Behavioral Modeling of Sequential Bargaining Games: Fairness and Limited Backward Induction

Xia Qu and Prashant Doshi
THINC Lab, Dept. of Computer Science
University of Georgia, Athens, GA 30602
quxia@uga.edu, pdoshi@cs.uga.edu

Abstract

Experiments show that in sequential bargaining games (SBG), subjects usually deviate from game-theoretic predictions. Previous explanations have focused on considerations of fairness in the offers, and social utility functions have been formulated to model the data. However, a recent explanation by Ho and Su (2013) for observed deviations from game-theoretic predictions in sequential games is that players engage in limited backward induction. A suite of computational models that integrate different choice models with utility functions are comprehensively evaluated on SBG data. These include DeBruyn and Bolton's recursive quantal response with social utility functions, those based on Ho and Su's dynamic level-k, and analogous extensions of the cognitive hierarchy with dynamic components. Our comprehensive analysis reveals that in extended SBG with 5 rounds, models that capture violations of backward induction perform better than those that model fairness. However, we did not observe this result for SBG with less rounds, and fairness of the offer remains a key consideration in these games.

Introduction

Sequential bargaining games (SBG) (Stahl, 1972; Rubinstein, 1982) typically consist of a finite number of rounds of bargaining over the partition of a pie between a proposer and a responder. If a partition from the proposer is accepted by the responder, the game ends and the pie is divided according to this partition; otherwise, both players exchange their roles in the next round and repeat the process until a partition is accepted or the predefined number of rounds elapse in which case both players get nothing. Discount factors are applied to the pie in subsequent rounds for both players in order to make a disagreement in bargaining costly. The game-theoretic solution is derived by applying backward induction (Stahl, 1972; Rubinstein, 1982).

Researchers have conducted several experiments on how humans behave in SBG (Binmore et al., 1985; Neelin et al., 1988; Ochs and Roth, 1989; Roth et al., 1991) and related negotiation games (Gal and Pfeffer, 2007). [1] Results from these experiments show that subjects primarily do not take the rational actions prescribed by backward induction. For example, in one-round bargaining games, also called ultimatum games, proposers in the experiments predominantly offer some amount that is less than half the pie but more than the rational offer to responders, and lower offers were rejected frequently (Roth et al., 1991). In two-round SBGs, most opening offers are between the equal split and subgame perfect equilibria (Binmore et al., 1985), though different discount factors for the proposers and responders affect the results slightly (Ochs and Roth, 1989). In the three-round and five-round extended games, we seldom observe subgame perfect equilibria amounts in the opening offers, and a majority of these are close to the second round pie size (Neelin et al., 1988).

One reason commonly explored to explain human behavioral deviation from rational play is that social factors such as fairness and reciprocity may affect a human player's utility (Johnson et al., 2002; Gal and Pfeffer, 2007). Based on this hypothesis, De Bruyn and Bolton (2008) investigate the role of fairness by employing two different utility functions and incorporating them into a quantal response equilibrium framework (McKelvey and Palfrey, 1995) for recursively computing the expected utility at each round. The quantal response model accounts for noise and experience. The two compared utility models include one with an equity-reciprocity competition (Bolton and Ockenfels, 2000) and the Fehr-Schmidt model (Fehr and Schmidt, 1999) that generalizes the previous utility model. The reported out-of-sample fits and model predictions on multiple data sets are consistent: the two models involving social factors exhibit better performance than the normative model, though the difference in performance varies for different data sets. Consequently, fairness is an important factor that needs to be considered while modeling bargaining.

Another reason that may potentially explain the deviation from backward induction in SBG is that humans may be exhibiting *limited backward induction* (Johnson et al., 2002; Ho and Su, 2013). Backward induction based solution is achieved under the assumption that players are rational and they believe that their opponents are rational who in turn believe that others are rational. However, this may not be

[1] The sequential negotiation game studied by Gal and Pfeffer (2007) involves a proposer who offers an exchange of colored chips to a responder who may accept or reject it. If the offer is rejected, the proposer and responder swap roles. While sharing important similarities with the SBG, it differs in a key aspect that the utility of the chips does not reduce over the finite rounds.

the case for human players as research suggests that humans have bounded recursive thinking power (Hedden and Zhang, 2002).

Ho and Su (2013) introduce a model that attributes level-k rules of varying levels to the other agent with Bayesian update of the distribution over the levels, allowing it to account for limited backward induction. This dynamic level-k model was fitted to the experimental data (McKelvey and Palfrey, 1992) on Centipede games (Rosenthal, 1981) giving a better likelihood compared to the default static level-k and backward induction. In a preliminary outline, Ho and Su proposed its applicability toward SBG as well but did not use it to fit any data. SBGs differ from Centipede games in having a much larger number of action choices for the proposer and role swapping, which complicates the construction of these models.

We construct a dynamic level-k model for SBG and a new model that generalizes the Poisson cognitive hierarchy (Camerer et al., 2004) with a Bayesian belief update to account for learning in repeated SBG. In contrast to the level-k model, a level-k player, $k > 0$, in the Poisson cognitive hierarchy chooses the action that optimizes its belief, which assumes a Poisson distribution over all of the lower levels. While the dynamic level-k model generally follows Ho and Su's paradigm, our construction differs in two ways: an agent at level 0 maximizes its expected utility instead of acting randomly and we integrate a quantal response function (McKelvey and Palfrey, 1995) to model decision errors.

Next, we perform a comprehensive and comparative analysis of the three different models for SBG data in combination with three different utility models two of which involve social factors. Given the new ways of thinking about how humans play sequential games, such an analysis is previously lacking and is needed crucially for SBG where data from multiple experiments is available. By combining the new level-based models with social utility models, we analyze the fit of two dominant theories on human behavior in two-player sequential games. For this analysis, we construct a large data set from 6 different experiments and perform out-of-sample testing. Interestingly, no one particular model performs the best on all the data sets. Rather, the level-based models perform better on data from an extended number of bargaining rounds while the recursive quantal response with social utility better models the behavioral data from SBG with less rounds.

In the following sections, we first talk about background of sequential bargaining games including data set. We then briefly talk about three social utility models. After that, we focus on three choice models which include the quantal response model (De Bruyn and Bolton, 2008) and two level-based models. We also compare the performance of these three choice models integrated with different utility models on experiment data set.

Background

SBG are widely studied (Stahl, 1972; Rubinstein, 1982; Binmore et al., 1985; Neelin et al., 1988; Ochs and Roth, 1989; Roth et al., 1991; Johnson et al., 2002; De Bruyn and Bolton, 2008) because of their widespread use and economic impact. We briefly describe SBG followed by an outline of the experimental data that is available.

Sequential Bargaining Games

In a multi-round SBG, two agents, i and j, bargain over how to split a pie of c units in n rounds. In odd rounds (initial round, t = 1), agent i proposes a split and offers x of the pie to agent j. If agent j accepts this proposal, the game ends and the pie is divided accordingly with i receiving $c - x$ and j receiving x. If agent j rejects the offer, the game continues to round $t + 1$ and the pie shrinks according to the discount factor, (δ_i, δ_j), where δ_i is a number between 0 and 1 representing agent i's cost of delay, and analogously for j. In even rounds, the two agents exchange their roles and agent j is the proposer. If no agreement is achieved after n rounds, the game ends and both players get nothing. An example of a 3-round game used by Neelin et al. (1988), with the pie size of $c = 5$, the discount factors, $\delta_i = 0.50$, $\delta_j = 0.50$, is shown in Fig. 1.

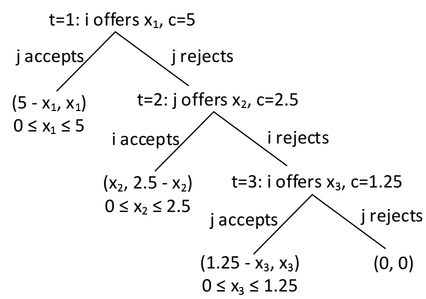

Figure 1: An example 3-round SBG from Neelin et al. (1988) with a pie size of 5 and discount factors, (0.50, 0.50). The split, (a,b), denotes that portion a is allocated to agent i and b is allocated to agent j.

The subgame perfect equilibrium of an n-round SBG is computed by backward induction. To illustrate, in 1-round (ultimatum) games, player j prefers some amount than nothing and would accept any offer; player i therefore should propose the smallest possible offer. In 2-round games, the proposer in the second round is j who would offer the smallest possible offer to player i and keep slightly less than $c\delta_j$. Therefore, in the first round, in order to make j accept the offer, proposer i needs to offer slightly larger than what j would get in round 2. Hence, i should offer $c\delta_j$ to j in the first round. In this way, we can compute the subgame perfect equilibrium for games with any number of rounds.

Behavioral Experiments on SBG

Extensive experimentation with human subjects playing SBG (Binmore et al., 1985; Güth and Tietz, 1986; Neelin et al., 1988; Ochs and Roth, 1989; Bolton, 1991; Roth et al., 1991; Güth and van Damme, 1998) has yielded much data that may inform the computational modeling of behavior in SBG. Among these experiments, we select data from those in which the SBGs exhibit more rounds. This allows the

Source	Data set	Pie Size	Number of Rounds	Discount factor	Number of subjects	Game repetitions	Mean opening	Game theory prediction
Neelin et al. (1988)	$NSS^{5,1}$	\$5	5	(0.34, 0.34)	80	1	0.343	0.250
	$NSS^{5,4}$	\$15	5	(0.34, 0.34)	30	4	0.359	0.250
Ochs and Roth (1989)	$OR^{3,10}$	\$30	3	(0.40, 0.40)	20	10	0.433	0.240
	$OR^{3,10}$	\$30	3	(0.60, 0.40)	20	10	0.450	0.160
	$OR^{3,10}$	\$30	3	(0.60, 0.60)	18	10	0.451	0.235
	$OR^{3,10}$	\$30	3	(0.40, 0.60)	18	10	0.466	0.350

Table 1: Experimental designs and collected data on SBG for our modeling. Values in the column, Game repetitions, indicate the number of games played by each subject in the pool. Values in the columns, Mean opening and Game theory prediction, are fractions of the opening pie size.

performance of level-based models such as level-k and Poisson cognitive hierarchy to be adequately tested and distinguished. Specifically, it facilitates attributing the other agent different levels and maintaining beliefs over these levels, which is otherwise precluded by 1- and 2-round games. Subsequently, we select data for SBGs with 3- and 5-rounds. Furthermore, experiments in which a small pool of subjects repeatedly play SBGs allow the applicability of the dynamic aspects of these models. This formed our second criteria for selecting the data sets.

Having searched a wide range of experiment data sets, we find six that satisfy our requirements: data from 2 experiments with 5-round games under different conditions by Neelin et al. (1988) and data from 4 experiments with 3-round games under different conditions by Ochs and Roth (1989). We summarize the experimental settings, results and the corresponding equilibrium predictions in Table 1. Observe that the averaged opening offers deviate substantially from equilibrium-based game theory predictions.

Social Utility Models

A well-known reason for behavioral deviations from rational play is that subjects may be influenced by social factors, and a factor especially relevant to SBG is fairness (or equitability). Human decision making in bargaining games and in negotiation (Gal and Pfeffer, 2007) not only considers the absolute payoff that is received from an offer but also the relative utility by comparing the received offer to the opponent's.

We consider two prominent utility models that include fairness considerations. One is the *equity-reciprocity competition* (ERC) (Bolton and Ockenfels, 2000), which is formally defined for a two-agent setting as:

$$U(\rho; b) = \begin{cases} c\left(\rho - \frac{b}{2}\left(\rho - \frac{1}{2}\right)^2\right) & \text{if } \rho < 1/2, \\ c\rho & \text{if } \rho \geq 1/2 \end{cases} \quad (1)$$

where ρ is the proportion of the pie of size c the agent – proposer or responder – *receives*, and b is a fairness factor measuring the importance of any inequitable allocation. The utility function defined in ERC when the proportion is less than half is nonlinear and consists of two components: the relative utility due to negative reciprocity reduced from the absolute payoff. Observe that the model is asymmetric because the utility is increasingly less when the agent is offered less than half the pie but it is the same as the payoff when the share is more than half the pie.

The second model is the *Fehr-Schmidt model* (FSC) (Fehr and Schmidt, 1999):

$$U(\rho; \alpha, \beta) = \begin{cases} c\left(\rho - \alpha\left(\frac{1}{2} - \rho\right)\right) & \text{if } \rho < 1/2, \\ c\left(\rho - \beta\left(\rho - \frac{1}{2}\right)\right) & \text{if } \rho \geq 1/2 \end{cases} \quad (2)$$

where ρ is as defined for Eq. 1, α and β represent the negative and positive reciprocity parameters, respectively.

The utility function in FSC is symmetric, a more wholesome measure of fairness and is a linear variant of the one used in ERC when a player receives less than half the pie. When the share is greater than half, relative utility due to positive reciprocity is reduced as well. [2]

For comparison, we also consider the *normative utility model* (NORM) in which the utility is simply defined as the payoff that each player receives.

$$U(\rho) = c\rho \quad (3)$$

These three social utility models, NORM, FSC, and ERC are incorporated into different choice models.

Quantal Response based Choice Model

De Bruyn and Bolton (2008) incorporate different social utility models into a recursive quantal response framework, denoted as RQR, which provides a distribution over the subject's actions. The quantal response choice model assigns probabilities to actions proportionally to their utilities. The cumulative probability for the responder accepting an offer proportion, $\sigma^t \in [0, 1]$, is the logit function:

$$P_{rsp}^{\sigma^t}(acc) = \frac{e^{\lambda \cdot U(\sigma^t)}}{e^{\lambda \cdot U(\varnothing)} + e^{\lambda \cdot U(\sigma^t)}} \quad (4)$$

where $U(\varnothing)$ is the utility of rejecting an offer, it is 0 when t is the last round, otherwise, it is the expected utility that the responder can get in the next round, $t + 1$, as a proposer. λ is the quantal parameter which controls how rational is the agent. The probability for the responder rejecting an offer proportion σ^t, $P_{rsp}^{\sigma^t}(rej)$ is defined similarly.

Let $P_{prp}(\sigma^t)$ be the probability of the proposer making an offer, σ^t, at round t. Then,

[2] Relative utilities due to negative and positive reciprocity may also be crudely characterized as guilt and envy (Ray et al., 2008).

$$P_{prp}(\sigma^t) = \frac{e^{\lambda \cdot E(U(1-\sigma^t))}}{\sum_{\sigma'} e^{\lambda \cdot E(U(1-\sigma'))}} \quad (5)$$

where $E(U(1 - \sigma^t))$ is the proposer's expected utility of offering σ. This takes the responder's probability of acceptance of this offer, $P_{rsp}^{\sigma^t}(acc)$, and rejecting this offer, $P_{rsp}^{\sigma^t}(rej)$, into consideration, $E(U(1 - \sigma^t)) = P_{rsp}^{\sigma^t}(acc) \cdot U(1 - \sigma^t) + P_{rsp}^{\sigma^t}(rej) \cdot E(U(\sigma^{t+1}))$.

Both ERC and FSC were integrated with RQR and compared by De Bruyn and Bolton (2008). The results showed that RQR+FSC performs relatively better than RQR+ERC.

Dynamic Level based Choice Models

Previous modeling of SBG data has explored the role of fairness with the conclusion that considering it provides an improved fit of the data compared to the normative model (De Bruyn and Bolton, 2008). Recently, Ho and Su (2013) illustrate that the constraint of *limited backward induction* helps explain behavioral data in multi-stage sequential games such as Centipede games, and possibly better compared to considerations of social factors such as positive and negative reciprocity as in Fehr and Schmidt (1999).

Multi-stage games exhibit the property of limited induction when a larger deviation from the backward induction based predictions is observed in extended stage games compared to games with lesser stages. In other words, limited induction causes the deviation to grow as the number of stages increases.

A dynamic level-k model (Ho and Su, 2013) captures this systematic violation of backward induction. Johnson et al. (2002) noted that subjects in a SBG study paid significantly less attention to the later rounds and that social factors do not explain the data, thereby providing preliminary evidence that subjects are engaging in limited induction.

Our first contribution is a construction of the dynamic level-k and Poisson cognitive hierarchy (Camerer, 2003) models for the multi-round SBG, in order to investigate whether limited backward induction capability provides an improved explanation of the behavioral data. Note that the latter represents violations of backward induction as well. We generalize these models to include belief-based learning for repeated SBG.

Level-k Model With Belief Update

Our dynamic level-k model (DLK) for predicting the actions of an agent, say i, in an n-round two-agent SBG consists of iterative decision rules and a belief distribution over these rules. Rules of levels ranging from 0 up to $n-1$ are attributed to the other agent, j, for predicting its possible action(s). A level l decision rule, $\theta_{j,l}$, where $0 \leq l \leq k-1$ and having l steps lookahead, is obtained by applying backward induction to an l-round subgame. An agent at level 0 acts as an utility maximizer believing she is the only agent in the game. Proposer i at level n initially has a *Poisson* distribution, parameterized by τ, over the different decision rules (probabilities are normalized over all possible levels):

$$b_i(\theta_{j,l}; \tau) \propto \frac{\tau^l e^{-\tau}}{l!} \quad 0 \leq l \leq k-1 \quad (6)$$

where $b_i(\theta_{j,l}; \tau)$ represents the probability that i believes j is a level l player. An example of i's belief in a 3-round SBG is given in Figure 2.

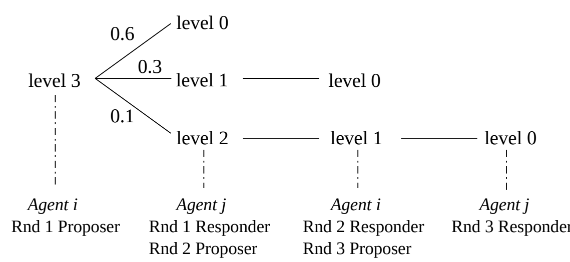

Figure 2: Level 3 belief in DLK for a 3-round SBG. Notice that agent j is a responder in round 1 and a proposer in round 2. A proposer's level l rule is a best response to a responder's level $l-1$ rule.

A *rational* responder at level 0 in the final round accepts an offer, σ^t, that maximizes her utility function, $U(\sigma^t)$. A responder at any higher level in the final round acts similarly. [3] A rational responder at level 0 or a higher level in an intermediate round computes the optimal proportion of the pie that she would offer in the next round (as a proposer), σ_{opt}^{t+1}, and accepts an offer whose utility, $U(\sigma^t)$, is greater than the utility of this portion, $U(1 - \sigma_{opt}^{t+1})$. This optimal portion computed by the responder is also used by a rational proposer at levels 1 or greater up to k in an intermediate round in order to decide her offer. The proposer offers a portion of the pie, σ^t, such that the proposer's share, $1 - \sigma^t$, maximizes the proposer's utility and the utility of the responder's share, σ^t, is greater than $U(1 - \sigma_{opt}^{t+1})$, for the responder. A level 0 proposer keeps the portion of the pie, $1 - \sigma^t$, of the current round that maximizes her utility, $U(1 - \sigma^t)$.

In DLK, the proposer at level k in the opening round keeps the portion, $1 - \sigma^1$, that both maximizes her utility and the utility of the responder's share σ^1 is greater than the *expectation* over the utilities of the responder's rules at levels 0 to $k-1$ in the next round. These utilities would differ because the optimal proportion is predicated on the level. The expectation uses the Poisson distribution over the lower levels. Additionally, analogous to RQR, we utilize a quantal response to model potential errors in the agents' behaviors. Given an offer in the current round, σ^t, a responder may accept this offer, denoted by acc, or reject it (rej). The utility of accepting the offer becomes $U(\sigma^t)$, while the utility of rejecting it in an intermediate round is the utility of next round's optimal portion that the responder would keep (as the proposer), $U(1 - \sigma_{opt}^{t+1,l})$. The responder's probability of

[3] Consequently, rules of levels higher than $n-1$ attributed to the other agent produce the same behavior as the rule of level $n-1$. Therefore, k need not be greater than n.

accepting an offer in round t, σ^t, at $l \geq 1$ is given as:

$$P_{rsp}^{\sigma^t,l}(acc;\lambda) = \frac{e^{\lambda \cdot U(\sigma^t)}}{e^{\lambda \cdot U(1-\sigma_{opt}^{t+1,l})} + e^{\lambda \cdot U(\sigma^t)}} \quad (7)$$

where λ parameterizes the degree to which the agent is playing rationally. The responder's probability of rejecting the offer, $P_{rsp}^{\sigma^t,l}(rej)$, is one minus the probability in Eq. 7.

In order to compute the probability of making an offer, σ^t, for a proposer at level, $1 \leq l \leq k-1$, we first need to decide whether the offer would be accepted or not. As we mentioned previously, an offer is accepted if the utility to the responder of her share, $U(\sigma^t)$, is greater than the utility of the optimal portion that the responder would keep in the next round, $U(1-\sigma_{opt}^{t+1})$. In this case, the utility of the offer to the proposer is simply $U(1-\sigma^t)$. On the other hand, if the offer is rejected, its utility to the proposer is the utility of the optimal portion that the responder offers in the next round, $U(\sigma_{opt}^{t+1})$. The probability of proposing an offer is:

$$P_{prp}(\sigma^t) = \frac{e^{\lambda \cdot \hat{U}(\sigma^t)}}{\sum_{\sigma'} e^{\lambda \cdot \hat{U}(\sigma')}} \quad (8)$$

where $\hat{U}(\sigma^t)$ is a piecewise function: it is $U(1-\sigma^t)$ if the offer is accepted, or $U(\sigma_{opt}^{t+1})$, otherwise.

At level k, the computation of $U(1-\sigma_{opt}^{t+1})$ is complicated due to agent i's Poisson distribution, b_i, over the lower level decision rules. Specifically, the decision rules of different levels lead to differing optimal portions as computed by the responder in the next round as the proposer. Consequently:

$$U(1-\sigma_{opt}^{t+1}) = \sum_{l=0}^{k-1} b_i(\theta_{j,l};\tau) U(1-\sigma_{opt}^{t+1,l})$$

Given the above modification, Eq. 8 now applies to a proposer at level k, where $U(\sigma_{opt}^{t+1})$ is computed analogously.

In the DLK model, agent i updates its distribution, b_i, between games on observing the action of the responder to an offer, σ^t, when it plays the SBG repeatedly. We may update this distribution using a simple Bayesian belief update given the observation of acc or rej,

$$b_i'(\theta_{j,l}|acc) \propto P_{rsp}^{\sigma^t,l}(acc) b_i(\theta_{j,l};\tau) \quad (9)$$

where $P_{rsp}^{\sigma^t,l}(acc)$ is as defined in Eq. 7. Analogously, the belief update on observing a rejection involves, $P_{rsp}^{\sigma^t,l}(rej)$. Note that the posterior may not remain a Poisson.

The level-k model described above differs from Ho and Su's dynamic model (2013) for SBG in two important ways: (i) While in our model, the proposer and responder at level 0 acts to maximize its own utility only, Ho and Su let the level 0 proposer and responder select a random threshold, which serves as the demand and the acceptance threshold. (ii) The initial belief over the rules of the responder assumes a Poisson distribution in our model. On the other hand, Ho and Su's construction places a probability 1 on a particular rule for the responder, which may vary for different participants.

We may integrate the social utility models defined in the previous section by substituting the normative utility function NORM with ERC and FSC.

Cognitive Hierarchy Model With Belief Update

Another behavioral model that could systematically capture violations of backward induction is the cognitive hierarchy model (Camerer, 2003; Camerer et al., 2004), which while sharing aspects with the level-k model also differs in key ways. In the past, this model has predominantly been utilized to model behavioral data on single-shot normal form games (Camerer, 2003; Wright and Leyton-Brown, 2010). We extend this model significantly to the context of repeated SBGs by including a belief update, and denote it as DCH.

While agent i at level k models the other agent, j, with decision rules at all levels from 0 up to $k-1$ where $k = n$ for an n-round SBG, DCH differs from DLK in that the proposer at any level, l, ascribes decision rules of levels ranging from 0 to $l-1$ to the other. An agent maintains a Poisson belief distribution over the rules of all lower levels. We illustrate the structure of DCH for a 3-round SBG in Fig. 3.

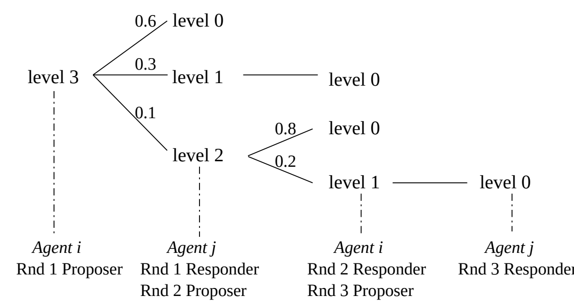

Figure 3: An agent at each level ascribes decision rules of all lower levels to the other, and maintains a belief distribution over them.

This difference from DLK in the structure of the model leads to changes in how the responder computes the optimal portion to offer in the next round as the proposer in that round. A responder at level 1 in any round t computes σ_{opt}^{t+1} similarly as in DLK since there is only one level below. However, a responder at level 2 computes the optimal portion that maximizes her expected utility where the expectation is due to the belief over her opponent in round $t+1$ being at level 1 or 0. We may apply this reasoning to a responder at any level, $1 < l \leq k-1$. Formally, the optimal portion to offer in the next round, $t+1$, for a responder, say j, of level, $L \geq 1$, (as the proposer) is, $\sigma_{opt}^{t+1,L} = \underset{\sigma}{argmax}\, U(1-\sigma)$, such that:

$$U(\sigma) \geq \sum_{l=0}^{L-1} b_j(\theta_{i,l};\tau) U(\sigma_{opt}^{t+1,l}) \quad (10)$$

where $b_j(\theta_{i,l};\tau)$ is the belief that j has as the proposer in the next round over the decision rules of the responder in that round, i, and $U(\sigma_{opt}^{t+1,l})$ is the utility of i's keep in round $t+1$ at level l.

Given the above sophistication in computing the optimal proportion that a responder at any level greater than 1 would offer in the next round, the probability distributions for the responder's and proposer's actions are computed identically to Eqs. 7 and 8, respectively. Furthermore, the proposer's belief over the decision rules of different levels ascribed to the responder is also updated similarly to Eq. 9, with one change from DLK: beliefs of the proposer at all levels from 2 to k are updated. Finally, we may integrate the social utility models as in DLK.

Performance Analysis

A second contribution of this paper is a comparative analysis of the different potential models on the SBG population data. We compare the previously prominent choice model – RQR in combination with the social utility models, NORM, ERC and FSC – with the new level-based choice models, DLK and DCH. The normative utility functions in the latter models may be substituted with those emphasizing fairness such as ERC and FSC. Because the level-based models are representative of a different behavioral explanation, we seek to answer the following important questions:

- Are there types of SBG where the recent hypothesis of limited backward induction offers a better explanation of the observed behavior, and, if so, why?
- Is a combination of being fair and limited backward induction influencing behavior in SBG thereby offering an improved explanation of SBG data in comparison to each individually?

We note that these are new questions whose answers could potentially reveal new insights into bargaining behavior.

Our methodology is to utilize one of the data sets in Table 1 for learning the parameters of the different models by minimizing the negative log likelihood of the models. We refer to the latter as the *fit* of the model and a greater negative log likelihood indicates a better fit. This is followed by an out-of-sample evaluation of the performance of the models with robustness checks. We evaluate the model performances in multiple ways: comparative fit based on the log likelihood, mean error in the predicted opening offers and a visual analysis of the distributions of opening offers. We select the $NSS^{5,4}$ data set for learning parameters in which each game is played for 5 rounds and a subject repeats it though not with the same opponent.

Learned Parameters

The choice model, RQR, is parameterized by the quantal response parameter, λ, and the update parameter, λ', that represents learning across the repeated games. Choice models, DLK and DCH, are both parameterized by the quantal response parameter, λ, and the Poisson distribution parameter, τ. Additionally, the social utility models, NORM, ERC and FSC, use zero, one and two parameters, respectively.

We learn the parameters of the different choice models in combination with the varying utility models as maximum likelihood point estimates on the $NSS^{5,4}$ data set. In Table 2, we report their values and the corresponding negative log likelihoods. For each choice model, we emphasize the highest log likelihood on using the different utility models, in bold. The overall best fit is also underlined.

Choice model	**Parameter**	**Social utility models**		
		NORM	ERC	FSC
RQR	λ	0.396	0.374	0.329
	λ'	-0.005	-0.016	-0.009
	b		12.293	
	α			1.813
	β			-1.631
log likelihood		-309.510	**-288.873**	-289.041
DLK	λ	0.601	0.931	0.697
	τ	0.313	17.070	12.269
	b		7.671	
	α			0.411
	β			-0.787
log likelihood		-263.358	-238.384	**-231.027**
DCH	λ	0.579	0.943	0.676
	τ	0.301	20.256	11.614
	b		7.170	
	α			0.391
	β			-0.880
log likelihood		-263.492	-238.426	**-230.713**

Table 2: Learned parameters for all combinations of the choice and utility models, obtained as maximum likelihood point estimates. A space indicates that the parameter is not applicable. Highest likelihood for each model is indicated in bold.

Observe that the level-based choice models, DLK and DCH, fit the data set better in comparison to RQR. Integrating the social utility models with the level-based choice models further improves the fit, with both DLK and DCH exhibiting negative log likelihood that is substantially greater than that of RQR. This indicates that in the longer 5-round SBG, considerations of limited backward induction positively impact the modeling of the data although fairness continues to play a role in the behavior as well. Notice that the value of τ increases from approximately 0.3 in the normative case to greater than 12 on integrating the social utility models. This indicates that the probability of modeling the other agent as a level 0 player drops from 0.72 to about 0. The former is consistent with the observation that opening offers in the $NSS^{5,4}$ data being between \$5 and \$6 about 73% of the times, where \$5.1 is the second round's pie size. Interestingly, social considerations provide an alternative explanation for this offer – mild levels of fairness – thereby allowing the weight on level 0 to reduce.

Model Predictions

We perform out-of-sample predictions using the parameters learned previously as shown in Table 2 on the remaining data sets of Table 1. These data include observations of subjects playing a 5-round SBG with no repetition and multiple 3-round SBGs with repetition. Out-of-sample testing differs from n-fold cross validation because the test data sets may represent distributions that are different from that of the training set. This could be due to subjects being drawn from different pools and games with much different parameters. Subsequently, such testing provides a robust performance evaluation and comparison.

The out-of-sample log likelihoods of the different choice models in combination with the social utility models are

Data Set	Observations	Random	RQR			DLK			DCH		
			NORM	ERC	FSC	NORM	ERC	FSC	NORM	ERC	FSC
NSS^{5,1}	0.343	0.5±0.013	0.423±0.012	0.404±0.010	0.374±0.011	0.494±0.010	0.435±0.006	0.452±0.008	0.495±0.005	0.435±0.003	0.450±0.003
NSS^{5,4}	0.359	0.5±0.025	0.348±0.012	0.420±0.008	0.392±0.008	0.419±0.006	0.397±0.003	0.387±0.004	0.423±0.010	0.396±0.006	0.386±0.008
OR$_1^{3,10}$	0.433	0.5±0.029	0.349±0.012	0.435±0.008	0.433±0.007	0.411±0.006	0.381±0.003	0.332±0.003	0.413±0.006	0.380±0.003	0.312±0.003
OR$_2^{3,10}$	0.450	0.5±0.029	0.309±0.012	0.419±0.010	0.404±0.009	0.608±0.006	0.458±0.003	0.495±0.004	0.610±0.007	0.458±0.003	0.485±0.004
OR$_3^{3,10}$	0.451	0.5±0.029	0.371±0.013	0.434±0.011	0.424±0.010	0.558±0.006	0.410±0.003	0.322±0.003	0.560±0.006	0.410±0.003	0.311±0.003
OR$_4^{3,10}$	0.466	0.5±0.029	0.437±0.013	0.458±0.009	0.459±0.009	0.361±0.006	0.301±0.004	0.192±0.003	0.363±0.006	0.291±0.003	0.192±0.003

Table 4: Mean opening offers (along with standard errors) for the different experiments as observed from the data and obtained from the model predictions. We include the predictions by a random model as well for comparison.

Choice model	**Parameter**	**Social utility models**		
		NORM	ERC	FSC
RQR	*NSS^{5,4}*	*-309.51*	*-288.87**	*-289.04**
	NSS^{5,1}	-241.85	-231.95*	-232.87*
	OR1^{3,10}	-387.84	-308.05	**-302.27**
	OR2^{3,10}	-442.92	**-335.43***	-338.56*
	OR3^{3,10}	-346.37	-303.40*	**-301.43***
	OR4^{3,10}	-313.80	-273.83	**-255.43**
DLK	*NSS^{5,4}*	*-263.36*	*-238.38*	***-231.03***
	NSS^{5,1}	-219.49*	-211.96*	**-211.54***
	OR1^{3,10}	-390.05*	-373.94*	-583.14
	OR2^{3,10}	-764.19*	-747.11*	613.92
	OR3^{3,10}	-717.59	-555.90*	-618.26*
	OR4^{3,10}	-750.41	-610.03	-1101.76
DCH	*NSS^{5,4}*	*-263.49*	*-238.43*	***-230.71***
	NSS^{5,1}	-220.24*	-212.48*	**-212.24***
	OR1^{3,10}	-383.32*	-386.05*	-667.83
	OR2^{3,10}	-746.28*	-747.79*	-587.11
	OR3^{3,10}	-699.48*	-591.50*	-711.28
	OR4^{3,10}	-730.91	-663.35	-1173.17

Table 3: Log likelihoods for the different choice models integrated with the three utility models. We include the fit on the training data, NSS^{5,4}, shown previously for completeness. Highest likelihood for each data set among the models is indicated in bold. * annotates likelihoods in a row whose difference is not significant.

shown in Table 3. They show distinct trends in the comparative performance of the different models: (a) DLK and DCH provide the best fit for NSS^{5,1} performing significantly better than RQR (Student's paired t-test, p-value = 0.002). Among the different utility models, ERC and FSC do not significantly improve on normative for both DLK and DCH on this data set. However, the two utility models do not cause a significant difference in the performance of DLK and DCH for NSS^{5,1}. (b) RQR in combination with either ERC or FSC provides the best fit for all the OR data sets. Furthermore, the performance of the level-based choice models degrades significantly for some of the OR data sets.

Based on the learned parameters, we may predict the mean opening offer by the proposer averaged across all participants and games. We report the models predictions in Table 4. Furthermore, taking the NSS^{5,1} data set as an example, we show in Fig. 4, the distributions of the opening offer proportions as predicted by the choice and utility models, in comparison with the observed distribution from the data.

The observed initial offers from the data in Table 4 are close to the size of the pie in the next round. From Table 4, we observe that surprisingly the mean offer proportion predicted by RQR is closer to the observed offer in comparison to the level-based choice models though the latter demonstrated better log likelihoods previously. However, we point out the higher standard error of RQR's offer predictions in comparison to the standard errors of the level-based choice models for the data set. Indeed, the distributions in Fig. 4 show that the two level-based choice models better simulate the distribution of opening offers compared to RQR, with minimal difference between the two. Additionally, the prediction of the opening offer by the random model that picks an offer at random is close to the observed data for the OR data sets, but removed from the opening offer for the NSS data sets. Because the pie size in the next round for the OR data sets is close to 50% of the first round's pie size, the close predictions of the mean random offer are coincidental.

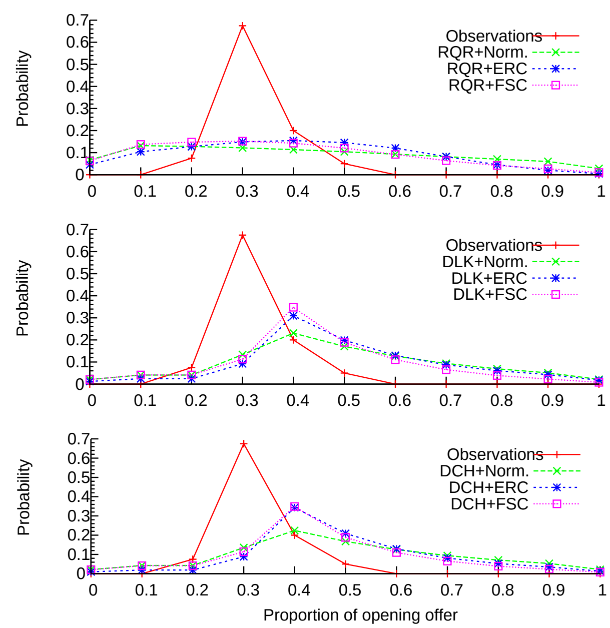

Figure 4: Distribution of opening offer proportions for test data set, NSS^5, 1, by the different models and from observed data. Notice that RQR's distribution is diffused and compares poorly with the observed distribution, whose support ranges from 0.1 to 0.6.

Conclusion

Several experiments on SBG have generated a large amount of behavioral data that is available for modeling. While previous models have focused on the relevance of fairness considerations while playing the games, we additionally explored whether considerations of limited backward induction could provide an improved model of the data. In this respect, we utilized two extended models that capture violations of backward induction for modeling SBG. While these models have traditionally been applied to single-shot normal-form games, their application to SBG is new. Our comprehensive empirical analysis with 9 different models provides evidence that in longer SBG with more rounds, limited backward induction plays a crucial role in how humans engage in SBG. Between the two models that capture violations of backward induction, we did not observe a significant difference indicating that the simpler of the two models, DLK, is sufficient. The modeling is further improved when combined with social utility indicating that fairness remains a consideration in the offer. However, for shorter rounds fairness of the offer remains the key consideration.

As future work, we are interested in exploring the effect of limited backward induction in other sequential problems. One such game is sequential negotiation as realized in the Colored Trail framework.

References

Binmore, K., Shaked, A., and Sutton, J. (1985). Testing noncooperative bargaining theory: A preliminary study. *American Economic Review*, 75(5):1178–80.

Bolton, G. (1991). A comparative model of bargaining: theory and evidence. Levine's Working Paper Archive 263, David K. Levine.

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

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction (Roundtable Series in Behaviorial Economics)*. Princeton University Press.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *Quarterly Journal of Economics*.

De Bruyn, A. and Bolton, G. E. (2008). Estimating the influence of fairness on bargaining behavior. *Management Science*, 54(10):1774–1791.

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

Gal, Y. and Pfeffer, A. (2007). Modeling reciprocal behavior in human bilateral negotiation. In *Twenty-Second Conference on Artificial Intelligence (AAAI)*, pages 815–820.

Güth, W. and Tietz, R. (1986). *Ultimatum bargaining for a shrinking cake: an experimental analysis*. Frankfurter Arbeiten zur experimentellen Wirtschaftsforschung. Professur für Volkswirtschaftslehre, insbesondere Verhaltensforschung, Johann-Wolfgang-Goethe-Univ.

Güth, W. and van Damme, E. (1998). Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study. *Journal of Mathematical Psychology*, 42(2/3):227–47.

Hedden, T. and Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, 85:1–36.

Ho, T. and Su, X. (2013). A dynamic level-k model in sequential games. *Management Science*, page in press.

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

McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica*, 60(4):803–836.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

Neelin, J., Sonnenschein, H., and Spiegel, M. (1988). A further test of noncooperative bargaining theory: Comment. *American Economic Review*, 78(4):824–36.

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

Ray, D., King-Casas, B., Montague, P. R., and Dayan, P. (2008). Bayesian model of behavior in economic games. In *Neural Information Processing Systems (NIPS)*, pages 1345–1352.

Rosenthal, R. (1981). Games of perfect information, predatory pricing and the chain store paradox. *Journal of Economic Theory*, 25:92–100.

Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and market behavior in jerusalem, ljubljana, pittsburgh, and tokyo: An experimental study. *American Economic Review*, 81(5):1068–95.

Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109.

Stahl, I. (1972). *Bargaining Theory*. Stockholm: Stockholm School of Eonomics.

Wright, J. R. and Leyton-Brown, K. (2010). Beyond equilibrium: Predicting human behavior in normal-form games. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, pages 901–907.