

How sensitive are bargaining outcomes to changes in disagreement payoffs?

Nejat Anbarci

School of Accounting, Economics and Finance
Deakin University
70 Elgar Road
Burwood VIC 3125, Australia
nejat.anbarci@deakin.edu.au

Nick Feltovich*

University of Aberdeen Business School
Edward Wright Building
Aberdeen AB24 3QY, United Kingdom
n.feltovich@abdn.ac.uk

February 14, 2011

Abstract

We use a human–subjects experiment to investigate how bargaining outcomes are affected by changes in the bargainers’ disagreement payoffs. Subjects play one of two bargaining games – a standard simultaneous–move Nash demand game, or a related unstructured bargaining game – against changing opponents. In both games, the disagreement outcome is asymmetric, and varies over plays of the game. Both bargaining parties are informed of both disagreement payoffs (and the cake size) prior to bargaining. We find that bargaining outcomes do vary with the disagreement outcome, but subjects underreact both to changes in their own disagreement payoff and to changes in the opponent’s disagreement payoff, relative to the risk–neutral prediction. This effect is observed in both games, and for two different cake sizes. We show theoretically that standard models of expected utility maximisation are unable to account for this effect – even when risk aversion is introduced – but a model of other–regarding preferences can explain it.

Journal of Economic Literature classifications: C78, C72, D81.

Keywords: Nash demand game, unstructured bargaining, disagreement, experiment, risk aversion, social preference, other–regarding behaviour.

*Corresponding author. Financial support from Deakin University’s Theoretical and Applied Economic Research fund is gratefully acknowledged. We thank John Boyd III, Emin Gahramanov, Lata Gangadharan and Randy Silvers for their suggestions and other invaluable help.

1 Introduction and background

Many negotiations – for example, between an employer and an employee, or the owner of a car and a potential buyer – involve a relation-specific surplus for the parties involved: if a used car is worth \$5,000 to the current owner and \$8,000 to the potential buyer, then a surplus of \$3,000 is available to be divided by the two parties. The fundamental role of bargaining in such decentralised markets has long been recognised (Edgeworth, 1881). However, until the 1950s, bilateral bargaining situations were deemed to lack a clear predicted outcome. The only prediction was that the division of the surplus would depend on the two parties’ relative bargaining power.

Nash (1950) proposed a framework which selected a unique feasible outcome – with certain desirable properties – as the solution of any bargaining situation that satisfies a few weak conditions.¹ Soon after, Nash (1953) proposed a non-cooperative game (now known as the *Nash Demand Game*, which we will abbreviate as NDG) in which two players simultaneously make demands, and where each player receives the payoff they demand if the demands are compatible; otherwise some default “disagreement” outcome is imposed. Both axiomatic and non-cooperative game-theoretic analyses of bargaining situations can serve as alternative but complementary ways of understanding the outcome of the bargaining process.²

Both analytical techniques provide testable implications for particular bargaining situations. As a simple example, consider the situation where players bargain over one unit of utility – specifically, a set of feasible agreements (a *bargaining set*) S equal to the convex hull of the points (0, 0), (1, 0) and (0, 1) and a disagreement outcome of (0, 0). (This means that the players are allowed to agree on any pair of payoffs $(x_1, x_2) \in S$, and if they fail to reach agreement, they each receive a payoff of 0.) In this case, all of S is *individually rational*: all payoff pairs in S yield to each party a payoff not worse than their payoffs under the disagreement outcome. The prominent axiomatic bargaining solutions, such as the Nash (1950) solution and the Kalai–Smorodinsky (1975) solution (see also Raiffa, 1953), make identical predictions in this case: agreement on the (0.5, 0.5) outcome. In addition, (0.5, 0.5) is the unique symmetric efficient Nash equilibrium in the NDG, as well as the outcome implied by risk-dominance (Harsanyi and Selten, 1988).

Now suppose that Player 1’s disagreement payoff increases from 0 to 0.5 and Player 2’s remains the same; that is, the disagreement point moves to (0.5, 0). Then the new individually rational bargaining set S' is the convex hull of (0.5, 0), (1, 0) and (0.5, 0.5) (see Figure 1), and both the Nash solution and the Kalai–Smorodinsky solution predict (0.75, 0.25) to be the outcome of this new bargaining situation. Moreover, the risk-dominant outcome and (if the bargaining parties focus only on individually rational outcomes) the symmetric efficient Nash equilibrium outcome also shift from (0.5, 0.5) in a Nash Demand Game with the first bargaining set to (0.75, 0.25) in the second.

Thus, most of the commonly used techniques for analysing bargaining situations agree on how players adjust to changes in their relative bargaining position (i.e., their disagreement payoff relative to the opponent’s). In the example above, the increase of 0.5 in Player 1’s disagreement payoff, with no change to Player 2’s disagreement payoff, led to an increase of 0.25 in Player 1’s payoff from bargaining, and a corresponding decrease of 0.25 in Player

¹Formally, a two-person cooperative (axiomatic) bargaining problem is described by a pair (S, d) where $S \subset \mathbf{R}^2$ is the *set of feasible agreements* with a *disagreement point* $d = (d_1, d_2) \in S$ being the allocation that results if no agreement is reached. Nash’s solution requires only that S is compact and convex, and that it contains some (x_1, x_2) with $x_1 > d_1$ and $x_2 > d_2$ (that is, the bargaining problem (S, d) is not “trivial”).

²As a matter of fact, the Nash Demand Game provides non-cooperative foundations for the Nash solution: Nash (1953) proved that the Nash solution outcome converges to the unique Nash equilibrium outcome of a “smoothed” Nash demand game in which a pair of incompatible demands may nonetheless be implemented with a small probability which goes to zero in the limit. See Binmore et al. (1993) for a bargaining experiment using a smoothed Nash demand game.




Figure 1: Bargaining problems and bargaining solutions (S and S' are sets of feasible agreements; d and d' are disagreement outcomes)

2's bargaining payoff. Given a bargaining set with an isosceles right triangular shape (like S or S' in Figure 1), any unit increase in one of the players' disagreement payoffs ought to lead to an increase in that player's ultimate bargaining payoff of exactly half a unit, along with a decrease in the other player's ultimate bargaining payoff of exactly half a unit. This implication is intuitively appealing, as it simply quantifies the likelihood that when a player's relative bargaining position improves, the outcome of bargaining becomes more favourable to her.

Whether this theoretically robust property holds in real bargaining situations is, of course, an empirical question. The goal of this paper is to examine whether and how bargaining outcomes actually are affected by changes to players' disagreement payoffs. We accomplish this by means of a human-subjects experiment, which allows us precise control over both the disagreement outcome and the total amount being bargained over (which, following standard bargaining terminology, we refer to as the "size of the cake"). We use two bargaining games, both of which capture essential features of real-life bargaining. One game is the NDG, described above. The other game is an unstructured variation of the NDG, which we call the *Unstructured Bargaining Game* (UBG). In the UBG, the bargaining set is the same, but instead of making simultaneous demands, players have a fixed, known amount of time available to negotiate a mutually-agreed division of the cake. Both players can make proposals, which have to be in the bargaining set, though they need not be individually rational or efficient. Either player can accept any opponent proposal; the first accepted proposal is implemented. If no proposal is accepted before the time limit, both players receive their disagreement payoffs.

In the experiment, subjects play one of these games (NDG or UBG) repeatedly against randomly chosen opponents, with randomly chosen disagreement payoffs. They play a set of rounds with low stakes (a cake size of £5) and a set with high stakes (£20). Our main finding is that while subjects do take into account their relative bargaining position – in the sense that increases in one's own disagreement payoff, and decreases in the opponent's disagreement payoff, translate into higher bargaining outcome payoffs – they are much less sensitive to changes in their bargaining position compared to the theoretical predictions described above. Specifically, when bargaining is successful, the sum of the magnitudes of the own-disagreement-payoff and opponent-disagreement-payoff effects is only around one-half, whereas the theoretical predictions imply that the sum should be one. This result is robust to which bargaining game was played, as well as to changes in the cake size and in the ordering in which subjects faced the cake sizes. We provide theoretical evidence, in Section 7, that this result cannot be explained solely by

subjects' aversion to risk. In Section 8, we show that while Fehr and Schmidt's (1999) model of other-regarding preferences also cannot explain our result, a slight modification of it can (though we note that other explanations are also possible).

2 The bargaining environment

We describe here the two-player bargaining problem underlying both games used in the experiment; see also Figure 2. There is a fixed sum of money (a *cake*) of size $\pounds M$ available to the players. The way bargaining occurs depends on the game, but in either case, the set of feasible agreements is the set of non-negative pairs totalling M or less. Also in both games, if bargaining is unsuccessful, the players receive *disagreement payoffs*. The disagreement outcome is asymmetric: the *favoured player* receives d_f and the *unfavoured player* receives d_u , with $d_f > d_u > 0$ and $d_f + d_u < M$.³ The values of M , d_f and d_u (along with which player is the favoured one) are assumed to be common knowledge. We use the term *surplus* to mean the portion of the cake remaining after subtracting the sum of the disagreement payoffs ($M - d_f - d_u$); this positive quantity represents the gains available from successful bargaining.

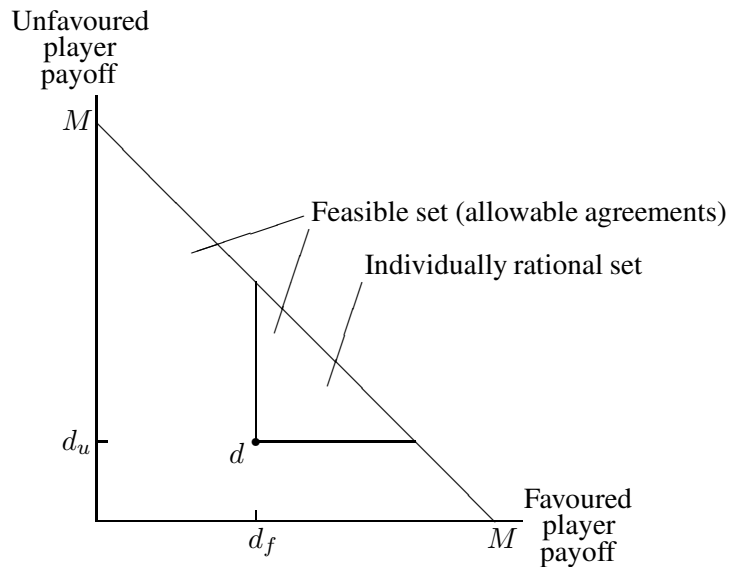


Figure 2: The bargaining environment

2.1 Nash demand game (NDG)

In the Nash demand game (Nash, 1953), bargaining consists of a single pair of simultaneously made demands x_f and x_u by the favoured and unfavoured players, respectively. If the demands are compatible ($x_f + x_u \leq M$), then each player receives the amount demanded (any remainder is left “on the table”). If the demands are incompatible ($x_f + x_u > M$), then both receive their disagreement payoffs.

The NDG is simple enough to be analysed by standard non-cooperative game theory, but the result is not a unique prediction. Rather, the game typically has a large number of Nash equilibria, including (1) efficient pure-

³In what follows, we will use female pronouns to refer to the favoured player, and male pronouns for the unfavoured player. In the experiment, of course, types were assigned irrespective of sex.

strategy equilibria in which $x_f \geq d_f$, $x_u \geq d_u$ and $x_f + x_u = M$, leading to equilibrium payoffs (x_f, x_u) ; (2) inefficient pure-strategy equilibria in which $x_f > M - d_u$ and $x_u > M - d_f$, with resulting equilibrium payoffs (d_f, d_u) ; and (3) inefficient mixed-strategy equilibria with expected payoffs totalling less than M but more than $d_f + d_u$.

Equilibrium selection criteria such as payoff dominance or efficiency can reduce the set of equilibria somewhat, eliminating the inefficient equilibria in (2) and (3) above. If an additional symmetry criterion is imposed, with symmetry defined relative to the individually rational set, then the unique prediction is for the players to split the surplus evenly: $x_f = \frac{1}{2}(M + d_f - d_u)$ and $x_u = \frac{1}{2}(M - d_f + d_u)$. This is also the prediction of risk dominance (Harsanyi and Selten, 1988).

2.2 Unstructured bargaining game (UBG)

In the unstructured bargaining game, players have a fixed, known amount of time available to negotiate a division of M . Either player can make proposals, which take the form (x_f, x_u) with $x_f, x_u \geq 0$ and $x_f + x_u \leq M$. There is no constraint (other than the time available) on the number of proposals that can be made, and the cake size remains the same until the time runs out, by contrast with Rubinstein's (1982) bargaining model. Either player can accept any opponent proposal; the first accepted proposal is implemented. (In case both players accept proposals at the same time, each is implemented with probability one-half.) If no proposal is accepted before the time limit, the disagreement outcome is imposed.

The UBG is far too complex to allow the use of standard non-cooperative game-theoretic methods for its analysis, without the imposition of additional assumptions.⁴ Instead, we make use of techniques from cooperative game theory. These techniques say little about the precise strategies used by the two players; rather, they have implications about what the outcome of bargaining is. The *core* predicts that the division of the cake corresponds to an efficient Nash equilibrium outcome ($x_f \geq d_f$, $x_u \geq d_u$ and $x_f + x_u = M$), but makes no sharper prediction. Axiomatic bargaining solution concepts can refine this multiplicity of predicted outcomes to a unique one; however, they require an assumption about the relationship between monetary payments and payoffs. If the relationship is proportional (risk neutrality), then the outcome of every well known axiomatic bargaining solution (including the Nash and Kalai-Smorodinsky solutions) coincides, with $x_f = \frac{1}{2}(M + d_f - d_u)$ and $x_u = \frac{1}{2}(M - d_f + d_u)$.

2.3 Theoretical predictions

The prediction of Nash equilibrium (with the additional assumptions of either efficiency and symmetry or risk dominance) for the NDG, and the predictions of the well-known axiomatic bargaining solutions for both the NDG and the UBG – discussed in the previous two sections – therefore imply the same outcome. In all cases, the players evenly share the the surplus (the remainder of the cake left over once both are paid their disagreement payoff). There is thus a sharp theoretical prediction concerning the relationship between the disagreement payoffs and the bargaining outcome in both games:

$$\frac{\partial x_f}{\partial d_f} = \frac{1}{2} = \frac{\partial x_u}{\partial d_u} \quad \text{and} \quad \frac{\partial x_f}{\partial d_u} = -\frac{1}{2} = \frac{\partial x_u}{\partial d_f}.$$

That is, an increase of £1.00 in a player's own disagreement payoff results in a £0.50 increase in that player's payoff resulting from bargaining, while an increase of £1.00 in the opponent's disagreement payoff results in a £0.50

⁴See Simon and Stinchcombe, 1989; Perry and Reny, 1993, 1994; and de Groot Ruiz et al., 2010 for non-cooperative game-theoretic analyses of unstructured bargaining using additional assumptions.

decrease in that player’s payoff from bargaining. Thus, the sum of the magnitudes of the two changes is equal to one:

$$\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = \left| \frac{\partial x_u}{\partial d_f} \right| + \left| \frac{\partial x_u}{\partial d_u} \right| = 1.$$

3 Related literature

While the NDG has the desirable feature of simplicity, one might criticise it as an excessive simplification of real-life bargaining. However, theorists have tended to defend it from this charge. Binmore (2007) points out that when bargainers can commit to demands, but neither has the ability to commit before the other, the NDG is the limiting case where both bargainers “rush to get a take-it-or-leave-it demand on the table first” (p. 496), resulting in simultaneous irrevocable demands.⁵ Moreover, Skyrms (1996) argues that in modelling the bargaining process, “[o]ne might imagine some initial haggling...but in the end each of us has a bottom line” (p. 4); focussing on these bottom lines results in the NDG. Our use of the UBG, by contrast, admits the possibility that not all important aspects of bargaining are captured by these final demands.

The literature on bargaining experiments is immense, and a review, even limiting consideration to those using games like the NDG and UBG, is far beyond the scope of this paper. (Surveys of bargaining experiments can be found in Roth, 1995 and Camerer, 2003, pp. 151–198.) Rather, we discuss the most closely related experiments to ours. Hoffman and Spitzer (1982) examined unstructured bargaining games with (in essence) a fixed, known cake size and one of two randomly chosen disagreement outcomes.⁶ Disagreement outcomes tended to be very asymmetric; for example, in their “Decision 1”, the two possible disagreement outcomes as portions of the cake were approximately (0.79, 0) and (0, 0.83). Hoffman and Spitzer found a substantial frequency of equal splits of the cake – irrespective of which disagreement outcome was chosen – even though this means that some bargainers were accepting payments that were well below their disagreement payoffs. This result may not have much implication for our experiment, however, as it is likely at least partly explained by Hoffman and Spitzer’s use of face-to-face bargaining (leading to a lack of subject anonymity). Hoffman and Spitzer (1985) reported a similar result in another experiment with face-to-face bargaining, but additional treatments provide an alternative explanation: that subjects placed randomly into a favourable bargaining position feel that they haven’t “earned” this position, and are thus reluctant to exploit it.⁷

More recently, Fischer, Güth and Pull (2007) examine bargaining in the ultimatum game and in a variant of the NDG. In this variant, players simultaneously submit an ambitious demand x_i and a (typically smaller) fallback demand g_i ; the players receive their ambitious demands if they total the cake size or less; if not, they each get their fallback demand if those total the cake size or less. If both pairs of demands total more than the cake size,

⁵The alternative case, when one bargainer can commit earlier than the other, gives rise to the ultimatum game. See Fischer et al. (2006) for an experiment that nests the ultimatum game and the NDG.

⁶In their setup, an agreement involved bargainers settling on one of a small number of payment pairs, but side-payments were allowed, making the bargaining set one with a fixed cake size. Rather than directly implementing disagreement outcomes, Hoffman and Spitzer assigned one of the bargainers the role of “controller”; in the case of disagreement, the controller unilaterally imposed one of the payment pairs. Assuming that controllers would always choose the most favourable payment pair, this was equivalent to randomly choosing one of two disagreement outcomes.

⁷Hoffman and Spitzer (1985) find that subjects fully exploit their bargaining position only when both (1) favourable position is seen to be earned, e.g. by scoring well on a test of general knowledge or cognitive skills; and (2) instructions are written to specifically encourage subjects to make use of their bargaining power (i.e., they are told this is acceptable behaviour). See Gächter and Riedl (2005) for another experiment using a quiz to allocate the favoured and unfavoured player roles.

each receives a disagreement payoff d_i .⁸ Fischer, Güth and Pull were primarily interested in whether behaviour was similar in the two bargaining games (i.e., whether bargainers failed to exploit the differences in structural bargaining power that exist between the games), but they also varied the disagreement outcome in both games. Subjects faced a total of eleven different disagreement payoff pairs: (0, 50), (5, 45), (10, 40), ..., (50, 0). Fischer, Güth and Pull’s design, with disagreement payoffs perfectly negatively correlated between opposing players, does not allow for a distinction between the effects on bargaining outcomes due to changes in own disagreement payoffs and those due to changes in opponent disagreement payoffs, but one can still compute the sum of these effects using their data. On average, the results they report imply that $|\partial x_i/\partial d_i| + |\partial x_i/\partial d_j| \approx 0.38$ and $|\partial g_i/\partial d_i| + |\partial g_i/\partial d_j| \approx 0.41$. That is, subjects’ demands increased as their bargaining position improved, but they were far from fully exploiting their bargaining power (which, as noted in Section 2.3, would have made these sums equal to one).⁹

We stress that the focus of our paper is limited to the effect of disagreement payoffs on bargaining outcomes; we use multiple games (NDG and UBG) and cake sizes (£5 and £20) purely to verify the robustness of the phenomena we observe.¹⁰ We note, however, that these other manipulations could serve as research topics in their own right, and indeed both have been addressed in previous work. Our use of the NDG and UBG games roughly parallels Feltovich and Swierzbinski’s (2011) “baseline” and “contracts” treatments, the former of which modified the NDG by giving one of the players an outside option (which could be chosen in lieu of bargaining), and the latter of which added a pre-play unstructured negotiation stage to this outside-option game. Feltovich and Swierzbinski found substantially higher agreement frequencies when pre-play negotiation was possible, and more surprisingly, they found differences between the treatments in the shares captured by the favoured and unfavoured players conditional on reaching agreement.¹¹ There is also a fair-sized literature examining the effect of the cake size in bargaining, usually using ultimatum games, and taken together, they have yielded fairly consistent results. When subjects are given opportunities to learn through repetition of the game, increasing the cake size raises the likelihood of a given demand (as a fraction of the cake) being accepted, and sometimes leads to higher demands (Slonim and Roth, 1998; Munier and Zaharia, 2003). However, in one-shot ultimatum games, no cake-size effect is typically discernible, even for quite large differences in cake sizes (Cameron, 1999).

4 Experimental design and procedures

All sessions lasted for forty rounds, split into two halves of twenty rounds each. The cake size was £5 in one half and £20 in the other half, with the order varied in an effort to control for any order effects. Thus, the ordering of cake sizes, as well as the game played (NDG or UBG) were varied between-subjects, while the cake size itself, player type (favoured or unfavoured) and the disagreement outcomes were varied within-subject.

The experimental sessions took place at the Scottish Experimental Economics Laboratory (SEEL) at the University of Aberdeen. Subjects were primarily undergraduate students from University of Aberdeen, and were recruited from a database of people expressing interest in participating in economics experiments. No one took part in this experiment more than once, nor did anyone take part who had participated in any previous bargaining experiments at SEEL.

⁸We alter their notation somewhat, to parallel the notation in the current paper.

⁹Harrison (1987) also varies disagreement payoffs in an unstructured bargaining game, but with perfect *positive* correlation between disagreement payoffs; his “Type 1 game” has a disagreement outcome of (0, 0), while in his “Type 3 game”, both players receive equal positive payments in case of disagreement.

¹⁰In this, we follow Roth et al. (1991), who vary stake sizes by a factor of three in some cells of their four-country experiment.

¹¹See also de Groot Ruiz et al. (2010) for a comparison of highly structured and less structured three-player bargaining games.

At the beginning of a session, subjects were seated in a single room and given written instructions for the first twenty rounds; these instructions described the bargaining environment, the sequence of events within a round of play, and the way the money payments they would receive were connected to their decisions.¹² They were informed then that the experiment would comprise two halves totalling forty rounds, but details of the second half were not announced until after the first half had ended. The instructions were also read aloud to the subjects, in an attempt to make the rules of the game common knowledge. Then, the first round of play began. After the twentieth round was completed, each subject was given a copy of the instructions for rounds 21–40. These instructions were also read aloud, before round 21 was played.

The experiment was run on networked personal computers, and was programmed using the z-Tree experiment software package (Fischbacher, 2007). Subjects were asked not to communicate with other subjects except via the computer program. Subjects were randomly matched in each round, with each other subject equally likely to be the opponent in a given round (a one-population matching protocol). Within each pair, roles were assigned randomly, so a given subject was equally likely to be the favoured or unfavoured player in that round.¹³ No identifying information was given about opponents (in an attempt to minimise incentives for reputation building and other supergame effects). Rather than using potentially biasing terms like “opponent” or “partner” for the other player, we used the neutral though somewhat cumbersome “player matched to you” and similar phrases.

Each round of the game began with a screen telling each subject the cake size and disagreement outcome (both own and opponent disagreement payoff) for that round. The disagreement payoff for a favoured player was drawn from a uniform distribution, from 25% to 45% of the cake; for an unfavoured player it was between 5% and 25% of the cake (both draws were rounded to the nearest £0.01). These draws were independent across rounds and pairs of subjects. After viewing their disagreement outcome, subjects in the NDG treatment were prompted to choose their demands. Demands were required to be whole-number multiples of £0.01, between zero and the cake size inclusive.¹⁴ After all subjects had chosen their demands and clicked to continue, they received end-of-round feedback: own demand, opponent demand, whether agreement was reached (i.e., whether demands totalled at most the cake size), own payoff and opponent payoff. A subject’s previous results were also collected into a history table at the top of the computer screen; these could be reviewed at any time. After all subjects clicked a button on the screen to continue, the session proceeded to the next round.

In the UBG cells, subjects were given a 90-second “negotiation stage” to reach agreement on a division of the cake. Figure 3 shows a sample screen viewed by subjects during this time. Subjects could make as many or as few proposals as they wished during the 90 seconds; a proposal consisted of a nonnegative multiple of £0.01 for the sender and one for the receiver, adding up to the cake size or less. Other than that, there were no constraints on proposals (e.g., there was no requirement that later proposals had to be more favourable to the receiver than earlier

¹²Sample instructions are shown in Appendix B. The remaining sets of instructions, as well as the raw data from the experiment, are available from the corresponding author upon request.

¹³Thus, with extremely high probability, a subject plays some rounds as favoured player and others as unfavoured player. Some researchers (for example, Binmore, Shaked and Sutton, 1985) have found that giving subjects experience in both bargaining roles can mitigate other-regarding preferences, though Bolton (1991) found no difference between sessions with changing roles and those with fixed roles.

¹⁴Our restriction of demands and disagreement payoffs to hundredths of a pound, necessitated by the discreteness of money, has at most minor effects on theoretical predictions. In particular, when the sum of disagreement payoffs is an odd number of pence, there is no longer a unique prediction according to symmetry, risk dominance and the axiomatic bargaining solutions; instead, there will be two distinct predictions, differing by one penny, and instead of each player receiving exactly half of the surplus, each receives half of the surplus plus/minus £0.005. For example, for a cake size of £5 and a disagreement outcome of (£1.00, £1.99), all of these concepts predict agreements of either (£2.01, £2.99) or (£2.00, £3.00). The discreteness of disagreement payoffs also meant that there was a small chance that both subjects in a pair would have the same disagreement payoff (25% of the cake), though this never actually happened in the experiment.

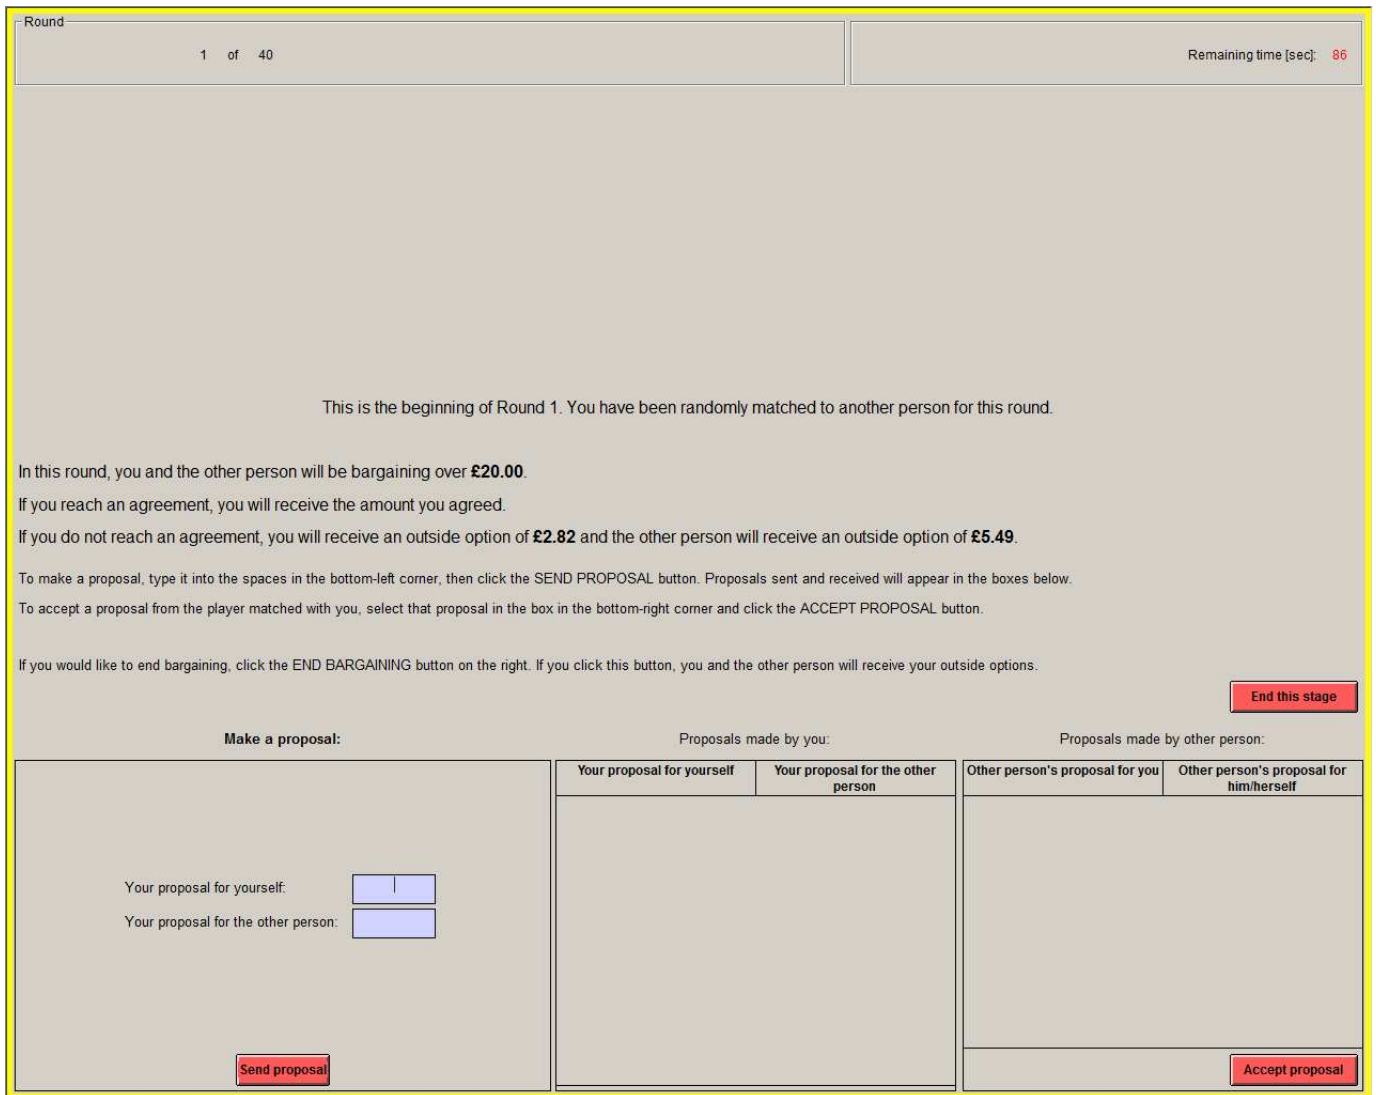


Figure 3: Screen-shot from negotiation stage of UBG treatment

ones). Proposals could not be withdrawn once made, and no messages were possible apart from the proposals.¹⁵ Both the subject's own proposals and the proposals of the opponent were shown on the subject's screen (in separate places), but it was not possible to view proposals for other pairs of subjects. As long as the negotiation stage hadn't ended, a subject could choose to accept any of the opponent's proposals, at which time that proposal would become binding. The opponent's proposals were listed in order of decreasing payoff to the subject, so there was almost no cognitive effort required to determine the most favourable opponent proposal (it was always at the top of the list), though of course a subject could accept a less favourable proposal if desired. The negotiation stage ended if a proposal was accepted, if either subject in a pair chose to end it (by clicking a button on the screen), or after the 90

¹⁵Our prohibition of cheap talk, and the restriction of negotiation to computers rather than face-to-face interaction, were intended to maintain anonymity between bargainers in the experiment. This is important, as removing this anonymity opens up the possibility of side-payments or threats outside the laboratory, after an experimental session has concluded. However, we acknowledge that lack of anonymity can be an important feature of some real bargaining situations. We also note that a side consequence of both of these design choices is they keep the level of social distance between the bargainers relatively high. Some research (e.g., Bohnet and Frey, 1999; Rankin, 2006) has found that lower levels of social distance are associated with a greater prevalence of other-regarding behaviour.

seconds had expired without an accepted proposal; in these latter two cases, the disagreement outcome was imposed.

In either game, at the end of the fortieth round, the experimental session ended and subjects were paid, privately and individually. For each subject, two rounds from each block of twenty were randomly chosen, and the subject was paid his/her earnings in those rounds. There was no show-up fee. Subjects' total earnings averaged about £20. NDG sessions typically lasted about 45 minutes, UBG sessions about 90 minutes.

5 Hypotheses

Our experiment was designed with several hypotheses in mind; these hypotheses will assist us in organising our analysis and discussion of the experimental results. The first four hypotheses concern the effect on payoffs from bargaining from changes to the disagreement outcome. As mentioned in Section 2.3, a player's payoff as a share of the cake size should increase by half of any change to her own disagreement outcome, and should decrease by half of any change to the opponent's disagreement outcome. By the same token, both players' payoffs – as shares of the surplus available – should be unaffected by changes to either player's disagreement payoff. We thus have:

Hypothesis 1 *In both treatments, for both player types and both cake sizes, a one-unit increase in a player's own disagreement payoff is associated with a one-half-unit increase in that player's payoff as a share of the cake size.*¹⁶

Hypothesis 2 *In both treatments, for both player types and both cake sizes, a one-unit increase in a player's opponent's disagreement payoff is associated with a one-half-unit decrease in that player's payoff as a share of the cake size.*

Hypothesis 3 *In both treatments, for both player types and both cake sizes, a player's payoff as a share of the surplus is unaffected by changes to the player's own disagreement payoff.*

Hypothesis 4 *In both treatments, for both player types and both cake sizes, a player's payoff as a share of the surplus is unaffected by changes to the opponent's disagreement payoff.*

A fifth hypothesis reflects the prediction of axiomatic bargaining solutions, as well as efficient Nash equilibrium and risk dominance, that agreement occurs with probability one, and is thus not affected by changes to the disagreement outcome – in contrast with some experimental results (e.g., Murnighan et al., 1988) that have found a negative correlation between disagreement payoffs and agreement frequencies.

Hypothesis 5 *In both treatments, for both player types and both cake sizes, the frequency of agreement is unaffected by changes to either player's disagreement payoff.*

6 Experimental results

The experiment comprised eight sessions – two for each combination of game (NDG or UBG) and cake-size ordering (increasing or decreasing) – with a total of 108 subjects (varying from 10–18 in a session). We begin the analysis of results in Section 6.1 with descriptive aggregate statistics; these will show the effects of some of our treatment variables (cake size, favoured versus unfavoured player) on bargaining outcomes. Later in the section, we

¹⁶To save space, we only state the null hypotheses. The corresponding alternative hypotheses should be clear.

will disaggregate the data somewhat, in order to examine how bargaining outcomes are affected by changes to the disagreement payoffs. Then, in Section 6.2 we use regressions to disentangle the effects due to the disagreement payoffs from effects due to changes in other variables.

Many of the results we examine will involve two statistics, which we define now in order to avoid confusion. A *demand as a portion of the cake* is a demand, normalised onto a scale from 0 to 1 so that a zero demand corresponds to 0 and a demand of the entire cake corresponds to 1:

$$\text{demand as portion of cake} = \frac{\text{demand}}{M}.$$

We divide by the cake size M in order to facilitate comparison of results with different cake sizes. A *demand as a portion of the surplus* is also normalised, but in such a way that a demand equal to the subject’s own disagreement payoff corresponds to 0, and a demand of the whole cake minus the opponent’s disagreement payoff corresponds to 1. That is,

$$\text{demand as portion of surplus} = \frac{\text{demand} - d_f}{M - d_f - d_u}$$

for the favoured player and

$$\text{demand as portion of surplus} = \frac{\text{demand} - d_u}{M - d_f - d_u}$$

for the unfavoured player. (Hence values less than zero or greater than one for this statistic are possible, though the former is weakly dominated and the latter is not rationalisable.) We will often normalise subjects’ payoffs in a similar way – as proportions (or sometimes as percents) of the cake and of the surplus.

6.1 Aggregate behaviour

Some aggregate data are presented in Tables 1 and 2. Table 1 shows results for the NDG treatment. For both cake sizes, and both for all rounds and for rounds 11–20 (the second half) of each cake size, the table shows the frequency of agreement and mean demands by both types of player (favoured and unfavoured), both as a percent of the cake size and as a percent of the surplus available to the bargainers. Also shown are the mean payoffs to both types of player conditional on agreement (thus identical to mean demands conditional on agreement), again as percents of the cake size and of the surplus. Table 2 shows corresponding results for the UBG treatment: agreement frequencies

Table 1: Aggregate statistics – NDG treatment

| | | £5 cake | | £20 cake | |
|--|----------------------------------|---------|-------|----------|-------|
| | | All | 11–20 | All | 11–20 |
| | Rounds: | | | | |
| | Agreement frequency (%) | 57.6 | 58.1 | 60.7 | 61.9 |
| Mean demand | favoured player (% of cake) | 58.0 | 58.1 | 57.0 | 58.1 |
| | unfavoured player (% of cake) | 47.6 | 47.9 | 47.1 | 47.2 |
| | favoured player (% of surplus) | 46.4 | 46.1 | 44.4 | 47.6 |
| | unfavoured player (% of surplus) | 65.3 | 66.2 | 63.9 | 63.5 |
| Mean payoff (conditional on agreement) | favoured player (% of cake) | 51.3 | 51.6 | 50.6 | 51.6 |
| | unfavoured player (% of cake) | 42.1 | 43.1 | 41.2 | 42.2 |
| | favoured player (% of surplus) | 33.3 | 33.6 | 32.1 | 35.0 |
| | unfavoured player (% of surplus) | 53.3 | 55.7 | 51.5 | 52.7 |

and mean payoffs for both types of player conditional on agreement.¹⁷

Table 2: Aggregate statistics – UBG treatment

| | | £5 cake | | £20 cake | |
|--|----------------------------------|---------|-------|----------|-------|
| | | All | 11–20 | All | 11–20 |
| Rounds: | | | | | |
| Agreement frequency (%) | | 83.3 | 85.2 | 83.5 | 86.3 |
| Mean payoff (conditional on agreement) | favoured player (% of cake) | 57.5 | 58.5 | 56.9 | 57.0 |
| | unfavoured player (% of cake) | 42.1 | 41.4 | 42.8 | 42.6 |
| | favoured player (% of surplus) | 44.9 | 47.1 | 42.4 | 42.5 |
| | unfavoured player (% of surplus) | 54.5 | 52.7 | 56.8 | 56.6 |

One clear result from Tables 1 and 2 is that the favoured player – the bargainer with the larger disagreement payoff – makes some, but only limited, use of her better bargaining position. In the NDG, favoured players demand on average roughly an extra tenth of the cake (with only slight variation depending on the cake size and which rounds we consider) compared to unfavoured players. Conditional on agreement in both games, favoured players’ average shares of the cake are also higher than those of the unfavoured players by about 10% of the cake. Nonparametric statistical tests find that these differences in shares are significant (Wilcoxon signed–ranks test, pooled NDG and UBG session–level data, $p \approx 0.004$ for both cake sizes).¹⁸ However, they are substantially smaller than the approximately 20% average difference in the disagreement payoffs themselves between favoured and unfavoured players (34.8% vs. 15.3% respectively in NDG and 35.3% vs. 15.0% in UBG).

The comparative lack of exploitation of bargaining position is further highlighted when we examine demands and payoffs as proportions of the available surplus. Favoured players’ average demands in the NDG correspond to just under half the available surplus, while unfavoured players demand nearly two–thirds of the available surplus. Similarly, conditional on agreement, favoured players’ average shares of the surplus are only about one–third in the NDG – compared to over half for unfavoured players – and the corresponding shares in the UBG are between 42% and 47% for favoured players and between 52% and 57% for unfavoured players. The differences observed between favoured and unfavoured players’ shares are also significant for both cake sizes (Wilcoxon signed–ranks test, pooled NDG and UBG session–level data, $p \approx 0.020$ for the £5 cake, $p \approx 0.027$ for the £20 cake).

Figure 4 presents some more disaggregated information about the relationship between bargaining outcomes and disagreement payoffs. To construct this figure, we first classified the outcome from each individual pair of subjects in every round according to (a) whether the difference between favoured and unfavoured players’ disagreement payoffs (as shares of the cake) fell into the interval $[0, 0.05]$, $[0.05, 0.1)$, ... or $[0.35, 0.4]$, and (b) whether the difference between favoured and unfavoured players in a particular statistic (demands in NDG; payoffs conditional on agreement in NDG and UBG) as a share of the cake was in $[-1, -0.15]$, $(-0.15, -0.05]$, $(-0.05, +0.05]$, ... , $(+0.35, +0.45]$, $(+0.45, +0.55]$ or $(+0.55, 1]$. Then, for each of those three statistics, we recorded the total number of times

¹⁷Notice that favoured and unfavoured players’ payoffs don’t add up to 100% of the cake, even in the UBG conditional on agreement. Out of 901 agreements in this treatment, 11 left positive amounts of money “on the table”.

¹⁸See Siegel and Castellan (1988) for descriptions of the nonparametric statistical tests used in this paper, as well as for tables of critical values. We note that in implementing these tests, we err on the side of conservatism in two ways. First, we use session–level data rather than more disaggregated data, so that we ignore the information that can be gained by looking at individuals separately. (While individuals within a session should not be assumed to be independent of each other, neither are they perfectly correlated.) Second, we pool data from the NDG and UBG treatments; to the extent that these data are different in any important way, this will add a source of variance that will reduce the apparent significance of our test statistics.

the outcome fell into each of the 72 possible interval pairs (e.g., disagreement payoff difference in [0.05, 0.1) and difference between demands in (+0.35, +0.45]). Finally, for each of these 72 interval pairs, we plotted a circle whose radius is proportional to the number of outcomes in that interval pair (so that larger circles correspond to outcomes that were observed more often). Also shown in each panel of the figure, for comparison, are the horizontal line

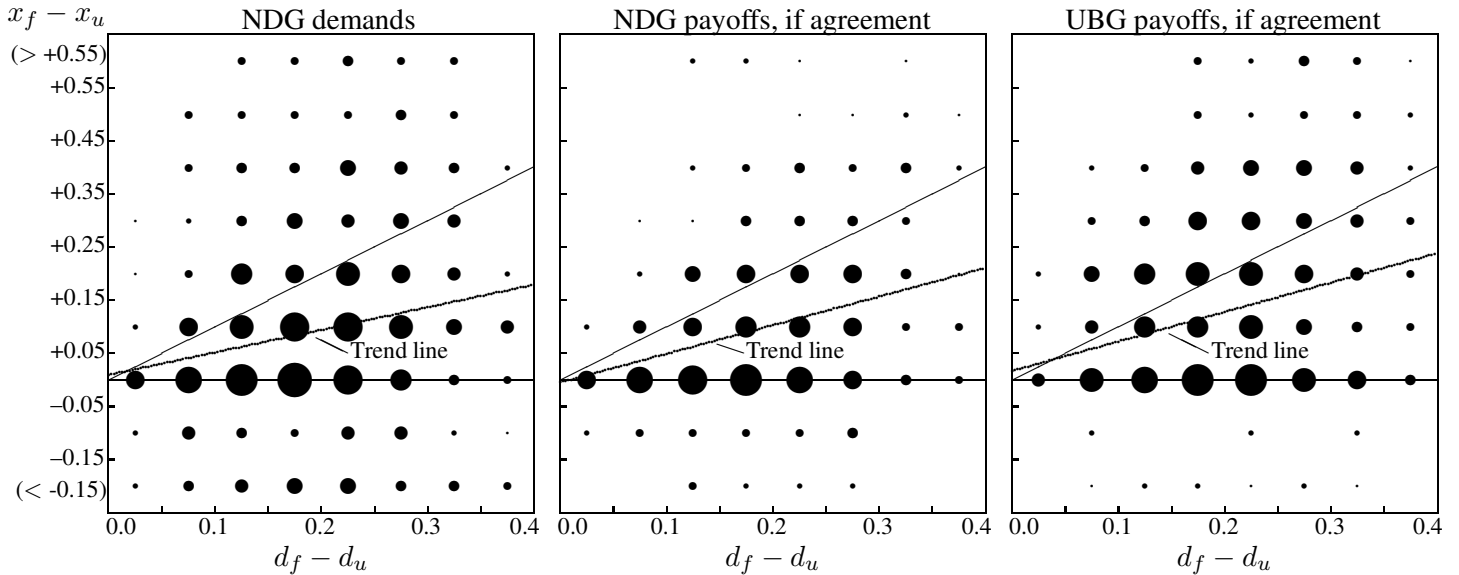


Figure 4: Bargaining outcomes as share of the cake, disaggregated by difference in disagreement payoffs (area of circle is proportional to number of outcomes)

Note: horizontal line represents equal split of the cake; diagonal solid line represents equal split of the surplus; diagonal dotted line represents linear least-squares fit to data

segment corresponding to an equal split of the cake and the diagonal segment corresponding to an equal split of the surplus.¹⁹ Additionally, each panel shows (as a dotted line) a least-squares trend line fitted to the data, to illustrate the association between changes in relative bargaining position and changes in bargaining outcomes.

As the figure illustrates, when neither player has a strongly advantageous position (the difference in disagreement payoffs is low), outcomes with approximately equal shares of the cake are most common, with most deviations in the direction favouring the player with the higher disagreement payoff. As the favoured player’s position improves ($d_f - d_u$ increases), there is an apparent tendency toward better outcomes for this player (as shown by the trend lines), but most outcomes continue to be between equal shares of the cake and equal shares of the surplus.

6.2 Parametric statistical analysis

We next use parametric methods to disentangle the effects of some of the factors that might influence bargaining outcomes in our two games. We begin by looking at subjects’ demands – as fractions of the cake or as fractions of the available surplus. For the former, we estimate Tobit models with zero and one as the endpoints; for the latter, we estimate linear models. In keeping with our hypotheses, our primary explanatory variables are the subject’s own disagreement payoff and that of the opponent. Additional right-hand-side variables are the player type (1=favoured

¹⁹Thus, circles below the horizontal line segment correspond to outcomes in which the unfavoured player received a larger absolute share of the cake (for example, if the favoured and unfavoured players capture 40% and 60% of the cake, respectively), while circles above the diagonal line segment correspond to outcomes with the favoured player capturing more than half of the available surplus (for example, if the disagreement payoffs are 30% and 10% of the cake, and the favoured and unfavoured players capture 80% and 20% respectively).

player), cake size (1=£20 cake), cake size ordering (1=increasing) and round number (1–20 for each cake size). All of the models were estimated using Stata (version 11), and incorporated individual–subject random effects.

Table 3 presents the results of these regressions: coefficient estimates and standard errors for each variable, and log likelihoods for each model. The main results are remarkably robust, changing little depending on whether we consider demands in the NDG or demands conditional on agreement in either game. Consistent with what was

Table 3: Regression results (coefficients and standard errors) – demands as proportions of the cake or of the surplus

| Dependent variable: Sample: | Demand, as fraction of cake | | | Demand, as fraction of surplus | | |
|---------------------------------------|-----------------------------|-----------------------|----------------------|--------------------------------|----------------------|----------------------|
| | NDG (all) | NDG (agreements) | UBG (agreements) | NDG (all) | NDG (agreements) | UBG (agreements) |
| constant | 0.513*** (0.024) | 0.434*** (0.017) | 0.481*** (0.015) | 0.483*** (0.049) | 0.454*** (0.035) | 0.467*** (0.031) |
| own disag. payoff (frac. of cake) | 0.235*** (0.044) | 0.318*** (0.038) | 0.280*** (0.034) | –0.468*** (0.090) | –0.528*** (0.080) | –0.457*** (0.071) |
| opp. disag. payoff (frac. of cake) | –0.218*** (0.044) | –0.217*** (0.038) | –0.287*** (0.034) | 0.654*** (0.090) | 0.397*** (0.080) | 0.427*** (0.071) |
| favoured player type | 0.011 (0.013) | –0.008 (0.011) | 0.031*** (0.011) | 0.022 (0.027) | –0.022 (0.024) | 0.061*** (0.023) |
| large cake | –0.008 (0.005) | –0.006 (0.004) | –0.000 (0.004) | –0.016 (0.010) | –0.013 (0.009) | –0.001 (0.008) |
| incr. cake–size order | 0.010 (0.026) | –0.013 (0.011) | 0.004 (0.009) | 0.022 (0.049) | –0.029 (0.022) | 0.008 (0.018) |
| round | 0.0003 (0.0004) | 0.0013*** (0.0003) | 0.0001 (0.0003) | 0.001 (0.001) | 0.003*** (0.001) | 0.000 (0.001) |
| <i>N</i> | 2160 | 1278 | 1802 | 2160 | 1278 | 1802 |
| $-\ln(L)$ | 1388.683 | 1389.439 | 1881.538 | 78.158 | 447.730 | 569.742 |

* (**, ***): Coefficient significantly different from zero at the 10% (5%, 1%) level.

seen in the descriptive statistics, demands as fractions of the cake size are sensitive to both a player’s own and the opponent’s disagreement payoff, but less sensitive than they should be according to the theoretical predictions. Instead of a £1 increase in one’s own disagreement option leading to the predicted £0.50 increase in one’s demand and payoff, the increase varies only from £0.23–0.32, depending on which statistic we are considering. Similarly, a £1 increase in the opponent’s own disagreement option should lead to a £0.50 decrease in one’s demand and payoff, but the actual decrease varies from £0.21–0.29. In all three of these models, differences between the coefficient for own disagreement payoff and +0.5, and differences between the coefficient for opponent disagreement payoff and –0.5, are significant at the 1% level or better (see Table 4). Moreover, chi–square tests find that the sum of these coefficients’ magnitudes is always significantly different from one at the 0.1% level or better. Additionally, we find weak evidence that subjects respond differently to changes in their own disagreement payoff than to changes in the opponent disagreement payoff, as in one of the three cases (agreements in the NDG), the magnitude of the own–disagreement–payoff effect is significantly larger than that of the opponent–disagreement–payoff effect, though only at the 10% level, and there is no significant difference in the other two cases. In sum, we are able to reject Hypotheses 1 and 2.

The comparative insensitivity of bargaining outcomes to changes in disagreement payoffs can also be seen on the right side of Table 3, which concentrates on demands as a fraction of the available surplus. As already noted, both

Table 4: Additional hypothesis test results from Table 3 regressions

| Dependent variable: | Demand, as fraction of cake | | |
|-------------------------------------|-----------------------------|-------------------|------------------|
| | Treatment: | NDG | NDG (agreements) |
| $\beta_{d_i} = +0.5$ | $p < 0.001$ | $p < 0.001$ | $p < 0.001$ |
| $\beta_{d_j} = -0.5$ | $p < 0.001$ | $p < 0.001$ | $p < 0.001$ |
| $ \beta_{d_i} + \beta_{d_j} = 1$ | $p < 0.001$ | $p < 0.001$ | $p < 0.001$ |
| $ \beta_{d_i} = \beta_{d_j} $ | $p \approx 0.79$ | $p \approx 0.056$ | $p \approx 0.88$ |

Notes: β_{d_i} = coefficient for own-disagreement-payoff variable; β_{d_j} = coefficient for opponent-disagreement-payoff variable

cooperative and non-cooperative bargaining solution techniques imply that these should be unaffected by changes to either player's disagreement payoff; however, the table shows a significant negative effect from the player's own disagreement payoff, and a significant positive effect from the opponent's disagreement payoff. That is, demands as a fraction of the surplus tend to decrease as one's own disagreement payoff increases, and increase as the opponent's disagreement payoff increases. These own-disagreement-payoff and opponent-disagreement-payoff variables are also jointly significant at the the 1% level or better in all three of these models. We can therefore also reject Hypotheses 3 and 4.

Lastly, we note that our other control variables have – for the most part – little apparent effect on bargaining outcomes. This includes the favoured-player dummy, which is significant only in the UBG, suggesting that the differences between the types seen in Tables 1 for the NDG can be explained by the sizes of their disagreement payoffs, rather than by being favoured or unfavoured per se. Also, the cake size seems to have little effect on demands, though this is not especially surprising in light of the fact that we vary it by a factor of only four.

Table 5 presents additional regression results, this time with agreement as the dependent variable and using a probit model with individual-subject random effects. This table shows little in the way of systematic results. In the NDG data, there is some evidence that players' disagreement payoffs have an effect on the frequency of agreement, as either decreasing the favoured player's disagreement payoff or increasing the unfavoured player's leads to a statistically significant increase in the likelihood of an agreement. (They are also jointly significant at the 5% level.) On the other hand, in the UBG data, neither player's disagreement payoff has a significant effect, nor are they jointly significant at conventional levels. We thus find mixed support for our Hypothesis 5.

7 Can risk aversion explain our main result? No.

One criticism that can be levelled at our experimental design, and interpretation of the results, is that bargaining in our experiment takes place over (expected) money amounts, while bargaining theory involves utilities. Treating these as equivalent is akin to assuming that bargainers are risk neutral, whereas there is substantial evidence that people are actually risk averse (see Holt and Laury, 2002, for evidence from a carefully designed experiment).²⁰

Of course, the pure-strategy Nash equilibria of the NDG (in particular the efficient equilibria, which include our prediction) are robust to assumptions about bargainers' risk attitudes, as long as utility is increasing in money for all players. However, it is well known that predictions arising from axiomatic bargaining solutions such as the

²⁰Some researchers have used the *binary lottery* mechanism (Roth and Malouf, 1979), in which players bargain over probabilities of winning a prize rather than monetary amounts, to control for risk aversion among expected utility maximising subjects.

Table 5: Probit regression results (coefficients and standard errors)

| Dependent variable: Treatment: | Agreement indicator | |
|-----------------------------------|---------------------|---------------------|
| | NDG | UBG |
| constant | 0.616** (0.312) | 0.433 (0.360) |
| d_f (fraction of cake) | -1.680** (0.725) | 0.176 (0.871) |
| d_u (fraction of cake) | 1.316* (0.708) | -0.427 (0.854) |
| large cake | 0.114 (0.083) | 0.036 (0.098) |
| increasing cake-size ordering | -0.269 (0.183) | 0.684*** (0.147) |
| round | 0.009 (0.007) | 0.029*** (0.008) |
| N | 1080 | 1080 |
| $-\ln(L)$ | 668.630 | 448.775 |

* (**, ***): Coefficient significantly different from zero at the 10% (5%, 1%) level.

Nash solution can differ under risk aversion compared to under risk neutrality; as an example, if bargainers differ in their level of risk aversion, the less risk averse bargainer will receive a larger share of the cake (Kannai, 1977; Roth, 1979).²¹ Also, the mixed-strategy equilibria of the NDG change when bargainers' risk attitudes change.

In this section, we examine the possibility that our main result, the under-sensitivity of bargaining outcomes to changes to disagreement payoffs, can be explained by relaxing the implicit assumption of risk neutrality: specifically, allowing bargainers to be risk averse. We will see that this is *not* the case; in fact, none of the commonly used classes of risk-averse expected-utility functions is able to explain this pattern of results.

To our knowledge, nearly all modelling of risk aversion uses one of two single-parameter families of expected-utility functions: those with *constant absolute risk aversion* (CARA) and those with *constant relative risk aversion* (CRRA). We begin by discussing CARA, which has the advantage (over CRRA and other expected-utility functions) that decision making under uncertainty is unaffected by the individual's current wealth level, which is nearly always unobservable to the researcher. The general form for a CARA utility function with risk aversion is $u(x) = -e^{-\alpha x}$, where x is the gain from bargaining and $\alpha > 0$ is a risk-aversion parameter.

Proposition 1 *If both bargainers are risk averse with (perhaps different) CARA utility functions, then the Nash bargaining solution implies $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = 1$.*²²

Proof: see Appendix A.

²¹As much of the literature does (e.g., Roth and Malouf, 1979; Rubinstein et al. 1992), we will abuse terminology somewhat by referring to "risk aversion" when we actually mean "diminishing marginal utility of money". Of course, the mathematics of the utility functions we use – and the results that derive from them – are unaffected by which of these interpretations of their curvature is used.

²²In this section and in the next, we assume that the utility functions of the bargainers are common knowledge, as is typical in this literature (see, e.g., Kannai, 1977 or Roth, 1979).

Note that an immediate corollary of the proposition is that $\left| \frac{\partial x_u}{\partial d_f} \right| + \left| \frac{\partial x_u}{\partial d_u} \right|$ is also equal to 1.²³ Proposition 1 tells us that even though the sensitivity of the payoff from bargaining to changes in own and opponent disagreement payoffs need not be $+\frac{1}{2}$ and $-\frac{1}{2}$ respectively, as they are in the case of risk neutrality, their magnitudes still must add up to one. By contrast, the corresponding sums in Table 3 are far less than one (they vary from about 0.45 to about 0.57). Thus, our results cannot be explained by risk aversion with CARA utility.

We next move to CRRA utility, which is even more widely used by experimental economists to model preferences of risk-averse subjects, despite the fact that CRRA implies that decisions under uncertainty are affected by unobserved wealth levels. The general form for a CRRA utility function is

$$u(w, x) = \begin{cases} \frac{1}{1-\alpha}(w+x)^{1-\alpha} & \text{with } \alpha > 0 \text{ and } \alpha \neq 1; \\ \ln(w+x) & \text{for } \alpha = 1; \end{cases}$$

where w is the individual's initial wealth and x is the gain from bargaining.

Proposition 2 *If both bargainers are risk averse with (perhaps different) CRRA utility functions, then the Nash bargaining solution implies $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.*

Proof: see Appendix A.

As with Proposition 1, an immediate corollary of Proposition 2 is that $\left| \frac{\partial x_u}{\partial d_f} \right| + \left| \frac{\partial x_u}{\partial d_u} \right| \geq 1$. Proposition 2 yields a slightly weaker result than Proposition 1, with weak inequality replacing equality. However, the inequality is in the wrong direction for explaining our result, leading to the same implication as before: CRRA utility does not account for the low values of $\left| \frac{\partial x_f}{\partial d_f} \right|$ and $\left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$ seen in the experimental data.²⁴

8 Other-regarding preferences

If risk aversion doesn't explain our results, what does? One possibility is that subjects have tastes for fairness that prevent them from making full use of their bargaining power, pushing outcomes toward 50–50 splits of the cake (as observed behaviour in dictator–game and ultimatum–game experiments seems to suggest; see Camerer, 2003, pp. 48–59 for a survey). There are now several models of such other-regarding preferences, and a full treatment of all of them is well beyond the scope of this paper. However, we show that a minor adaptation of the most widely used model – that of Fehr and Schmidt (1999) – is sufficient to explain the underreaction of bargaining outcomes to changes in disagreement payoffs.

In the Fehr–Schmidt (1999) model, players have utility functions that depend on both own and opponent money payments. Specifically, for Player $i = 1, 2$ in a two–player game,

$$U_i(x) = x_i - \alpha_i \cdot \text{Max}|x_j - x_i, 0| - \beta_i \cdot \text{Max}|x_i - x_j, 0|,$$

for $i = f, u$, with $0 \leq \beta_i < 1$ and $\alpha_i \geq \beta_i$. The first term is the money payment itself; the second term captures dislike for unfavourable inequality, which will be relevant for the unfavoured player in our setup; and the third term captures aversion to favourable inequality, relevant for the favoured player. Note that both types of disutility are

²³The (binding) constraint $x_f + x_u = M$ implies $\frac{\partial x_u}{\partial x_f} = -1$ and hence $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = \left| \frac{\partial x_u}{\partial d_f} \right| + \left| \frac{\partial x_u}{\partial d_u} \right|$ from the chain rule.

²⁴Similar methods to those used in the proof of Proposition 2 can be used to prove that when one bargainer has CARA utility and the other has CRRA utility, the result $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$ continues to hold. In addition, numerical techniques suggest that this property holds for general continuous and concave utility functions. However, we have thus far failed to find a direct proof of this latter claim.

linear in the magnitude of the inequality, and that standard own–payoff–maximising preferences are obtained when $\alpha = \beta = 0$.

Proposition 3 *If both bargainers have Fehr–Schmidt (1999) preferences, then the Nash bargaining solution implies $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right|$ is generically either 0 or 1.²⁵*

Proof: see Appendix A.

Intuitively, this model allows for two possibilities. If the players dislike inequality greatly (α_u or β_f is relatively large) or if the disagreement outcome is fairly equitable ($d_f - d_u$ is small), then the Nash bargaining solution yields an equal split, and $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = 0$. Otherwise, the Nash solution gives the favoured player strictly more than half of the cake, and $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = 1$.

Thus, while the basic Fehr–Schmidt model can yield a value of $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right|$ less than one, it does not yield values like those seen in our experiment. In order to get these, we must make a small change to the model. We do this by making the disutility of unfavourable inequality convex (rather than linear) in the magnitude of the inequality. Under this modification, the utility function becomes:

$$U_i(x) = x_i - \alpha_i \cdot (\text{Max}|x_j - x_i, 0|)^2 - \beta_i \cdot \text{Max}|x_i - x_j, 0|,$$

for $i = f, u$ and with $0 \leq \alpha_i, \beta_i < 1$.²⁶

Given these utility functions, the bargaining problem has the form of the one in Figure 5, as long as β_f and β_u are strictly less than one–half.²⁷ If either β_f or β_u is strictly positive, the Pareto frontier will be kinked at the equal split point $\left(\frac{M}{2}, \frac{M}{2} \right)$, and when $\alpha_f > 0$ (resp. $\alpha_u > 0$), the upper (lower) segment of the Pareto frontier will be bowed away from the origin.

As in the basic Fehr–Schmidt model, when $d_f > d_u$, the Nash bargaining solution will either yield an equal split (in this case, when $\beta \geq \frac{d_f - d_u}{d_f - 3d_u + M}$) or a division favourable to the favoured player (when $\beta < \frac{d_f - d_u}{d_f - 3d_u + M}$). If the latter is true, the favoured player receives

$$\begin{aligned} x_f &= \frac{-1}{12\alpha(2\beta - 1)} \left\{ -1 + 2\beta + 4\alpha d_f - 4\alpha\beta(d_f - d_u) + 4\alpha M(1 - 3\beta) \right. \\ &+ \frac{1}{2} \left(-48\alpha(2\beta - 1)[(\beta - 1)(4\alpha m - 1)d_f + (3\beta - 4\alpha\beta m - 1)d_u + m(1 - 3\beta - \alpha m + 6\alpha\beta m)] \right. \\ &\left. \left. + (2 - 8\alpha(d_f + m) + 4\beta(2\alpha(d_f - d_u + 3m) - 1))^2 \right)^{1/2} \right\}, \end{aligned}$$

and the sum of own–disagreement–payoff and opponent–disagreement–payoff effects is given by

$$\frac{2 - 4\beta + 4\alpha d_f - 4\alpha\beta(d_f - d_u) - 2\alpha m + \sqrt{K}}{3\sqrt{K}},$$

²⁵Generically, because there is an additional knife–edge case where $\left| \frac{\partial x_f}{\partial d_f} \right| > 0$ and $\left| \frac{\partial x_f}{\partial d_u} \right| = 0$, as noted in the appendix. Of course, this case also cannot characterise our experimental results, since we find that $\left| \frac{\partial x_f}{\partial d_u} \right|$ is well above zero in all treatments.

²⁶A more general utility function with both linear and quadratic terms for both favourable and unfavourable inequality would also give the result we obtain here, and would have the additional advantage of symmetric treatment of both types of inequality; however the version we use has the advantage of having the same number of free parameters as the basic Fehr–Schmidt model, as well as mathematical tractability.

²⁷If $\beta_f \geq \frac{1}{2}$, increases in the favoured player’s payoff beyond 50% of the cake (ceteris paribus) do not increase her utility, so that the lower segment in Figure 5 would be positively sloped (vertical in the case of $\beta_f = \frac{1}{2}$). Then the Nash bargaining solution would yield an equal split for any $d_f \geq d_u$, and $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = 0$. Similarly, if $\beta_u \geq \frac{1}{2}$, the upper segment in Figure 5 would be positively sloped (horizontal in the case of $\beta_u = \frac{1}{2}$, though the Nash solution would be unaffected as long as $d_f \geq d_u$).




Figure 5: Example of bargaining set under variation of Fehr–Schmidt preferences

where

$$\begin{aligned}
K &= 1 + 4[4\alpha^2(d_f - d_u)^2 + 2\alpha(d_f - 7d_u + 3m) + 1]\beta^2 \\
&\quad - 4[4(d_f - d_u)(2d_f - m)\alpha^2 + (3d_f - 13d_u + 5m)\alpha + 1]\beta \\
&\quad + 4\alpha[\alpha(m - 2d_f)^2 + d_f - 3d_u + m].
\end{aligned}$$

An illustration of how this expression depends on α and β is given by Figure 6. Each panel shows, for a particular disagreement outcome (d_f, d_u) , the region of the (α, β) unit square where $\left|\frac{\partial x_f}{\partial d_f}\right| + \left|\frac{\partial x_f}{\partial d_u}\right| = 0$, and “iso-effect” curves where $\left|\frac{\partial x_f}{\partial d_f}\right| + \left|\frac{\partial x_f}{\partial d_u}\right| = 0.4, 0.5$ and 0.6 (values similar to what we observed in the experiment).




Figure 6: Selected values of $\left|\frac{\partial x_f}{\partial d_f}\right| + \left|\frac{\partial x_f}{\partial d_u}\right|$ under modified Fehr–Schmidt preferences (£5 cake, three disagreement outcomes)

As the figure shows, values of $\left|\frac{\partial x_f}{\partial d_f}\right| + \left|\frac{\partial x_f}{\partial d_u}\right|$ in the range of what we saw in the experiment are easily achievable using this modified Fehr–Schmidt model. Moreover, Table 6 shows that the divisions of the cake according to these parameters are also very similar to typical divisions observed in the experiment. This table shows, for the £5 cake and for the three disagreement outcomes used in Figure 6, the minimum and maximum value of x_f implied by all

Table 6: Favoured player shares of cake: ranges implied by modified Fehr–Schmidt preferences, and observed means from experiment (£5 cake)

| | | Disagreement outcome | | |
|--------------------|---|----------------------|----------------|----------------|
| | | (2.00, 0.50) | (1.75, 0.75) | (1.50, 1.00) |
| Model implications | $\frac{\partial x_f}{\partial d_f} + \frac{\partial x_f}{\partial d_u} = 0.4$ | (0.501, 0.549) | (0.504, 0.535) | (0.502, 0.520) |
| | $\frac{\partial x_f}{\partial d_f} + \frac{\partial x_f}{\partial d_u} = 0.5$ | (0.507, 0.567) | (0.508, 0.547) | (0.503, 0.521) |
| | $\frac{\partial x_f}{\partial d_f} + \frac{\partial x_f}{\partial d_u} = 0.6$ | (0.515, 0.587) | (0.510, 0.554) | (0.511, 0.522) |
| Experimental data | NDG | 0.547 | 0.505 | 0.490 |
| | UBG | 0.599 | 0.566 | 0.536 |

parameterisations of our modified Fehr–Schmidt model that yield values of $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right|$ equal to 0.4, 0.5 or 0.6 (as in Figure 6). Also shown are the corresponding mean observed payoffs for the favoured player (conditional on agreement), where the disagreement outcome was within £0.25 for each player.²⁸

9 Discussion and concluding remarks

The standard theoretical techniques used for analysing bargaining situations – both axiomatic solutions and non-cooperative game-theoretic methods – make sharp, testable predictions for bargaining situations involving a fixed, known cake and a known disagreement outcome. For each unit one’s own disagreement payoff increases, or alternatively for each unit the opponent’s disagreement payoff decreases, one’s own payoff from bargaining increases by one-half of a unit.

We conduct a human–subjects experiment to test empirically whether this property actually holds. Subjects play asymmetric bargaining games repeatedly against changing opponents, with disagreement payoffs chosen randomly in each round for both favoured and unfavoured players. In the interest of robustness, we vary the particular bargaining game played – in the Nash Demand Game (NDG), bargaining consists only of a pair of simultaneous demands, while in the Unstructured Bargaining Game (UBG), subjects can freely make proposals and counter-proposals over a specified period of time – as well as the stake size (a £5 cake versus a £20 cake) and the order in which these stake sizes were faced. Our design is novel, as there has been very little previous study of the effects of disagreement payoffs on bargaining outcomes, and (to our knowledge) no study that attempts to disentangle the effects of one’s own disagreement payoff from the effects of the opponent’s disagreement payoff.

Our main finding is that while bargaining outcomes do vary with changes to subjects’ bargaining positions, they vary substantially less than predicted by the theory. This is true for both bargaining games (NDG and UBG), for both low and high stakes, and for both orderings of stake sizes. Specifically, we find that a one–unit increase in a subject’s disagreement payoff translates to an increase of only 0.24 units in that subject’s demand in the NDG, while a one–unit increase in the opponent’s disagreement payoff in that game translates to a decrease of only 0.22 units, in contrast to theoretical predictions of 0.5 units in both cases. If we focus on outcomes where bargaining was successful, results are broadly similar: a one–unit increase in a subject’s own disagreement payoff is associated with payoff increases of 0.32 in the NDG and 0.28 in the UBG, while a one–unit increase in the opponent’s disagreement payoff is associated

²⁸For example, the means for the column “(2.00, 0.50)” were calculated from the observations where the disagreement outcome gave amounts in (1.875, 2.125) to the favoured player and amounts in (0.375, 0.625) to the unfavoured player.

with payoff decreases of 0.22 in the NDG and 0.29 in the UBG, again compared to predicted changes of 0.5 in each case. For the most part, subjects underreact equally to changes in their own and their opponents' disagreement payoffs, though in one case, we find weak evidence that subjects are more sensitive to their own disagreement payoff than to the opponent's.

One common criticism of experiments in which subjects bargain over money amounts (such as our experiment, as well as most other bargaining experiments including those of Hoffman and Spitzer, 1982 and 1985, and Fischer et al., 2007, while Harrison, 1987, used the binary lottery technique only in the event of agreement) is that axiomatic bargaining solutions involve utility amounts, not money amounts, so that results that seem to be inconsistent with these solutions (when they are applied to money amounts) might simply be showing that utility cannot be identified with monetary payments (that is, subjects are not risk-neutral expected-utility maximisers). However, we show in Section 7 that if bargainers are risk averse, with utility functions that satisfy either of the two widely used models of risk-averse preferences (constant absolute risk aversion or constant relative risk aversion), the theoretical implication of the Nash bargaining solution is almost as strong: while it does not imply that the magnitudes of own-disagreement-payoff effect and the opponent-disagreement-payoff are each 0.5, it still implies that their sum is at least 1. Hence, we conclude that our experimental results cannot be accounted for by subjects' risk aversion on its own.

Another explanation for seemingly anomalous results in bargaining experiments involves other-regarding preferences; indeed, several such models have been developed at least partly in order to explain such results (e.g., Rabin, 1993; Fehr and Schmidt, 1999). Fehr and Schmidt's (1999) model of inequity aversion is probably the most widely used model of other-regarding preferences, combining substantial explanatory power and mathematical simplicity. While we show that Fehr and Schmidt's basic model is also unable to account for our main result, we also demonstrate that a slight adaptation to their model can account for this result.

We hasten to acknowledge that our illustration that other-regarding preferences can explain our result does not constitute proof that it is the sole cause, even after being able to rule out risk aversion as an alternative explanation.²⁹ There may be still other explanations; for example, it may be that subjects are affected by the framing of the bargaining problem in our experiment. While the theory predicts that subjects completely internalise the disagreement payments, so that bargaining occurs only over the remainder of the cake (the individually rational portion of the bargaining set), some subjects might fail to do so, instead concentrating on the entire feasible bargaining set.³⁰ This would also push outcomes toward the 50–50 split, and decrease sensitivity to the disagreement point.

Our experiment was not designed to distinguish between other-regarding preferences and other competing (though not mutually exclusive) explanations, so at best, we could hope to find indirect evidence in favour of one of them. On the face of it, the fact that very similar results were observed under both low stakes and high stakes might speak against the other-regarding preferences explanation, since one might expect subjects to be less willing to express tastes for equity as they become more costly (that is, as the cake size increases). Such intuition is found not only in some theories of other-regarding preferences (e.g., Rabin, 1993), but also in some experimental results

²⁹In particular, we certainly do not claim based on this that people are not risk averse; it is easy to show, for example, that a model that combined inequity aversion and risk aversion can also explain the results observed in the experiment. All that we conclude is that risk aversion on its own is neither necessary nor sufficient to explain these results. We also note that even if the other-regarding-preferences explanation is correct, the particular inequity-aversion model we use is not the only one consistent with our results, though it might be the simplest such model.

³⁰One potential cause of such failure to internalise is that subjects may have been reluctant to exploit a favourable bargaining position that they considered to be "unearned", along the lines of Hoffman and Spitzer's (1982) result, mentioned in Section 3. Future experiments might allow favoured/unfavoured status, and the size of the disagreement payoffs, to be assigned based on the result of a "real effort" task.

(Slonim and Roth, 1995; Cameron, 1999). However, these supporting experimental results have typically involved quite large changes in stake sizes (payoff ratios of 50 and 40 in the two aforementioned experiments respectively), so not observing a difference with stakes raised only by a factor of four is likely not conclusive evidence. Moreover, not all theories of fairness predict such changes in behaviour as stake sizes increase; for example, Fehr and Schmidt's (1999) basic model predicts no stake-size effect at all (though the variation we consider in Section 8 does predict an effect: as the cake becomes larger, the division moves away from a 50–50 split in absolute terms, but closer to it in relative terms).

We would like to encourage other experimental researchers to replicate our results and attempt to distinguish amongst the alternative explanations described above, and others. Based on our results, we would also like to encourage theorists, when constructing models involving bargaining, to consider whether limiting attention to the individually rational portion of the bargaining set is as innocuous as it's usually assumed to be.

Finally, we would also like to point out that even though our main results are at odds with the standard theory, there are some silver linings in our results for axiomatic bargaining theory. The fact that behaviour appears robust to differences in the cake size suggests that subjects are able to normalise the cake size when they face a common scale factor in payoffs; that is, they do not violate the “homogeneity” axiom (Kalai, 1977), which all known axiomatic solutions satisfy. Moreover, since additionally the sets of disagreement payoffs in our experiment scale up proportionally to the cake size, the subjects do not seem to violate the “origin invariance” (OI) component of the “scale and origin invariance” axiom (SOI) either.³¹

References

- Binmore, K. (2007), *Playing for Real: a Text on Game Theory*, Oxford University Press, Oxford.
- Binmore, K., A. Shaked, and J. Sutton (1985), “Testing noncooperative bargaining theory: a preliminary study,” *American Economic Review* 75, pp. 1178–1180.
- Binmore, K., J. Swierzbinski, S. Hsu, and C. Proulx (1993), “Focal points and bargaining,” *International Journal of Game Theory* 22, pp. 381–409.
- Bohnet, I and B. Frey (1999), “Social distance and other-regarding behavior in dictator games: comment,” *American Economic Review* 89, pp. 335–339.
- Bolton, G.E. (1991), “A comparative model of bargaining: theory and evidence,” *American Economic Review* 81, pp. 1096–1136.
- Camerer, C.F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press, Princeton, NJ.
- Cameron, L. (1999), “Raising the stakes in the ultimatum game: experimental evidence from Indonesia,” *Economic Inquiry* 37, pp. 47–59.

³¹SOI was first coined as the “Independence of Equivalent Utility Representations” axiom by Nash (1950) but later became to be known as SOI. Other experimenters, such as Nydegger and Owen (1974), have observed violations of the “scale invariance” (SI) part of SOI in lab experiments. We note also that Kalai's (1977) homogeneity is a weaker property than SI; the former requires only invariance when the bargaining set and disagreement outcome are scaled by a common factor for all players, while the latter requires invariance even to scaling by different factors for different players.

- de Groot Ruiz, A., R. Ramer, and A. Schram (2010), "Structure in legislative bargaining", working paper, CREED, University of Amsterdam.
- Edgeworth, F.Y. (1881), *Mathematical Psychics: An Essay on the Application of Mathematics to the Moral Sciences*, Kegan Paul, London.
- Fehr, E., and K.M. Schmidt (1999), "A theory of fairness, competition and cooperation", *Quarterly Journal of Economics* 114, pp. 817–868.
- Feltovich, N. (2005), "Critical values for the robust rank–order test", *Communications in Statistics—Simulation and Computation* 34, pp. 525–547.
- Feltovich, N. and J. Swierzbinski (2010), "The role of strategic uncertainty in games: an experimental study of cheap talk and unstructured bargaining in the Nash demand game", forthcoming, *European Economic Review*.
- Fischbacher, U. (2007), "z–Tree: Zurich toolbox for ready–made economic experiments", *Experimental Economics* 10, pp. 171–178.
- Fischer, S., W. Güth, W. Müller, and A. Stiehler (2006), "From ultimatum to Nash bargaining: theory and experimental evidence", *Experimental Economics* 9, pp. 17–33.
- Fischer, S., W. Güth, and K. Pull (2007), "Is there as–if bargaining?" *Journal of Socio–Economics* 36, pp. 546–560.
- Gächter, S. and Riedl, A. (2005), "Moral property rights in bargaining with infeasible claims", *Management Science* 51, pp. 249–263.
- Harrison, G., (1987), "Risk aversion and the Nash solution in stochastic bargaining experiments", *Economics Letters* 24, pp. 321–326.
- Harsanyi, J.C. and R. Selten (1988), *A General Theory of Equilibrium Selection in Games*, MIT Press, Cambridge, MA.
- Hoffman, E. and M.L. Spitzer (1982), "The Coase theorem: some experimental tests", *Journal of Law and Economics* 25, pp. 73–98.
- Hoffman, E. and M.L. Spitzer (1985), "Entitlements, rights, and fairness: an experimental examination of subjects' concepts of distributive justice," *Journal of Legal Studies* 14, pp. 259–297.
- Holt, C.A. and S.K. Laury (2002), "Risk aversion and incentive effects", *American Economic Review* 92, pp. 1644–1655.
- Kalai, E. (1977), "Proportional solutions to bargaining situations: interpersonal utility comparisons", *Econometrica* 45, pp. 1623–1630.
- Kalai, E. and M. Smorodinsky (1975), "Other solutions to Nash's bargaining problem", *Econometrica* 43, pp. 513–518.
- Kannai, Y. (1977), "Concavifiability and constructions of concave utility functions", *Journal of Mathematical Economics* 4, pp. 1–56.

- Munier, B. and C. Zaharia (2003), “High stakes and acceptance behavior in ultimatum bargaining”, *Theory and Decision* 53, pp. 187–207.
- Murnighan, J.K., A.E. Roth, and F. Schoumaker (1988), “Risk aversion in bargaining: an experimental study”, *Journal of Risk and Uncertainty* 1, pp. 101–124.
- Nash, J. (1950), “The bargaining problem”, *Econometrica* 18, pp. 155–162.
- Nash, J. (1953), “Two–person cooperative games”, *Econometrica* 21, pp. 128–140.
- Nydegger, R. V. and Owen, G. (1974), “Two–person bargaining: an experimental test of the Nash axioms”, *International Journal of Game Theory* 3, pp. 239–249.
- Perry, M. and P.J. Reny (1993), “A non–cooperative bargaining model with strategically timed offers”, *Journal of Economic Theory* 59, pp. 50–77.
- Perry, M. and P.J. Reny (1994), “A noncooperative view of coalition formation and the core”, *Econometrica* 62, pp. 795–817.
- Rabin, M. (1993), “Incorporating fairness into game theory and economics”, *American Economic Review* 83, pp. 1281–1302.
- Raiffa, H. (1953), “Arbitration schemes for generalized two–person games”, in *Contributions to the Theory of Games. II: Annals of Mathematics Studies*, H.W. Kuhn and A.W. Tucker, eds., Princeton University Press, Princeton, NJ, pp. 361–387.
- Rankin, F.W. (2006), “Requests and social distance in dictator games”, *Journal of Economic Behavior & Organization* 60, pp. 27–36.
- Roth, A.E. (1979), *Axiomatic Models of Bargaining* Springer, Berlin and New York.
- Roth, A.E. (1995), “Bargaining experiments”, in *Handbook of Experimental Economics*, J.H. Kagel and A.E. Roth, eds., Princeton University Press, Princeton, NJ, pp. 253–348.
- Roth, A.E. and M.W.K. Malouf (1979), “Game–theoretic models and the role of information in bargaining”, *Psychological Review* 86, pp. 574–594.
- Roth, A.E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir (1991), “Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study”, *American Economic Review* 81, pp. 1068–1095.
- Rubinstein, A. (1982), “Perfect equilibrium in a bargaining model”, *Econometrica* 50, pp. 97–109.
- Rubinstein, A., Z. Safra, and W. Thomson (1992), “On the interpretation of the Nash bargaining solution and its extension to non–expected utility preferences”, *Econometrica* 60, pp. 1171–1186.
- Siegel, S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences*, McGraw–Hill, New York.
- Simon, L.K. and M.B. Stinchcombe (1989), “Extensive form games in continuous time: pure strategies”, *Econometrica* 57, pp. 1171–1214.

Skyrms, B. (1996), *Evolution of the Social Contract*, Cambridge University Press, Cambridge.

Slonim, R. and A.E. Roth (1998), “Learning in high stakes ultimatum games: an experiment in the Slovak Republic”, *Econometrica* 66, pp. 569–596.

A Proofs of Propositions 1, 2 and 3

A.1 Proposition 1: CARA utility

Suppose both bargainers are risk averse, with (perhaps different) CARA utility functions. We wish to show that the Nash bargaining solution implies $\left| \frac{dx_f}{dd_f} \right| + \left| \frac{dx_f}{dd_u} \right| = 1$.

CARA utility implies that the bargainers’ utility functions can be written (if necessary, by taking an affine transformation) as

$$\begin{aligned} u_f(x) &= -e^{-\alpha x} \\ u_u(x) &= -e^{-\beta x} \end{aligned}$$

for the favoured and unfavoured players respectively, with $\alpha, \beta > 0$. The feasible bargaining set can then be written

$$S = \{(u_f(x_f), u_u(x_u)) : x_f + x_u \leq M\},$$

and the disagreement outcome is $d = (u_f(d_f), u_u(d_u))$.

For this bargaining problem, the outcome implied by the Nash solution is the pair (x_f, x_u) that maximises the Nash product $[u_f(x_f) - u_f(d_f)][u_u(x_u) - u_u(d_u)]$ such that $x_f + x_u \leq M$. Plugging in the players’ utility functions gives us the constrained optimisation problem solved by the Nash bargaining solution:

$$\begin{aligned} \text{Maximise} \quad & [(-e^{-\alpha x_f}) - (-e^{-\alpha d_f})][(-e^{-\beta x_u}) - (-e^{-\beta d_u})] \\ \text{subject to} \quad & x_f + x_u \leq M. \end{aligned}$$

Since both bargainers’ utility functions are strictly increasing in money, the cake-size constraint will be binding: $x_f + x_u = M$. The optimisation problem thus has an implicit solution for x_f and x_u in terms of parameters (along with $x_f + x_u = M$):

$$(\alpha - \beta)e^{-[\alpha x_f + \beta x_u]} = \alpha e^{-[\alpha x_f + \beta d_u]} - \beta e^{-[\alpha d_f + \beta x_u]}. \quad (1)$$

To find the effect on x_f and x_u of changes to the disagreement payoffs, totally differentiate Equation 1 to yield

$$(\alpha - \beta)e^{-[\alpha x_f + \beta x_u]} \cdot (\beta - \alpha)dx_f = \alpha e^{-[\alpha x_f + \beta d_u]} \cdot (-\alpha \cdot dx_f - \beta \cdot dd_u) - \beta e^{-[\alpha d_f + \beta x_u]} \cdot (-\alpha \cdot dd_f - \beta \cdot dd_u),$$

and collecting terms gives us

$$\begin{aligned} & \left[\beta^2 e^{-[\alpha d_f + \beta x_u]} - (\alpha - \beta)^2 e^{-[\alpha x_f + \beta x_u]} + \alpha^2 e^{-[\alpha x_f + \beta d_u]} \right] dx_f \\ & = \alpha \beta e^{-[\alpha d_f + \beta x_u]} dd_f - \alpha \beta e^{-[\alpha x_f + \beta d_u]} dd_u. \end{aligned} \quad (2)$$

Squaring both sides of Equation 1, and substituting into the middle term in the top line of Equation 2, allows us to simplify:

$$\left(e^{-[\alpha d_f + \beta x_u]} + e^{-[\alpha x_f + \beta d_u]} \right) dx_f = e^{-[\alpha d_f + \beta x_u]} dd_f - e^{-[\alpha x_f + \beta d_u]} dd_u,$$

so that

$$dx_f = \frac{e^{-[\alpha d_f + \beta x_u]}}{e^{-[\alpha d_f + \beta x_u]} + e^{-[\alpha x_f + \beta d_u]}} dd_f - \frac{e^{-[\alpha x_f + \beta d_u]}}{e^{-[\alpha d_f + \beta x_u]} + e^{-[\alpha x_f + \beta d_u]}} dd_u.$$

This last equation implies that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = 1$, proving Proposition 1.

A.2 Proposition 2: CRRA utility

Suppose both bargainers are risk averse, with (perhaps different) CRRA utility functions. We wish to show that the Nash bargaining solution implies $\left| \frac{dx_f}{dd_f} \right| + \left| \frac{dx_f}{dd_u} \right| \geq 1$.

CRRA utility implies that the bargainers' utility functions can be written (if necessary, by taking an affine transformation) as

$$\begin{aligned} u_f(x) &= \frac{1}{1-\alpha}(w_f+x)^{1-\alpha} \text{ with } \alpha \neq 1, \text{ or } u_f(x) = \ln(w_f+x) \text{ for } \alpha = 1 \\ u_u(x) &= \frac{1}{1-\beta}(w_u+x)^{1-\beta} \text{ with } \beta \neq 1, \text{ or } u_u(x) = \ln(w_u+x) \text{ for } \beta = 1 \end{aligned}$$

for the favoured and unfavoured players respectively, with $\alpha, \beta > 0$ (a value of 0 implies risk neutrality), and where w_f and w_u are their initial (non-negative) wealth levels. The outcome implied by the Nash solution is the pair (x_f, x_u) that maximises the Nash product $[u_f(x_f) - u_f(d_f)][u_u(x_u) - u_u(d_u)]$ such that $x_f + x_u \leq M$. As with CARA utility, strict monotonicity of CRRA utility implies that $x_f + x_u = M$ at the solution.

Demonstrating that $\left| \frac{dx_f}{dd_f} \right| + \left| \frac{dx_f}{dd_u} \right| \geq 1$ for all versions of CRRA utility requires breaking up the space of (α, β) pairs into nine subsets, according to whether α and β are greater than, less than or equal to one. Below are three of the nine possible cases; the others proceed analogously and are left out for space reasons, but can be obtained from the corresponding author upon request.

Case 1: $\alpha, \beta < 1$

The resulting constrained optimisation problem is

$$\begin{aligned} \text{Maximise} \quad & \left[\frac{1}{1-\alpha}(w_f+x_f)^{1-\alpha} - \frac{1}{1-\alpha}(w_f+d_f)^{1-\alpha} \right] \left[\frac{1}{1-\beta}(w_u+x_u)^{1-\beta} - \frac{1}{1-\beta}(w_u+d_u)^{1-\beta} \right] \\ \text{subject to} \quad & x_f + x_u \leq M. \end{aligned}$$

Solving yields the Nash condition

$$(1-\beta) \left[w_f + x_f - (w_f + x_f)^\alpha (w_f + d_f)^{1-\alpha} \right] = (1-\alpha) \left[w_u + x_u - (w_u + x_u)^\beta (w_u + d_u)^{1-\beta} \right].$$

Totally differentiating the Nash condition yields

$$\begin{aligned} \left[2 - \alpha - \beta - \alpha(1-\beta) \left(\frac{w_f+x_f}{w_f+d_f} \right)^{\alpha-1} - \beta(1-\alpha) \left(\frac{w_u+x_u}{w_u+d_u} \right)^{\beta-1} \right] dx_f \\ = (1-\alpha)(1-\beta) \left[\left(\frac{w_f+x_f}{w_f+d_f} \right)^\alpha dd_f - \left(\frac{w_u+x_u}{w_u+d_u} \right)^\beta dd_u \right]. \end{aligned}$$

Define $y_f = \frac{w_f+x_f}{w_f+d_f}$ and $y_u = \frac{w_u+x_u}{w_u+d_u}$; note that both are greater than or equal to one, since $x_f \geq d_f$ and $x_u \geq d_u$.

Then the above simplifies to

$$\left[2 - \alpha - \beta - \alpha(1-\beta)y_f^{\alpha-1} - \beta(1-\alpha)y_u^{\beta-1} \right] dx_f = (1-\alpha)(1-\beta) \left[y_f^\alpha dd_f - y_u^\beta dd_u \right],$$

so that

$$\begin{aligned} \frac{\partial x_f}{\partial d_f} &= \frac{(1-\alpha)(1-\beta)y_f^\alpha}{2 - \alpha - \beta - \alpha(1-\beta)y_f^{\alpha-1} - \beta(1-\alpha)y_u^{\beta-1}} \\ &= \frac{(1-\alpha)(1-\beta)y_f^\alpha}{(1-\alpha)(1-\beta y_u^{\beta-1}) + (1-\beta)(1-\alpha y_f^{\alpha-1})}, \end{aligned} \tag{3}$$

and similarly

$$\frac{\partial x_f}{\partial d_u} = -\frac{(1-\alpha)(1-\beta)y_u^\beta}{(1-\alpha)(1-\beta y_u^{\beta-1}) + (1-\beta)(1-\alpha y_f^{\alpha-1})}. \quad (4)$$

Note that whenever $\alpha < 1$, both $1-\alpha$ and $1-\alpha y_f^{\alpha-1}$ are positive (since $y_f, y_u \geq 1$), and whenever $\alpha > 1$, both are negative, and similarly for β . This means that when α and β are both larger or both smaller than 1, the numerators and denominators of Equations 3 and 4 are positive, and when α and β are on opposite sides of 1, both numerators and denominators are negative. Since for this case we are assuming that $\alpha, \beta < 1$, we have

$$\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = \frac{(1-\alpha)(1-\beta)(y_f^\alpha + y_u^\beta)}{2-\alpha-\beta-\alpha(1-\beta)y_f^{\alpha-1}-\beta(1-\alpha)y_u^{\beta-1}}.$$

Let Num and Den be the numerator and denominator of the right-hand-side expression above:

$$\begin{aligned} Num &= (1-\alpha)(1-\beta)(y_f^\alpha + y_u^\beta) > 0 \\ \text{and } Den &= 2-\alpha-\beta-\alpha(1-\beta)y_f^{\alpha-1}-\beta(1-\alpha)y_u^{\beta-1} > 0, \end{aligned}$$

and let $D = Num - Den$. We want to establish that $D \geq 0$ and thence that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.

We can write

$$D = D(\alpha, \beta, y_f, y_u) = -2 + \alpha + \beta + (1-\alpha)(1-\beta)[y_f^\alpha + y_u^\beta] + \alpha(1-\beta)y_f^{\alpha-1} + \beta(1-\alpha)y_u^{\beta-1}. \quad (5)$$

Now,

$$\begin{aligned} \frac{\partial D}{\partial y_f} &= \alpha(1-\alpha)(1-\beta)y_f^{\alpha-1} - \alpha(1-\alpha)(1-\beta)y_f^{\alpha-2} \\ \text{and } \frac{\partial D}{\partial y_u} &= \beta(1-\alpha)(1-\beta)y_u^{\beta-1} - \beta(1-\alpha)(1-\beta)y_u^{\beta-2}, \end{aligned}$$

and it is easy to show that $\frac{\partial D}{\partial y_f} = 0$ when $y_f = 1$ and $\frac{\partial D}{\partial y_u} = 0$ when $y_u = 1$; that is, $(y_f, y_u) = (1, 1)$ is a stationary point of D . Also, note that $D(\alpha, \beta, y_f = 1, y_u = 1) = 0$ for any α and β .

Finally, define $\hat{\alpha} = \alpha(1-\alpha)(1-\beta)$ and $\hat{\beta} = \beta(1-\alpha)(1-\beta)$, and note that $\hat{\alpha}, \hat{\beta} > 0$. Then we have

$$\begin{aligned} \frac{\partial^2 D}{\partial y_f^2} &= \hat{\alpha}(\alpha-1)y_f^{\alpha-2} - \hat{\alpha}(\alpha-2)y_f^{\alpha-3} \\ &= \hat{\alpha}y_f^{\alpha-3}[(\alpha-1)y_f - (\alpha-2)] > 0 \end{aligned}$$

(since $y_f \geq 1$), and similarly

$$\frac{\partial^2 D}{\partial y_u^2} = \hat{\beta}y_u^{\beta-3}[(\beta-1)y_u - (\beta-2)] > 0.$$

Since $\frac{\partial^2 D}{\partial y_u \partial y_f} = \frac{\partial^2 D}{\partial y_f \partial y_u} = 0$, D reaches a global minimum when $y_f = y_u = 1$, and as we have shown, $D = 0$ there.

This means that $D(\alpha, \beta, y_f, y_u) \geq 0$, and thus that $Num \geq Den$, and thus that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.

Case 2: $\alpha < \beta = 1$

In this case, the resulting constrained optimisation problem is

$$\begin{aligned} \text{Maximise} & \quad \left[\frac{1}{1-\alpha}(w_f + x_f)^{1-\alpha} - \frac{1}{1-\alpha}(w_f + d_f)^{1-\alpha} \right] [\ln(w_u + x_u) - \ln(w_u + d_u)] \\ \text{subject to} & \quad x_f + x_u \leq M, \end{aligned}$$

and the implicit solution is

$$\left[w_f + x_f - (w_f + x_f)^\alpha (w_f + d_f)^{1-\alpha} \right] = (w_u + x_u) [\ln(w_u + x_u) - \ln(w_u + d_u)].$$

Then, following the Case 1 steps up to Equation 5 yields

$$D = D(\alpha, 1, y_f, y_u) = -2 + (1 - \alpha)y_f^\alpha + \alpha y_f^{\alpha-1} + y_u - \ln(y_u).$$

As in Case 1, $D(\alpha, 1, y_f = 1, y_u = 1) = 0$ for any α , and $\frac{\partial D}{\partial y_f} = \alpha(1 - \alpha)y_f^{\alpha-1} - \alpha(1 - \alpha)y_f^{\alpha-2}$ and $\frac{\partial D}{\partial y_u} = 1 - \frac{1}{y_u}$, so both first derivatives are zero when $y_f = y_u = 1$. Also, it is easy to show that $\frac{\partial^2 D}{\partial y_f^2}$ and $\frac{\partial^2 D}{\partial y_u^2}$ are positive, and $\frac{\partial^2 D}{\partial y_u \partial y_f} = \frac{\partial^2 D}{\partial y_f \partial y_u} = 0$, so as in Case 1, D reaches a global minimum of 0 when $y_f = y_u = 1$, again entailing that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.

Case 3: $\alpha < 1 < \beta$

In this case, the constrained optimisation problem and solution are as in Case 1, and following the steps up to Equations 3 and 4 again yields

$$\frac{\partial x_f}{\partial d_f} = \frac{(1 - \alpha)(1 - \beta)y_f^\alpha}{(1 - \alpha)(1 - \beta y_u^{\beta-1}) + (1 - \beta)(1 - \alpha y_f^{\alpha-1})} \quad (6)$$

$$\text{and } \frac{\partial x_f}{\partial d_u} = -\frac{(1 - \alpha)(1 - \beta)y_u^\beta}{(1 - \alpha)(1 - \beta y_u^{\beta-1}) + (1 - \beta)(1 - \alpha y_f^{\alpha-1})}. \quad (7)$$

In this case, however, the numerators and denominators of both fractions are negative, so that

$$\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = \frac{(1 - \alpha)(1 - \beta) [y_f^\alpha + y_u^\beta]}{2 - \alpha - \beta - \alpha(1 - \beta)y_f^{\alpha-1} - \beta(1 - \alpha)y_u^{\beta-1}},$$

noting that both numerator and denominator are negative. So, defining Num , Den and D as in Case 1, we have

$$D = D(\alpha, \beta, y_f, y_u) = -2 + \alpha + \beta + (1 - \alpha)(1 - \beta) [y_f^\alpha + y_u^\beta] + \alpha(1 - \beta)y_f^{\alpha-1} + \beta(1 - \alpha)y_u^{\beta-1},$$

but since Num and Den are negative, we need to show that $D \leq 0$ in order to find that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.

As in Case 1, $D(\alpha, \beta, y_f = 1, y_u = 1) = 0$ for any α and β , and the first derivatives are

$$\begin{aligned} \frac{\partial D}{\partial y_f} &= \alpha(1 - \alpha)(1 - \beta)y_f^{\alpha-1} - \alpha(1 - \alpha)(1 - \beta)y_f^{\alpha-2} \\ \text{and } \frac{\partial D}{\partial y_u} &= \beta(1 - \alpha)(1 - \beta)y_u^{\beta-1} - \beta(1 - \alpha)(1 - \beta)y_u^{\beta-2}, \end{aligned}$$

with both equal to zero when $y_f = y_u = 1$. Finally, taking second derivatives shows that $\frac{\partial^2 D}{\partial y_f^2}$ and $\frac{\partial^2 D}{\partial y_u^2} < 0$, and again $\frac{\partial^2 D}{\partial y_u \partial y_f} = \frac{\partial^2 D}{\partial y_f \partial y_u} = 0$, so that D reaches a global maximum at 0 when $y_f = y_u = 1$. This means that $D(\alpha, \beta, y_f, y_u) \leq 0$, and thus that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| \geq 1$.

These three cases, along with the other six (which are proved analogously), complete the proof of Proposition 2.

A.3 Proposition 3: Fehr–Schmidt (1999) preferences

Suppose both players have utility functions as in the Fehr–Schmidt (1999) model:

$$U_i(x) = x_i - \alpha_i \cdot \text{Max}|x_j - x_i, 0| - \beta_i \cdot \text{Max}|x_i - x_j, 0|,$$

with $0 < \beta_i < 1$ and $\alpha_i \geq \beta_i$. We wish to show that $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right|$ is generically equal to zero or one.³²

Then, since $d_f \geq d_u$, we have $U_f(d) = d_f - \beta_f(d_f - d_u) = (1 - \beta_f)d_f + \beta_f d_u$, and $U_u(d) = d_u - \alpha_u(d_f - d_u) = (1 + \alpha_u)d_u - \alpha_u d_f$; note that $U_f(d) \geq d_u \geq U_u(d)$. Also, as long as $x_f \geq x_u$, we will have $U_f(x) = (1 - \beta_f)x_f + \beta_f x_u$, and $U_u(x) = (1 + \alpha_u)x_u - \alpha_u x_f$.

We begin by noting that irrespective of α_f , α_u , β_f and β_u , adding the same amount to both x_f and x_u always makes both players strictly better off (increasing the first term of the utility function, leaving the other two terms unchanged), so the Nash bargaining solution implies $x_f + x_u = M$.

Next, we prove a result about (x_f, x_u) .

Lemma 1 *If $d_f \geq d_u$, the Nash bargaining solution implies $x_f \geq x_u$.*

Proof: Consider the level curves of the Nash bargaining solution, given by

$$[U_f(x) - U_f(d)][U_u(x) - U_u(d)] = K.$$

Each of these curves has a slope of -1 along the ray $U_f(x) - U_f(d) = U_u(x) - U_u(d)$ (or equivalently $U_f(x) - U_u(x) = U_f(d) - U_u(d)$), is steeper (slope less than -1) to the left of this ray (i.e., where $U_f(x) - U_u(x) < U_f(d) - U_u(d)$) and is flatter to the right of it. Since $U_f(d) \geq U_u(d)$, these level curves must therefore have slope less than -1 when $U_u(x) > U_f(x)$.

Now, suppose by contradiction that the Nash solution implies $x_f < x_u$. Then $U_f(x) = x_f - \alpha_f(x_u - x_f) < x_f$, and $U_u(x) = x_u - \beta_u(x_u - x_f) = (1 - \beta_u)x_u + \beta_u x_f > x_f$, so that $U_u(x) > U_f(x)$. This means that one of the Nash solution level curves is tangent to the upper segment of the Pareto frontier at (x_f, x_u) with $x_f < x_u$. However, this segment is linear, with endpoints $\left(\frac{M}{2}, \frac{M}{2}\right)$ (where each player gets $M/2$) and $(-\alpha_f M, (1 - \beta_u)M)$ (where the unfavoured player gets the entire M), so its slope is

$$\frac{(1 - \beta_u)M - \frac{M}{2}}{-\alpha_f M - \frac{M}{2}} = - \left(\frac{1 - 2\beta_u}{1 + 2\alpha_f} \right) > -1$$

(since the fraction in parentheses has a numerator less than or equal than one, and denominator greater than or equal to one). Since the slope of this segment is greater than -1 , it cannot be tangent to any Nash solution level curve at (x_f, x_u) with $x_f < x_u$, completing the proof of the lemma.

From Lemma 1, we need not be concerned with α_f and β_u , so we can simplify notation by dropping the subscripts for α and β : $\alpha \equiv \alpha_u$ and $\beta \equiv \beta_f$.

Then, the Nash bargaining solution solves the constrained optimisation problem

$$\begin{aligned} \text{Maximise} \quad & [(1 - \beta)(x_f - d_f) + \beta(x_u - d_u)][(1 + \alpha)(x_u - d_u) - \alpha(x_f - d_f)] \\ \text{subject to} \quad & x_f + x_u \leq M \text{ and } x_f \geq x_u. \end{aligned}$$

There are two possible solutions, depending on whether the constraint $x_f \geq x_u$ is binding (see Figure 7).

$$x_f = \text{Max} \left\{ \frac{1 + 3\alpha - \beta - 4\alpha\beta}{2 + 4\alpha - 4\beta - 8\alpha\beta} d_f + \frac{1 + \alpha - 3\beta - 4\alpha\beta}{2 + 4\alpha - 4\beta - 8\alpha\beta} (M - d_u), \frac{M}{2} \right\}.$$




Figure 7: Nash bargaining solution outcomes under Fehr–Schmidt preferences, when $d_f \geq d_u$

Case 1: $(1+2\alpha)(1-2\beta)(d_f-d_u) > (\alpha+\beta)(M-d_f-d_u)$. Then, $x_f = \frac{1+3\alpha-\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta}d_f + \frac{1+\alpha-3\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta}(M-d_u) > \frac{M}{2}$, so that $x_f > x_u$. In this case,

$$\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right| = \frac{1+3\alpha-\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta} + \frac{1+\alpha-3\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta} = 1.$$

Case 2: $(1+2\alpha)(1-2\beta)(d_f-d_u) < (\alpha+\beta)(M-d_f-d_u)$. Then, $\frac{1+3\alpha-\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta}d_f + \frac{1+\alpha-3\beta-4\alpha\beta}{2+4\alpha-4\beta-8\alpha\beta}(M-d_u) < \frac{M}{2}$, so $x_f = x_u = \frac{M}{2}$. Then $\left| \frac{\partial x_f}{\partial d_f} \right| = \left| \frac{\partial x_f}{\partial d_u} \right| = 0$, so their sum is zero as well.

Thus, generically we have $\left| \frac{\partial x_f}{\partial d_f} \right| + \left| \frac{\partial x_f}{\partial d_u} \right|$ equal to zero or one. (In the knife-edge case where $(1+2\alpha)(1-2\beta)(d_f-d_u) = (\alpha+\beta)(M-d_f-d_u)$, we have $\left| \frac{\partial x_f}{\partial d_f} \right| > 0 = \left| \frac{\partial x_f}{\partial d_u} \right|$, which also is not consistent with our experimental results.)

³²Small modifications to the proof show that the result continues to hold when the α s and β s can be zero.

B Sample instructions

Below is the text of instructions from our cell with the NDG and increasing cake sizes, followed by that from our cell with the UBG and decreasing cake sizes. The other two sets of instructions are available from the corresponding author upon request.

Instructions: first part of experiment [NDG, increasing cake sizes]

You are about to participate in a decision making experiment. Please read these instructions carefully, as the amount of money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment.

This experiment consists of two parts, each made up of 20 rounds. These instructions are for the first half; you will receive instructions for the second half after this half has ended. Each round in this half consists of one play of a simple bargaining game, played between two people via the computer. In every round, you are randomly matched to another participant, with whom you will play this bargaining game. You will not be told the identity of the person you are matched with in any round, nor will they be told your identity – even after the end of the session.

The bargaining game is as follows. You and the person matched to you bargain over a £5.00 prize. You and the other person make simultaneous claims for shares of this prize.

- If your claims add up to the amount of the prize or less, you receive your claim, and the other person receives his/her claim.
- If your claims add up to more than the amount of the prize, you receive an “outside option”, and the other person receives a different “outside option”.

These outside options are chosen randomly by the computer, and vary from round to round and from person to person. In each round, you and the person matched to you are informed of both of your outside options before choosing your claims.

Sequence of Play: The sequence of play in a round is as follows.

- (1) The computer randomly matches you to another participant, and randomly determines your outside option and the outside option of the other person. Your computer screen will display both your outside option and that of the other person.
- (2) You choose a claim for your share of the £5.00 prize. The other person chooses a claim for his/her share of the prize. Your claim can be any multiple of 0.01, between zero and 5.00 inclusive. Both of you choose your claim before being informed of the other's.
- (3) The round ends. You receive the following information: your own choice, the choice made by the person matched with you, your own payoff for the round, the payoff of the person matched with you.

After this, you go on to the next round.

Payments: At the end of the experimental session, two rounds from this half will be chosen randomly for each participant. You will be paid the total of your earnings in those two rounds. In addition, there will be opportunities for payments in the second half of the session. Payments are made privately and in cash at the end of the session.

Instructions: second part of experiment

The procedure in this part of the experiment is nearly the same as that in the first part. You will play the same bargaining game as before, for 20 additional rounds. The participant matched with you will still be chosen randomly in every round, and your outside options will also be chosen randomly in every round.

The difference from the first part of the experiment is that the prize is now worth £20. So, you and the other person are now choosing shares of £20 instead of £5. Your claim – and that of the other person – can now be any multiple of 0.01, between zero and 20.00 inclusive.

As before, if your claims add up to the amount of the prize or less, you receive your claim, and the other person receives his/her claim. If your claims add up to more than the amount of the prize, you both receive your respective outside options.

At the end of the experimental session, two rounds from this half will be chosen randomly for each participant. You will be paid the total of your earnings in those two rounds. Your earnings from this part of the experiment will be added to your earnings from the previous part.

Instructions: first part of experiment [UBG, decreasing cake sizes]

You are about to participate in a decision making experiment. Please read these instructions carefully, as the amount of money you earn may depend on how well you understand them. If you have a question at any time, please feel free to ask the experimenter. We ask that you not talk with the other participants during the experiment.

This experiment consists of two parts, each made up of 20 rounds. These instructions are for the first half; you will receive instructions for the second half after this half has ended. Each round in this half consists of one play of a simple bargaining game, played between two people via the computer. In every round, you are randomly matched to another participant, with whom you will play this bargaining game. You will not be told the identity of the person you are matched with in any round, nor will they be told your identity – even after the end of the session.

The bargaining game is as follows. You and the person matched to you bargain over a £20.00 prize. You do this by sending and receiving proposals for dividing the prize during a “negotiation stage” of the game. Below is an example of how the bottom portion of your computer screen will look during the negotiation stage.

To make a proposal, type it into the spaces in the bottom-left corner, then click the SEND PROPOSAL button. Proposals sent and received will appear in the boxes below.

To accept a proposal from the player matched with you, select that proposal in the box in the bottom-right corner and click the ACCEPT PROPOSAL button.

If you would like to end bargaining, click the END BARGAINING button on the right. If you click this button, you and the other person will receive your outside options.

End this stage

| Make a proposal: | Proposals made by you: | | Proposals made by other person: | |
|--|----------------------------|------------------------------------|---------------------------------|---|
| | Your proposal for yourself | Your proposal for the other person | Other person's proposal for you | Other person's proposal for him/herself |
| Your proposal for yourself: <input type="text"/> | | | | |
| Your proposal for the other person: <input type="text"/> | | | | |
| Send proposal | | | | Accept proposal |

To send a proposal to the other person, type the amounts for yourself and the other person in the “Make a proposal” box, then click “Send proposal”. The amounts you enter must be between zero and the amount of the prize (inclusive), and can have 0, 1 or 2 decimal places. The two amounts together must add up to the amount of the prize, or less. All of your proposals will appear in the box in the bottom-centre of your screen, and all of the proposals made by the other person will appear in the box in the bottom-right. The person matched to you will see these proposals as well, but no one else will be able to see your proposals, nor will you be able to see theirs.

You may accept any one of the proposals from the person matched to you, or none of them. To accept a proposal, highlight the one you wish to accept and click “Accept proposal”. If either you or the other person accepts a proposal, then you have reached an agreement, and the prize is divided according to the accepted proposal.

The negotiation stage lasts for up to 90 seconds; you may send as many or as few proposals as you wish during that time. You may end the negotiation stage before the 90 seconds are over, by clicking on the button labelled “End this stage” on the right of your screen. Once you or the person matched with you has clicked this button, it is not possible to send or accept proposals.

If you or the other person ends the negotiation stage early, or if the time available for proposals ends without you reaching an agreement, then you receive an “outside option”, and the other person receives a different “outside option”. These outside options are chosen randomly by the computer, and vary from round to round and from person to person. In each round, you and the person matched to you are informed of both of your outside options at the beginning of the negotiation stage.

Sequence of Play: The sequence of play in a round is as follows.

- (1) The computer randomly matches you to another participant, and randomly determines your outside option and the outside option of the other person. Your computer screen will display both your outside option and that of the other person.
- (2) The negotiation stage begins. You can send proposals for dividing the £20.00 prize. The other person can also send proposals for dividing the £20.00 prize; you can accept one of these proposals or none of them.
- (3) The round ends. You receive the following information: whether or not you reached an agreement, your own payoff, the payoff of the person matched with you.
After this, you go on to the next round.

Payments: At the end of the experimental session, two rounds from this half will be chosen randomly for each participant. You will be paid the total of your earnings in those two rounds. In addition, there will be opportunities for payments in the second half of the session. Payments are made privately and in cash at the end of the session.

Instructions: second part of experiment

The procedure in this part of the experiment is nearly the same as that in the first part. You will play the same bargaining game as before, for 20 additional rounds. The participant matched with you will still be chosen randomly in every round, and your outside options will also be chosen randomly in every round.

The difference from the first part of the experiment is that the prize is now worth £5. So, you and the other person are now sending and receiving proposals for dividing £5 instead of £20. The amounts you propose for yourself – and for the other person – can now be any multiple of 0.01, between zero and 5.00 inclusive, and they must add up to 5.00 or less.

As before, if you or the other person ends the negotiation stage early, or if the time available for proposals ends without reaching agreement, then you both receive your respective outside options.

At the end of the experimental session, two rounds from this half will be chosen randomly for each participant. You will be paid the total of your earnings in those two rounds. Your earnings from this part of the experiment will be added to your earnings from the previous part.