Strategic ambiguity and risk in alternating pie-sharing experiments

Anna Conte[1] · Werner Güth[2] · Paul Pezanis-Christou[3]

Abstract

We experimentally study a class of pie-sharing games with alternating roles from a decision-making perspective. For this, we consider a variant of a two-stage alternating-offer game which introduces an imbalance in the protagonists' bargaining powers. This game class enables us to investigate how exposure to risk and strategic ambiguity affects one's bargaining behaviour. Two structural econometric models of behaviour, a *naïve* and a *sophisticated* one, capture remarkably well the observed deviations from the game-theoretic benchmark. Our findings indicate, in particular, that a higher exposure to strategic ambiguity leads to a behaviour that is less responsive to the game's parameters and to distorted, yet consistent, beliefs about other's behaviour. We also find evidence of a backward-reasoning whereby first-stage decisions relate to the second-stage ones but which do not call for the counterfactual reasoning that characterises rationality in such settings.

Keywords Strategic ambiguity · Ambiguity aversion · *a*-insensitivity · Risk · Probability distortion · Alternating-offer bargaining · Impunity

JEL classification: C33 · C34 · C78 · C91 · D80

✉ Anna Conte
anna.conte@uniroma1.it

Werner Güth
gueth@coll.mpg.de

Paul Pezanis-Christou
paul.pezanis-christou@adelaide.edu.au

[1] Department of Statistical Sciences, Sapienza University of Rome, Viale Regina Elena 295G, Rome 00161, Italy

[2] Max Planck Institute for Collective Goods (Bonn), Kurt-Schumacher-Str. 10, Bonn 53113, Germany

[3] School of Economics and Public Policy, University of Adelaide, Pulteney St. 10, Adelaide 5000, South Australia, Australia

🖄 Springer

1 Introduction

Consider a lottery yielding a positive prize N with probability p and a 0-payoff otherwise. The monetary expectation of this lottery, Np, depends symmetrically on N and p. While the literature has extensively reported on how people perceive probabilities and evaluate monetary gains and losses in such lottery contexts (see, e.g., Wakker, 2010), little is known about the way people approach N and p in strategic contexts, and in particular how they trade them off in a multi-stage setting. In this paper, we investigate these questions in the context of alternating-offer bargaining games, *à la* Ståhl (1972) and Rubinstein (1982), which have been extensively studied in the laboratory (see Roth, 1995 for an early survey). These games typically feature a finite number of stages of bargaining to share a pie according to Ultimatum Game (UG, henceforth) rules with alternating roles of proposer and responder. Rubinstein (1982) solves such games via backward induction and predicts a first-stage agreement which, in a two-stage setting with a second-stage pie of size N and a probability p that the second stage occurs, consists in offering (and accepting) the monetary expectation of the above lottery example.

Despite the simplicity of this prediction, the experimental evidence on these games reports systematic deviations (i.e., larger-than-expected offers and frequent first-stage disagreements) which have been explained in terms of payoff-based social preferences (Ochs & Roth, 1989; Bolton, 1991; Goeree & Holt, 2000; Cooper & Kagel, 2016), reputation (Embrey et al., 2015), or failures to recognise that sequential interaction requires backward induction reasoning (Binmore et al., 1985; Neelin et al., 1988; Binmore et al., 1988). The latter argument has been further investigated by Binmore et al. (2002) and Johnson et al. (2002) who show that even though players may have fairness concerns, their failure to backward induct mostly results from a limited cognition that impedes game-theoretic reasoning. Ho & Su (2013) propose a dynamic *level-k* model to rationalise these failures and attribute them to players' limited inductive reasoning and/or their inability to update their beliefs about others' depth of reasoning as the game unravels. A common feature of the above rationales is their reliance on some form of expected utility maximisation.

In this article, we forego the notion of expected utility maximisation altogether and experimentally investigate whether behaviour in two-stage bargaining games can be rationalised in terms of risk (induced by a known probability and resolved by nature) and strategic ambiguity (induced by ignorance of the other's intentions and the probability of each alternative being chosen, and resolved by the other's decision). In the setting we study, "risk" is captured by the probability p that a second stage occurs if no first-stage agreement is reached. There is ample evidence that, besides nonlinear evaluations of monetary payoffs, people's probability misperceptions can affect their decisions and expectations when confronting risky decisions (see e.g., Tversky & Kahneman, 1992; Quiggin, 1982; Barberis, 2012). In this regard, the players' misperceptions of the probability p and attitudes to risk may well affect their expectations and drive behaviour in

our bargaining games. As for "strategic ambiguity," it results from the lack of some critical knowledge to entertain beliefs about others' behaviour. Its effect is typically put in perspective by comparing behaviour in treatments that manipulate players' ambiguity-generating information and/or by eliciting their preferences or cognitive skills. Although this type of strategic uncertainty has proven very useful to organise behaviour in experiments on coordination, labour-markets, public goods, asset markets, and one-stage bargaining games (see, e.g., Heinemann et al., 2009; Cabrales et al., 2010; Di Mauro & Finocchiaro Castro, 2011; Li et al., 2017; Akiyama et al., 2017; Greiner, 2018), it remains unclear whether it results from players' aversion to ambiguous events (which is emotional and reflects people's dislike of ambiguous situations, see Li, 2017, p. 245) or from their beliefs about others' behaviour (which is cognitive). Li et al. (2019) address this question in the context of trust games and propose belief-free measures of ambiguity aversion and of a(mbiguity generated)-insensitivity, i.e., an insensitivity to likelihood changes. They show, in particular, that strategic ambiguity has a two-fold effect on behaviour: it can trigger ambiguity-aversion, which leads players to prefer using safe strategies, and/or a-insensitivity, which makes them less likely to act based on their beliefs (see also Li et al., 2020). Although these measures do not suit the analysis of games with a richer interaction like bargaining games with alternating roles, the identified two-fold effect provides a useful grasp to study the roles of strategic ambiguity and risk in a simplified version of these games.

To best highlight the effect of strategic ambiguity, we assign dictatorial power to the second-stage proposer, i.e., her/his payoff is not affected by the second-stage responder's decision, as in an Impunity Game (Bolton, 1991, IG henceforth). By doing so, we leave the game-theoretic prediction unchanged and we remove all cognitive difficulty for the second-stage proposer to decide what to offer. That is, the first-stage responder only has to deal with the strategic ambiguity of the first stage and the uncertainty (risk) of a second stage of interaction, whereas the first-stage proposer has to cope with the additional strategic ambiguity of the second stage.

Next, we assess the roles of ambiguity aversion and a-insensitivity in these games with two structural models that are inspired from a descriptive analysis of the observed behaviour and that forego expected utility maximisation. The simplest, *naïve*, model assumes that players perceive the two stages as being more or less dependent on the second-stage parameters (p and N) so their strategies are belief-free (or safe). This model only aims at capturing ambiguity aversion concerns. The *sophisticated* model, which allows for probability misperceptions and for a link between the first- and second-stage decisions, aims at capturing both ambiguity aversion and a-insensitivity, and it nests the game-theoretic model with backward induction as a special case.

Finally, we document these behavioural traits in a relatively large game class by confronting participants to multiple (N, p)-constellations that characterise a range of dilemmas that spans from a UG-like game (i.e., when both N and p are small) to Rubinstein-like games (i.e., when p is close to one). For each of these constellations, we use the strategy method to elicit dyadic decisions (i.e., an offer for the first stage and an acceptance threshold for the second stage or *vice versa*) and provide no

feedback on the stage or round outcome to avoid the confounding effect of the players' repeated interactions on our conclusions.

We summarise our findings in the following three points. First, observed behaviour is best fit by a variant of the *sophisticated* model which supports a "backward reasoning" for the determination of first-stage decisions. Unlike the fully rational backward induction that underlies the game-theoretic prediction, this reasoning is based on what players expect to receive or to offer given their respective roles in each stage (proposer or responder), the likelihood of a second stage (p), and the second-stage pie (N). The model fits the first-stage behaviour remarkably well and indicates that proposers tend to offer more than what responders request. This holds when participants with an almost invariant behaviour are discarded from the estimations and thus suggests the use of safe strategies that are in line with the ambiguity aversion induced by the players' roles. First-stage proposers are also found to be substantially more a-insensitive than first stage responders, so the participants' first-stage behaviour is consistent with their respective exposures to strategic ambiguity. Yet, first-stage disagreements occur but these are mostly observed when both N and p are large, i.e., about 40% of the time.

Second, the model's estimates indicate that the higher the first-stage offers (requests), the higher the second-stage requests (offers). This pattern is positively related to the players' a-insensitivities, and thus, in line with their exposures to strategic ambiguity. Requesting or offering positive amounts in the second stage is clearly suboptimal but in line with the "emotional commitment" rationale of Yamagishi et al. (2009) for such behaviour in single-stage IG experiments: players do so because they emotionally commit to what they consider to be a fair share of the pie to request or to offer.[1] In this regard, the players' backward reasoning organises their first-stage decisions and supports the role of strategic ambiguity in explaining the observed misbehaviour.

Third, the estimated model predicts that on average second-stage interactions end in disagreement as is observed for 60% of the time after a first-stage disagreement when both N and p are large. However, when the estimations discard participants with an almost invariant behaviour in the first stage, the model predicts virtually identical second-stage offers and requests. Since players received no feedback on the stage- or round-outcomes, this suggests that despite the players' asymmetric exposures to strategic ambiguity (which affects their a-insensitivity), their beliefs about each other's behaviour are consistent on average and that those with an almost invariant behaviour in the first stage failed to match the other's choices because they were most ambiguity averse.

The game class we study and the strategic ambiguity considerations we pursue in this study are presented in the next section. Section 3 discusses the experimental protocols and procedures. Section 4 reports a descriptive analysis of observed behaviour.

---

[1] It also follows from Yamagishi et al. (2009) that this argument cannot be confounded with signalling anger (by responders) or generosity (by proposers) because participants were not informed of eventual rejections.

Section [5] spells out the *naïve* and *sophisticated* structural models and reports the estimation results, and Section [6] concludes.

2 The game class

The game class we study involves two players, A and B, who bargain over a prize of 100 and who both know that a second-stage of bargaining over a prize of $N < 100$ might occur with probability p if no agreement is reached in the first stage. In what follows, we describe the sequence of individual decisions to be taken by each player at each stage.

The first-stage proposer (henceforth A) is asked to offer an amount y, with $0 < y < 100$, to the first-stage responder (henceforth B) who is asked to determine an acceptance threshold $\underline{y}$, with $0 < \underline{y} < 100$. If $y \geq \underline{y}$, then there is a first-stage agreement that B earns the amount y and A the residual $100 - y$. If $y < \underline{y}$, then there is no agreement and both parties earn zero payoffs with probability $(100 - p)$, where $0 < p < 100$ stands for the probability that a second stage occurs (expressed in percentage for convenience). With probability p, A and B reverse their roles and engage in a second stage to share $0 < N < 100$ as B decides: B can offer any amount x with $0 < x < N$ to A and keeps the remainder $N - x$ for herself, whereas A can only choose an acceptance threshold $\underline{x}$, with $0 < \underline{x} < N$. If $x \geq \underline{x}$, then A gets x, and B keeps $N - x$. Any offer x smaller than $\underline{x}$ is lost for both parties.

Denoting by $\mathbb{1}_{(y \geq \underline{y})}$ and $\mathbb{1}_{(x \geq \underline{x})}$ the indicator functions, taking the value 1 for $y \geq \underline{y}$ and $x \geq \underline{x}$, respectively, and 0 otherwise, the players' payoffs are defined as:

$$\mathbb{1}_{(y \geq \underline{y})} \times (100 - y) + \left[1 - \mathbb{1}_{(y \geq \underline{y})}\right] \times \begin{cases} \mathbb{1}_{(x \geq \underline{x})} \times x, & \text{with probability } p \\ 0, & \text{with probability } (100 - p) \end{cases} \quad \textit{for A}$$

and

$$\mathbb{1}_{(y \geq \underline{y})} \times y + \left[1 - \mathbb{1}_{(y \geq \underline{y})}\right] \times \begin{cases} (N - x), & \text{with probability } p \\ 0, & \text{with probability } (100 - p) \end{cases} \quad \textit{for B.}$$

The determination of the players' optimal (equilibrium) decisions invokes expected-profit maximisation and backward induction. Thus, in the second stage, A should accept any positive offer x and set $\underline{x} = 0$, whereas B should keep as much as possible of N so that B's (second-stage) expected payoff Np determines B's acceptance threshold in the first stage. Moving (backwards) to the first stage, A foresees her second-stage payoff of 0 and therefore offers $y = Np$ in the first stage. Any offer $y < Np$ would be rejected by B and would drag A into a probable second stage that yields a 0 payoff, whereas any offer $y > Np$ would be overly generous (to B), i.e., A could further increase her payoff by offering some y closer to Np. This defines the players' optimal agreement which is to be reached in the first-stage. Thus, given (N, p) the optimal offer and acceptance threshold are:

$$\underline{x}^{opt} = x^{opt} = 0 \qquad \text{and} \qquad \underline{y}^{opt} = y^{opt} = Np.$$

2.1 Strategic ambiguity considerations

The literature offers a profusion of definitions for "strategic ambiguity" and we will summarily define it as characterising an interactive decision context where ambiguity emanates from the fact that the probability distributions of others' decisions are unknown and resolved by others' decisions. Such strategic ambiguity can be game-theoretically resolved by assuming common(ly) known rationality, but it emerges peremptorily when that assumption is relaxed.

Strategic ambiguity is inherent to our context and leads to two broad types of modelling approaches. The first maintains the assumption of expected utility maximisation and relaxes full rationality by allowing for inconsistent beliefs about others' behaviour. This applies, for example, to the case of reinforcement-learning models (see Nawa et al. 2002), and of the *level-k* models of Stahl & Haruvy (2008) (for UG-like settings) and of Ho & Su (2013) (for sequential games) which assume that players best-respond to protagonists who are less rational than themselves, i.e., whose depth-of-reasoning is one step lower than their own.[2]

The second, which we pursue here, is to forego the assumption of expected utility maximisation and revert to the use of simple heuristics like *(i)* treating the decisions to be taken at each stage as being more or less independent of each other, *as if* players opted for a "complexity reduction," see Kovářák et al. (2016), *(ii)* opting for an equal-split of the pie, as Binmore et al. (1985) and Stahl & Haruvy (2008) report in their experiments, or *(iii)* emotionally responding to the game's parameters (*N* and *p*), as in Yamagishi et al. (2009). Alternatively, one could revert to the formulation of beliefs or predictions about others' decisions which may be biased by probability misperception. These heuristic- and belief-based alternatives are not necessarily mutually exclusive and may complement each other in many possible ways. For example, although the use of heuristics (or safe strategies) reflects a behaviour that maintains ambiguity aversion, it does not prevent the formation of possibly distorted beliefs about others' behaviour or probabilistic events which would reflect a player's *a*-insensitivity. The experiments aim at checking whether these two effects help organising behaviour in these games.

The players' respective dilemmas in this game class can be summarised as follows. As *p* increases, the second stage becomes a less risky alternative to be considered by first-stage responders especially when *N* is large, whereas for first-stage proposers it becomes more likely that their payoff will entirely depend on the good-will of the second-stage proposers if a second stage occurs. Thus, the players' incentives to reach a first-stage agreement are not aligned, especially when both *p* and *N* become large; and this leads first-stage proposers to trade off their strategic ambiguity with ultimatum power (in the first stage) for risk and strategic ambiguity with no

---

[2] Crawford et al. (2013) review more elaborate *level-k* models (for one-stage games) in which the players' lower depths-of-reasoning are Poisson distributed. It is also worth noting that the recursive nature of these models render their predictions sensitive to how one defines the behaviour of *level-0* players, i.e., those with a zero depth of reasoning – the latter are typically assumed to display a uniformly random behaviour over the set of available choices.

Table 1 First-stage benchmark predictions

				N (in ECU)		
		5	25	50	75	95
	5	0.2	1.2	2.5	3.7	4.7
	25	1	6	12.3	18.5	23.5
p (in %)	50	2	12	24.5	37	47
	75	3	18	36.8	55.5	70.5
	95	3.8	22.8	46.6	70.3	89.3

veto power (in the second) whereas first-stage responders simply need to trade off strategic ambiguity (in the first stage) for risk (in the second).

3 Experimental procedures

The experiment implements our game class by confronting participants to multiple (N, p)-constellations to assess how they react to changes in N (holding p constant), in p (holding N constant), and in both N and p (holding Np constant).

Table 1 displays the (N, p)-constellations used in the experiments and the corresponding first-stage agreement predictions $(N - 1)p$, where N is expressed in Experimental Currency Units (ECU) and p in percentages. In each round of play, participants are confronted to a randomly drawn (N, p)-constellation, and we use the strategy method to elicit an offer and an acceptance threshold for each stage (see the Instructions and screenshots in Online Resource A). Thus, A specifies an offer $y \in \{1, 2, \ldots, 99\}$ for the first stage and an acceptance threshold $\underline{x} \in \{1, 2, \ldots, N - 1\}$ for the second, whereas B specifies an acceptance threshold $\underline{y} \in \{1, 2, \ldots, 99\}$ for the first stage and an offer $x \in \{1, 2, \ldots, N - 1\}$ for the second. Since participants could only enter positive integer amounts, the optimal second- and first-stage choices are amended as:

$$\underline{x}^{opt} = x^{opt} = 1 \qquad \text{and} \qquad \underline{y}^{opt} = y^{opt} = (N - 1)p,$$

and the payoffs are defined as:

- if $y \geq \underline{y}$ then A earns $100 - y$ and B earns y (both in ECUs);
- if $y < \underline{y}$ then B earns $N - x$ with probability p and 0 probability $(100 - p)$, whereas A earns x in case of $x \geq \underline{x}$ with probability p and 0 otherwise.

To best assess how participants deal with the strategic ambiguity and risk of this game class and to prevent confounding effects from their repeated play and accumulated experience, participants were given no end-stage or end-of-round feedback. They were also explained that their rewards for participating in the experiments would be determined by matching their choices for the 25 rounds with those of another randomly chosen participant (of the other type, A or B) and by accumulating

the payoffs from these 25 interactions. The resulting amount was then converted into Euros at the exchange rate of €0.005 per ECU and added to a show-up fee of €5.

The experiment's only treatment variation relates to the sequence in which participants' choices for each stage were elicited. In a "Forward" (control) treatment, their first- and second-stage choices are elicited in sequence as outlined in the previous section, whereas in a "Backward" treatment, second-stage choices are elicited first to induce participants to think about their respective roles in that stage and document how this affects their respective first-stage choices.

The experiments were conducted at the Max Planck Laboratory of Experimental Economics in Jena (Germany) with a total of 716 participants recruited online from a pool of students in Business Administration, Sciences, and Engineering from the University of Jena.[3],[4] The sessions took no longer than two and a half hours to complete (including the time needed to read instructions) and participants earned on average €10.55 (s.d. 6.00) in the Forward treatment and €10.65 (s.d. 6.10) in the Backward treatment in addition to their show-up fee.

4 Descriptive Data Analysis

We start with providing an overview of A- and B-players' behaviour in each stage of the Forward and Backward treatments to unravel prominent behavioural traits and document the outcomes of their interactions. In this regard, we first check for significant differences in the behaviour of A- and of B-players in the Forward and Backward treatments by conducting, for each player type and role (Proposer and Responder), a constellation-specific Wilcoxon rank-sum test that checks for a significant treatment difference. This generates 25 test outcomes (p-values) for each type of player and role which we use to test the joint null hypothesis that all 25 null hypotheses hold against the alternative that at least one of the null hypotheses is rejected with the sequential Holm-Bonferroni procedure (HB, henceforth, see Holm, 1979). We report the details of test outcomes and of the HB procedure in Online Resource B. As we find no significant difference to report, we pooled the data of the Forward and Backward treatments.

---

[3] The experiments originally involved a total of 308 participants. However, we conducted additional sessions with "perturbed" (N, p)-constellations, i.e., sessions where each of the 25 (N, p)-constellations of Table 1 had symmetric deviations of size −2, 0, or +2 in N and in p, to confront participants to an even richer environment with 75 possible constellations. These sessions took about half an hour longer to complete and involved an additional 408 participants who received a show-up fee of €2.5. Since the data of (N, p)-constellations with a 0-deviation displayed no significant difference with the original data on sessions with 25 ("unperturbed") constellations, we included it in our analysis, see Table in Online Resource B for test outcomes on the samples' comparison.

[4] The experiments were programmed in *z-tree* (Fischbacher, 2007) and participants were recruited using ORSEE (Greiner, 2015).
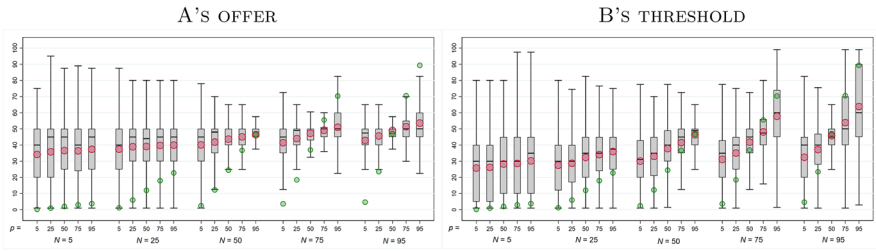
Fig. 1 A-players' offers and B-players' thresholds in stage 1.
Note: Box plots of A-player's offer and B-player's threshold by N and p. Red (large) and green (small) circles represent means and benchmark $((N-1)p)$ predictions, respectively

4.1 First-stage behaviour

Figure 1 displays box plots with the superimposed means (red/large circles) of A-players' offers and B-players' thresholds for each cell of Table 1 when constellations are sorted by N and p. The plots document three salient behavioural patterns.

First, both mean offers and thresholds closely match predictions in only two constellations: in (50, 95), where deviations are not significant for B-players only (*t* test: *p*-value = 0.183), and in (95, 50) where they are not significant for both player types (*t* test: *p*-values = 0.782 for A and = 0.838 for B). Note that these constellations characterise settings where "the stakes are full ($N = 95$) and the chances are half ($p = 50$)," or *vice versa* and for which the solution is an almost equal division of the original pie of 100 (47), so this finding could simply obtain from the use of prominent reference measures (i.e., "maximum and half") that "naturally lead" to an (almost) equal split of 100.

Second, mean offers and thresholds are always greater than predicted when $(N-1)p < 50$ and always lower otherwise. The deviations from equilibrium predictions are also inversely related to the benchmark predictions when $(N-1)p < 50$ and they increase otherwise, which suggests that both player types share some behavioural traits that render the (50, 95)- and (95, 50)-constellations pivotal. This is supported by the finding that, on average, offers are typically higher than thresholds when $(N-1)p < 50$ and lower (implying disagreement) otherwise, and the "offer-threshold" absolute differences are inversely related to the benchmark predictions when $(N-1)p < 50$ and are increasing otherwise.

Third, B-players are more responsive to changes in N and p than A-players, and especially to changes in p when $N \geq 50$. In this case, the bulk of offers is around 50 no matter N or p, whereas for a given N, thresholds sharply increase with p. Further evidence on this asymmetric responsiveness to the game's parameters can be found in the players' frequencies of invariant choices reported in Table 2. The figures reveal that 12% of A-players appear to always make the same offer (11% always offer 50), whereas only 6% of B-players always request the same amount (2% always request 50). This holds when considering participants with at least 80% of invariant choices since 23% of A-players and 15% of B-players then still do so, with 20% offering 50 and 4% requesting 50 so that A-players are more prone to overlook important information about the games they play than B-players. Such invariant

Table 2 Invariance of behaviour to N and p in stage 1

	A's offer		B's threshold	
Invariance % *	100%	≥ 80%	100%	≥ 80%
# participants (%)	42/358 (11.7%)	82/358 (22.9%)	21/358 (5.9%)	55/358 (15.4%)
Modal y and $\underline{y}$				
40	–	–	42.86%	32.73%
50	95.24%	89.02%	33.33%	29.09%
other numbers	4.76%	10.98%	23.81%	38.18%

*: % of times a participant submits/requests the same amount

behaviour is consistent with the loads of strategic ambiguity and risk induced by the players' roles in each stage. Indeed, recall that while both player types face risk (via p), A-players must also deal with strategic ambiguity in both stages, whereas B-players only do so in the first stage. Thus, following Li et al. (2019), this invariant behaviour can be expected to relate to the players' insensitivity to the game's parameters, N and p, that is induced by their exposure to strategic ambiguity. In this regard, since B-players experience less strategic ambiguity than A-players, they are also expected to react more to the game's parameters, as we find.

4.2 Second-stage behaviour

Figure 2 displays box plots with superimposed means (red/large circles) of A-players' thresholds and B-players' offers against N and p. The plots indicate that both requests and offers are greater than the benchmark prediction of 1 and that they both significantly increase with N. However, conditional on N, the deviations from predictions are largely invariant to p so the players' behaviour is not contingent on irrelevant information like p. Furthermore, A-players' average thresholds are all higher than B-players' average offers so that "money is typically left on the table", and this is more likely as N increases.[5]

In terms of invariant behaviour, participants can be sorted into three categories: those who accept anything and offer 1 whatever N (14.8% of A-players and 28.72% of B-players); those who always request or offer more than 1 ECU (34.1% and 23.5%, respectively); and the remainder who request or offer 1 ECU when N is small and more otherwise (51.1% and 47.8%, respectively). We also checked for "equal-splitters" who only offer or request $N/2$ but found no compelling evidence of such behaviour; and the few who did hardly ever did so in the first stage, as Stahl & Haruvy (2008) conjecture.

In sum, over two-thirds of participants display a second-stage behaviour that is inconsistent with expected utility maximisation and suggests that A-players'

---

[5] See Tables C.3 and C.4 in Online Resource for detailed (constellation-specific) statistics.
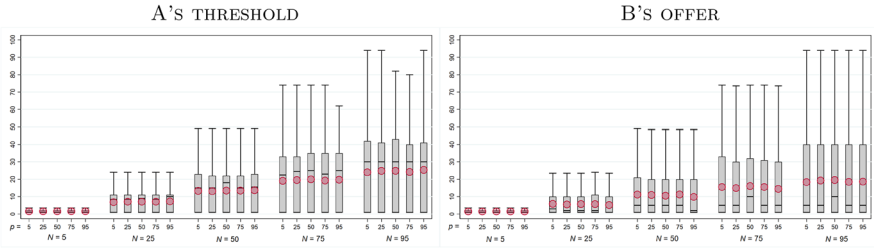
Fig. 2 A's thresholds and B's offers in stage 2.
Note: Box plots of A-player's offer and B-player's threshold by *N* and *p*. Red circle represent means. The benchmark prediction is 1 for all constellations

thresholds reveal (or relate to) their beliefs about what they expect to receive from B-players.

4.3 Disagreement rates

An important economic aspect of the game class we study is the extent to what the (N, p)-constellations considered yield inefficient outcomes, i.e., a first-stage disagreement. To assess disagreement rates, we recombined participants' choices for each (N, p)-constellation 25,000 times and report in Table 3 the mean and standard deviation of the proportion of matches under each constellation that ended in conflict.

The outcomes for the first-stage interactions indicate an overall disagreement rate of 32.2% and suggest three distinct clusters of constellations: those with $N \geq 75$ and $p \leq 25$ which generate a low disagreement rate of about 24%, those with *N* and

Table 3 Disagreement rates

		First stage						Unconditional second stage			
				N						*N*		
		5	25	50	75	95		5	25	50	75	95
	5	0.320	0.284	0.263	0.249	0.235	5	0.306	0.440	0.451	0.460	0.489
		(.015)	(.014)	(.014)	(.014)	(.014)		(.013)	(.014)	(.014)	(.014)	(.013)
	25	0.298	0.263	0.247	0.235	0.235	10	0.319	0.459	0.447	0.486	0.487
		(.014)	(.014)	(.014)	(.014)	(.013)		(.013)	(.013)	(.014)	(.014)	(.013)
p	50	0.307	0.308	0.251	0.272	0.314	*p* 50	0.320	0.438	0.460	0.477	0.483
		(.014)	(.013)	(.013)	(.013)	(.013)		(.013)	(.013)	(.014)	(.014)	(.014)
	75	0.326	0.323	0.302	0.400	0.467	75	0.326	0.448	0.454	0.470	0.483
		(.015)	(.014)	(.013)	(.014)	(.014)		(.012)	(.014)	(.014)	(.014)	(.013)
	95	0.329	0.331	0.369	0.558	0.566	95	0.333	0.483	0.479	0.493	0.486
		(.014)	(.014)	(.013)	(.014)	(.014)		(.013)	(.013)	(.013)	(.013)	(.013)

Figures are calculated by randomly re-matching A- and B-participants for 25,000 times and by taking the average of the resulting frequencies that the (N, p)-constellation considered ended in conflict. Standard deviations are reported in parentheses. Second-stage unconditional disagreement rates assume that the second stage is always reached

$p \geq 75$ which yield about 50% of disagreements, and those with $N \leq 50$ that generate about 31% of disagreements.[6] Interestingly, the 32% chance of a disagreement in the $(5, 5)$-constellation, whose "insignificant" second stage makes it most similar to a UG setting, is also in line with the one-third rejection rate reported by Güth et al. (1982) for their UG experiments.[7]

As for the second-stage interactions (where a disagreement implies a zero payoff for A-players only), we find a lowest disagreement rate of about 31% when $N = 5$ no matter p. This is in line with the 30% reported by Yamagishi et al. (2009) for their one-stage IG experiments. Otherwise, a second-stage disagreement is observed between 40% and 50% of the time, which is remarkably high.[8]

5 Econometric modelling

Having highlighted some salient characteristics of A- and B-players' behaviour in each stage and their consequences on the outcomes, we proceed with proposing two structural econometric models to organise the observed behaviour in the light of the strategic ambiguity and risk. For this, we consider two simultaneous equation models of the players' behaviour in either role, with one equation for the first-stage choices and one equation for the second-stage choices. These models characterise a "naïve" and a "sophisticated" way of dealing with strategic ambiguity and risk, and each of them is separately estimated for each type of player.

In what follows, let $i \in S$ denote the generic A- or B-player in sample S and Ω the set of (N, p)-constellations. Recall that for each (N, p)-constellation, A-players submit an offer y_i for the first stage and an acceptance threshold $\underline{x}_i$ for the second, whereas B-players submit an acceptance threshold $\underline{y}_i$ for the first stage and an offer x_i for the second stage, so for expository convenience we simplify these notations by denoting A-players' choices for first and second stages by $Y_i(N, p)$ and $X_i(N, p)$; and those of B-players for the first and second stages by $Y_i(N, p)$ and $X_i(N, p)$.

5.1 The *naïve* model

This model characterises a "naïve" approach in that it assumes players to perceive each stage in isolation and to possibly take account of the parameters N and/or p in their

---

[6] In their study of legal case outcomes in the US, Ashenfelter & Currie (1990) report an average disagreement rate of 40% which is the one observed in our setting when both the stakes and the chances of a second stage are relatively high, namely $N = p = 75$.

[7] The mean offer and mean threshold for this constellation are also in keeping with what is typically reported in UG settings, see Güth & Kocher (2014).

[8] We also looked at conditional probabilities of second-stage "disagreement" (i.e., conditional on having a disagreement in the first stage and the second stage occurs) and found that they are typically smaller than the unconditional ones when $N < 50$ and typically larger otherwise, what suggests that second-stage "disagreements" are more likely when the stakes are large, see Online Resource D.

first- and second-stage choices. Player i's first-stage choice $Y_i(N,p)$, i.e., an A-player's offer or a B-player's threshold, is modelled as a function of an individual-specific intercept η_i and a polynomial $l(N,p)$ that is common to all players of the same type and that captures the players' possible reactions to the second-stage parameters N and p. It also involves a stochastic term $c_{i(N,p)}$ that captures noise in i's first-stage decisions and that has mean 0 and standard deviation $\sigma_c = \check{\sigma}_c\sqrt{\exp(\kappa Np)}$, where κ captures the extent of heteroscedasticity as a function of Np.

$$\begin{cases} Y_i(N,p) = \eta_i + l(N,p) + c_{i(N,p)} \\ X_i^*(N,p) = \widetilde{X_i(N)} + d_{i(N,p)} = \theta_i + m(N) + d_{i(N,p)} \end{cases} \quad \forall(N,p) \in \Omega \quad (1)$$

$$X_i(N,p) = \max\left[X_i^*(N,p), 1\right] \quad (2)$$

$$\begin{pmatrix} \eta_i \\ \theta_i \end{pmatrix} \sim \text{NID}\left[\begin{pmatrix} \mu_\eta \\ \mu_\theta \end{pmatrix}, \begin{pmatrix} \sigma_\eta^2 & \rho_{\eta\theta}\sigma_\eta\sigma_\theta \\ \rho_{\eta\theta}\sigma_\eta\sigma_\theta & \sigma_\theta^2 \end{pmatrix}\right] \quad (3)$$

$$\begin{pmatrix} c_{i(N,p)} \\ d_{i(N,p)} \end{pmatrix} \sim \text{NID}\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_c^2 & \rho_{cd}\sigma_c\sigma_d \\ \rho_{cd}\sigma_c\sigma_d & \sigma_d^2 \end{pmatrix}\right] \quad (4)$$

$$\sigma_c = \check{\sigma}_c\sqrt{\exp(\kappa Np)} \quad (5)$$

$$\sigma_d = \check{\sigma}_d\sqrt{\exp(\lambda N)} \quad (6)$$

The second stage choice $X_i^*(N,p)$ is modelled via the individual-specific intercept θ_i and a polynomial $m(N)$ that is common to all players of the same type and that accounts for players' eventual nonlinear evaluations of N. Most importantly, the model assumes that $\widetilde{X_i(N)}$ is indirectly observed via a latent function $X_i^*(N,p) = \widetilde{X_i(N)} + d_{i(N,p)}$, which is inferred from the data through the observation rule $X_i(N,p) = \max\left[X_i^*(N,p), 1\right]$. This latent function captures a player's expectation for the second stage and assumes that this expectation does not affect the player's first-stage choice. We treat A- and B-players' choices for the second stage as being left-censored at 1. Noise in second-stage choices is modelled via $d_{(N,p)}$ with mean 0 and standard deviation $\sigma_d = \check{\sigma}_d\sqrt{\exp(\iota N)}$, where $\exp(\iota N)$ controls for heteroscedasticity as a function of N via the parameter ι. Finally, the intercepts η_i and θ_i have a joint normal distribution, like the stochastic terms $c_{i(N,p)}$ and $d_{i(N,p)}$ and these joint normal distributions (with μ, σ and ρ standing for the mean, standard deviation and correlation coefficient, respectively) are estimated along with the other parameters of the model.[9]

---

[9] The joint distribution of the model's structural parameters has been estimated as in Conte et al. (2011) who study individual choice under risk. We refer the interested reader to Online Resource E for further estimation details, including the derivation of the likelihood function.

5.1.1 The *naïve* model: Outcomes

The left panel of Table 4 collects the results of the *naïve* model when the estimations refer to the whole data set. The reported specifications have been selected via likelihood-ratio tests for nested alternatives and Akaike Information Criteria (AIC) for non-nested ones.

First-stage behaviour The estimates of μ_η, the mean of the first-stage random intercept, reflect the means of A-players' offers and B-players' thresholds and confirm the finding that offers tend to be significantly higher than requests, cf. Fig. 1. The estimated standard deviations, σ_η, vouch for a substantial heterogeneity in behaviour, especially in B-players' choices. The estimated polynomial $l(N, p)$ reveals a positive and significant effect of N and Np that is similar for both roles, whereas the cumulative effect of p and p^2 indicates that as the likelihood of a second stage increases, offers increase less (and are capped at $p = 0.75$) whereas thresholds increase more. Further, a Wald test for the joint significance of the parameters in $l(N, p)$ overwhelmingly rejects the null for both roles and both samples and thus suggests that, on average, participants took account of the second-stage parameters N and p in their first stage decisions.

Second-stage behaviour The means of the second-stage random intercepts, μ_θ, are significantly negative for both roles and their standard deviations, σ_θ, also support a substantial heterogeneity, especially in A-players' choices. As for the estimated polynomial $m(N)$, it is increasing with an almost concave shape in N that is similar across roles. These estimates confirm the patterns reported in Section 4.2, namely that the average A-player always requests more than what the average B-player offers. They also define critical values of N for each player-type, $N_A^* = 20$ and $N_B^* = 40$, at which the average A-player expects to receive and the average B-player intends to give the minimal amount (1 ECU). For N larger than these critical values, the average A-player expects to receive positive amounts that increase with N, what is in keeping with the emotional argument of Yamagishi et al. (2009), and the average B-player intends to offer positive amounts that also increase with N. For N smaller than these critical values, the negative intercept estimates suggest that the average A-player requests more than N and the average B-player offers less than 1 (or equivalently requires a larger pie for her/himself). We could read this as the two players deeming the second-stage pie too small and requesting a compensation if their choice alternatives were not censored at 1.[10]

Stage-to-stage behaviour The significant correlations between the random intercept of the first- and second-stage equations, $\rho_{\eta\theta}$, suggest that, for both roles, offers are on average positively related to thresholds. Thus, the more (less) A-players offer in the

---

[10] See e.g. List (2007) and Bardsley (2008) who both study taking vs. giving behaviour in dictator experiments.

Table 4 *Naïve* model: Maximum simulated likelihood estimates

Description	Structural components	Estimated parameter	All Data		Trimmed Data	
			A's model	B's model	A's model	B's model
Random intercepts	$(\eta_i, \theta_i)'$	μ_η	34.283 [32.784, 35.781]	26.539 [25.465, 27.613]	28.666 [27.453, 29.878]	23.077 [21.807, 24.347]
		μ_θ	-11.467 [-12.633, -10.300]	-17.614 [-18.763, -16.465]	-14.519 [-15.553, -13.484]	-17.942 [-19.566, -16.318]
		σ_η	11.951 [10.995, 12.907]	14.441 [13.55646 15.325]	17.893 [16.839, 18.947]	14.117 [13.137, 15.097]
		σ_θ	19.589 [18.537, 20.642]	15.278 [14.536, 16.020]	10.940 [10.379, 11.502]	21.115 [19.547, 22.683]
		$\rho_{\eta\theta}$	0.384 [0.338, 0.430]	0.157 [0.127, 0.187]	0.064 [0.008, 0.119]	0.204 [0.143, 0.265]
First-stage polynomial		N	0.086 [0.074, 0.097]	0.047 [0.035, 0.060]	0.110 [0.095, 0.125]	0.058 [0.003, 0.006]
		p	4.559 [1.522, 7.596]	-2.066 [-5.410, 1.276]	5.542 [1.732, 9.351]	-2.042 [-5.891, 1.807]
		p^2	-3.237 [-6.032, -0.442]	4.288 [15.161, 17.156]	-3.979 [-7.482, -0.475]	4.952 [1.399, 8.504]
		Np	0.113 [0.091, 0.135]	0.327 [0.302, 0.351]	0.150 [0.123, 0.178]	0.384 [0.356, 0.412]
Second-stage polynomial	$m(N)$	N	0.638 [0.591, 0.686]	0.597 [0.533, 0.662]	0.644 [0.585, 0.704]	0.618 [0.541, 0.694]
		$N^2/10^2$	-0.555 [-0.688, -0.422]	-0.461 [-0.631, -0.292]	-0.564 [-0.728, -0.400]	-0.538 [-0.739, -0.337]
		$N^3/10^4$	0.266 [0.169, 0.363]	0.165 [0.045, 0.284]	0.256 [0.137, 0.376]	0.210 [0.069, 0.352]

Table 4 (continued)

			All Data		Trimmed Data	
Description	Structural components	Estimated parameter	A's model	B's model	A's model	B's model
Error terms	$\left(c_{i(N,p)} \ d_{i(N,p)}\right)'$	$\check{\sigma}_c$	9.640 [9.429, 9.852]	10.153 [9.925, 10.382]	10.649 [10.384, 10.913]	10.894 [10.630, 11.158]
		$\check{\sigma}_d$	3.100 [2.926, 3.275]	4.843 [4.538, 5.147]	3.585 [3.357, 3.813]	5.025 [4.639, 5.410]
		ρ_{cd}	-0.066 [-0.094, -0.038]	-0.170 [-0.199, -0.142]	-0.035 [-0.066, -0.003]	-0.141 [-0.173, -0.108]
		κ	0.012 [0.011, 0.013]	0.017 [0.016, 0.018]	0.011 [0.009, 0.012]	0.014 [0.013, 0.016]
		λ	0.034 [0.032, 0.036]	0.025 [0.023, 0.027]	0.032 [0.030, 0.034]	0.025 [0.022, 0.027]
# observations			8950	8950	6825	7450
# subjects			358	358	273	298
Log-likelihood			-55886.05	-55061.72	-42873.27	-45432.12
AIC			111806.11	110157.45	85780.53	90898.25

μ, σ, and ρ represent mean, standard deviation, and correlation, respectively, with suffix indicating the related random variable(s). The vector of random intercepts, η_i and θ_i for the first- and second-stage behavioural equation, respectively, follows a bivariate normal distribution. $l(N, p)$ is a polynomial in N, p, p^2, and Np. $m(N)$ is a polynomial in N, N^2, and N^3. The vector of error terms of the first- and second-stage behavioural equation, c_i and d_i, where κ and λ control for heteroscedasticity, respectively, follows a bivariate normal distribution. The "Trimmed Data" discards participants who submit the same first-stage offer or threshold more than 80% of the time (see Table 2). 95% confidence intervals in brackets. Confidence intervals for all the σs and ρs are derived using the delta method (Oehlert, 1992). All estimated coefficients are significantly different from 0 with p-values < 0.05, except for p in B's model (All Data: p-value = 0.226; Trimmed Data: p-value = 0.298)

first stage, the more (less) they expect to receive in the second, whereas the more (less) B-players request in the first stage, the more (less) they intend to offer in the second. Such behaviour alludes again to the emotional commitment argument of Yamagishi et al. (2009) and suggests a consistency between the players' behaviour in the first stage and their expectations for the second. To this extent, A-players would emotionally commit to what they expect to be a fair share of the second stage-pie, whereas B-players would emotionally commit to their first-stage requests.

As a robustness check of our findings, we trimmed off the data of "almost invariant" participants, i.e., who submitted the same first-stage choice at least 80% of the time, and report the estimation results in the left panel of Table 4. The gap in the players' intercepts pertaining to the first stage gets smaller but remains significant, and thus, confirms an intrinsic difference as to how "regular" A- and B-players perceive the first stage, with the former tending to submit overly generous offers no matter N or p. These positive intercepts estimates actually indicate the use of safe strategies, what Li et al. (2019) observed to be a characteristic of ambiguity-averse individuals (in trust games), and their difference can be explained here in terms of A-players' higher exposure to strategic ambiguity, cf. Fig. 1. Otherwise, the estimated polynomial $l(N, p)$, the second-stage intercepts, and the estimated polynomial $m(N)$ hardly change and confirm the reported differences between A- and B-players with regard to N and/or p.

5.2 The *sophisticated* model

This model builds on the *naïve* one by *(i)* accounting for features that characterise individual decision making under risk and ambiguity that may bias their decisions, like (individual) probability misperceptions as defined in Cumulative Prospect Theory (Tversky & Kahneman, 1992), and *(ii)* allowing for a possible backward reasoning that takes account of the second-stage choice to determine the first-stage one.

The first stage offers or thresholds $Y_i(N, p)$ in Eq. (7) are modelled as an affine function of the second-stage outcome with a random intercept term α_i and a stochastic term meant to capture noise in participant i's first-stage decisions, $e_{i(N,p)}$. The latter follows a normal distribution with mean 0 and standard deviation $\sigma_e = \breve{\sigma}_e \sqrt{\exp(\delta N p)}$, where the parameter δ controls for heteroscedasticity as a function of Np.

$$\begin{cases} Y_i(N, p) = \alpha_i + w_i(p)\left[N - \widetilde{X_i(N)}\right] + e_{i(N,p)} \\ X_i^*(N, p) = \widetilde{X_i(N)} + u_{i(N,p)} = \beta_i + z(N) + u_{i(N,p)} \end{cases} \quad \forall (N, p) \in \Omega \qquad (7)$$

$$X_i(N, p) = \max\left[X_i^*(N, p), 1\right] \qquad (8)$$

$$w_i(p) = \exp\left(-\varphi(-\log(p))^{\gamma_i}\right) \qquad \varphi, \gamma_i > 0 \qquad (9)$$

$$\begin{pmatrix} \alpha_i \\ \beta_i \\ \log(\gamma_i) \end{pmatrix} \sim \text{NID} \left[\begin{pmatrix} \mu_\alpha \\ \mu_\beta \\ \mu_\gamma \end{pmatrix}, \begin{pmatrix} \sigma_\alpha^2 & \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta & \rho_{\alpha\gamma}\sigma_\alpha\sigma_\gamma \\ \rho_{\alpha\beta}\sigma_\alpha\sigma_\beta & \sigma_\beta^2 & \rho_{\beta\gamma}\sigma_\beta\sigma_\gamma \\ \rho_{\alpha\gamma}\sigma_\alpha\sigma_\gamma & \rho_{\beta\gamma}\sigma_\beta\sigma_\gamma & \sigma_\gamma^2 \end{pmatrix} \right] \tag{10}$$

$$\begin{pmatrix} e_{i(N,p)} \\ u_{i(N,p)} \end{pmatrix} \sim \text{NID} \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_e^2 & \rho_{eu}\sigma_e\sigma_u \\ \rho_{eu}\sigma_e\sigma_u & \sigma_u^2 \end{pmatrix} \right] \tag{11}$$

$$\sigma_e = \check{\sigma}_e \sqrt{\exp(\delta N p)} \tag{12}$$

$$\sigma_u = \check{\sigma}_u \sqrt{\exp(\iota N)} \tag{13}$$

Players are assumed to misperceive the probability p of a second stage via Prelec's (1998) Probability Weighting Function (PWF) $w_i(p) = \exp\left(-\varphi(-\log(p))^{\gamma_i}\right)$, with $\varphi, \gamma_i > 0$. This PWF has two key parameters, γ_i which is individual-specific and φ which is common to all players; and it implies no probability distortion when $\gamma_i = \varphi = 1$. γ_i captures the under- or over-weighting of probabilities: the PWF has an S-shape when $\gamma_i > 1$ so that small (large) probabilities are underweighted (overweighted) and it has an inverted-S shape when $\gamma_i < 1$ so that small (large) probabilities p are overweighted (underweighted). As $\gamma_i \to 0$, the PWF becomes a step function that is flat everywhere except at the edges of the probability interval and characterises what Tversky & Fox (1995) called "ambiguity-generated insensitivity" or "a-insensitivity", i.e., people's inability to distinguish between different likelihoods or to "understand the ambiguous decision situation from a cognitive perspective" (Li, 2017, p. 242). The coefficient φ determines the PWF's elevation: given γ_i, a larger φ implies a lower $w_i(p)$ for each p or equivalently, a less "attractive" gamble (Gonzalez & Wu, 1999).[11]

This model can be seen as a normal-form representation of a two-stage game where the second-stage outcome is characterised by $\left[N - \widetilde{X_i(N)} \right]$. Here, the second-stage offer function $\widetilde{X_i(N)}$ is defined by $\beta_i + z(N)$, with β_i standing for an individual-specific intercept term and $z(N)$ for a polynomial in N that is common to all players of the same type and that possibly accounts for a nonlinear evaluation of N. Note that this polynomial affects first-stage decisions via the latent function and $w_i(p)$, so it captures risk aversion concerns *à la* CRRA-CARA on the players' first- and second-stage decisions without imposing a particular structure. This specification can also represent subject i's preferences for how to split N in risky/ambiguous conditions, including an equal-split (as reported in Binmore et al., 1985) or a fully rational behaviour.

As in the *naïve* model, $\widetilde{X_i(N)}$ is indirectly observed via a latent function $X_i^*(N,p) = \widetilde{X_i(N)} + u_{i(N,p)}$, inferred from the data via the observation rule $X_i(N,p) = \max\left[X_i^*(N,p), 1\right]$. We treat A's and B's second-stage choices as being

---

[11] For an extensive discussion on PWF curvature and elevation and their meanings, see Abdellaoui et al. (2010).

left-censored at 1. The additive error term $u_{i(N,p)}$ follows a normal distribution with mean 0 and standard deviation $\sigma_u = \breve{\sigma}_u \sqrt{\exp(\iota N)}$, where $\exp(\iota N)$ controls for heteroscedasticity as a function of N via the parameter $_\iota$. The model also assumes that α_i, β_i, and $\log(\gamma_i)$, are individual-specific and follow the joint normal distribution in Eq. (10) and that the error terms $e_{i(N,p)}$ and $u_{i(N,p)}$ follow the bivariate normal distribution in Eq. (11).[12] Finally, note that when the distributions of α_i, β_i, and γ_i are degenerate with point mass one at zero, zero, and one, respectively, φ is set to one and all the coefficients in $Z_{(N)}$ are set to zero, the model collapses the game-theoretic benchmark.

Let us clarify here the interpretation of A-players' behaviour in this model and what it brings in addition to the *naïve* model. Recall that the fully rational approach calls for backward induction and thus for an A-player to solve the B-player's decision problem and act accordingly to reach a first-stage agreement. To this end, A must anticipate the offer that s/he will receive from B in the second stage (while B determines her/his own second-stage payoff directly). The proposed *sophisticated* model relaxes, without discarding, the full rationality assumption and posits that A-player's second-stage threshold actually reveals this crucial information, whereas the *naïve* model posits that it only is a belief-free request. To this extent, a comparison of the models' goodness-of-fit performances would vindicate (or not) this assumption regarding how A-players' second-stage thresholds reveal their beliefs about B's second-stage offers.

To summarise, this model includes elements inspired from Game Theory, like a backward reasoning that encompasses the game-theoretic backward induction as a special case, and from empirical/experimental findings, like probability distortion and its implications in terms of strategic ambiguity. In addition, by comparing its goodness-of-fit to that of the *naïve* model, it can document or confirm the roles of ambiguity aversion and *a*-insensitivity in explaining behaviour in this game class.

5.2.1 The *sophisticated* model: Outcomes

The estimation results are reported in Table 5 and complement the behavioural insights from the *naïve* model in terms of strategic ambiguity and risk.

First-stage behaviour The means of the first-stage random intercepts, μ_α, confirm that on average, A-players offer more than what B-players request. This holds when the estimations discard "almost invariant" participants and thus confirms the use of safe strategies that are in keeping with the ambiguity aversion induced by the exposure of each player's type to strategic ambiguity. The estimated PWFs for A- and B-players display the well-documented inverse-S shape pattern as the means of γ_i are substantially smaller than 1 for both player roles and smaller for A-players than for B-players, i.e., 0.1703 (std. err. 0.0123) and 0.3890 (std. err. 0.0245), respectively.[13]

---

[12] See Online Resource E for further details.

[13] Recall that γ_i follows a log-normal distribution $LN(\mu_\gamma, \sigma_\gamma^2)$ so its mean is $\exp\left(\mu_\gamma + \sigma_\gamma^2/2\right)$.

Table 5 *Sophisticated* model: Maximum simulated likelihood estimates (please place table 5 on one page or split between two pages)

Description	Structural components	Estimated parameter	All Data		Trimmed Data	
			A's model	B's model	A's model	B's model
Random coefficients		μ_α	30.374 [29.568, 31.179]	23.490 [22.683, 24.296]	26.092 [25.214, 26.970]	18.060 [17.159, 18.960]
		μ_β	-11.483 [-12.504, -10.463]	-15.971 [-17.025, -14.917]	-18.180 [-19.309, -17.052]	-16.766 [-18.189, -15.343]
		μ_γ	-3.031 [-3.279, -2.784]	-1.640 [-1.742, -1.539]	-2.517 [-2.698, -2.335]	-1.405 [-1.504, -1.305]
		σ_α	12.342 [11.725, 12.959]	16.141 [15.563, 16.720]	12.399 [11.859, 12.939]	13.382 [12.826, 13.938]
	$(\alpha_i, \beta_i, \log(\gamma_i))'$	σ_β	17.780 [16.686, 18.874]	16.158 [15.161, 17.156]	16.482 [15.550, 17.414]	16.701 [15.777, 17.624]
		σ_γ	1.588 [1.400, 1.776]	1.180 [1.044, 1.316]	0.916 [0.786, 1.045]	1.022 [0.922, 1.122]
		$\rho_{\alpha\beta}$	0.377 [0.329, 0.425]	0.384 [0.348, 0.419]	0.396 [0.349, 0.444]	0.476 [0.430, 0.521]
		$\rho_{\alpha\gamma}$	-0.374 [-0.432, -0.315]	-0.372 [-0.415, -0.328]	-0.550 [-0.598, -0.501]	-0.140 [-0.206, -0.075]
		$\rho_{\beta\gamma}$	-0.331 [-0.378, -0.285]	-0.368 [-0.421, -0.315]	-0.167 [-0.270, -0.064]	-0.300 [-0.395, -0.205]
PWF Elevation		φ	1.666 [1.610, 1.721]	1.465 [1.419, 1.512]	1.427 [1.371, 1.482]	1.292 [1.247, 1.337]
Second-stage polynomial		N	0.638 [0.592, 0.684]	0.594 [0.535, 0.653]	0.646 [0.589, 0.704]	0.617 [0.545, 0.689]
	$Z(N)$	$N^2/10^2$	-0.554 [-0.682, -0.425]	-0.508 [-0.667, -0.348]	-0.607 [-0.766, -0.448]	-0.607 [-0.796, -0.418]
		$N^3/10^4$	0.266 [0.172, 0.359]	0.200 [0.086, 0.315]	0.297 [0.181, 0.413]	0.269 [0.136, 0.403]

Table 5 (continued)

Structural components	Estimated parameter	All Data		Trimmed Data	
Description		A's model	B's model	A's model	B's model
Error terms	$\check{\sigma}_e$	10.241 [10.008, 10.474]	10.842 [10.600, 11.085]	11.226 [10.948, 11.504]	13.382 [12.826, 13.938]
	$\check{\sigma}_u$	3.212 [3.034, 3.390]	4.156 [3.891, 4.421]	3.472 [3.247, 3.697]	4.660 [4.306, 5.015]
$(e_{i(N,p)}\; u_{i(N,p)})'$	ρ_{eu}	-0.065 [-0.094, -0.037]	-0.158 [-0.187, -0.129]	-0.034 [-0.065, -0.002]	-0.153 [-0.185, -0.121]
	δ	0.004 [0.003, 0.005]	0.005 [0.004, 0.007]	0.004 [0.003, 0.006]	0.005 [0.003, 0.006]
	ι	0.033 [0.031, 0.035]	0.030 [0.028, 0.032]	0.033 [0.031, 0.035]	0.027 [0.025, 0.030]
# observations		8950	8950	6825	7450
# subjects		358	358	273	298
Log-likelihood		-55767.69	-54578.08	-42768.67	-45081.05
AIC		111571.38	109192.16	85573.34	90198.10

μ, σ, and ρ represent mean, standard deviation and correlation, respectively, with suffix indicating the related random variable(s). The vector of random coefficients, which includes the intercepts α_i and β_i of the first- and second-stage behavioural equations, respectively, and $\log(\gamma_i)$, the logarithm of the parameter that shapes the PWF, follows a trivariate normal distribution. φ determines the PWF elevation. $Z(N)$ is a polynomial in N, N^2, and N^3. The vector of error terms of the first- and second-stage behavioural equation, e_i and u_i, where δ and ι control for heteroscedasticity, respectively, follows a bivariate normal distribution. Trimmed Data discards participants who submit the same first-stage offer or threshold more than 80% of the time (see Table 2). 95% confidence intervals in brackets. Confidence intervals for all the σs and ρs are derived using the delta method (Oehlert, 1992). All estimated coefficients are significantly different from 0 with p-values < 0.01, except for ρ_{eu} in A's model when using Trimmed Data (p-value = 0.036)

A-PARTICIPANT   B-PARTICIPANT
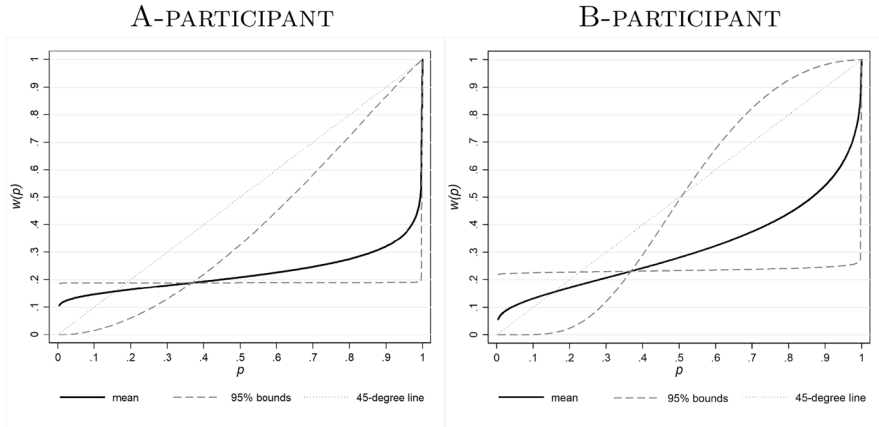
Fig. 3 Estimated Prelec's Probability Weighting Function

As for the estimated φ parameters, they are substantially larger than 1 for both roles and larger for A-players than for B-players, so that the prospect of reaching a second stage is more "unattractive" for A-players than for B-players. It thus follows from the estimated PWFs that A-players are more *a*-insensitive than B-players; see Fig. 3 for a display of these PWFs with their bounds encompassing 95% of the estimated distributions of γ_i. These behavioural biases are confirmed when discarding "almost invariant" participants from the estimations as A-players then appear even more *a*-insensitive (and less heterogeneous) than B-players, see Fig. F.1 in Online Resource F, and their intercept estimates still suggest that they are also more ambiguity averse.

Second-stage behaviour The negative mean of the random intercepts of the offer and threshold (latent) functions, μ_β, are significantly larger for A- than for B-players and the estimates characterising $z(N)$ indicate a nonlinear increasing relationship between the players' choices and N that is role-invariant. While this is in keeping with the inferences made from the *naïve* model, the estimation outcomes pertaining to the Trimmed Data reveal that the intercepts, like the polynomial estimates, are no more significantly different between A- and B-players, *as if* the expectations to receive from B-players and to concede to A-players of "regular" average A- and B-players are virtually identical. This vindicates our assumption that A-players' second-stage choices reveal their predictions about B-players' behaviour, and shows that the reported gap in the intercepts when estimating the model with all participants is mostly due to the presence of "almost invariant" A-players. As such, these participants would be most affected by the strategic ambiguity bored by the games they play and, consequently, they are unable to form unbiased expectations about B-players' behaviour.[14]

---

[14] According to Satisficing Theory (Simon, 1956), the B-players' requests can be seen as revealing a minimal aspiration to reach a first-stage agreement, whereas A-players' requests can be seen as revealing what they deem satisficing if a second stage is reached. See Online Resource G for more on this.

Backward reasoning The *sophisticated* specification outperforms the *naïve* one in terms of goodness-of-fit (measured by AIC), and thus, indicates that participants did use some "backward reasoning" to determine their first-stage choices.[15] Unlike the game-theoretic backward induction that hinges on expected utility maximisation and calls for counterfactual thinking, this reasoning refers to a multi-stage thinking based on what one expects to receive (for A-players) or to concede (for B-players) given N and p. To this extent, this finding indicates that the emotional commitment that characterises the players' second-stage behaviour can be actually traced-back to their first-stage choices in our experiments, what gives a grounding to the emotional commitment rationale of Yamagishi et al. (2009) for single-stage IG settings.

Stage-to-stage behaviour A Wald test of the joint null that the restrictions on the *sophisticated* model that turn it into the game-theoretic benchmark model hold against the alternative that they do not is overwhelmingly rejected for both roles, so the observed stage-to-stage behaviour is inconsistent with what game theory predicts. The significantly positive correlation between the intercepts α_i and β_i, $\rho_{\alpha\beta}$, and their significantly negative correlations with $\log(\gamma_i)$, $\rho_{\alpha\gamma}$, and $\rho_{\beta\gamma}$, unveil the central role of the players' *a*-insensitivity (partly captured by their γ_i parameter) in explaining the observed behaviour. Indeed, these estimated correlations indicate that a more *a*-insensitive A-player will offer more in the first stage and expect to receive more in the second and, likewise, a more *a*-insensitive B-player will request more in the first stage and expect to concede more in the second. This pattern becomes salient for A-players when the "almost invariant" participants are discarded from the estimations since their estimated $\rho_{\alpha\gamma}$ coefficient is then significantly different and thrice larger (in absolute value) than the one for B-players. We interpret this as evidence that A-players' exposure to strategic ambiguity compounds their ambiguity aversion and *a*-insensitivity more than what B-players' exposure does.

The other side of the coin revealed by the estimated correlations indicates that a less *a*-insensitive A-(B-)player will offer (request) less in the first stage and expect to receive (concede) less in the second. This points to behaviour much closer to the benchmark prediction.

Figure 4 puts in perspective the above findings for the Trimmed Data by displaying the estimated offers and thresholds for each stage and player's type as functions of N and p, along with the benchmark predictions (darkest surfaces) and average data for each (N, p)-constellation (dots). The surface-plots in the upper panels indicate that the *sophisticated* model fits the data remarkably well – there only is a mild underestimation of A-players' average first-stage offers in the (50,50)- and (75,50)-constellations and a mild overestimation of B-players' average thresholds when p equals 95. The upper-right panel further indicates that the model predicts a disagreement when $p > 50$ and that the latter increases with N, which is in keeping with the simulated disagreement rates of Table 3.[16]

---

[15] The "best" model according to these criteria, which includes a penalisation for the number of parameters estimated, is the one with the smallest AIC value.

[16] See Online Resource F for the surface-plots pertaining to the full dataset (All Data).
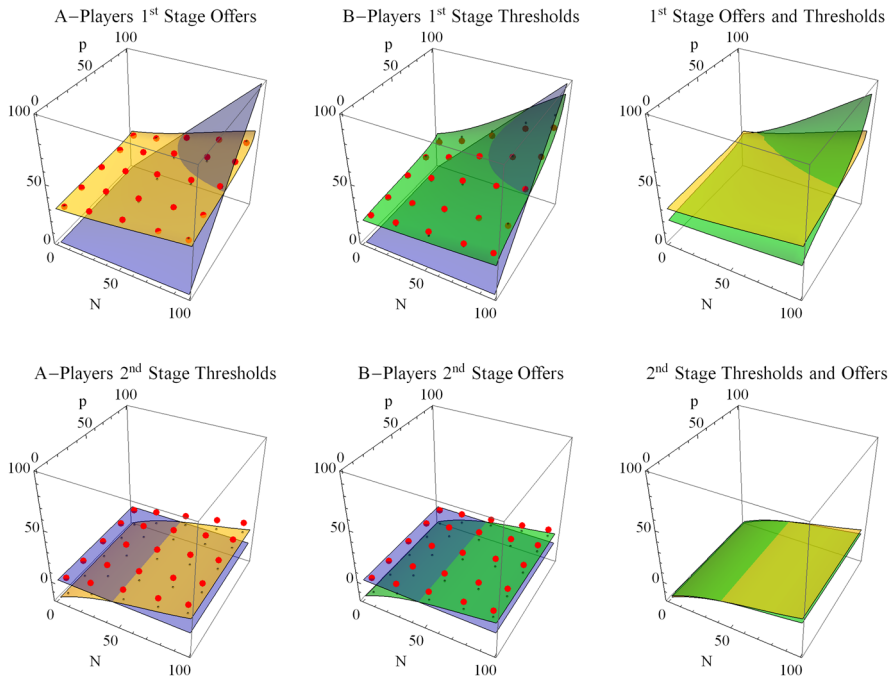
Fig. 4 Estimated offer and threshold functions for the first and second stages, Trimmed Data.
Note: The darkest surfaces (red dots) characterise the benchmark predictions (average data). The yellow surfaces for A-player and the green surfaces for B-player stand for the *sophisticated* model predictions for an average player type and are based on Trimmed Data estimates in Table 5

As for the players' second-stage average decisions, they are overestimated by censoring. The surface-plots in the lower panels predict that on average, A- and B-players request and offer (respectively) the minimum of 1 ECU for $N < 40$ and they otherwise request and offer very similar amounts that increase with N (cf. lower-right panel). The latter finding suggests that A- and B-players have a similar understanding of how to behave in this game class. As already noted above, the trimming of "almost invariant" players mostly concerns A-players whose use of safe strategies suggests that they were most ambiguity averse. To this extent, the similarity of the predicted second-stage behaviour of "regular" A- and B-players is most surprising given A-players' higher *a*-insensitivity and the non-provision of end-of-stage or end-of-round information feedback, *as if* players' distorted beliefs about each other's average behaviour were indeed consistent.

6 Conclusion

We experimentally investigate the role of strategic ambiguity and risk on behaviour in a class of pie-sharing games with alternating roles. The games we consider characterise a wide range of bargaining contexts and share the common features that

the second stage of interaction follows the rules of an Impunity Game (i.e., the proposer's offer cannot be vetoed by the responder but the latter can reject a positive offer without affecting the responder's payoff) and is uncertain (i.e., it occurs with a known probability p if no agreement is reached in the first stage).

The experimental evidence on similar two-stage bargaining games reports systematic deviations from the benchmark predictions that point to a limited cognition that would impede rational (game-theoretic) reasoning. Yamagishi et al. (2009) show in particular that rejected offers in single-stage Impunity Games can be rationalised in terms of responders' emotional commitment to what they believe to be a fair share of the pie for them to receive. As such rejections revert to "leaving money on the table," it follows that no matter one's beliefs about others' lower rationality (as in the *level-k* model of Ho & Su, 2013), such behaviour can hardly be rationalised in terms of expected utility maximisation. Furthermore, since proposers in these experiments received no end-of-stage or end-of-round information feedback, such rejections cannot be attributed to anger or moral disgust (since the responder's decision would not be disclosed to the proposer) neither to other-regarding preferences (since rejections exacerbate rather than reduce inequality, see Yamagishi et al., 2009).

Therefore, we propose to study these games through the lens of strategic ambiguity and risk without referring to expected utility maximisation or beliefs about others' rationality. Strategic ambiguity refers to the players' ignorance of the other's intentions and of the likelihoods of the other's possible actions, whereas risk refers to the known probability p of a second stage of interaction if no first-stage agreement is reached. The game class we consider introduces an asymmetry in the players' bargaining power, and thus, in their overall exposure to strategic ambiguity: a first-stage responder is exposed to strategic ambiguity in the first stage and to the risk of a second stage, whereas the first-stage proposer is additionally exposed to the strategic ambiguity of the second stage.

To assess the role of strategic ambiguity on the players' behaviour, we follow Li et al. (2019), who show that in a (trust) game context, it can trigger ambiguity aversion, i.e., an emotional response reflecting one's dislike of ambiguous situations and which calls for the use of safe strategies, and/or a-insensitivity, i.e., which is cognitive and makes one less likely to use belief-based strategies. As for risk, we assume that players may misperceive its role as in Cumulative Prospect Theory.

We investigate the relevance of these traits to organise behaviour with two structural models. The first is *naïve* in that it assumes that the players' first-stage decisions are not affected by their second-stage choices, although they may factor in details of the second-stage interaction (N and/or p). Since this model only allows the use of belief-free (or safe) strategies for the first and second stages, it only captures the players' ambiguity aversion. The second is *sophisticated* and builds on the *naïve* one by additionally allowing for probability distortions and for the first-stage decisions to be affected by the second-stage ones. Since it allows the use of safe and/ or belief-based strategies, it can capture both ambiguity aversion and the players' a-insensitivities – and it nests the game-theoretic benchmark model.

The estimations indicate that the *sophisticated* specification with a backward reasoning fits the data best. The model fits remarkably well the players' average behaviour in the first stage as well as their stage-to-stage behaviour. We find in particular

that first-stage offers are hardly affected by the game's parameters and that they tend to be higher than requests (as long as N and p are not too large), which is in keeping with the use of safe strategies that characterise ambiguity aversion and with the higher exposure of first-stage proposers to strategic ambiguity. These players also distort probabilities more than first-stage responders, and in a way that expresses more a-insensitivity, i.e., an insensitivity to likelihood changes. Further, the average second-stage decisions (offers or requests) are positively related to the average first stage ones (requests or offers). Such a stage-to-stage pattern is captured by the backward reasoning and suggests that the players' emotional commitments can actually be traced-back to their first-stage decisions. Interestingly, the saliency of this pattern increases with the players' a-insensitivity and thus with their exposure to strategic ambiguity.

These patterns are confirmed when we discard participants who display an almost invariant behaviour (i.e., the most ambiguity averse) from the estimations. In this case, besides the model's improved ability to explain first-stage decisions, it also predicts virtually identical decisions for the second-stage average proposer and average responder. This indicates that the least ambiguity averse players predict what the other expects to receive or intends to give almost perfectly on average, despite the higher a-insensitivity of "regular" A-players. Finally, since participants received no end-of-stage or end-of-round feedback, such matching predictions suggest that the players' beliefs about their protagonists were actually consistent and that most of the second-stage "disagreements" result from their heterogeneous play.

Deciding as a proposer or a responder in alternating-offer bargaining is a complex task which has drawn a lot of attention in the economics literature. The settings we study and the structural models we propose provide new insights into the analysis of strategic behaviour that highlight the role of strategic ambiguity in a multi-stage context.

References

Abdellaoui, M., L'Haridon, O., & Zank, H. (2010). Separating curvature and elevation: A parametric probability weighting function. *Journal of Risk and Uncertainty, 41*, 39–65.

Akiyama, E., Hanaki, N., & Ishikawa, R. (2017). It is not just confusion! Strategic uncertainty in an experimental asset market. *The Economic Journal, 127*, F563–F580.

Ashenfelter, O., & Currie, J. (1990). Negotiator behavior and the occurrence of dispute. *The American Economic Review, 80*(2), 416–420.

Barberis, N. (2012). A model of casino gambling. *Management Science, 58*(1), 35–51.

Bardsley, N. (2008). Dictator game giving: Altruism or artefact? *Experimental Economics, 11*, 122–133.

Binmore, K., McCarthy, J., Ponti, G., Samuelson, L., & Shaked, A. (2002). A backward induction experiment. *Journal of Economic Theory, 104*(1), 48–88.

Binmore, K., Shaked, A., & Sutton, J. (1985). Testing noncooperative bargaining theory: A preliminary study. *American Economic Review, 75*(5), 1178–1180.

Binmore, K., Shaked, A., & Sutton, J. (1988). A further test of noncooperative bargaining theory: Reply. *American Economic Review, 78*(4), 837–839.

Bolton, G. E. (1991). A comparative model of bargaining: Theory and evidence. *The American Economic Review, 81*(5), 1096–1136.

Cabrales, A., Miniaci, R., Piovesan, M., & Ponti, G. (2010). Social preferences and strategic uncertainty: An experiment on markets and contracts. *American Economic Review, 100*(5), 2261–2278.

Conte, A., Hey, J., & Moffatt, P. (2011). Mixture models of choice under risk. *Journal of Econometrics, 162*(1), 79–88.

Cooper, D., & Kagel, J. (2016). Other regarding preferences: A selective survey of experimental results. *The Handbook of Experimental Economics*, 2.

Crawford, V., Costa-Gomes, M., & Iriberri, N. (2013). Strutural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature, 51*(1), 5–62.

Di Mauro, C., & Finocchiaro Castro, M. (2011). Kindness, confusion, or ... ambiguity? *Experimental Economics, 14*, 611–633.

Embrey, M., Fréchette, G. R., & Lehrer, S. F. (2015). Bargaining and reputation: An experiment on bargaining in the presence of behavioural types. *Review of Economic Studies, 82*(2), 608–631.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics, 10*(2), 171–178.

Goeree, J., & Holt, C. (2000). Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *European Economic Review, 44*, 1079–1089.

Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology, 38*(1), 129–166.

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association, 1*(1), 114–125.

Greiner, B. (2018). Strategic uncertainty aversion in bargaining - experimental evidence. mimeo.

Güth, W., & Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior & Organization, 108*, 396–409.

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

Heinemann, F., Nagel, R., & Ockenfels, P. (2009). Measuring strategic uncertainty in coordination games. *The Review of Economic Studies, 76*(1), 181–221.

Ho, T. -H., & Su, X. (2013). A dynamic level-k model in sequential games. *Management Science, 59*(2), 452–469.

Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics, 6*(2), 65–70.

Johnson, E., Camerer, C., Sen, S., & Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory, 104*(1), 16–47.

Kovářák, J., Levin, D., & Wang, T. (2016). Ellsberg paradox: Ambiguity and complexity aversions compared. *Journal of Risk and Uncertainty, 52*(1), 47–64.

Li, C. (2017). Are the poor worse at dealing with ambiguity? *Journal of Risk and Uncertainty, 54*(3), 239–268.

Li, C., Turmunkh, U., & Wakker, P. P. (2019). Trust as a decision under ambiguity. *Experimental Economics, 22*(1), 51–75.

Li, C., Turmunkh, U., & Wakker, P. P. (2020). Social and strategic ambiguity versus betrayal aversion. *Games and Economic Behavior, 123*, 272–287.

Li, Z., Loomes, G., & Pogrebna, G. (2017). Attitudes to uncertainty in a strategic setting. *The Economic Journal, 127*(601), 809–826.

List, J. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy, 115*(3), 482–493.

Nawa, N. E., Shimohara, K., & Katai, O. (2002). On fairness and learning agents in a bargaining model with uncertainty. *Cognitive Systems Research, 3*(4), 555–578.

Neelin, J., Sonnenshein, H., & Spiegel, M. (1988). A further test of noncooperative bargaining theory: Comment. *American Economic Review, 78*(4), 824–836.

Ochs, J., & Roth, A. (1989). An experimental study of sequential bargaining. *American Economic Review, 79*(3), 355–84.

Oehlert, G. W. (1992). A note on the delta method. *The American Statistician, 46*(1), 27–29.

Prelec, D. (1998). The probability weighting function. *Econometrica, 66*(3), 497–527.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization, 3*(4), 323–343.

Roth, A. E. (1995). Bargaining experiments. In J. Kagel & A. E. Roth (Eds.), *Handbook of Experimental Economics, Handbook of Experimental Economics* (pp. 253–348). Princeton University Press.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica, 50*, 97–109.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review, 63*, 129–138.

Stahl, D. O., & Haruvy, E. (2008). Level-n bounded rationality in two-player two-stage games. *Journal of Economic Behavior & Organization, 65*(1), 41–61.

Ståhl, I. (1972). *Bargaining Theory*. Stockholm: Stockholm School of Economics.

Tversky, A., & Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review, 102*, 269–283.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297–323.

Wakker, P. (2010). *Prospect Theory: For Risk and Ambiguity*. Cambridge, UK: Cambridge University Press.

Yamagishi, T., Horita, Y., Takagishi, H., Shinada, M., Tanida, S., & Cook, K. (2009). The private rejection of unfair offers and emotional commitment. *Proceedings of the National Academy of Sciences of the United States of America, 106*, 11520–11523.