions. Regarding model fit, we expect that the model with the "lenient" fairness conceptualization will fit the data better than the one with the "strict" conceptualization. Regarding model parameters, we expect the f parameter to be lower in the standard version of the paradigm (Experiment 1) than in the third-party version (Experiment 2). Similarly, we expect the a parameter to be lower in the third-party version in comparison to the standard version.

The third goal (c) is achieved by evaluating how different models (with "strict" or "lenient" fairness conceptualization) fit observed data across all three experiments and, in particular, in Experiment 3. In Experiment 3, participants receive the explicit instruction of "not taking any side" before a Third-Party Ultimatum Game. Therefore, it is possible that both our models show unsatisfactory fit, as an indifferent decision-maker cannot maximize relative gain. If this is the case, Experiment 3 may provide insights on the boundary conditions under which our model does not apply.

2. Experiment 1

In Experiment 1, we provide a proof of concept for our models. Specifically, we test our hypothesis that Ultimatum Game responses can be modeled relying on three processes, namely fairness, relative gain maximization, and absolute gain maximization. Moreover, we aim at probing potential fairness boundaries by testing the "strict" and "lenient" fairness conceptualizations embedded in the two models presented above. Data and materials for Experiment 1 are available at the OSF repository (https://osf.io/tzbpd/, accessed on 26 December 2024).

2.1. Procedure

In Experiment 1, participants were involved in a standard Ultimatum Game. The experiment was conducted online, and prior to the beginning of the procedure, participants provided their informed consent to take part in the experiment. The experiment began with the collection of socio-demographic information about the participant, then instructions were presented. Participants were informed that they were about to take part in an Ultimatum Game. Specifically, participants were informed that some of them would be assigned to the role of the proposer and that their task was to come up with split proposals of a total of 10 Euro. The remaining participant would be assigned to the role of the receiver, and their task was to decide whether to accept or reject each offer. Moreover, participants were informed that the proposer/receiver pairs would be constant throughout the experiment, meaning that they would always interact with the same partner. All participants were assigned to the receiver role. A translation of the final instructions received by the participant is available in the online repository (https://osf.io/3xt4y, accessed on 26 December 2024) and Appendix A.

To increase the ecological validity of the experimental situation and avoid asking the participants hypothetical choices, the Ultimatum Game was incentivized. Participants were informed that the total amount of money collected from the game would be converted into raffle tickets to win a 20 Euro voucher for a local bookstore. In total, four prizes were awarded. Specifically, we informed participants that for each Euro collected, they would receive one raffle ticket and, as it follows logically, collecting more tickets was linked to increased chances of winning the voucher. Similarly, participants were informed that the total amount of money collected by the proposer would be converted in the same way and that the proposer would enter the raffle too. The English translation of the instructions the participants received is available at the online repository (https://osf.io/tzbpd/).

Before the beginning of the experimental task, comprehension checks were implemented. Specifically, participants were presented with a random offer, and they were asked to indicate how much money the proposer and the receiver would receive if the receiver accepted the offer. Similarly, participants were asked to indicate the same quantities if the receiver rejected the offer. Immediately after submitting their responses, participants were provided feedback about their response's correctness. Correct amounts were highlighted in green and incorrect amounts were highlighted in red. In the case of incorrect answers, participants were asked to respond again until the correct answers were provided.

After the comprehension checks, the experimental task began. Participants were presented with nine offers in total, one at a time in random order. They were asked to evaluate one offer per trial and accept or reject the offer by clicking the corresponding button on the screen. At the end of the experiment, participants were informed of the total amount of money they collected, which was converted into raffle tickets. After the last participant took part in the experiment, four winners were drawn from the raffle, and the vouchers were provided.

2.2. Participants

The total sample size required for Experiment 1 was determined based on response frequency simulations and power analysis. Our power analysis was based on the joint model estimating parameters separately for the standard (Experiment 1) and third-party (Experiment 2) conditions. This provided us with more power than needed for Experiment 1 but granted us the statistical power required for the joint analysis comparing the two experiments. The target effect size was obtained from the existing literature (Biella & Sacchi, 2018; Experiment 1b, which most closely resembled the intended setup of our experiment). In the first step, we estimated how many participants would be necessary to ensure that no more than 5% of expected cell frequencies would be below 5, such that model misfits would likely be detected. This suggested 200 participants per experiment to provide adequate expected cell frequencies. In the second step, we tested for 80% power to detect the difference between the two experiments in parameters a, $r1$, and $r2$, a difference that we found in our re-analysis of Ruessmann and Topolinski (2020, Experiment 1b). These analyses indicate minimum sample sizes of 264, 281, and 355 participants per experiment. To ensure a sufficient sample size after excluding participants that failed the comprehension checks, we interrupted the advertisement of the data collection after recruiting the 500th participant. We excluded one participant that took the comprehension checks more than 10 times. All the others answered the checks correctly immediately (N = 415), after one instance of feedback (N = 74) or after additional feedback (N = 32). The final sample size consisted of N = 521 participants (162 males, 354 females, 4 diverse, and 1 undisclosed), aged between 18 and 70 (M = 25.60, SD = 7.33).

2.3. Results

Participants' acceptance rate for each offer level (by experiment) can be found in Figure 2.

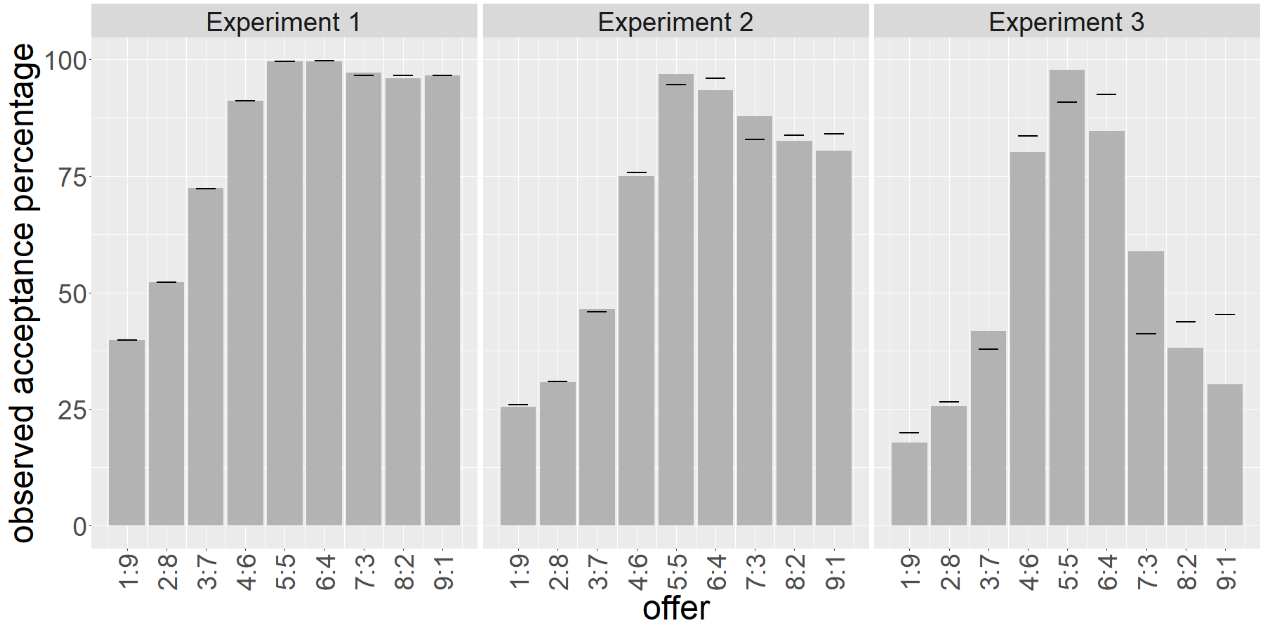

Figure 2. Observed acceptance percentage by offer level across experiments. Horizontal lines represent predictions from best-performing models (lenient fairness).

The first analysis involved fitting the MPT model with the "strict" conceptualization of fairness. Acceptance and rejection frequencies were aggregated for every offer level and used to fit the MPT model (Figure 1A) using the MPTinR package (version 1.14.1) in R (version 4.4.2). Model fit was assessed using the two canonical indexes G^2 and Cohen's w. The first model exhibited unsatisfactory fit, $G^2(3) = 22.06$, $p < 0.001$. However, the effect size of this deviation is below the threshold for what is generally considered a small effect, w = 0.069, which indicates the deviation from fit to be negligible. Specifically, this might occur even for models with satisfactory fit, as the goodness of fit index G^2 is known to be oversensitive to larger sample sizes (Foldnes & Henning Olsson, 2015; Powell & William, 2001), which is the case for this experiment. Next, we fitted the second model involving the "lenient" conceptualization of fairness. This second model fitted the data better, reaching satisfactory fit, $G^2(3) = 1.21$, $p = 0.75$, w = 0.016. Once a model with satisfactory fit was obtained, we proceeded to investigate the individual.

Using the "lenient" fairness model, we estimated the individual contributions of each process. Specifically, we obtained one f parameter encoding fairness' contribution, four r parameters ($r1$, $r2$, $r3$, $r4$) encoding relative gain maximization's contribution at different levels, and one a parameter encoding absolute gain maximization's contribution. The estimations and confidence intervals for each parameter are available in Table 1 and are represented graphically in Figure 3.

To test for the need of the specific process parameter, we fitted alternative models by constraining the parameters to 0 (or to 0.5 in the case of the last parameter in the hierarchy) and compared the alternative model's fit against the fit of the unconstrained model. This analysis suggested that all the parameters for each process are required in order to properly model the data, G^2s > 63.93, ps > 0.001, ws > 0.12. The results of the comparison of the alternative and unconstrained model's fits are available in Table 1.

Table 1. Parameter estimates, confidence intervals, and model fit when the parameter is constrained to 0 for individual models across all three experiments.

Parameter	Study	Unconstrained Model		Constrained Model	
		Estimation	95% CI	G^2	p
f	Standard [Exp. 1]	0.03	0.02, 0.04	63.93	<0.001
	Third-Party [Exp. 2]	0.14	0.12, 0.26	267.95	<0.001
	Fair Instruction [Exp. 3]	0.50	0.47, 0.53	978.30	<0.001
r_1	Standard [Exp. 1]	0.09	0.06, 0.11	80.85	<0.001
	Third-Party [Exp. 2]	0.23	0.19, 0.28	123.01	<0.001
	Fair Instruction [Exp. 3]	0.18	0.10, 0.26	19.56	<0.001
r_2	Standard [Exp. 1]	0.25	0.21, 0.29	221.28	<0.001
	Third-Party [Exp. 2]	0.43	0.37, 0.48	219.59	<0.001
	Fair Instruction [Exp. 3]	0.07	−0.05, 0.19	1.25	0.26
r_3	Standard [Exp. 1]	0.46	0.41, 0.50	534.07	<0.001
	Third-Party [Exp. 2]	0.62	0.57, 0.67	434.48	<0.001
	Fair Instruction [Exp. 3]	0.35	0.23, 0.46	30.33	<0.001
r_4	Standard [Exp. 1]	0.59	0.54, 0.63	779.90	<0.001
	Third-Party [Exp. 2]	0.68	0.63, 0.73	521.66	<0.001
	Fair Instruction [Exp. 3]	0.51	0.41, 0.61	63.40	<0.001
a	Standard [Exp. 1]	1.00	0.99, 1.00	2995.62	<0.001
	Third-Party [Exp. 2]	0.94	0.92, 0.95	1196.71	<0.001
	Fair Instruction [Exp. 3]	0.81	0.78, 0.85	269.32	<0.001

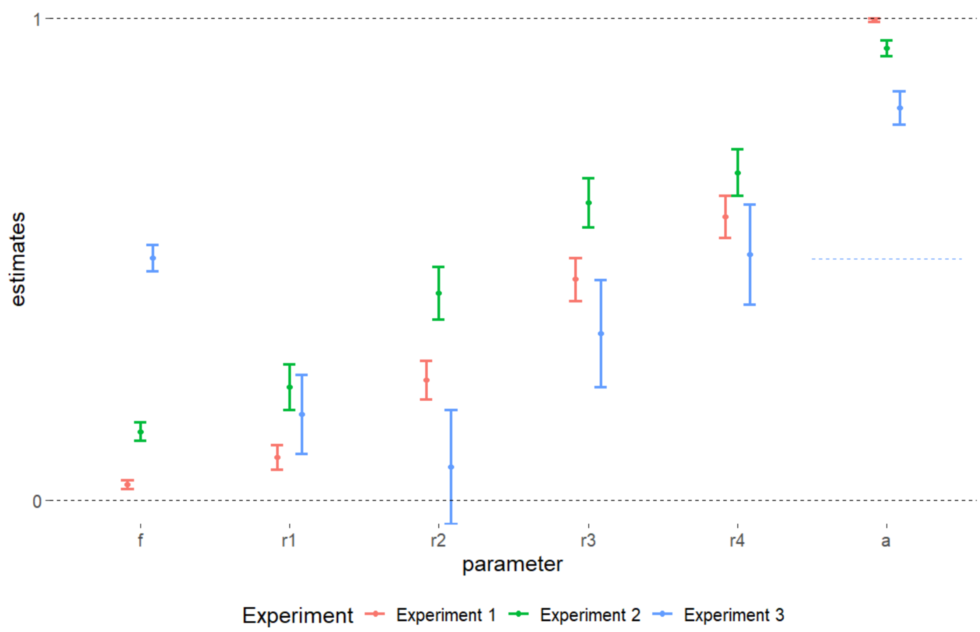

Figure 3. Parameter estimates and confidence intervals for all processes estimated using the "Lenient" model separately for each experiment.

Finally, we tested whether individual r parameters are required in order to properly model the empirical data. To do so, we fitted alternative models by constraining each r parameter to be equal to the subsequent one (i.e., $r1 = r2$, $r2 = r3$, and $r3 = r4$). Once again, the comparison of constrained models against the unconstrained one suggested that each r parameter needs to be unconstrained in order to reproduce the empirical data, G^2s > 17.62, ps > 0.001, ws > 0.06.

To fully explore our data, we probed the predictive capabilities of the best-performing model, the one encoding the lenient operationalization of fairness. The best performance was obtained for the 1:9 offer, for which the model predicted 207.01 acceptance responses out of the 521 trials. We observed 207 acceptance responses for a total error of 0.01 out of the 521 trials. The worst performance was obtained for the 7:3 offer, for which the model predicted 502.60 acceptance responses against the observed 506. The overall error was 3.40 out of 521 trials. The observed (column height) and predicted (horizontal bars) acceptance are shown in Figure 2.

2.4. Discussion

In this first experiment, the MPT model bounded to the Ultimatum Game paradigm proved to be satisfactory, at least in its "lenient" fairness formulation. Modeling the Ultimatum Game's responses based on the three processes, namely fairness, relative gain maximization, and absolute gain maximization, seems to be possible. The model embedding "strict" fairness remains questionable, as the G^2 is not satisfactory. However, the Cohen's w suggests the opposite conclusion. Being within a reasonable threshold, the Cohen's w suggests that the model embedding the "strict" conceptualization of fairness does not deviate from empirical data. Differently, the model based on the "lenient" fairness conceptualization was satisfactory regardless of how fit was assessed. Moreover, this second model allowed us to test for the need of each individual parameter. The results support the hypothesis that all three processes contributed to producing the responses in the Ultimatum Game paradigm. This first experiment yielded encouraging results in our endeavor to model the Ultimatum Game's responses using an MPT model.

3. Experiment 2

In order to fully validate our MPT model, the parameters in the model must be susceptible to being manipulated in a theoretically meaningful way, yielding theoretically meaningful changes. Indeed, we developed the second experiment to prove that a theoretically meaningful manipulation can affect the individual parameters in predictable ways. As fairness is of primary concern for the present work, we decided to use the third-party version of the Ultimatum Game (Civai et al., 2013) in order to produce the theoretically expected increase in the f parameter. In such a version, participants are asked to evaluate the proposer's offers on the behalf of the receiver. In this situation, it is possible that participants "take the side" of the receiver (Civai et al., 2013); however, as in the Third-Party Ultimatum Game, the decision-maker is no longer receiving any monetary utility, so they are expected to be more guided by inequality aversion (Bolton, 1991; Bolton & Ockenfels, 2000) than in the standard Ultimatum Game (i.e., Experiment 1). Therefore, Experiment 2 aims at demonstrating that, in the Third-Party Ultimatum Game, our model fits the data adequately and reflects the expected decrease in the f parameter in comparison with Experiment 1. Again, we will test both models with the "strict" and "lenient" fairness conceptualization, we will investigate both model fit and the individual parameters as in Experiment 1, and we will compare parameters across the two experiment with the same procedure used so far, namely, comparing unrestricted versus restricted models.

3.1. Procedure

The procedure of Experiment 2 was identical to the procedure of Experiment 1, with only two modifications. First, participants were informed about the role of the decision-maker, a participant required to evaluate the proposer's offer without receiving any money. Second, the comprehension checks included estimating how much money the decision-maker would receive in the case they accept or reject the offer. Obviously, the correct answer was zero in both cases. As before, participants were informed that the proposers' and the receivers' accumulated money would be converted into raffle tickets, while decision-makers would be endowed a single ticket as a participation fee, regardless of their decisions. To avoid recruiting participants who took part in the first experiment, we added an additional item asking the participants if they recalled taking part in a similar experiment.

3.2. Participants

An a priori power analysis guided us in the determination of the required sample size. However, we included two additional practical criteria for interrupting data collection. First, data collection would have been interrupted if participant recruitment dropped below 10 per day. Second, data collection would have been interrupted if it exceeded 27 days. After 8 days, data collection dropped below 10; therefore, data collection was interrupted.

Participants that declared they had participated in the previous experiment were excluded (N = 21). Moreover, only participants that understood the task correctly according to our comprehension checks were retained. N = 437 participants completed the comprehension checks on the first attempt, N = 48 answered correctly after a single instance of feedback, and N = 14 answered correctly after an additional piece of feedback. No participant made more than seven attempts. The final sample was N = 499 participants (166 males, 327 females, 4 diverse, and 2 undisclosed), aged between 18 and 68 (M = 24.52, SD = 7.20).

3.3. Results

As in Experiment 1, we began our analysis with the global fit of both models. The fit of the model implementing the "strict" fairness conceptualization was not satisfactory, $G^2(3) = 55.73$, $p < 0.001$. In contrast to Experiment 1, Cohen's w indicated that the size of the deviation was not negligible, w = 0.11. As before, we proceeded with assessing the model implementing the "lenient" fairness conceptualization. This model exhibited a sub-optimal fit based on the G^2, which is known to be overly sensitive to large sample sizes (Foldnes & Henning Olsson, 2015; Powell & William, 2001), $G^2(3) = 27.40$, $p < 0.001$. Crucially, the size of the deviation was negligible, w = 0.08, suggesting that the model predicts the data well and that the deviation of optimal fit resulted from the large sample size.

As before, the individual parameters quantifying the contribution of fairness, relative gain maximization, and absolute gain maximization were estimated (Figure 3), and models constraining each parameter to zero were compared against the unconstrained model (Table 1). These analyses supported, once again, the need for each individual parameter for predicting the responses in the Third-Party Ultimatum Game, as removal of any parameter significantly reduced model fit, G^2s > 123.01, ps < 0.001, ws > 0.18.

Again, we tested the need for multiple r parameters by constraining two r parameters to be equal at a time. All constrained models showed diminished fit, G^2s > 24.40, ps > 0.001, ws > 0.11, except the last one. When the $r3$ and $r4$ parameters were set to be equal, the constrained model showed similarly satisfactory fit as the unconstrained model, $G^2 = 3.05$, $p = 0.08$, w = 0.08.

To fully explore our data, we probed the predictive capabilities of the best-performing model, the one encoding the lenient operationalization of fairness. The best performance was obtained for the 2:8 offer, for which the model predicted 153.63 acceptance responses out of the 499 trials. We observed 153 acceptance responses with a total error of 0.63 out of the 499 trials. The worst performance was again obtained for the 7:3 offer, for which the model predicted 412.91 acceptance responses against the observed 438. The overall error was 25.09 out of 499 trials. The observed (column height) and predicted (horizontal bars) acceptance are shown in Figure 2.

3.4. Third-Party Manipulation

To fully validate our MPT model, we tested the effect of the third-party manipulation, comparing the results of Experiment 1 and those of Experiment 2. We developed a joint model containing pairs of parameters, one for the standard Ultimatum Game (Experiment 1) and one for the third-party Ultimatum Game (Experiment 2), reflecting the contributions of fairness, relative gain maximization, and absolute gain maximization. Once again, we set the pairs of parameters to be equal across conditions and compared the fit of the constrained model against the fit of the unconstrained model to assess the need for distinct parameters. As the model embedding the "lenient" fairness conceptualization showed satisfactory fit in both experiments, we used that model as a base to develop the joint model.

The joint model's overall fit was satisfactory, $G^2(6) = 28.61$, $p < 0.001$, w = 0.06. We consider the G^2, which is overly sensitive to large sample size (Foldnes & Henning Olsson, 2015; Powell & William, 2001), especially unreliable in this case, as the sample size of this analysis is the sample size of both experiments combined. Moreover, the Cohen's w showed a remarkably low level of w = 0.06, which increased our confidence in the satisfactory fit of the MPT model. Therefore, we estimated the parameters for both the standard and the third-party version of the model (Table 2).

Table 2. The estimation, confidence intervals, significance, and comparisons of individual parameters in the joint model involving responses from the standard and Third-Party Ultimatum Game. In the comparison portion of the table, the values and Cohen's w refer to the tests in which the parameter was constrained to be equal in both conditions.

	Standard Ultimatum Game					Third-Party Ultimatum Game					Comparison		
	Est.	**95% CI**	**G^2**	***p***	**w**	**Est.**	**95% CI**	**G^2**	***p***	**w**	**G^2**	***p***	**w**
f	0.03	0.02, 0.04	63.93	<0.001	0.10	0.14	0.12, 0.16	267.95	<0.001	0.18	104.59	<0.001	0.12
r_1	0.09	0.06, 0.11	80.85	<0.001	0.11	0.23	0.19, 0.28	123.00	<0.001	0.13	29.82	<0.001	0.08
r_2	0.25	0.21, 0.29	221.28	<0.001	0.16	0.43	0.37, 0.48	219.59	<0.001	0.16	25.35	<0.001	0.08
r_3	0.46	0.41, 0.50	534.07	<0.001	0.25	0.62	0.57, 0.67	434.48	<0.001	0.22	20.27	<0.001	0.07
r_4	0.59	0.54, 0.63	770.90	<0.001	0.30	0.68	0.63, 0.73	521.66	<0.001	0.24	7.37	0.01	0.06
a	1.00	0.99, 1.00	2995.62	<0.001	0.57	0.94	0.92, 0.95	1196.70	<0.001	0.36	61.07	<0.001	0.10

Moreover, we constrained each parameter to be equal to its counterpart between the standard and the third-party conditions and tested the fit of the constrained model against the fit of the unconstrained joint one.

As evidenced by the analysis, the fairness process affects the two conditions differently, $G^2 = 104.59$, $p < 0.001$, w = 0.12. Specifically, participants in the third-party condition, f = 0.14, CI95% = [0.12, 0.16], are more concerned with fairness than participants in the standard condition, f = 0.03, CI95% = [0.02, 0.04]. Reasonably, absolute gain maximization was marginally more central for participants in the standard condition, a = 1.00, CI95% = [0.99, 1.00] than for participants in the third-party condition, a = 0.94, CI95% = [0.92, 0.95], although this result must be taken with caution, given both its tenuous magnitude and the position of the a parameter at the bottom of the hierarchy. Relative gain maximization differed only descriptively between the two conditions, G^2s = 7.37, ps < 0.01, ws < 0.08.

3.5. Discussion

Taken together, the findings of Experiment 2 support the hypothesis that our MPT model properly accounts for Ultimatum Game responses in a theoretically sound way. On one hand, the overall model supported the notion that fairness, relative gain maximization, and absolute gain maximization drive responses in this paradigm, and on the other hand, the joint model yielded meaningful predictions in suggesting that fairness is more of a concern in the third-party than in the standard condition. Moreover, the results seem to indicate that relative gain maximization is not affected by our manipulation. This result, together with the tenuous difference in the *a* parameter, makes sense, as participants in the third-party condition did not receive any monetary utility and might only marginally take the side of the receiver (Civai et al., 2013).

4. Experiment 3

In Experiment 3, we aim to validate the model even further. We implemented a manipulation of the Third-Party Ultimatum Game that instructs participants to decide based only on their sense of fairness. Indeed, we expect the f parameter to be higher than the same parameter in both previous experiments. As before, we will proceed one step at a time. First, the global fit of the model will be assessed. Second, if the global fit is adequate, we will constrain every parameter to be zero, assessing the need for the parameter to be in the model. Third, we will run an integrative analysis, allowing us to compare joint models including data from Experiment 3 and the data from the previous experiments in order to compare the f parameter across experiments. If the latter analysis is inconclusive, we will conclude that the f parameter is unaffected by the manipulation in Experiment 3. If any of the model's parameters prove to not be necessary, or if the global fit of the model is not satisfactory, we will conclude that the model we developed cannot describe responses produced in a situation in which participants are asked to base their responses on one process only (fairness in this case) but have no stake in the game. If all analyses are successful and predictions are met, this third experiment will validate our model even further; however, if poor model fit is obtained, the third experiment will be considered a test of boundary conditions outside of which our model cannot be applied.

It is worth noting that our instructions to participants can be seen as controversial. If a participant is instructed to behave in a certain way, the resulting behavior cannot be generalized outside of the laboratory setting, where the instructions are not imposed. However, generalizing participants' behavior to other situations is not the goal of this experiment. The rather artificial situation we created in the lab can provide insights into whether participants can regulate their behavior (i.e., base their responses on fairness alone) if they are instructed to do so. Such an artificial situation would be unacceptable if the goal of the experiment was, for example, to test participants' innate disposition to comply with fairness. However, as the experiment is meant to stress-test the model, the artificiality of the experimental situation is less problematic.

4.1. Procedure

The procedure of Experiment 3 was identical to Experiment 2, with one minor modification. In Experiment 3, participants were asked to decide as a "fair judge". That is, participants were instructed to avoid "taking sides" and base their judgment only on their sense of what is fair. The series of experimental trials, comprehension checks, and inquiry as to whether the participants recalled taking part in a similar experiment were identical to Experiments 2 and 1.

4.2. Participants

The required sample size was determined similarly to Experiment 2. Again, data collection would have been interrupted if it exceeded 27 days, or if participant recruitment fell below 10 participants per day.

After five days, the minimal threshold of recruitment frequency was reached; however, as at that time the total sample was composed of only N = 197, without excluding those who took part in the previous experiments, we continued data collection while refraining from analyzing the unreliable sample. Recruitment frequency remained low, probably due to the data collection occurring outside the regular semester. Therefore, on day 13, we sent a reminder to the same mailing list. The reminder worked, as on the day on which it was sent, we recruited N = 142 new participants. Three days after the reminder, recruitment frequency dropped below the threshold again, but this time, the total sample size (without exclusion) was sufficiently large, N = 415.

A total of N = 19 participants declared that they took part in a similar experiment and were excluded. A total of N = 353 participants completed the comprehension checks immediately, N = 36 after one instance of feedback, and N = 7 after additional feedback. No participant took more than six attempt to complete the comprehension checks. The final sample was N = 396 participants (131 males, 256 females, 8 diverse, and 1 undisclosed), aged between 18 and 74 (M = 25.68, SD = 8.39).

4.3. Results

As in our previous studies, the analytical strategy started with assessing the global fit of models encoding both the "strict" and "lenient" fairness conceptualizations. However, in this case, both models required one additional assumption. Indeed, as the participant is instructed to avoid "taking sides", predictions based on the relative gain maximization are somewhat misleading, as the process subsumes asymmetrical preferences. Therefore, for fitting purposes only, we assumed that relative gain maximization prescribes acceptance of offers favoring the receiver and rejection for those favoring the proposer. With this in mind, model fit can be investigated. Regarding the first model encoding the "strict" fairness conceptualization, both fit measures pointed toward poor fit, $G^2(3) = 418.89$, $p < 0.001$, w = 0.34. Similarly, the model encoding the "lenient" fairness conceptualization did not fit the data adequately, $G^2(3) = 158.51$, $p < 0.001$, w = 0.21. Given the lack of fit of both models, proceeding to the estimation and interpretation of individual parameters as well as the integrative analysis comparing Experiment 3 with the previous two is unwarranted. The interested reader who wishes to explore such analyses is directed toward the online repository containing the raw data and required scripts.

4.4. Discussion

The results of Experiment 3, although negative, are still informative. Indeed, the lack of model fit can guide us to conclude that, in the context of Third-Party Ultimatum Games, the MPT model proposed cannot explain data in which the decision-maker is totally indifferent between favoring the proposer and the receiver. More speculatively, the models' misfit can be imputed to the theoretical shortcoming of relative gain maximization predictions. Here, it is theoretically impossible to decide a priori which side of the offer's continuum is preferred; therefore, the model's predicted responses cannot be informed by relative gain maximization.

In sum, the results of Experiment 3 suggest that our MPT model, and its underlying processes, do not apply to situations in which the decision-maker has no stake in the game. Therefore, these results can inform us of the scope of our model and its boundary conditions. Limiting the scope of our model based on empirical evidence and theoretical constraints, however, is desirable as it makes our theory more precise.

5. General Discussion

Since its first development, the Ultimatum Game unveiled the discrepancy between the normative behavior prescribed by standard economic theory and actual behavior documented by empirical observation. Many accounts have been proposed to explain such a discrepancy, from the "wounded pride/spite model" (Pillutla & Murnighan, 1996) to negative reciprocity (Rabin, 1993) and the norm of fairness (Messick & Schell, 1992; Messick, 1995). The debate on which account is most suited for such an explanatory purpose is still ongoing, and it generally divides behavior drivers into two categories, namely cognitive heuristics and emotional reactions (Civai, 2013). In the present paper, we overcome such a dichotomy, addressing the issue by relying on several processes that partially belong to both categories. Crucially, the present investigation is capable of quantifying each process's contribution to the final response produced. The multinomial processing tree models investigated combine the roles of fairness, relative gain maximization, and absolute gain maximization and assign to each process part of the responsibility for the final decision.

Based on three studies, we initially tested the models' ability to account for the processes at play, demonstrating that a theoretically meaningful manipulation (Third-Party Ultimatum Game) leads to predictable changes in the relative importance of each process, and provides boundary conditions delimiting the scope of the models. Moreover, moving monetary utility away from the decision-maker (i.e., Experiment 2) successfully shifted its concerns toward preserving a fair state of affairs. Such a successful manipulation further validates our model.

5.1. Contibutions to the Bundary Conditions of the Existing Literature

Several theoretical accounts provide reliable explanations for the empirical observations on cooperative and competitive games. For example, Fehr and Schmidt's work (Fehr & Schmidt, 1999) provided a broad perspective on the determinants of decision-making. They show that cooperation is maintained even though competitive behavior might lead to greater monetary gain. Similarly, fairness is always part of the equation. Other works, focusing on the role of reciprocity, demonstrate that this is not a marginal construct (Bolton, 1991; Rabin, 1993; Thaler, 1988). Our work extensively builds upon this literature and does not aim to criticize these theories. On the contrary, we embedded these theoretical accounts into our model. The key contributions from our work attempt to consolidate our support for the reviewed theoretical accounts and to probe how far they can go to explain decision-making in the Ultimatum Game. For example, our investigation provides insight into the theories' boundary conditions. For example, the decision-makers' indifference causes theoretical shortcoming highlighted by the impossibility of the mode to reach a satisfactory fit. It seems that the scope of existing theories does not extend to situations in which the decision-maker is guided by fairness alone. However, this conclusion is based on a negative finding and caution is warranted. It is advisable that further research explores other potential boundary conditions highlighting the scope of existing theories.

5.2. Fairness (and Other Processes) in Competitive Games

Another contribution of the present research relates to fairness in the context of competitive games. Specifically, our investigation supports the notion that fairness is not as clear-cut as classical economic accounts suggest. Indeed, fairness boundaries are wider than expected, as the model with the "lenient" fairness conceptualization always outperformed the one embedding a more "strict" concept of fairness. Therefore, our results are in line with a conceptualization of fairness that is more liquid than what perfectly distributive justice would prescribe (Hardin, 1968; Schroeder et al., 2003). Specifically, we speculate that decision-makers are willing to tolerate mild deviance from a perfectly even offer (i.e., 5:5), as the boundaries of what is "fair" include uneven splits. Moreover, toleration for such mild deviance has also been documented in one-shot games (Suleiman, 1996), suggesting that such a tolerance is present even without the expectation to be compensated in the future. However, building expectations that span across several (future) interactions could be an additional process in repeated games. As several Ultimatum Game rounds are played, a reasonable decision-maker might accept an offer deviating slightly from an even split, expecting a compensatory offer in later rounds. This explanation, which is ad hoc for repeated games, allows for construing fairness in a more interactive way that is spread over time (i.e., over multiple rounds). Such a theorization resonates more with social norms and an implicit form of procedural/restorative justice (Schroeder et al., 2003). However, further research is required to explicitly investigate this account.

5.3. Limitations and Future Directions

As with any research, the present investigation is not free from limitations. For example, one of such limitations concerns the nature of the rewards used. Even though the number of resources accumulated by participants was indeed converted into real chances of winning the raffle, allowing participants to "compete" for real money might have triggered alternative motivations. Moreover, the total amount of money at stake in our experiment is quite small. It remains unclear how the processes we investigate might unfold when more resources are at stake. Further research might involve greater amounts of money directly provided to participants to test if our results hold under such conditions. Additionally, our experimental situation placed participants, both receivers and decision-makers, at a "safe" distance from the proposers. Indeed, embodying proposers with real participants, physically present in the same room as test subjects, might have unexpected consequences. Being in the presence of a social other actively violating the norm of fairness might trigger stronger reactance and consequently increased rejection. Similarly, the same situation might lead to increased acceptance due to participants' intimidation caused by bold proposers that do not hesitate to violate fairness if it serves their own purpose. Which of the two opposing reactions is more likely to happen is an empirical question that cannot be addressed by an online interaction, which lacks the simultaneous presence of both the proposer and the receiver/decision-maker in the same physical space.

These limitations, however, are quite common in the reviewed literature and fall outside of the scope of the present research. Addressing these limitations is undoubtedly a viable avenue to extend our findings, which, however, already contribute to the existing literature on economic decision-making.

Author Contributions: Conceptualization, M.B., M.H. and L.O.; methodology, M.B. and M.H., formal analysis, M.B.; resources, M.B.; data curation, M.B.; writing—original draft preparation, M.B.; writing—review and editing, M.H. and L.O., visualization, M.B.; supervision, M.H.; project administration, M.B. and L.O.; funding acquisition, M.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the European Association of Social Psychology that awarded a Seedcorn Grant (2020) to Marco Biella. The APC was funded by the Publication Fund of the University of Basel for Open Access.

Data Availability Statement: Data and materials (R scripts to reproduce the analysis) are available on the OSF repository (https://osf.io/tzbpd/).

Acknowledgments: The authors thank Nihels Kukken for the input on the analysis.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Translation of Participants' Instructions

Experiment 1

Your task is to decide on the Proposer's offer. If you accept the offer, you will both be credited with the amount of money specified in the proposal. If you reject the offer, neither of you will be credited with any money in this round. The Proposer will only receive feedback about the money they received at the end of the entire game and not after each round.

Your chances of winning the raffle at the end of the study will be calculated based on the amount of money accumulated. For every euro you receive in the study, you will receive one ticket for the raffle.

Experiment 2

Your task is to decide on the Proposer's offer on behalf of the Receiver. If you accept the offer, the Proposer and the Receiver will both be credited with the amount of money specified in the proposal. If you reject the offer, neither the Proposer nor the Receiver will be credited with any money in this round. The Proposer will only receive feedback about the money they received at the end of the entire game and not after each round.

The Receiver's and the Proposer's chances of winning the raffle at the end of the study will be calculated based on the amount of money accumulated. For every euro the Receivers and the Proposer obtain in the study, they will receive one ticket for the raffle.

Experiment 3

Your task is to decide as a "fair judge". If you accept the offer, the Proposer and the Receiver will both be credited with the amount of money specified in the proposal. If you reject the offer, neither the Proposer nor the Receiver will be credited with any money in this round. The Proposer will only receive feedback about the money they received at the end of the entire game and not after each round.

The Receiver's and the Proposer's chances of winning the raffle at the end of the study will be calculated based on the amount of money accumulated. For every euro the Receivers and the Proposer obtain in the study, they will receive one ticket for the raffle.

References

Aina, C., Battigalli, P., & Gamba, A. (2020). Frustration and anger in the Ultimatum Game: An experiment. *Games and Economic Behavior*, *122*, 150–167. [CrossRef]

Alexander, R. D. (1987). *The biology of moral systems*. Routledge.

Biella, M., Rebholz, T. R., Holthausen, M., & Hütter, M. (2023). The interaction game: A reciprocity-based minimal paradigm for the induction of social distance. *Journal of Applied Social Psychology*, *53*(8), 796–814. [CrossRef]

Biella, M., & Sacchi, S. (2018). Not fair but acceptable... for us! Group membership influences the tradeoff between equality and utility in a Third Party Ultimatum Game. *Journal of Experimental Social Psychology*, *77*, 117–131. [CrossRef]

Boehm, C. (1999). *Hierarchy in the forest: Egalitarianism and the evolution of human altruism*. Harvard University Press.

Bolton, G. E. (1991). A comparative model of bargaining: Theory and evidence. *The American Economic Review*, *81*, 1096–1136.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, *90*(1), 166–193. [CrossRef]

Camerer, C., & Thaler, R. H. (1995). Anomalies: Ultimatums, dictators and manners. *Journal of Economic Perspectives*, *9*(2), 209–219. [CrossRef]

Civai, C. (2013). Rejecting unfairness: Emotion-driven reaction or cognitive heuristic? *Frontiers in Human Neuroscience*, *7*, 126. [CrossRef] [PubMed]

Civai, C., Corradi-Dell'Acqua, C., Gamer, M., & Rumiati, R. I. (2010). Are irrational reactions to unfairness truly emotionally-driven? Dissociated behavioural and emotional responses in the Ultimatum Game task. *Cognition*, *114*(1), 89–95. [CrossRef]

Civai, C., Rumiati, R. I., & Rustichini, A. (2013). More equal than others: Equity norms as an integration of cognitive heuristics and contextual cues in bargaining games. *Acta Psychologica*, *144*(1), 12–18. [CrossRef] [PubMed]

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. Routledge.

Conway, P., & Gawronski, B. (2013). Deontological and utilitarian inclinations in moral decision making: A process dissociation approach. *Journal of Personality and Social Psychology*, *104*(2), 216. [CrossRef]

Darwin, C. (1871). *The descent of man*. D. Appleton.

Debove, S., Baumard, N., & André, J. B. (2016). Models of the evolution of fairness in the ultimatum game: A review and classification. *Evolution and Human Behavior*, *37*(3), 245–254. [CrossRef]

Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, *114*(3), 817–868. [CrossRef]

Foldnes, N., & Henning Olsson, U. (2015). Correcting too much or too little? The performance of three chi-square corrections. *Multivariate Behavioral Research*, *50*(5), 533–543. [CrossRef]

Forsythe, R., Horowitz, J. L., Savin, N. E., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic Behavior*, *6*(3), 347–369. [CrossRef]

Güth, W., Schmittberger, R., & Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, *3*(4), 367–388.

Hardin, G. (1968). The tragedy of the commons: The population problem has no technical solution; it requires a fundamental extension in morality. *Science*, *162*(3859), 1243–1248. [CrossRef] [PubMed]

Harrison, G. W., & McCabe, K. A. (1996). Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory*, *25*, 303–327. [CrossRef]

Haruvy, E., & Roth, Y. (2022). On the impact of an intermediary agent in the ultimatum game. *Games*, *13*(3), 43. [CrossRef]

Hennig, M., & Hütter, M. (2020). Revisiting the divide between deontology and utilitarianism in moral dilemma judgment: A multinomial modeling approach. *Journal of Personality and Social Psychology*, *118*(1), 22. [CrossRef]

Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, *7*(3), 346–380. [CrossRef]

Hu, X., & Batchelder, W. H. (1994). The statistical analysis of general processing tree models with the EM algorithm. *Psychometrika*, *59*(1), 21–47. [CrossRef]

Hütter, M., & Klauer, K. C. (2016). Applying processing trees in social psychology. *European Review of Social Psychology*, *27*(1), 116–159. [CrossRef]

Hütter, M., & Sweldens, S. (2018). Dissociating controllable and uncontrollable effects of affective stimuli on attitudes and consumption. *Journal of Consumer Research*, *45*(2), 320–349. [CrossRef]

Hütter, M., Sweldens, S., Stahl, C., Unkelbach, C., & Klauer, K. C. (2012). Dissociating contingency awareness and conditioned attitudes: Evidence of contingency-unaware evaluative conditioning. *Journal of Experimental Psychology: General*, *141*(3), 539. [CrossRef]

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1986). Fairness and the assumptions of economics. *Journal of Business*, *59*, S285–S300. [CrossRef]

Kamas, L., & Preston, A. (2012). Distributive and reciprocal fairness: What can we learn from the heterogeneity of social preferences? *Journal of Economic Psychology*, *33*(3), 538–553. [CrossRef]

Klauer, K. C., Stahl, C., & Voss, A. (2011). Multinomial models and diffusion models. In K. C. Klauer, A. Voss, & C. Stahl (Eds.), *Cognitive methods in social psychology* (pp. 367–390). The Guilford Press.

Kravitz, D. A., & Gunto, S. (1992). Decisions and perceptions of recipients in ultimatum bargaining games. *The Journal of Socio-Economics*, *21*(1), 65–84. [CrossRef]

Messick, D. M. (1995). Equality, fairness, and social conflict. *Social Justice Research*, *8*(2), 153–173. [CrossRef]

Messick, D. M., & Schell, T. (1992). Evidence for an equality heuristic in social decision making. *Acta Psychologica*, *80*(1–3), 311–323. [CrossRef]

Messick, D. M., & Thorngate, W. B. (1967). Relative gain maximization in experimental games. *Journal of Experimental Social Psychology*, *3*(1), 85–101. [CrossRef]

Pillutla, M. M., & Murnighan, J. K. (1996). Unfairness, anger, and spite: Emotional rejections of ultimatum offers. *Organizational Behavior and Human Decision Processes*, *68*(3), 208–224. [CrossRef]

Powell, D. A., & William, D. (2001). The robustness of the likelihood ratio chi-square test for structural equation models: A meta-analysis. *Journal of Educational and Behavioral Statistics*, *26*(1), 105–132. [CrossRef]

Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, *83*, 1281–1302.

Riefer, D. M., & Batchelder, W. H. (1988). Multinomial modeling and the measurement of cognitive processes. *Psychological Review*, *95*(3), 318. [CrossRef]

Ruessmann, J. K., & Topolinski, S. (2020). Economic decisions for others are more favorable for close than distant clients. *Personality and Social Psychology Bulletin*, *46*(3), 393–407. [CrossRef] [PubMed]

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, *300*(5626), 1755–1758. [CrossRef] [PubMed]

Schroeder, D. A., Steel, J. E., Woodell, A. J., & Bembenek, A. F. (2003). Justice within social dilemmas. *Personality and Social Psychology Review*, *7*(4), 374–387. [CrossRef] [PubMed]

Singmann, H., & Kellen, D. (2013). MPTinR: Analysis of multinomial processing tree models in R. *Behavior Research Methods*, *45*(2), 560–575. [CrossRef]

Suleiman, R. (1996). Expectations and fairness in a modified ultimatum game. *Journal of Econonomic Psychology*, *17*, 531–554. [CrossRef]

Suleiman, R. (2017). Economic harmony: An epistemic theory of economic interactions. *Games*, *8*(1), 2. [CrossRef]

Suleiman, R. (2022). Economic harmony—A rational theory of fairness and cooperation in strategic interactions. *Games*, *13*(3), 34. [CrossRef]

Thaler, R. H. (1988). Anomalies: The ultimatum game. *Journal of Economic Perspectives*, *2*(4), 195–206. [CrossRef]

Van't Wout, M., Kahn, R. S., Sanfey, A. G., & Aleman, A. (2006). Affective state and decision-making in the ultimatum game. *Experimental Brain Research*, *169*(4), 564–568. [CrossRef] [PubMed]