EURO Journal on Decision Processes 11 (2023) 100028

Contents lists available at ScienceDirect

EURO Journal on Decision Processes

journal homepage: www.elsevier.com/locate/ejdp

Confidence in bargaining processes and outcomes: Empirical tests of a conceptual model

Rudolf Vetschera[a,1,*], Luis C. Dias[b,2]

[a] Dept. of Business Decisions and Analytics University of Vienna, Vienna, Austria
[b] University of Coimbra CeBER, Faculty of Economics Coimbra, Portugal

ARTICLE INFO

Keywords:
Negotiation
Zeuthen-Hicks bargaining
Asymmetric Nash bargaining solution
Experiment

ABSTRACT

The relationship between negotiation processes and outcomes is a challenging problem for theoretical and empirical analyses. In this paper, we study whether a dynamic bargaining model that incorporates a notion of negotiator confidence in the process and that predicts the asymmetric Nash bargaining solution as the outcome is compatible with observations in negotiation experiments. This requires establishing how the compatibility between the model and the actual bargaining process can be assessed, without knowing a key parameter in the model. We find that the model is largely compatible with the observed bargaining process, but that actual agreements tend to be more balanced than the solution predicted by the model. We also find a close relationship between the parameter representing negotiator confidence in the model and the negotiator's (independently ascertained) aspiration levels, thus providing additional evidence for the model's external validity.

1. Introduction

Bargaining is a complex, dynamic process that leads to a clearly defined outcome, either an agreement that resolves all the issues that formed the subject of the negotiation, or the situation that the parties did not reach an agreement. Clearly, in real life the outcome depends on the bargaining process. However, the relationship between bargaining process and outcomes is a challenging topic both for theoretical and empirical research.

Game theory literature distinguishes between axiomatic and strategic bargaining models. Axiomatic bargaining models are frequently associated with concepts of cooperative game theory (Kibris, 2021). Starting with the seminal work of Nash (1950), these models aim to characterize outcomes of bargaining processes using normatively appealing axioms. In contrast, strategic bargaining models have their roots in non-cooperative game theory (Chatterjee, 2021) and try to model the behavior of negotiators during the bargaining process. Well-known examples of strategic bargaining models are the Zeuthen-Hicks model (Bishop, 1964; Harsanyi, 1956) or Rubinstein's bargaining model (Rubinstein, 1982). Axiomatic and strategic models have largely developed separately, even though there are some connections between them, e.g., it can be shown that the Zeuthen-Hicks model leads to the Nash bargaining solution. Consequently, the fact that axiomatic bargaining models describe a solution but do not prescribe steps to reach that solution has been seen as a major limitation of these models when it comes to supporting negotiators (Munier and Rulliére, 1993). Even very recently, Engler and Page (2022) noted a lack of conceptual knowledge about the bargaining process.

Empirical negotiation research has studied negotiation processes from different perspectives (Vetschera, 2013). Already in the 1960s, researchers begun to study the relationship between bargaining processes and outcomes (Benton et al., 1972; Chertkoff and Conley, 1967; Hinton et al., 1974). A problem encountered in this type of empirical research is that negotiations consist of multiple steps (offers, concessions etc.) embedded in a dynamic process, but reach one single outcome. Early research thus mainly considered aggregate measures such as total or average concessions. Later on, other constructs like types of bargaining steps (Filzmoser and Vetschera, 2008; Hindriks et al., 2007; Kersten et al., 2013) were used. However, these steps also need to be aggregated across the negotiation process, e.g., by considering the fraction of times a specific type of steps such as a concession step occurs among all bargaining steps. This concept still does not take into account the dynamics of the process. An overall representation of the entire bargaining process is a concession curve, which describes the position taken

* Corresponding author.
E-mail addresses: rudolf.vetschera@univie.ac.at (R. Vetschera), lmcdias@fe.uc.pt (L.C. Dias).
[1] ORCID: 0000-0003-2809-8989.
[2] ORCID: 0000-0002-1127-1071.

https://doi.org/10.1016/j.ejdp.2023.100028
Received 26 September 2022; Received in revised form 28 November 2022; Accepted 16 January 2023

(the utility claimed) by a negotiator over time. However, this is still not a compact representation of the bargaining process that can easily be related to outcomes. Researchers therefore used parametric representations of the concession curve, which allow to distinguish convex or concave curves (Carbonneau and Vahidov, 2014) to classify different process patterns. Other approaches used machine learning to simultaneously consider many structural properties of concession curves (Nastase, 2006). Although these approaches provide an empirical description of concessions in a bargaining process, they lack a theoretical basis in analytical bargaining models that would allow to connect the process to characteristics of the negotiators.

Theoretical bargaining models typically represent bargainers via their utility functions. Several empirical studies have attempted to connect properties of the utility functions to concession making and bargaining outcomes (Mumpower, 1991; Northcraft et al., 1995; 1998; Vetschera, 2007). These studies uncovered some plausible relationships, but bargaining behavior might depend on other factors in addition to a negotiator's preferences. Empirical bargaining theory has identified different bargaining strategies such as conceding or compromising behavior (Kilmann and Thomas, 1977), or negotiator toughness and kindness (Engler and Page, 2022), which reflect properties of the negotiator beyond their preferences for outcomes. These concepts cannot be represented easily in game theoretic bargaining models.

This paper contributes to reconcile these different streams of literature. It develops an approach to study the connection between bargaining process and bargaining outcomes based on a theoretical model that allows to characterize different bargaining strategies. The theoretical model we use for this purpose is a two-party bargaining model with the risk of an unfavorable outcome in case the negotiation breaks down. This model was analyzed in detail by Dias and Vetschera (2022) and can be regarded as a generalization of the Zeuthen-Hicks bargaining model. It takes into account not only the bargainers' preferences, but also a parameter that represents their confidence level in relation to the final outcome of their bargaining process, as it is known that beliefs about outcomes play an important role in bargaining processes (Dickinson, 2009). For a given confidence level, this model allows making predictions about the bargaining process as well as its outcomes, as we will explain in section two.

We test empirically whether this model describes actual bargaining processes, based on data from a negotiation experiment conducted at four universities. These data include information about multiple bargaining processes (the offers exchanged), information about the outcomes (the final agreement), and the self-reported aspiration and reservation levels (i.e., the utility the parties aim to achieve and the minimum utility they will accept). The parameter that represents the confidence level of each bargainer, however, is not known, which makes it more challenging to assess the compatibility between the model and the experimental data.

For a model which contains parameters representing individual characteristics, such an assessment can be performed in different ways. We do not expect that this model (or any other model) can precisely represent all possible bargaining processes. Therefore, we can determine in how many cases the actual process is compatible with the model and how frequently it is not. An observed process, as well as an observed outcome, imply certain values of the model's parameters. Since the model predicts both process and outcome of a negotiation, we can test whether the same parameter values are compatible with the observed process and outcome. Thus, our empirical research aims to answer the following main research questions:

RQ1 To what extent are observed negotiation *processes* compatible with the model?

RQ2 To what extent are observed negotiation *outcomes* (agreements) compatible with the predictions of the model and how do actual and predicted agreements differ?

RQ3 Does the model predict that the observed processes lead to the observed outcomes ?

Finally, the confidence parameter in the model has a clear substantial interpretation in terms of negotiator characteristics. We can therefore also test whether other characteristics of the negotiator match the estimated parameter in a plausible way, thus addressing an additional research question:

RQ4 Do confidence parameters estimated from observed behavior exhibit theoretically plausible relationships to aspiration and reservation levels?

This paper therefore makes a threefold contribution: we introduce a compact representation of negotiator behavior in the bargaining process based on a theoretical bargaining model and having a clear interpretation; we envisage how to assess the compatibility between the conceptual model and experimental data when a key model parameter cannot be observed; and we apply this approach in an empirical study to obtain new insights into the relationship between bargaining process and outcome.

The remainder of this paper is structured as follows. In section two, we provide an overview of the bargaining model. In section three, we present the empirical evaluation strategies and the experimental setting. Section four presents the empirical results and section five summarizes and concludes the paper.

2. Overview of a model with confidence parameters

The present work considers a model studied by Dias and Vetschera (2022), which includes parameters to explicitly model the confidence of each party in relation to what it expects to obtain as a negotiated agreement. This model considers two parties bargaining over a single issue, by successively exchanging offers. As a source of attrition, negotiations might break down with some probability at each round. For instance, a breakdown might result from the intervention of a higher authority (an impatient principal who can dismiss its agent, or some authority with power of arbitration who might punish the parties for not being able to show progress). We adopt this model as a behavioral model for parties that actually exchange offers, as occurs in real-life bargaining processes, rather than a normative model that parties would use to instantaneously reach an equilibrium agreement following game-theoretical considerations.

Without loss of generality, we refer to two bargaining parties as the "buyer" and the "seller", and we consider the single issue they are bargaining over to be the "price". We adopt the following notation (Dias and Vetschera, 2022):

- s and b represent the offers of the seller and the buyer, respectively. When used as subscripts, these letters also refer to the seller and the buyer.
- $u_s(x)$ and $u_b(x)$ represent the utility of an offer x for the seller and for the buyer.
- d represents the outcome if the parties fail to reach an agreement.

The two parties have completely opposing objectives: if the seller prefers offer x to option y, then the reverse occurs for the buyer, and vice-versa. Thus, the "price" can be a package involving multiple attributes (as in the empirical study of the present paper). For simplicity we refer here to the price as a single issue, which the seller seeks to maximize and the buyer seeks to minimize, i.e., $u'_s(x) > 0$ and $u'_b(x) < 0$. We consider utilities are normalized, i.e, $u_s(x), u_b(x) \in [0, 1]$. Outcome d is unattractive for both parties yielding $u_s(d) = 0$ and $u_b(d) = 0$.

From the moment when the initial positions of the parties are known until an agreement is reached, there are two offers on the table, s and b. We assume that these offers are better than d for both parties, so that they are interested in bargaining with each other. Naturally, $u_s(s) > u_s(b)$ and $u_b(s) < u_b(b)$. In each bargaining step, a party can insist on its own

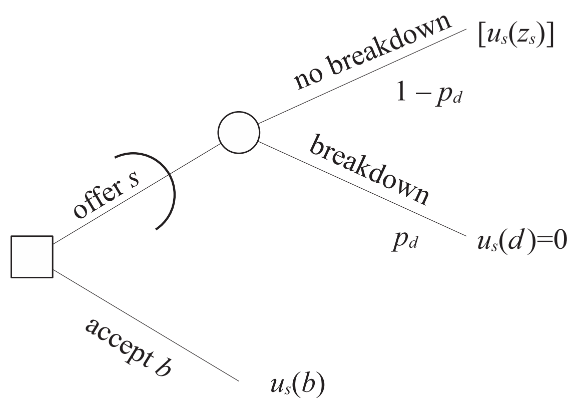

Fig. 1. Framework for one negotiation round (Seller's decision).

offer, accept the other party's offer, or make a concession by improving its offer for the other party.

Let us consider at time t it is the seller's turn to respond to an offer (for the buyer the reasoning is similar, with the necessary adjustments). The decision faced by the seller is outlined in Fig. 1. The lower branch in this decision tree corresponds to accepting the buyer's offer b, thus ending the bargaining process getting $u_s(b)$. The upper branch represents all the different offers s the seller could make, including insisting on its previous offer. If the seller does not accept the buyer's offer and insists on its previous offer causing a stalemate, then the negotiation might break down with probability p_d, and the final outcome is therefore uncertain.

This model assumes the parties do not know what will the final outcome be if they do not accept the offer from the other party. Each of them will seek to maximize its expected utility based on its subjective expectations. Let $u_s(z_s)$ represent the subjective expected utility that the seller foresees (at time t, omitted as a subscript) to obtain as the final outcome if the negotiation does not break down in the ensuing rounds, given offers s (an offer the seller can adjust) and b. This subjective expected utility encompasses not only the current offers, but also the expectation about future offers and the possibility of breakdown at a later round. While other models also take into account expectations, they mostly deal with expectations about the disagreement outcome (Dickinson, 2009). In contrast, our model focuses on expectations about bargaining outcomes, which can be strongly affected by negotiator characteristics such as the overconfidence bias (Caputo, 2013). Therefore, we expect that a parameter related to each party's confidence is able to capture an important behavioral influence on the negotiation process.

The outcome estimated by the seller for the case the negotiation does not break down will reasonably lie in-between the two offers currently on the table, i.e. the seller will not expect to get more than the offer it proposes, or less than the offer the buyer proposes (otherwise it would be better to accept that offer). Therefore, for some value $\gamma_s \in (0, 1]$ representing the confidence of the seller, we can write the seller's certainty equivalent of the final outcome as:

$$z_s = u_s^{-1}(\gamma_s u_s(s) + (1 - \gamma_s) u_s(b)) \tag{1}$$

(And, respectively, $z_b = u_b^{-1}(\gamma_b u_b(b) + (1 - \gamma_b) u_b(s))$, with $\gamma_b \in (0, 1]$ representing the confidence of the buyer).

Since the seller seeks to maximize expected utility, its decision will depend on the interplay between the estimated outcome and the estimated breakdown probability:

$$\text{offer } s \succ \text{accept } b \Leftrightarrow (1 - p_d) u_s(z_s) > u_s(b)$$
$$\Leftrightarrow p_d < \frac{u_s(z_s) - u_s(b)}{u_s(z_s)} \tag{2}$$

And since the buyer also seeks to maximize expected utility, from its perspective the preference is:

$$\text{offer } b \succ \text{accept } s \Leftrightarrow (1 - p_d) u_b(z_b) > u_b(s)$$
$$\Leftrightarrow p_d < \frac{u_b(z_b) - u_b(s)}{u_b(z_b)} \tag{3}$$

Therefore, the parties will hold to their current offers s and b if their estimate of the breakdown risk p_d is very low. But a stalemate situation increases this risk. Thus, at some point in time this risk becomes sufficiently large to warrant a concession. The model predicts that the seller will be the one conceding if $\frac{u_s(z_s)-u_s(b)}{u_s(z_s)} < \frac{u_b(z_b)-u_b(s)}{u_b(z_b)}$, i.e.,

$$u_s(z_s) u_b(s) < u_s(b) u_b(z_b), \tag{4}$$

and the buyer will be the one conceding if this inequality is reversed (and a simultaneous concession would occur in the event of an equality).

A concession does not imply accepting the other party's offer. It can be sufficiently small to make the uncertain branch of continuing bargaining (the top branch in Fig. 1) preferred to the lower branch. One can note that each party is able to recognize the need of a concession without information about the utility function and expectations of the other party, but such information, at least concerning a vicinity of offers on the table, would be required to compute the minimum concession needed to reverse the inequality (4) (in the case of the seller) or to enforce this inequality (in the case of the buyer).

This process of two parties successively enforcing and then reversing inequality (4) until eventually reaching an agreement (if no breakdown takes place) is reminiscent of Zeuthen's principle, which according to Harsanyi (1977) is the only rule consistent with subjective probabilities that rational players can entertain about each other's behavior. This model will coincide with a standard Zeuthen-Hicks model (Bishop, 1964; Harsanyi, 1956) if $\gamma_s = \gamma_b = 1$, i.e., if the two parties have the maximum level of confidence. Thus, the main difference of the framework used in this work (Fig. 1) is the introduction of the parameters γ_s and γ_b acknowledging that parties can realistically expect that the final outcome of the process will give them less utility than what they are offering at time t. It also explicitly acknowledges each party has many options (all potential concessions) besides insisting on its own offer or accepting the other party's offer.

The process of successively reversing inequalities in a standard Zeuthen-Hicks model leads the parties to obtain the Nash Bargaining Solution (NBS) (Nash, 1950), although this is guaranteed only if both parties have concave utility functions (Dias and Vetschera, 2019). In turn, for the model considered here, Dias and Vetschera (2022) show that under the assumption that the two parties are risk neutral or risk averse (i.e., they have concave utility functions), the model predicts the parties will obtain the nonsymmetric (or asymmetric) NBS, i.e.,

$$x^* = \arg\max_x u_s(x)^{\gamma_s} u_b(x)^{\gamma_b} \tag{5}$$

This corresponds to the solution obtained in Nash's framework when the symmetry axiom is dropped (see, e.g., Muthoo, 1999; Roth, 1979), considering here confidence as the source of asymmetry.

The process and the solution depends only on the ratio of these parameters, $r = \gamma_s / \gamma_b$, rather than their absolute values (for details, see Dias and Vetschera, 2022). Indeed, Eq. (4) can be transformed into

$$r(u_s(s) - u_s(b)) u_b(s) < (u_b(b) - u_b(s)) u_s(b) \tag{6}$$

and Eq. (5) is equivalent to

$$x^* = \arg\max_x u_s(x)^r u_b(x). \tag{7}$$

3. Materials and methods

3.1. Compatibility assessment

The model that we have reviewed in the previous section makes predictions about the offers that each party makes in each step of the negotiations and the outcome of the negotiation, which depend on the values of the confidence parameters of both parties. As a working hypothesis, we consider these confidence parameters do not change during the negotiation process. For an empirical test of the model, in a situation where we cannot observe the confidence factors directly, we propose a two-step approach. First, we estimate the ratio r of the two confidence parameters (which, according to the analytical results, determines both process and outcomes) from observed offers and outcomes. In a second step, we can then test whether the values of r which we have estimated exhibit a consistent pattern.

Next, we explain how one can estimate confidence parameters from negotiation processes and negotiation outcomes, and then how to assess the agreement between these estimates, considering a general model as well as the model outlined in the previous section.

The bargaining process consists of offers made by both sides in an alternating way. For simplicity, we index the offers sequentially, so that an offer made by one party at time $t-1$ is followed by an offer made by the other party at t, for $t \in \{1, \dots, T\}$. This notation is used for convenience, it does not imply that the time difference between two subsequent offers is constant.

A general model characterizing the successive the offers of the parties as a function of their characteristics r (here considered as a single parameter, but the reasoning could be extended to multiple parameters) can be stated as follows:

- If s_t is an offer of the seller at time t, then it satisfies inequality $f_s(b_{t-1}, s_t, r) > 0$, where f_s is the function that defines the process model and b_{t-1} is the last offer from the buyer.
- If b_t is an offer of the buyer at time t, then it satisfies inequality $f_b(s_{t-1}, b_t, r) > 0$, where f_b is the function that defines the process model and s_{t-1} is the last offer from the seller.

A general model characterizing the outcome of the bargaining process as a function the the bargainers' characteristics r can state that the agreement x^* is such that

$$x^* = arg\max\{g(x,r) : x \in X\} \tag{8}$$

where g is the function characterizing the outcome model and X is the set of potential agreements.

Moreover, a general model characterizing the outcome can also characterize the sequence of offers (i.e., the process) as a succession that converges to the maximum of g:

$$\dots < g(s_{t-3}, r) < g(b_{t-2}, r) < g(s_{t-1}, r) < g(b_t, r) < \dots \le g(x^*, r) \tag{9}$$

3.1.1. Estimating the confidence ratio from offers

If the characteristics of the bargainers r are unknown, but data about the process (successive offers) is available, then it is possible to infer a subset of the unknown parameter space for each offer. One possibility is based on the model characterization given by f_s and f_b:

$$r \in \Omega_t^f = \begin{cases} \{r : f_s(b_{t-1}, s_t, r) > 0\}, & \text{if the seller made an offer at } t \\ \{r : f_b(s_{t-1}, b_t, r) > 0\}, & \text{if the buyer made an offer at } t \end{cases} \tag{10}$$

Applying this reasoning to the model presented in Section 2, if an offer b_{t-1} made by the buyer at time $t-1$ is on the table, the seller will make an offer s_t at time t so that the condition

$$r(u_s(s_t) - u_s(b_{t-1}))u_b(s_t) > (u_b(b_{t-1}) - u_b(s_t))u_s(b_{t-1}) \tag{11}$$

is fulfilled. This provides a lower bound for r as

$$r \in \Omega_t^f \Leftrightarrow r > lo_t = \frac{(u_b(b_{t-1}) - u_b(s_t))u_s(b_{t-1})}{(u_s(s_t) - u_s(b_{t-1}))u_b(s_t)} \tag{12}$$

Similarly, if the buyer makes an offer at time t, this provides an upper bound on r as

$$r \in \Omega_t^f \Leftrightarrow r < up_t = \frac{(u_b(b_t) - u_b(s_{t-1}))u_s(b_t)}{(u_s(s_{t-1}) - u_s(b_t))u_b(s_{t-1})} \tag{13}$$

Rational bargainers will make offers which are better for them than their counterpart's current offer, so we can assume that $u_s(s_t) > u_s(b_{t-1})$ and $u_b(b_t) > u_b(s_{t-1})$. As long as these conditions are fulfilled, both bounds will be positive.

Another possibility to infer a subset of the unknown parameter space for each offer, derived from (9), is based on the model characterization given by g:

$$r \in \Omega_t^g = \begin{cases} \{r : g(s_t, r) > g(b_{t-1}, r)\}, & \text{if the seller made an offer at } t \\ \{r : g(b_t, r) > g(s_{t-1}, r)\}, & \text{if the buyer made an offer at } t \end{cases} \tag{14}$$

The model presented in Section 2 predicts that the offers of both parties will move the process toward an agreement that corresponds to the asymmetric Nash bargaining solution. This solution is the agreement x which maximizes the function

$$g(x, r) = u_s(x)^r u_b(x) \tag{15}$$

where x is an offer from any side. If the seller makes an offer which contributes to that maximization, it must fulfill the condition

$$u_s(s_t)^r u_b(s_t) > u_s(b_{t-1})^r u_b(b_{t-1}) \tag{16}$$

from which we can obtain the condition

$$r \in \Omega_t^g \Leftrightarrow r > lo_t = \frac{\ln u_b(b_{t-1}) - \ln u_b(s_t)}{\ln u_s(s_t) - \ln u_s(b_{t-1})} \tag{17}$$

Similar to the previous conditions, we obtain the analogous upper bound from offers made by the buyer:

$$r \in \Omega_t^g \Leftrightarrow r < up_t = \frac{\ln u_b(b_t) - \ln u_b(s_{t-1})}{\ln u_s(s_{t-1}) - \ln u_s(b_t)} \tag{18}$$

As a further benchmark, we also consider the standard Zeuthen-Hicks bargaining model, which posits that negotiations converge to the standard (symmetric) Nash bargaining solution. We therefore also check whether offers increase the unweighted function $g(x, 1) = u_s(x) \cdot u_b(x)$.

During a negotiation, several offers are made by both sides, which result in several upper and lower bounds. If we assume that both parties exactly follow the process described by the model for a given value of r (denoted as r^{true}), then the condition

$$\max_t lo_t \le r^{true} \le \min_t up_t \tag{19}$$

must hold. However, we do not assume this model (or any other one) fits all processes perfectly, i.e., that subjects make "correct" offers according to the true value of r all the time. Therefore, we formulate a weaker condition to assess compatibility. We consider a process to be compatible with the model if the average bounds that are obtained from considering all offers from each side form a non-empty interval:

$$\frac{1}{T}\sum_t lo_t \le \frac{1}{T}\sum_t ub_t \tag{20}$$

This assessment can be based on the Ω_t^f and on the Ω_t^g perspectives.

3.1.2. Estimating bounds from outcomes

We consider a multi-issue negotiation with a finite number of options in each issue. Therefore, the total number of possible packages (agreements) is also finite. The model presented in Section 2 predicts that the agreement of the negotiation is the package which maximizes (15), which characterizes the outcome model as in (8). A package x maximizes that function if for any other package y,

$$u_s(x)^r u_b(x) \ge u_s(y)^r u_b(y) \tag{21}$$

which can be rewritten as

$$r(\ln u_s(x) - \ln u_s(y)) \ge \ln u_b(y) - \ln u_b(x) \tag{22}$$

(22) yields an upper or lower bound on r, depending on the sign of $\ln u_s(x) - \ln u_s(y)$. Condition (21) needs to be checked (and bounds are computed) only for Pareto-optimal packages y (let PO denote this set).

If another package z is dominated by y, the condition $u_s(z)^r u_b(z) < u_s(y)^r u_b(y)$ holds and therefore (21) is trivially also fulfilled for z. This also implies that a bound obtained from z will be less tight than a bound obtained from y, so dominated packages can be ignored in calculating bounds on r. Therefore, the subset of the unknown parameter space inferred from the outcome is:

$$r \in \Omega_*^g \Leftrightarrow \{r : r(\ln u_s(x) - \ln u_s(y)) \geq \ln u_b(y) - \ln u_b(x), \forall y \in PO\} \quad (23)$$

The above condition places a challenge, as it is well known from empirical literature that negotiators often fail to reach an efficient agreement (Gettinger et al., 2016; Kersten and Mallory, 1999; Metcalfe, 2000). Therefore, actual agreements will often not be maximizers of (15). This has also been observed for the symmetric Nash bargaining solution, which is sometimes missed because bargainers prefer a more balanced, even if inefficient, agreement (Bone et al., 2014). In those cases, we cannot estimate an interval for r from the outcome of the negotiation. However, we still can address the compatibility between the outcome model and the actual outcome by considering the agreement that would be predicted by the process model, i.e., by considering the value of r that is inferred from the successive negotiation offers.

From Eq. (22), we can assign to each possible weighted Nash solution an interval of r that would make this solution the optimal one. Similarly, we obtain an interval of r from the process by taking either the average or extreme upper and lower bounds estimated from offers. It is therefore possible that several intervals associated with potential asymmetric Nash solutions overlap with the interval estimated from the process. To obtain a unique prediction of the outcome of the negotiation in these cases, we use the maximizer that is most compatible with the range of r obtained from the process. The conditions (21) define an interval of r for which that package maximizes (15). Denote the bounds of that interval for maximizing package m by lo_m and up_m and the bounds obtained from the process, e.g. by averaging across all offers according to (20), by lo_p and up_p. We then select as a surrogate Pareto-optimal outcome the package m which maximizes

$$\frac{\min(up_m, up_p) - \max(lo_m, lo_p)}{up_p - lo_p} \quad (24)$$

i.e. the package whose range covers the largest part of the interval implied by the process.

3.2. Data

For the empirical analysis we used data from a negotiation experiment conducted at four universities in four countries in Europe and South America (Pfeiffer, 2021). The main aim of that experiment was to study the effect of cultural differences on the use of different negotiation tactics. The analysis we perform here is a secondary analysis of the data set generated in that study.

The negotiation case used was adapted from literature (Pesendorfer et al., 2007). It was a mixed motive multi-issue buyer-seller negotiation between a representative of a farmers' association and a representative of a pharmaceutical company about the sale of vaccines against Mad Cow Disease (BSE). The case involved five issues, for which between four and seven options were available for a total of 1,680 possible packages (potential agreements combining the levels for the multiple issues). Additive multi-attribute utility functions were provided to participants in the form of scoring tables showing participants only the utility values of their own side for each option. The case contained 32 efficient (Pareto optimal) packages, out of which 16 were possible asymmetric Nash bargaining solutions for some values of r.

In total, 216 students (138 female, 72 male, 6 undisclosed, mean age 25.4 years) participated in the experiment as part of negotiation classes taught at their home universities. Participants were assigned to roles in the case in a way that all students from the same university played the same role, to minimize the risk of spreading confidential role-specific information such as payoff values to participants playing the other role. 108 dyads performed the negotiation, out of which 101 dyads reached an agreement. All negotiations were conducted electronically using the Negotiation Support System NEGOISST (Schoop et al., 2003) and took place in November and December 2020.

The experiment was conducted in several stages. First, an initial (pre-briefing) questionnaire was administered to participants, in which their demographic data, their experience in negotiations and with negotiation support systems and their knowledge of English (the language in which the negotiations were conducted) were measured. Results of this survey are briefly summarized in Table 1. On a one to five scale, users rated their knowledge of English on average between good (3) and very good (4), but had only little to medium experience in negotiations and the majority had never used a negotiation support system before.

In the subsequent week, the course instructors at their respective universities introduced participants to the negotiation support system and participants had the opportunity to perform a trial negotiation with a software agent to familiarize themselves with the system. Subsequently, the negotiation case was made available to participants and a pre-negotiation questionnaire was administered to them. This questionnaire consisted of two main parts. One part measured cultural variables that were needed for the main purpose of the experiment. The second part, which is mainly relevant for this analysis, elicited the expectations of subjects about the upcoming negotiation. In this part, subjects where asked about their expectations about the atmosphere of the negotiation, the importance they assigned to achieving a good outcome for themselves as well as reaching a fair agreement, and their aspiration and reservation levels in each issue of the negotiation. Aspiration levels refer to the level a negotiator wants to achieve in the negotiation, and reservation levels indicate the levels below which a negotiator would consider the outcome of a negotiation as unacceptable. Both concepts were clearly explained to subjects in the questionnaires.

Subsequently, the actual negotiations were conducted. Subjects had five days to reach an agreement, but were also informed that they could terminate the negotiations without agreement at any time and that negotiations would automatically be terminated without agreement after the deadline. Finally, a post-negotiation questionnaire was administered to participants in which their satisfaction with the negotiation, their outcome, and the behavior of their opponent as well as their evaluation of the system were measured.

3.3. Research questions

After the preceding presentation, we can now revisit in more detail the research questions outlined in the introduction:

RQ1 *To what extent are observed negotiation processes compatible with the model?*
We analyze this question by assessing the compatibility between the bounds inferred from all the offers made from time 1 to time T, considering the process model $(\Omega_1^f, \ldots, \Omega_T^f)$ (Eqs. (12)-(13)) as well as considering the assumed convergence to the outcome $(\Omega_1^g, \ldots, \Omega_T^g)$ (Eqs. (17)-(18)). The latter will be analyzed considering not only the weighted Nash model $g(x, r)$, but also the NBS model $g(x, 1)$ as an additional benchmark. Compatibility at time t is established if the set Ω_t is not empty. For each offer made by one side, we consider the process to be compatible with the model if the bound implied by that offer is compatible with the bound implied by the previous offer of the other side. For example, if an offer from the seller implies a lower bound lo_t, and the preceding offer of the buyer an upper bound up_{t-1}, we consider the process is compatible with the model at time t if $lo_t \leq up_{t-1}$.
As a follow-up question, we analyze whether information becomes more precise over time, i.e., whether the width of the interval, $up_{t-1} - lo_t$, decreases as the negotiation proceeds.

Table 1
Language proficiency and negotiation experience of subjects.

	Language proficiency	Negotiation experience	Used an NSS
Scale from (1)	Not good	No experience	Never
Scale to (5)	Native speaker	Regular business negotiator	More than 10 times
Mean	3.433	2.367	1.186
Median	4	2	1
SD	0.799	0.832	0.535

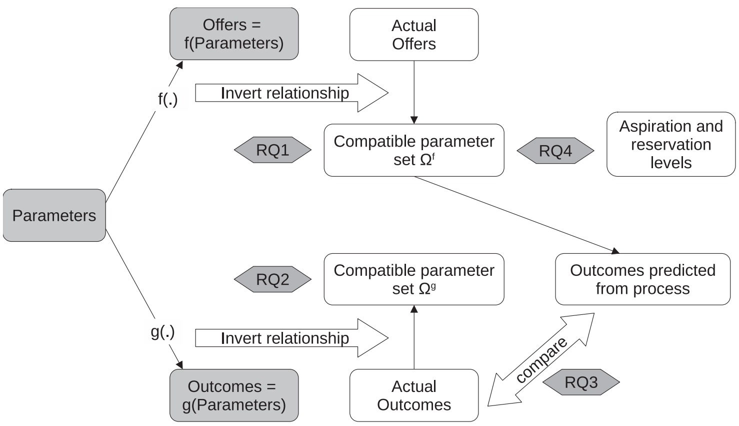

Fig. 2. Overview of the research approach and research questions.

RQ2 *To what extent are observed negotiation outcomes (agreements) compatible with the predictions of the model and how do actual and predicted agreements differ*?
We analyze the first part of this question by observing how many agreements fulfil the maximization condition (23) for some value of r. For the second part, we compare the actual outcomes of negotiations to the asymmetric Nash solutions which are closest to the actual agreements in utility space, where we define closeness in terms of Euclidean distance. To examine the second part of this question from a different perspective, we also consider the asymmetric Nash solutions that best match the range of r estimated from the process, i.e., having the largest overlap per Eq. (24). These predictions, possibly different from the solution minimizing the Euclidean distance, are also compared to the actual agreements,

RQ3 *Does the model predict that the observed processes lead to the observed outcomes*?
To analyze this question, we consider the outcome that the model would predict based on the parameter r estimated from the process and compare that predicted outcome to the actual outcome of the negotiation. Since the range of possible parameter values that can be inferred from the offers might extend beyond the interval that is compatible with a particular outcome, we again use the outcome which has the largest overlap with process parameters according to Eq. (24).

RQ4 *Do confidence parameters estimated from observed behavior exhibit theoretically plausible relationships to aspiration and reservation levels*?
As mentioned before, data collected before the negotiations elicited the aspiration and the reservation levels of the subjects. To analyze this question, we study the relationship between the bounds for r inferred from the process and the differences in aspiration and reservation levels between the two parties.

Figure 2 provides a graphical overview of our approach and how the research questions relate to it.

4. Results

4.1. Offers

We first analyze whether individual offers are compatible with the model, in terms of process, to answer **RQ1**. In total, 769 offers were made during the experiments. Table 2 classifies these offers according to their compatibility with the process model. For the first offer in each negotiation, there is no offer from the other side on the table. We can assume that the implicit offer of the other side is the best package for the opponent. However, this approach is only useful in the case of the symmetric Nash solution. The opponent's best offer has a utility of zero for the focal negotiator, which would lead to taking the logarithm of zero in (17). For comparability, we also excluded these cases in the calculation of bounds from Eqs. (13) and (12). Furthermore, if the offers from both sides are still at the extreme positions, the denominator in these equations becomes zero. All these cases that do not provide any information about r are counted as N/A (Not Applicable) in Table 2.

Since we had to exclude several offers in evaluating the model, the total number of compatible offers for our model is smaller than for the symmetric Nash bargaining solution. However, if one considers only the offers for which a bound can be calculated, it becomes clear that the fraction of compatible offers, which exceeds 80%, is much higher for the asymmetric than for the symmetric Nash bargaining solution. This difference is significant according to a χ^2-test ($\chi^2 = 23.353, p < 0.1\%$).

Considering offers from the two parties separately, the fact that most negotiations were started by the buyer is reflected in larger number of missing data points from the buyer (for which no previous offer from the seller was available). The fraction of compatible offers is higher for the seller than the buyer (85.8 vs. 78.8% for the first criterion), this difference is significant at $p = 3\%$. However, this difference should not be interpreted as causality that buyers are more inclined to make offers that are inconsistent with the model than sellers. Compatibility means that e.g. an upper bound calculated from an offer of the buyer is larger than the lower bound calculated from the preceding offer from the seller. If an offer from the buyer results in incompatibility, this could be caused

Table 2
Compatibility of individual offers with the model.

Criterion	Role		Total	N/A	Incomp.	Compatible
$r \in \Omega_i^f$ (12)-(13)	All	N	769	128	113	528
		% of total		16.6%	14.7%	68.7%
		% of valid			17.6%	82.4%
	Buyer	N	408	106	64	238
		% of total		26.0%	15.7%	58.3%
		% of valid			21.2%	78.8%
	Seller	N	361	22	49	290
		% of total		6.1%	13.6%	80.3%
		% of valid			14.5%	85.5%
$r \in \Omega_*^{g(x,r)}$ (asymmetric Nash)	All	N	769	128	115	526
		% of total		16.6%	15.0%	68.4%
		% of valid			17.9%	82.1%
	Buyer	N	408	106	65	237
		% of total		26.0%	15.9%	58.1%
		% of valid			21.5%	78.5%
	Seller	N	361	22	50	289
		% of total		6.1%	13.9%	80.1%
		% of valid			14.7%	85.3%
$r \in \Omega_*^{g(x,1)}$ (symmetric Nash)	All	N	769		224	545
		% of total			29.1%	70.9%
	Buyer	N	408		96	312
		% of total			23.5%	76.5%
	Seller	N	361		128	233
		% of total			35.5%	64.5%

Table 3
Statistics on changes in lower and upper bounds over time.

	Change in bound	
	lower	upper
Mean	0.1041	-0.5442
Median	0.1129	-0.2580
SD	1.6304	2.1419
Min	-10.3882	-8.8795
Max	8.3457	11.4199
N	251	195
p	0.0055	0.0002

Table 4
Compatibility of offers across the entire negotiation.

	Extreme bounds		Average bounds	
Total	108	100.0%	108	100.0%
Incompatible	76	70.4%	12	11.1%
Compatible	32	29.6%	95	88.0%
- with an interval below 1	9	8.3%	19	17.6%
- with an interval including 1	9	8.3%	45	41.7%
- with an interval above 1	14	13.0%	31	28.7%

by a rather low upper bound calculated from the buyer's offer, but also from a rather high lower bound calculated from the previous offer of the seller that was just below the previous upper bound. Thus, incompatibility cannot be causally connected to an offer from either one side.

Next, for the RQ1 follow-up, we analyze the evolution of the estimated bounds over time. Since negotiators learn about each other's preferences during the negotiation, we expect them to make offers that fit the model better over time. Consequently, the range of possible values of r should shrink, i.e., lower bounds should increase and upper bounds decrease over time.

To test this relationship, we first consider the differences in estimated bounds on r between subsequent offers from the same party as shown in Table 3.

Both the mean and median changes have the expected signs, lower bounds increase and upper bounds decrease over time, although as the extreme values indicate, there are also changes in the opposite direction. A nonparametric Wilcoxon test confirms that changes in lower bounds are significantly larger than zero and changes in upper bounds are significantly lower than zero.

We extend this analysis by applying the SIPA method of Vetschera and Filzmoser (2012). This method uses linear interpolation between observations made at discrete points in time (the time offers are actually made) to approximate the value of a variable describing the negotiation process at identical points (such as the end of each quarter of the negotiation) in all negotiations. Here, we apply this method to the estimated upper and lower bounds of parameter r.

Figure 3 shows the results of this analysis graphically. Again, it can be observed that as the negotiation proceeds, lower bounds increase and upper bounds decrease. Only at the beginning of the negotiation, the trend of decreasing upper bounds is not observed. However, most negotiations were started by the buyer, and therefore this part of the data is based on a small number of observations. Most differences are statistically significant: only values at the end of the negotiation are not significantly different from the situation after three quarters of the negotiation for both types of bounds, and also not significantly different from the situation halfway through the negotiation for lower bounds.

Aggregating bounds across the entire negotiation yields the results shown in Table 4. Considering the highest lower and the lowest upper bound, only a third of all negotiations are fully compatible with the model. However, as we have already argued, this requirement is probably too strict. Negotiators do not know the utility function and expectations of their opponent and therefore can easily make concessions that are either too small or too large. Considering the average bound across all offers of each party, as shown in the two rightmost columns of Table 4, therefore provides a more realistic picture. In Table 4 we also provide a more detailed analysis of those negotiations which are compatible with the model according to the estimated range of the confidence ratio r. A ratio $r = 1$ would imply that both parties have the same confidence parameter and thus act with the same level of confidence during the negotiation. In about 40% of the negotiations, the interval we have estimated for r includes that value. The intervals which are strictly above or strictly below one indicate that one party has a higher level of confidence than the other party. These asymmetric situation show a difference between roles. In 31 negotiations, the seller acted in a more confident way than the buyer, whereas the opposite occurred in only 19 negotiations. However, this distribution is not significantly different from an equal split ($\chi^2 = 2.42, p = 11.98\%$).

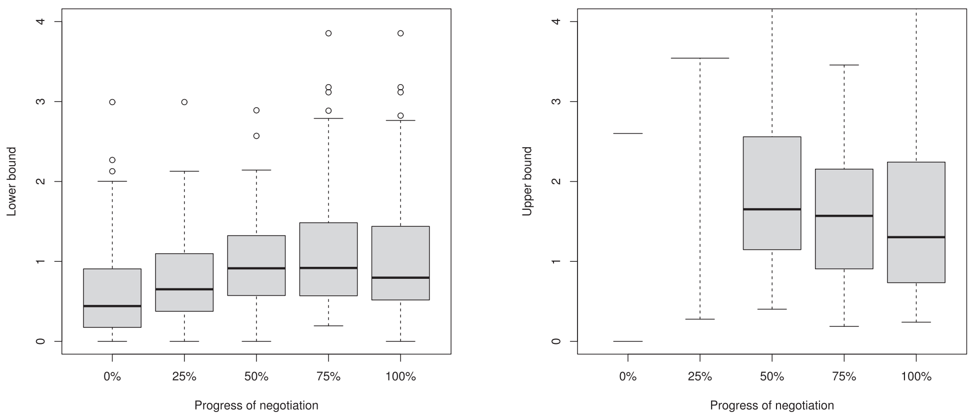

Fig. 3. Development of lower and upper bounds over time.

Table 5
Correlation between estimated confidence ratio r and differences in aspiration and reservation levels.

Difference in		Extreme bounds		Average bounds	
		lower	upper	lower	upper
Aspirations	Correlation	0.2449	0.2411	0.2391	0.3066
	p	0.0330	0.0359	0.0375	0.0071
Reservations	Correlation	0.0756	-0.1702	-0.0789	-0.1256
	p	0.5165	0.1415	0.4983	0.2798

These results describe behavior during the entire negotiation and therefore provide an opportunity to analyze Research Question RQ4. This RQ relates the value of the confidence parameter to aspiration and reservation levels of negotiators, which were elicited before the actual negotiation started. Both high confidence and high aspirations/reservation levels can lead to a tougher bargaining style. We therefore expect that negotiators who have high aspiration or reservation levels for the negotiation would also exhibit a higher value of the confidence parameter. However, our parameter r only measures the ratio of the two confidence parameters, not the individual values of the two parties. Consequently, if in a negotiation one party has higher aspiration or reservation level than the other party, we expect the value of r to be in favor of that party.

Table 5 shows the correlation between the difference in aspiration (reservation) levels (seller minus buyer) and the estimated value of r (which is the ratio of seller's confidence divided by buyer's confidence). High values in both variables therefore indicate a situation that is favourable for the seller. There is a significant positive correlation between the difference in aspiration levels and the estimated value of r. The more the seller's aspiration level exceeds that of the buyer, the more confidently the seller acts compared to the buyer. No significant correlation exists between the difference in reservation levels and parameter r. In retrospect, this is not surprising: aspiration levels influence behavior throughout the negotiation, reservation levels become relevant only if a negotiator is already close to his or her reservation level. We can therefore conclude that our estimate of parameter r is indeed a useful representation of behavior during the bargaining process, that exhibits expected relationships to aspiration levels and therefore external validity.

4.2. Outcomes

4.2.1. Compatibility of agreements

We now turn to **RQ2**. As already indicated, it is quite unlikely that negotiators exactly reach the asymmetric Nash solution predicted by the model.

Figure 4 shows the location of agreements in utility space. Filled circles (green) indicate the utilities of agreements. Squares (blue) mark efficient solutions (in some cases, there seem to be two efficient solutions at the same utility value of the seller, in fact the one that is worse for the buyer is slightly better for the seller), and, among the efficient solutions, the potential asymmetric Nash bargaining solutions are additionally marked with an "X" (magenta). Out of the total 101 negotiations that reached an agreement, the agreement was inefficient (Pareto dominated in utility space by another possible contract) in 81 negotiations. Seven of the 20 efficient agreements also correspond to an asymmetric Nash bargaining solution for some value of r, the remaining 13 do not. We therefore have to answer the first part of RQ2 negatively.

To study the second part of RQ2, we compare the actual outcomes of negotiations to the asymmetric Nash solutions which are closest to the actual agreements in utility space, where we define closeness in terms of Euclidean distance.

Table 6 gives an overview of all 16 asymmetric Nash solutions that were possible in the negotiation case, and the number and average properties of actual agreements closest to them in utility space. Since r is the ratio of the seller's confidence divided by the buyer's confidence, solutions most favorable to the seller correspond to high values of r. Actual agreements are located near only six out of the 16 possible asymmetric Nash solutions. Most agreements were quite balanced in terms of the utilities of the two parties, so the solutions in which one party would be much better off were typically quite far from the actual agreements. Considering all negotiations, 57 agreements were located close to a solution that slightly favors the seller (solution 11, $u_b = 0.6300, u_s = 0.7005$) and where we therefore estimate the confidence parameter of the seller to be between 1.25 and 1.82 times larger than that of the buyer. Twenty agreements were close to the most balanced solution (number 10, $u_b = 0.6900, u_w = 0.6615$), where the interval of r includes the symmetric value of 1. This solution therefore also corresponds to the symmetric Nash bargaining solution. Comparing the (average) utilities of actual agreements with those of the theoretical solutions, it becomes apparent that actual agreements are inefficient and we also notice a tendency towards more equal payoffs, in particular in cases where the theoretical solution would be quite unbalanced.

4.2.2. Solutions predicted from the negotiation process

To answer **RQ3**, we consider those asymmetric Nash solutions which most closely match the process by providing the largest overlap between the interval of r implied by the solution and the interval implied by the process. Following the approach in Eq. (24), we assign to each negotiation the asymmetric Nash solution that has the largest overlap in terms of the intervals estimated from the process for parameter r.

Table 7 provides an overview of the solutions assigned to negotiations by that method similarly to Table 6. Compared to the approach minimizing the Euclidean distance, a larger number of negotiations is mapped to asymmetric Nash solutions which favor the seller (e.g., 20 ne-

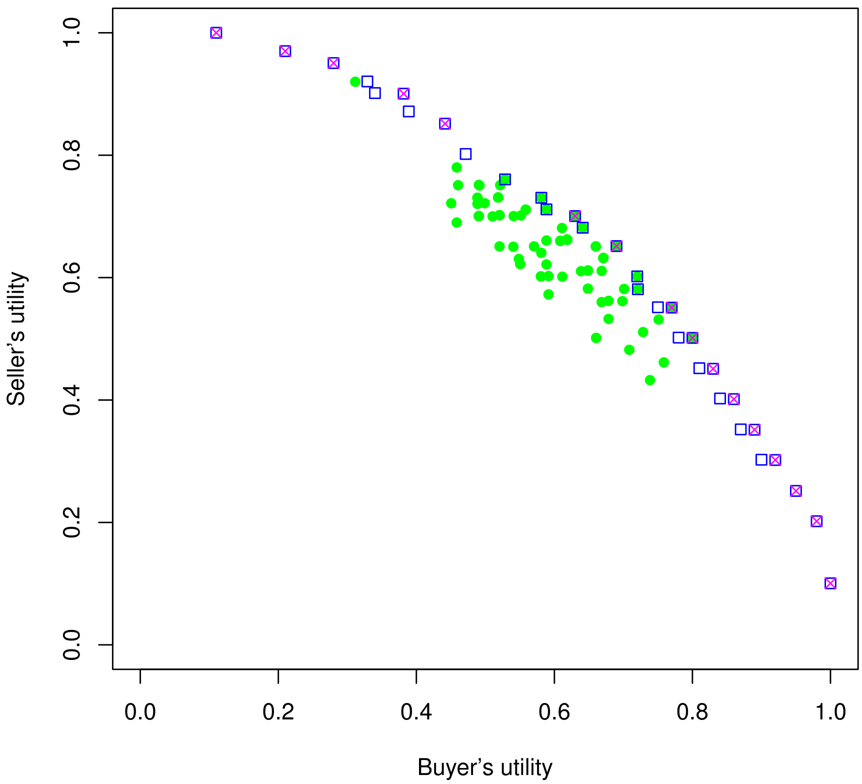

Fig. 4. Actual agreements, efficient solutions and asymmetric Nash bargaining solutions in utility space.

Table 6
Overview of possible asymmetric Nash solutions mapped to agreements (Distance corresponds to the Euclidean distance in the utility space).

Asymmetric Nash solutions					Closest actual agreements			
	Utility		Bounds on r		Average Utility			Average
Id	Buyer	Seller	Lower	Upper	Buyer	Seller	Number	Distance
1	1.0000	0.1005	-Inf	0.0289				
2	0.9800	0.2020	0.0289	0.1419				
3	0.9500	0.2515	0.1419	0.1754				
4	0.9200	0.3020	0.1754	0.2184				
5	0.8900	0.3515	0.2184	0.2578				
6	0.8600	0.4015	0.2578	0.3054				
7	0.8300	0.4510	0.3054	0.3469				
8	0.8000	0.5015	0.3469	0.4060	0.7657	0.4652	3	0.0500
9	0.7700	0.5510	0.4060	0.6548	0.7197	0.5323	11	0.0582
10	0.6900	0.6515	0.6548	1.2545	0.6607	0.6105	20	0.0562
11	0.6300	0.7005	1.2545	1.8214	0.5713	0.6831	57	0.0761
12	0.4415	0.8515	1.8214	2.6106	0.4751	0.7362	9	0.1213
13	0.3815	0.9005	2.6106	5.7241				
14	0.2800	0.9505	5.7241	14.1660	0.3115	0.9200	1	0.0438
15	0.2100	0.9700	14.1660	21.2293				
16	0.1100	1.0000	21.2293	Inf				

gotiations are mapped to solution number 12, in which the seller's utility is almost twice as large as that of the buyer, compared to 9 agreements mapped to that solution using Euclidean distance). Thus it seems that process parameters are more in favor of the seller than the actual outcomes. On average, the intervals of r implied by the asymmetric Nash solutions overlap more than 50% of the interval implied by the offers, as indicated in the last column of Table 7.

Figure 5 illustrates this assignment in utility space. Each part of the figure shows the actual agreements assigned to one of these six asymmetric Nash solutions, which is indicated by the solid dot. Empty circles show the utilities of all other symmetric Nash solutions and the crosses the utility values of actual agreements mapped to that particular asymmetric Nash solution. Clearly, the solutions predicted by the process parameters are not always the ones which are closest to the actual outcome in utility space.

To further study the second part of **RQ2**, we compare properties of the asymmetric Nash solutions predicted from process parameters to the solutions which are closest to the actual agreements in utility space and the actual agreements themselves. We apply three measures in that comparison: The first is the joint utility (the sum of utility values of both parties), which measures the efficiency of outcomes. The other two measures are related to fairness. We consider the utility of the weaker player and the contract imbalance (the difference in utilities).

Table 8 shows the results of this comparison. Since most of the actual agreements are not efficient, it is not surprising that the predicted solutions would have achieved a higher joint utility than the actual

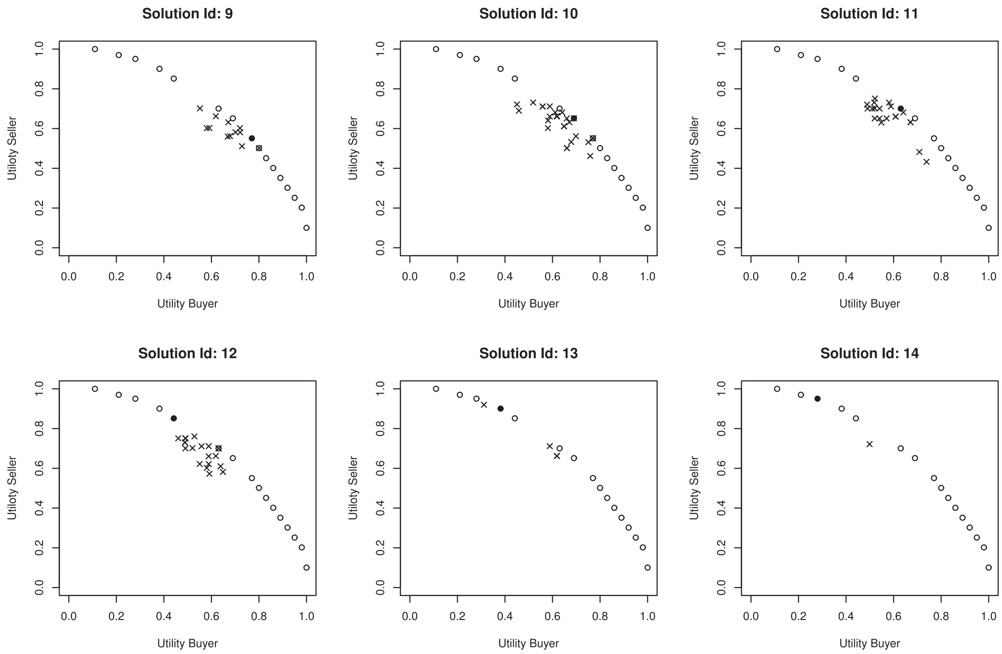

Fig. 5. Actual agreements mapped to asymmetric Nash solutions according to process parameters.

Table 7
Asymmetric Nash solutions predicted from the process.

	Asymmetric Nash solutions				Actual agreements	
	Utility		Bounds on r			Average
Id	Buyer	Seller	Lower	Upper	Number	Overlap
9	0.7700	0.5510	0.4060	0.6548	12	0.5581
10	0.6900	0.6515	0.6548	1.2545	31	0.6770
11	0.6300	0.7005	1.2545	1.8214	21	0.6735
12	0.4415	0.8515	1.8214	2.6106	20	0.5080
13	0.3815	0.9005	2.6106	5.7241	4	0.4336
14	0.2800	0.9505	5.7241	14.1660	1	1.0000

agreements. The picture is less clear concerning fairness. The utility of the worse-off player is lower in the actual agreements than in the closest asymmetric Nash solution as well as the solution that would be predicted from the process parameters. Thus, the lack of efficiency at least partially also affects the weaker player. The contract imbalance exhibits an interesting pattern. While the average contract imbalance is considerably higher in the predicted than the actual solution, the converse holds for the median. Thus the distribution of this variable is quite asymmetric. All solutions are significantly different from each other in terms of joint utility and the utility of the worse off player. The contract imbalance is significantly different between the actual and the closest asymmetric Nash solution, but not between the actual solution and the solution predicted from the process.

To study this effect further, we consider the individual asymmetric Nash solutions in Table 9. Here it becomes clear that the low contract imbalance of the predicted solutions results mainly from the two solutions number 10 and 11, which are predicted for a large number of negotiations (31 and 21, respectively). These are the two asymmetric Nash solutions which provide the most similar payoff to the two players. The actual negotiations (and agreements) for which these two theoretical solutions are predicted exhibit a considerable variation in payoff structures as can be seen in Fig. 5. In contrast, in all other predicted solutions, the contract imbalance of the predicted solution is much higher than in the actual agreements, leading to a high average contract imbalance of predicted solutions. Figure 5 also shows that most negotiations for which solution number 12 is the predicted outcome actually have led

Table 8
Characteristics of actual agreements, closest asymmetric Nash solutions in utility space and asymmetric Nash solutions predicted from process.

		Joint Utility	Utility worse off	Contract imbalance
Actual	Mean	1.2513	0.5597	0.1319
	Median	1.2530	0.5705	0.1230
	SD	0.0493	0.0600	0.0986
Closest	Mean	1.3284	0.6051	0.1181
	Median	1.3305	0.6300	0.0705
	SD	0.0148	0.0687	0.1230
Process	Mean	1.3213	0.5694	0.1826
	Median	1.3305	0.6300	0.0705
	SD	0.0226	0.0978	0.1735
Actual vs. Closest	W	390	1819	4653
	p	0.0000	0.0000	0.0422
Actual vs. Process	W	724	3027.5	3855.5
	p	0.0000	0.0064	0.7599

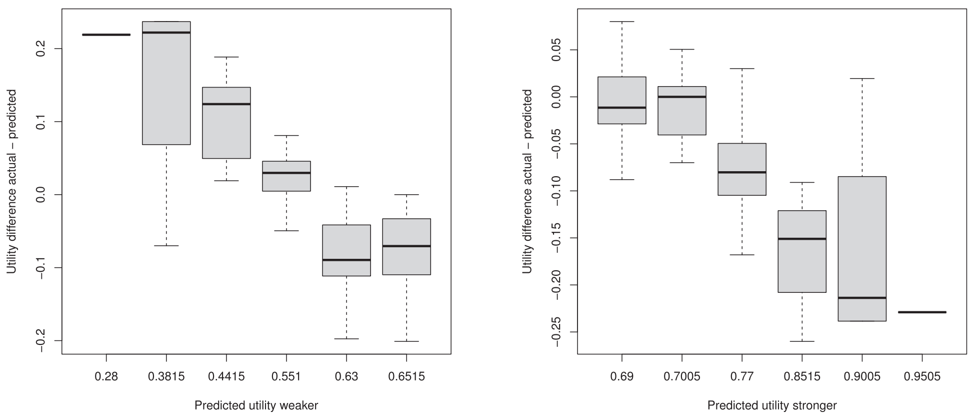

Fig. 6. Utility differences between predicted and actual outcomes for weaker (left) and stronger (right) player by predicted utility.

Table 9
Comparison of contract imbalance of predicted asymmetric Nash solutions and averages of actual agreements assigned to it.

		Predicted utility		Contract imbalance	
Id	N	Buyer	Seller	Predicted	Actual
9	12	0.7700	0.5510	0.2190	0.1153
10	31	0.6900	0.6515	0.0385	0.1162
11	21	0.6300	0.7005	0.0705	0.1441
12	20	0.4415	0.8515	0.4100	0.1345
13	4	0.3815	0.9005	0.5190	0.2046
14	1	0.2800	0.9505	0.6705	0.2225

to a considerably more balanced agreement, while in solution number 12, the seller receives almost twice the utility of the buyer.

These results indicate that in negotiations in which the process would have predicted a very balanced outcome (solutions 10 and 11), the actual outcomes are typically more imbalanced, while in the other cases, actual outcomes are more balanced than the predicted one. A more balanced payoff can be achieved if the weaker player receives a higher utility. Figure 6 shows the distribution of differences in utility between actual and predicted agreement for the weaker and the stronger player for the six predicted solutions. The left part shows the payoff difference for the weaker player. Clearly, in cases in which the predicted solution would provide only a very low payoff to the weaker player (such as 0.3815 in solution number 13 or 0.4415 in solution number 12), that player receives a substantially higher utility in the actual agreement and for almost all cases, the difference between actual and predicted agreement is positive. However, in the more balanced situations where the weaker player would receive a comparatively high utility (0.63 for the buyer in solution 11 and 0.6515 for the seller in solution 10) the difference is mostly negative. A symmetric pattern can be observed for the utility differences of the stronger player. Taken together, the two parts of this figure clearly show that compared to the predicted outcomes, subjects seemed to avoid extremely unbalanced solutions even if their behavior in the process would have predicted such an outcome.

5. Conclusions

In the present paper, we address the relationship between negotiation processes and outcomes using a model based on the decisions made in each bargaining step. We study whether the theoretical predictions of the model are reflected in the actual bargaining steps that negotiators make, and in the outcomes they reach.

Concerning **RQ1**, our results indicate that about 80% of the bargaining steps in the negotiation processes we experimentally observed are compatible with the model. This high level of compatibility is noteworthy given that the parties had no information about their opponents accept the offers they made. Concerning **RQ2**, the number of agreements that correspond exactly to the theoretically predicted one is very low. This is not surprising, because the model predicts an efficient agreement and the literature shows negotiators very often fail to achieve efficiency. Nevertheless, we could also show that actual solutions come close to the performance of the predicted solutions both in terms of efficiency and in terms of fairness. Concerning **RQ3**, a comparison of the actual agreements to those predicted by the model based on the preceding process revealed an interesting pattern: in cases in which the model would predict a very imbalanced solution, the actual solution is much more balanced. On the other hand, the actual agreement was somewhat less balanced in cases in which the predicted solution would have provided almost equal utility to both parties. Thus, actual agreements are more concentrated at a moderate level of inequality than the theoretically predicted solutions. Finally, we also considered the relationship between the estimated confidence parameters of our model and the aspiration levels of negotiators, a variable reported by them before bargaining begun (**RQ4**), We found a strong empirical evidence for the relationship between the imbalance in aspiration levels and the imbalance in the confidence parameters (both estimating the latter from the processes and estimating the latter from the outcomes).

Our research therefore makes both conceptual and empirical contributions. The ratio of confidence parameters is a compact way of representing bargaining behavior or, more precisely, the asymmetry of bargaining behavior between negotiators. Compared to existing measures such as the shape (concavity or convexity) of concession curves, we consider this measure to have two main advantages. First, it is closely linked to a theoretical perspective of the bargaining process, rather than being a purely descriptive concept. Furthermore, it is a measure at the dyad level rather than a characteristic of the behavior of one individual negotiator. This makes it easier to relate this measure to outcome dimensions, which are also the result of the behavior of both sides. Thus using this measure allows to avoid many conceptual problems of interdependence that occur when negotiation data is analyzed at the level of individual negotiators (Turel, 2010).

The application of this new measure led to some interesting empirical results. Actual agreements in our experiment exhibited a moderate level of imbalance, thus deviating both from more and from less imbalanced theoretically predicted agreements. Furthermore, we are able to demonstrate a clear and consistent relationship between the initially self-reported aspiration levels and our measure of confidence which, in addition to adding external validity to our measure, also provides a clear connection between a negotiator's attitudes before the negotiation and behavior during the negotiation.

However, our study also showed the limits of this new concept of representing negotiation processes. Our experiments still contained a number of bargaining steps which are not compatible with the model. This might be due to the fact that our subjects lacked information about their opponent's preferences and thus might have made concessions that would be insufficient according to the model. But it could also mean that the model needs to be extended or that more robust methods for estimating the model's confidence parameters need to be developed.

Furthermore, our analysis is based on one single experiment with students, using only one negotiation case. Although there is empirical evidence that student subjects do not behave differently than experienced negotiators (Herbst and Schwarz, 2011), the use of student subjects is clearly a limitation of this research and also generalization to other cases and other settings (involving for example more, or less distributive problems of different complexity) is clearly required. Such broader empirical studies will allow to delineate the possibilities, but also the limitations of our approach to model negotiation processes more clearly.

Funding

L. Dias acknowledges CeBER's funding by national funds, through FCT, Portuguese Science Foundation, under grant UIDB/05037/2020.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A: Negotiation case

The negotiation case involved a buyer-seller negotiation between the representative of a farmers' association (NCA) in the fictitious country of Litulia and the representative of a pharmaceutical company InnoVax about the sale and use of an experimental vaccine against BSE (Bovine spongiform encephalopathy, commonly known as "Mad Cow Disease"). The background story explained that there was an outbreak of BSE in Lithulia, that InnoVax had recently developed a vaccine against that disease, and that the two parties had started negotiations about the possibility to carry out a field test of the new vaccine on the cattle herds of NCA. The instructions to both sides clearly explained that either side had outside options (other countries in which the vaccine could be tested and another company at a similar stage of development), but that coming to an agreement would be beneficial for both sides.

The case involved five issues, with predefined options for each issue. The issues were:

1. Price of the vaccine, to be set in steps of 100 Dollars between 100 $ and 500 $.
2. Delivery time between one and seven months.
3. Guarantees by the supplier about the effectiveness and safety of the vaccine, ranging from no guarantee at all, a narrow guarantee that would cover only gross negligence on the side of the seller, a wide guarantee covering all possible damages from the vaccine up to a full guarantee, in which the seller would also compensate the buyer for damages from the disease in case the vaccine was not effective.
4. Publicity about the test being carried out, ranging from keeping the test secret to publications in medical journals and full information of the general public.
5. The possibility of future testing of other vaccines on NCA's cattle, with the options from no additional test, testing only improved versions for the BSE vaccine (minimal), testing of other vaccines already approved by authorities (standard) to unlimited testing of any new products.

Both parties were provided with a payoff table (data in Table A.1) that contained all the elements of an additive utility function for each side. Each party only received the payoff table for its own side and was instructed to treat this information confidentially.

Table A.1
Payoff tables used in the experiment.

	Attribute weights			Ratings	
Attribute	Buyer	Seller	Option	Buyer	Seller
Delivery time	35%	25%	1 month	100	0
			2 months	86	12
			3 months	71	24
			4 months	43	40
			5 months	29	80
			6 months	14	88
			7 months	0	100
Price	20%	35%	100	100	0
			200	90	29
			300	60	49
			400	30	86
			500	0	100
Guarantee	25%	15%	None	0	100
			Narrow	40	80
			Wide	68	67
			Total	100	0
Publicity	10%	10%	None	100	0
			Narrow	40	100
			Wide	0	60
Testing	10%	15%	None	0	0
			Minimal	70	40
			Standard	100	67
			Unlimited	70	100

References

Benton, A.A., Kelley, H.H., Liebling, B., 1972. Effects of extremity of offers and concession rate on the outcomes of bargaining. J. Pers. Soc. Psychol. 24 (1), 73–83. doi:10.1037/h0033368.

Bishop, R.L., 1964. A Zeuthen-Hicks theory of bargaining. Econometrica 32, 410–417. doi:10.2307/1913045.

Bone, J., Hey, J.D., Suckling, J., 2014. What price compromise? Theory Decis. 77 (3), 359–376. doi:10.1007/s11238-014-9465-9.

Caputo, A., 2013. A literature review of cognitive biases in negotiation processes. Int. J. Confl. Manage. 24 (4), 374–398. doi:10.1108/IJCMA-08-2012-0064.

Carbonneau, R.A., Vahidov, R.M., 2014. A utility concession curve data fitting model for quantitative analysis of negotiation styles. Expert Syst. Appl. 41 (9), 4035–4042. doi:10.1016/j.eswa.2013.12.029.

Chatterjee, K., 2021. Non-cooperative bargaining theory. In: Kilgour, D.M., Eden, C. (Eds.), Handbook of Group Decision and Negotiation. Springer, Cham, pp. 531–544. doi:10.1007/978-3-030-49629-6_9.

Chertkoff, J.M., Conley, M., 1967. Opening offer and frequency of concession as bargaining strategies. J. Pers. Soc. Psychol. 7 (2), 181–185. doi:10.1037/h0024997.

Dias, L.C., Vetschera, R., 2019. Multiple local optima in Zeuthen-Hicks bargaining: an analysis of different preference models. EURO J. Decis. Process. 1-2, 33–53. doi:10.1007/s40070-018-0089-0.

Dias, L.C., Vetschera, R., 2022. Two-party bargaining processes based on subjective expectations: a model and a simulation study. Group Decis. Negotiat. 31, 843–869. doi:10.1007/s10726-022-09786-x.

Dickinson, D.L., 2009. The effects of beliefs versus risk attitude on bargaining outcomes. Theory Decis. 66 (1), 69–101. doi:10.1007/s11238-007-9075-x.

Engler, Y., Page, L., 2022. Driving a hard bargain is a balancing act: how social preferences constrain the negotiation process. Theory Decis. 93, 7–36. doi:10.1007/s11238-021-09835-y.

Filzmoser, M., Vetschera, R., 2008. A classification of bargaining steps and their impact on negotiation outcomes. Group Decis. Negotiat. 17 (5), 421–443. doi:10.1007/s10726-008-9106-1.

Gettinger, J., Filzmoser, M., Koeszegi, S.T., 2016. Why cant we settle again? Analysis of factors that influence agreement prospects in the post-settlement phase. J. Bus. Econ. 86 (4), 413–440. doi:10.1007/s11573-016-0809-5.

Harsanyi, J.C., 1956. Approaches to the bargaining problem before and after the theory of games: a critical discussion of Zeuthen's, Hicks', and Nash's theories. Econometrica 24 (2), 144–157. doi:10.2307/1905748.

Harsanyi, J.C., 1977. Rational Behavior and Bargaining Equilibrium in Games and Social Situations. Cambridge University Press, Cambridge.

Herbst, U., Schwarz, S., 2011. How valid is negotiation research based on student sample groups? New insights into a long-standing controversy. Negotiat. J. 27 (2), 147–170. doi:10.1111/j.1571-9979.2011.00300.x.

Hindriks, K., Jonker, C.M., Tykhonov, D., 2007. Negotiation dynamics: analysis, concession tactics, and outcomes. In: Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, pp. 427–433. doi:10.1109/IAT.2007.50.

Hinton, B.I.L., Hamner, W.C., Pohlen, M.F., 1974. The influence of reward magnitude, opening bid and concession rate on profit earned in a managerial negotiation game. Behav. Sci. 19, 197–203. doi:10.1002/bs.3830190305.

Kersten, G.E., Mallory, G.R., 1999. Rational inefficient compromises in negotiations. J. Multi-Criteria Decis. Anal. 8 (2), 106–111. doi:10.1002/(SICI)1099-1360(199903)8:2<106::AID-MCDA243>3.0.CO;2-7.

Kersten, G.E., Vahidov, R., Gimon, D., 2013. Concession-making in multi-attribute auctions and multi-bilateral negotiations: theory and experiments. Electron. Commerce Res. Appl. 12, 166–180. doi:10.1016/j.elerap.2013.02.002.

Kibris, O., 2021. Negotiation as a cooperative game. In: Kilgour, D.M., Eden, C. (Eds.), Handbook of Group Decision and Negotiation. Springer Cham, pp. 545–568. doi:10.1007/978-3-030-49629-6_10.

Kilmann, R.H., Thomas, K.W., 1977. Developing a forced-choice measure of conflict-handling behavior: the "mode" instrument. Educ. Psychol. Meas. 37 (2), 309–325.

Metcalfe, D., 2000. Rethinking Pareto-efficiency and joint feasibility. Negotiat. J. 16. doi:10.1023/A:1007590219594. 29–22

Mumpower, J.L., 1991. The judgment policies of negotiators and the structure of negotiation problems. Manage. Sci. 37 (10), 1304–1324. doi:10.1287/mnsc.37.10.1304.

Munier, B.R., Rulliére, J.-L., 1993. Are game theoretic concepts suitable negotiation support tools? From Nash equilibrium refinements toward a cognitive concept of rationality. Theory Decis. 34 (3), 235–253. doi:10.1007/BF01075191.

Muthoo, A., 1999. Bargaining Theory with Applications. Cambridge University Press, Cambridge.

Nash, J.F., 1950. The bargaining problem. Econometrica 18 (2), 155–162. doi:10.2307/1907266.

Nastase, V., 2006. Concession curve analysis for inspire negotiations. Group Decis. Negotiat. 15, 185–193. doi:10.1007/s10726-006-9028-8.

Northcraft, G.B., Brodt, S.E., Neale, M.A., 1995. Negotiating with nonlinear subjective utilities: Why some concessions are more equal than others. Organ. Behav. Hum. Decis.Process. 63 (3), 298–310. doi:10.1006/obhd.1995.1081.

Northcraft, G.B., Preston, J.N., Neale, M.A., Kim, P.H., Thomas-Hunt, M.C., 1998. Nonlinear preference functions and negotiated outcomes. Organ. Behav. Hum. Decis.Process. 73 (1), 54–75. doi:10.1006/obhd.1998.2750.

Pesendorfer, E.-M., Graf, A., Koeszegi, S., 2007. Relationship in electronic negotiations: tracking behavior over time. Zeitschrift fõr Betriebswirtschaft 77 (12), 1315–1338. doi:10.1007/s11573-007-0327-6.

Pfeiffer, M.-C., 2021. The influence of cultural differences on computer-based negotiations: a four-country analysis. University of Vienna, Vienna, Austria Master thesis.

Roth, A.E., 1979. Axiomatic Models of Bargaining. Springer Verlag, Berlin.

Rubinstein, A., 1982. Perfect equilibrium in a bargaining model. Econometrica 50 (1), 97–109. doi:10.2307/1912531.

Schoop, M., Jertila, A., List, T., 2003. Negoisst: a negotiation support system for electronic business-to-business negotiations in e-commerce. Data Knowl. Eng. 47 (3), 371–401. doi:10.1016/S0169-023X(03)00065-X.

Turel, O., 2010. Interdependence issues in analyzing negotiation data. Group Decis. Negotiat. 19, 111–125. doi:10.1007/s10726-008-9118-x.

Vetschera, R., 2007. Preference structures and negotiator behavior in electronic negotiations. Decis. Support Syst. 44 (1), 135–146. doi:10.1016/j.dss.2007.03.007.

Vetschera, R., 2013. Negotiation processes: an integrated perspective. EURO J. Decis. Process. 1 (1–2), 135–164.

Vetschera, R., Filzmoser, M., 2012. Standardized interpolated path analysis of offer processes in e-negotiations. In: International Conference on Electronic Commerce ICEC 2012, pp. 134–140. doi:10.1145/2346536.2346560.