

Multiple local optima in Zeuthen-Hicks bargaining: An analysis of different preference models

Luis C. Dias · Rudolf Vetschera

Received: date / Accepted: date

Abstract Zeuthen-Hicks bargaining provides a dynamic model that explains how two parties in a negotiation make concessions to reach the Nash bargaining solution. However, it is not clear whether this process will always reach the global optimum corresponding to the Nash bargaining solution, or could end at a local optimum, or even in disagreement. In this paper, we analyze different types of utility functions, both analytically and in a computational study, to determine under which circumstances convergence to the Nash bargaining solution will be achieved. We show that non-standard preferences, involving e.g. reference point effects, might indeed lead to multiple local optima of the Nash bargaining objective function and thus failure of the bargaining process. This occurs more often if expectations of parties are mutually incompatible.

Keywords Zeuthen-Hicks · Bargaining · Negotiation process · Utility function · Reference point · Nash solution

1 Introduction

Many definitions of negotiations take a process perspective. For example, Kilgour and Eden (2010, p. 2) define negotiations as “... a process in which two or more independent, concerned parties make a choice”. Although the process perspective is thus seen as constituent of the concept of negotiations, actual models that describe the negotiation process are rare (Vetschera, 2013). This holds for many dimensions of a negotiation process, and in particular for the substantive level (Filzmoser et al, 2016) of offers and concession in a negotiation process.

L. C. Dias
University of Coimbra, Faculty of Economics, CeBER and INESC
Coimbra, Portugal
E-mail: lmedias@fe.uc.pt

R. Vetschera
University of Vienna, Faculty of Business, Economics and Statistics,
Vienna, Austria
E-mail: rudolf.vetschera@univie.ac.at

Formal models of bargaining based on economics and game theory are frequently classified into axiomatic and strategic models (e.g., Sutton, 1986). The origins of *axiomatic* bargaining models go back to Nash’s seminal work on the bargaining problem (Nash, 1950). They are linked to cooperative game theory by the assumption that any agreement reached by the bargainers is mutually binding (Kíbris, 2010). Axiomatic approaches define properties of a bargaining solution which can be interpreted either descriptively as properties that the outcome of a bargaining process between rational parties typically will fulfill, or from a prescriptive point of view as properties of bargaining outcomes that are socially desirable (Kíbris, 2010; Thompson, 1994). Although Nash’s axioms have been criticized in literature and many different axiomatic solutions fulfilling alternative sets of axioms have been proposed in the meantime (for a survey, see e.g., Kíbris, 2010), the Nash bargaining solution still remains a widely used concept.

Axiomatic models are only concerned with properties of the outcome of a bargaining process, but not with the process by which this outcome is reached. This question is at the focus of *strategic* models of bargaining, which are typically based on non-cooperative game theory (Chatterjee, 2010; Sutton, 1986). One of the most popular strategic models is the bargaining model by Rubinstein (1982), which assumes that two parties can make alternating offers about the distribution of a fixed asset. Each party can either accept the opponent’s offer, or, after some delay, make a counteroffer. Delays between offers are an important element in this model, since the model assumes that each party has a time discount rate and thus prefers to receive outcomes sooner rather than later. Consequently, the model predicts that already the first offer will be accepted, and that the allocation of outcomes will reflect the parties’ discount rates. Only extensions of the model considering e.g. incomplete information lead to bargaining processes that actually involve multiple rounds (Chatterjee, 2010).

There is a close relationship between axiomatic and strategic models. An important question concerning strategic models is whether the bargaining process leads to an outcome corresponding to some axiomatic solution. This relationship between bargaining processes and solutions is also the focus of the current paper. If a bargaining process consists of several steps (unlike the basic version of the Rubinstein model), the structure of the search space becomes an important issue that so far has been rarely addressed in literature. Axiomatic bargaining solutions are typically characterized as solutions which maximize some function, in the case of the Nash bargaining solution, the product of both parties’ utilities. If this function has multiple local optima, a bargaining process might end in a local rather than the global optimum. So even if a process can be shown to maximize a function of which the (global) maximum corresponds to an axiomatic bargaining solution, one cannot be sure that the process actually leads to the axiomatic solution if it may converge to a local optimum.

In the present paper, we address this question for the Nash bargaining function and study for which types of utility functions multiple local optima of the Nash objective function might exist that could lead a bargaining process to miss the Nash bargaining solution. Although this question might affect any process leading to the Nash bargaining solution, we specifically motivate our study by a particular process model, the Zeuthen-Hicks bargaining model (Harsanyi, 1956; Bishop, 1964). This model, which we will explain in detail in section 2, actually precedes the Rubinstein (1982) model and formalizes even earlier descriptions of

the bargaining process. One specific property of this model is that it describes a process consisting of several steps, that ultimately maximize the product of both parties' utilities (or more precisely the product of increases in utilities above the disagreement point). Thus, it will converge to the Nash bargaining solution if a unique optimum exists. Although this model makes clear predictions about which party makes a concession as well as the size of concessions, it has been rarely tested empirically (for one exception, see Fandel, 1985) and it has also not been much extended in theoretical research. For example, the initial paper by Harsanyi (1956) is cited less than 200 times according to Web of Science, and only few of these references deal with the actual bargaining process that the model describes. Besides the assumption of a specific and mechanical psychological process of the parties, other limitations of this model include the assumption that both parties are expected utility maximizers who have perfect information about the probability that an offer is rejected, which is unlikely in practice.

Nevertheless, the Zeuthen-Hicks bargaining model is one of the few models for which the question of convergence has to some extent already been discussed in literature. Since the model involves a comparison of offers to the disagreement outcome, the process might fail to converge to a unique solution if the two parties' perceptions of the disagreement outcome are different (Crawford, 1980). Saraydar (1971) pointed out the impact that different starting offers have on the process: If parties lack information of their opponent's preferences, the choice of a starting offer is not trivial and it could even happen that parties choose a starting offer that is worse to them than the Nash solution. In that case, they actually should make reverse concessions (increase their demands) to reach the Nash solution.

Characteristics of the utility functions and the resulting bargaining set have been studied by Saraydar (1971), who explicitly refers to the assumption that the utility functions of both parties are concave, and Coddington (2010), who showed that if the efficient frontier of the bargaining set in utility space is not convex, the process might fail to converge or the solution may be indeterminate. However, the question of how multiple local optima of the Nash objective function are related to characteristics of the bargainer's utility functions has not yet been addressed. In our view, this question is particularly important, since the existence of multiple local optima might not only affect the Zeuthen-Hicks model, but also other strategic bargaining models that supposedly converge to the Nash bargaining solution such as a sequential demand game (Thompson, 1994), or modifications of the Rubinstein game (Chatterjee, 2010).

We therefore study the existence of different local optima for different types of preferences. In particular, we focus on preferences that deviate from the classical assumption of decreasing marginal utility, by considering for example reference point effects as proposed in Prospect Theory (Kahneman and Tversky, 1979). Since there is considerable empirical evidence that negotiators do not always have concave utilities (Vetschera, 2006, 2007), studying the effects of non-standard preferences on existence of multiple local optima and thus the viability of the Zeuthen Hicks bargaining model or similar models is important for the applicability of such models to actual negotiations. This obviously not only concerns the suitability of these models as descriptive models of the bargaining process, but also their usefulness as tools to support negotiators, e.g. as part of electronic negotiation support systems (Kersten and Lai, 2010).

The remainder of the paper is structured as follows: In Section 2, we describe the Zeuthen-Hicks bargaining model in more detail and illustrate the existence of multiple local optima. Section 3 provides analytical results for some specifications of utility functions, which are complemented by a simulation study using more general classes of functions in Section 4. A discussion of results and an outlook on future research in Section 5 conclude the paper.

2 Model framework

2.1 The Zeuthen-Hicks two-party bargaining model

The Zeuthen-Hicks bargaining model is based on an analysis of the successive offers that two parties make during a bargaining process. We refer to the two parties as Buyer (B) and Seller (S), but the reasoning applies to other two-party negotiations over a single issue, provided that the parties have opposite preferences (e.g., management vs. union bargaining over salary increase). The outcome of the bargaining process should be a value x_{final} (the final ‘‘price’’). The Buyer prefers lower values (its utility function $u_B(\cdot)$ is monotonically decreasing) and the Seller prefers higher values (its utility function $u_S(\cdot)$ is monotonically increasing), but we do not make further assumptions about differentiability or other properties of the utility functions:

$$u_S(x + \delta) \geq u_S(x) \wedge u_B(x + \delta) \leq u_B(x), \forall \delta > 0 \quad (1)$$

Let us examine the problem from the Buyer’s perspective. Within each bargaining step, the Buyer can decide whether to accept the Seller’s offer x_S or make a counteroffer x_B . If the opponent’s offer is accepted, the bargaining process ends with $x_{final} = x_S$. If a counteroffer is made, there is the risk that the opponent terminates the bargaining process, leading to a disagreement outcome d . In the Zeuthen-Hicks bargaining model, this uncertainty about the opponent’s reaction is expressed by a probability. Let $p_S(x_B)$ denote the probability (estimated by the Buyer) that the Seller will reject an offer x_B . Since the Seller prefers a higher price, $p_S(x_B)$ is a monotonically decreasing function in x_B . The Buyer will make offer x_B only if the expected utility of this gamble is greater than the utility of accepting the Seller’s offer x_S :

$$p_S(x_B)u_B(d) + (1 - p_S(x_B))u_B(x_B) > u_B(x_S) \quad (2)$$

where u_B denotes the Buyer’s utility function for outcomes. Note that this formulation assumes that rejection of one party’s offer leads to a termination of the negotiation. Here we do not consider extensions of the model that explicitly distinguish between that case and a stalemate in which both parties hold out on their offers for a prolonged time (Wagner, 1979). From (2), we can compute a critical rejection probability for the Buyer as

$$p_S^* = \frac{u_B(x_B) - u_B(x_S)}{u_B(x_B) - u_B(d)} \quad (3)$$

As long as $p_S(x_B) < p_S^*$, the Buyer is willing to take the risk of making a counteroffer. Therefore, a higher p_S^* indicates a stronger bargaining position of the Buyer.

Similarly, a critical probability p_B^* concerning the Buyer's reaction to offer x_S can be computed from the Seller's perspective:

$$p_B^* = \frac{u_S(x_S) - u_S(x_B)}{u_S(x_S) - u_S(d)} \quad (4)$$

The central assumption of Zeuthen-Hicks bargaining is that the party with the higher critical probability can afford to run more risks and therefore is less inclined to accept the offer from the counterpart (Harsanyi (1956) notes this assumption can be derived from a set of assumptions that include perfect knowledge and expected-utility maximization, which can be viewed as practical limitations of this model; this model is also not considering costs incurred for elapsed time in the negotiation). According to the model, the party who has a lower critical rejection probability will adjust its offer so that its own critical probability exceeds that of the opponent. In other words, the Buyer will try to establish

$$p_S^* > p_B^* \quad (5)$$

and the Seller vice versa. After substituting (3) as well as (4) and some simplifications, (5) is equivalent to

$$(u_B(x_B) - u_B(d))(u_S(x_B) - u_S(d)) > (u_B(x_S) - u_B(d))(u_S(x_S) - u_S(d)) \quad (6)$$

Once the Buyer has the higher critical probability, the Seller will adjust its offer to gain the lead again, so in the course of the negotiation, the actions of both parties will lead to maximization of

$$n(x) = (u_B(x) - u_B(d))(u_S(x) - u_S(d)) \quad (7)$$

which is the function that is also maximized in the Nash bargaining solution. Note that in order to achieve convergence to the Nash bargaining solution, it is not necessary that both parties have full information about each other's utility functions, the model only requires the local information about utilities that is necessary to determine the concession to be made (Wagner, 1979).

The main questions addressed in this work are whether the successive adjustments of offers to increase the Nash product always lead to regular concessions in which one party reduces its own utility in favor of the other party's utility and whether this process will converge to a global optimum or end in a local optimum. We assume that a trade is beneficial, i.e., the Buyer's valuation of the good being traded is higher than the Seller's. Without loss of generality, we set the Seller's valuation to zero and the Buyer's to one, which means that the two parties basically bargain about how to split the surplus of one that is created by the trade. If bargaining fails, no trade takes place and no value is created, and therefore we consider the disagreement outcome for both parties to be zero ($u_B(d) = u_S(d) = 0$), which allows us to omit the disagreement outcome in (7).

The bargaining process will converge to a global maximum, and both parties will in their respective steps make regular concessions if the function

$$n(x) = u_S(x) \cdot u_B(x) \quad (8)$$

has a unique maximum at some price x_N , and is monotonically increasing for $x < x_N$ and monotonically decreasing for $x > x_N$, i.e., if $n(x)$ is quasiconcave. In

such conditions, when a party makes an offer that increases the product in (8), then this party is sure that it is moving in the direction of the global maximum. Otherwise, if $n(\cdot)$ is not quasiconcave, situations like those depicted in Figure 1 might arise. Suppose that, in the situation depicted on the left part of Figure 1, the Seller was initially asking for price $x_0 = 1$ and a Buyer eager to close a deal offered $x_1 = 0.45$. Then, $p_S^* > p_B^*$ and to reverse this inequality the Seller would have to at best offer $x_2 = 0.54$. To restore $p_S^* > p_B^*$ by making a regular concession, the Buyer would now offer $x_3 = 0.47$, and so on, until they would eventually reach an agreement at the local optimum $x^* = 0.50$. But the reader may note that this is not the global maximum. In such situations, moving towards the global maximum would not necessarily imply a concession: instead of moving from $x_1 = 0.45$ to $x_3 = 0.47$, the Buyer might have restored $p_S^* > p_B^*$ by *decreasing* the price offered. But this party might not be aware of this, even if it had perfect information about both utility functions in the neighborhood of its offer.

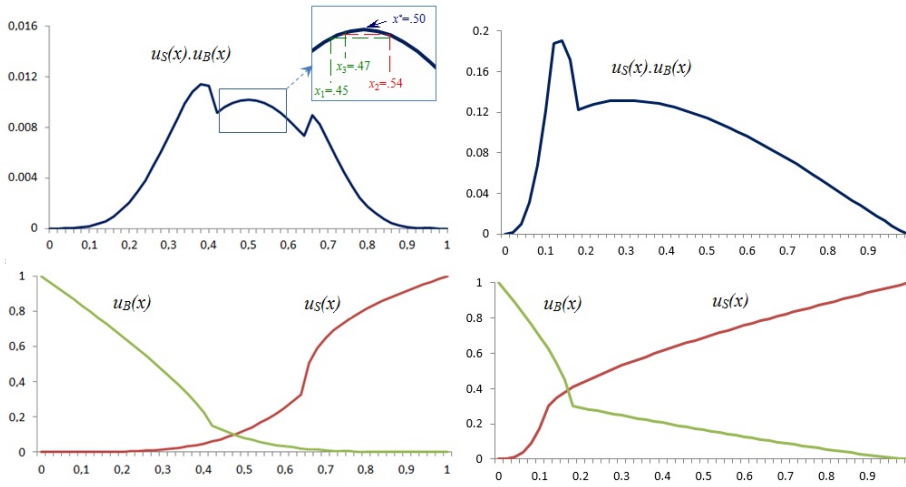


Fig. 1 Illustrative examples with multiple local optima

3 Analysis of specific functions

This section presents results concerning the quasiconcavity of the Nash product $n(x) = u_S(x)u_B(x)$ for some families of functions. First, we show that $n(x)$ is quasiconcave if the Buyer's and the Seller's utilities are concave. Next, we show that two well-known continuous and convex utility functions, the exponential and the power functions, also lead to a quasiconcave $n(x)$. Finally, we show that if the utilities are piece-wise linear then $n(\cdot)$ might not be a quasiconcave function, providing a characterization for the cases where this property holds.

3.1 Concave utility functions

Let us analyze first the case where u_B and u_S are both concave. This is the case of well-known utility functions such as the logarithmic function or the negative exponential function (e.g., $u_S(x) = \frac{1-e^{-\alpha x}}{1-e^{-\alpha}}$, with $\alpha > 0$) and corresponds to the common economic assumption of decreasing marginal utilities.

We will show that in these conditions the product of the utilities (8) is a quasiconcave function. As a preliminary step, note that $\log n(x) = \log u_S(x) + \log u_B(x)$ is a concave function in the domain where $u_S(x) > 0$ and $u_B(x) > 0$, since

- a) both $\log u_S(x)$ and $\log u_B(x)$ are concave when u_B and u_S are concave (the logarithm being a concave function as well);
- b) the sum of two concave functions is a concave function.

To show that $n(\cdot)$ is quasiconcave amounts to demonstrate that:

$$n(x) > n(y) \Rightarrow n(\lambda x + (1 - \lambda)y) \geq n(y), \forall \lambda \in [0, 1] \quad (9)$$

As we are working with non-negative utilities, if either $u_S(y) = 0$ or $u_B(y) = 0$ then $n(y) = 0$ and (9) trivially holds. In the domain where both utilities are strictly positive, $n(\lambda x + (1 - \lambda)y) < n(y)$ would imply

$$\log n(\lambda x + (1 - \lambda)y) < \log n(y) < \log n(x), \quad (10)$$

as the logarithm is a strictly increasing function. However this cannot be the true because $\log n(\cdot)$ is concave. Therefore, $n(\cdot)$ is quasiconcave.

3.2 Exponential functions

Let us now consider the utilities of the Seller and the Buyer in the domain $[0, 1]$ are exponential as defined next, for some positive values $s, b > 0$:

$$u_S(x) = (e^{sx} - 1)/c_S \quad (11)$$

$$u_B(x) = (e^{b(1-x)} - 1)/c_B \quad (12)$$

with $c_S = e^s - 1$ and $c_B = e^b - 1$ denoting two positive constants. Thus, for any $s > 0$, the Seller's utility increases from $u_S(0) = 0$ to $u_S(1) = 1$, whereas for any $b > 0$ the Buyer's utility decreases from $u_B(0) = 1$ to $u_B(1) = 0$.

We can write the product of utilities as:

$$n(x) = u_B(x)u_S(x) = (e^{b(1-x)} - 1)(e^{sx} - 1)/(c_S c_B)$$

This is a differentiable function with derivative

$$n'(x) = \frac{1}{c_S c_B} \left[s e^{sx} (e^{b(1-x)} - 1) - b e^{b(1-x)} (e^{sx} - 1) \right] = \frac{1}{c_S c_B} \left[s e^{sx} u_B(x) - b e^{b(1-x)} u_S(x) \right]$$

Since $e^{sx} = c_S u_S(x) + 1$ and $e^{b(1-x)} = c_B u_B(x) + 1$ we can rewrite:

$$\begin{aligned} n'(x) &= \frac{1}{c_S c_B} [s(c_S u_S(x) + 1)u_B(x) - b(c_B u_B(x) + 1)u_S(x)] \\ &= \frac{1}{c_S c_B} [s c_S u_S(x) u_B(x) + s u_B(x) - b c_B u_B(x) u_S(x) - b u_S(x)] \\ &= \frac{1}{c_S c_B} [s c_S n(x) + s u_B(x) - b c_B n(x) - b u_S(x)] \end{aligned}$$

We will now show that $n(x)$ is quasiconcave for any $s, b > 0$, i.e., even though both $u_B(x)$ and $u_S(x)$ are convex functions. First, we can note that $n'(0) = \frac{s}{c_S c_B} > 0$ and $n'(1) = \frac{-b}{c_S c_B} < 0$. Function $n(x)$ is initially increasing (positive derivative) then it attains a maximum and becomes decreasing (negative derivative). We need to show, however, that there cannot be multiple local maxima, i.e., that once $n(x)$ starts decreasing it will not increase again. To show this let us analyze the following function in the domain $]0, 1[$ where $n(x) > 0$:

$$\frac{n'(x)}{n(x)} = \frac{1}{c_S c_B} \left[s c_S + \frac{s}{u_S(x)} - b c_B - \frac{b}{u_B(x)} \right]$$

The above function is decreasing with x , because $u_S(x)$ is increasing and $u_B(x)$ is decreasing. Therefore when $\frac{n'(x)}{n(x)}$ becomes negative, it cannot become positive for larger values of x . But since $n(x)$ is strictly positive in the domain $]0, 1[$, this means that when $n'(x)$ becomes negative, it cannot become positive for larger values of x . Therefore, this demonstrates the quasiconcavity of $n(x)$: once $n(x)$ starts decreasing, it cannot increase for larger values of x .

3.3 Power functions

Consider the case of two individuals with utility function of the shape $u(x) = x^\alpha$ (which can be convex or concave depending on whether α is greater to or lower than 1). For any $\alpha > 0$, such a function grows from $u(0) = 0$ to $u(1) = 1$. Let us then define the utilities of the Seller and the Buyer in the domain $[0, 1]$ as follows, for some positive values $s, b > 0$:

$$\begin{aligned} u_S(x) &= x^s \\ u_B(x) &= (1-x)^b \end{aligned}$$

We can write the product of utilities as:

$$n(x) = u_B(x)u_S(x) = x^s(1-x)^b$$

This is a differentiable function in $]0, 1[$ with derivative

$$\begin{aligned} n'(x) &= s x^{s-1} (1-x)^b - b (1-x)^{b-1} x^s \\ &= x^{s-1} (1-x)^{b-1} [s(1-x) - bx] \end{aligned}$$

We will now show that $n(x)$ is quasiconcave, for any $s, b > 0$, i.e., even when both $u_B(x)$ or $u_S(x)$ are convex functions. As in the previous section, we need to show that there cannot be multiple local maxima, i.e., that once the $n(x)$ starts

decreasing it will not increase again. In other words, if $n'(x_0)$ is negative for some $x_0 \in]0, 1[$, then $n'(x)$ is still negative for $x > x_0$. To see this, let us note that $x^{s-1}(1-x)^{b-1} > 0$ for $x \in]0, 1[$, which means that the sign of $s(1-x) - bx$ will determine the sign of $n'(x)$. Therefore, if $n'(x_0)$ is negative for some $x_0 \in]0, 1[$, then $s(1-x_0) - bx_0$ must also be negative. Since $s, b > 0$, $s(1-x) - bx$ decreases with x , implying that $s(1-x) - bx < 0$ and consequently $n'(x)$ is negative for $x > x_0$.

3.4 Piecewise linear functions

Let us now analyze the case where $u_S(x)$ and $u_B(x)$ are piecewise linear functions, one increasing in x (Seller's utility), $u_S(x)$, and one decreasing in x (Buyer's utility).

Let us consider a point $x = b$, possibly a breakpoint in $u_B(x)$ and/or $u_S(x)$. Let $u_1 = u_S(b)$ and let $u_2 = u_B(b)$. Since $u_S(x)$ and $u_B(x)$ are piecewise linear functions, in the vicinity of b we can write, for some slopes $l_S, r_S, l_B, r_B \geq 0$:

$$u_S(x) = \begin{cases} u_1 - l_S(b-x), & \text{if } x \leq b \\ u_1 + r_S(x-b), & \text{if } x \geq b \end{cases} \quad (13)$$

$$u_B(x) = \begin{cases} u_2 + l_B(b-x), & \text{if } x \leq b \\ u_2 - r_B(x-b), & \text{if } x \geq b \end{cases} \quad (14)$$

Then, up to b we have:

$$n(x) = [u_1 - l_S(b-x)] \cdot [u_2 + l_B(b-x)] \quad (15)$$

$$n'(x) = l_S[u_2 + l_B(b-x)] - l_B[u_1 - l_S(b-x)] \quad (16)$$

$$n''(x) = -2l_Sl_B < 0 \quad (17)$$

From b onwards:

$$n(x) = [u_1 + r_S(x-b)] \cdot [u_2 - r_B(x-b)]$$

$$n'(x) = r_S[u_2 - r_B(x-b)] - r_B[u_1 + r_S(x-b)]$$

$$n''(x) = -2r_Sr_B < 0$$

This brief analysis allows seeing that the $n(x)$ is quasiconcave everywhere except possibly at certain breakpoints. Indeed, there is no minimum in any segment between two breakpoints, where $n(x)$ is the product of two positive linear functions, leading to $n''(x) < 0$. To verify if a product of piecewise linear functions, $n(\cdot)$, is quasiconcave therefore amounts to check only the breakpoints: observing the one-sided derivatives at each breakpoint $x = b$, on the left side we have $n'(b) = l_Su_2 - l_Bu_1$ and on the right side of b we have $n'(b) = r_Su_2 - r_Bu_1$. There exists a local minimum of $n(x)$ at b if and only if $l_Su_2 - l_Bu_1 \leq 0 \wedge r_Su_2 - r_Bu_1 \geq 0$ (one of the inequalities being strict). Thus, $n(x)$ is quasiconcave at b if and only if:

$$l_Su_2 > l_Bu_1 \vee r_Su_2 < r_Bu_1 \vee (l_Su_2 = l_Bu_1 \wedge r_Su_2 = r_Bu_1) \quad (18)$$

Let us note that if $l_S \geq r_S$ and $l_B \leq r_B$ then (18) holds. Therefore, for $b \in]0, 1[$ to be a local minimum requires $l_S < r_S$ ($u_S(x)$ convex at b) or $l_B > r_B$ ($u_B(x)$ convex at b). This means that not all breakpoints need to be checked: to verify if a product of piecewise linear functions, $n(\cdot)$, is quasiconcave it is sufficient to check if condition (18) holds for all breakpoints in $u_S(x)$ where this function is convex and all breakpoints in $u_B(x)$ where the latter is convex.

4 Simulation analysis

4.1 Motivation

The analytical models in the preceding section have provided several cases in which a unique global optimum of $n(x)$ exists. However, the proofs given there cannot be extended to other types of utility functions. For example, the proof in subsection 3.1 depends on the assumption that both utility functions are concave and does not necessarily hold otherwise. Likewise, we could show only for a few convex utility function that $n(x)$ is quasiconcave. This leaves two questions open:

1. Do multiple (local) optima exist for all functions that do not fall into the classes studied in Section 3?
2. If not, which properties of the utility functions make occurrence of multiple local optima more likely?

To answer these questions, we resort to a simulation study, in which we analyze the occurrence of multiple local optima for different, randomly generated utility functions. In this study, we still focus on specific classes of functions, which are particularly relevant as models of preferences. Obviously, in a buyer-seller setting, utility functions are monotonically increasing (for the seller) or decreasing (for the buyer) functions of price. Furthermore, we consider two specific shapes.

The first class are convex functions. Although risk aversion (and thus concave utility functions) is a common phenomenon, there are also decision makers who are risk seeking even for gains (Schoemaker, 1990; Pennings and Smidts, 2003) and thus have a convex utility function.

Furthermore, we also consider s-shaped functions, as proposed for example in Prospect Theory (Kahneman and Tversky, 1979). Prospect Theory postulates that utility functions (which are called value functions in Prospect Theory) are convex for outcomes below and concave for outcomes above a reference level. There is considerable empirical evidence also with respect to expected utility theory that decision makers are risk seeking in the domain of gains, and risk averse in the domain of losses (Schoemaker, 1990; Rieger et al, 2015). In contrast to expected utility theory, Prospect Theory defines “gains” and “losses” not in terms of positive or negative changes to final wealth, but relative to a reference point. The former interpretation would mean that in the present problem, all outcomes are seen as gains since there is no change if no trade takes place, and, by limiting prices to the zero-one interval, no party is worse off if trade takes place either. However, we consider it more realistic to assume that each party, at the beginning of the bargaining process, has some expectations about the outcome, for example a price that was reached in previous exchanges of the same or a similar good. For the Seller, a price below that expectation might be considered as a loss (relative to the

reference point), and a price above the reference point as a gain. For the Buyer, the opposite holds.

Prospect Theory not only assumes that utility functions are convex below and concave above the reference level, it also makes more specific assumptions about the shape of the utility function. In particular, it is assumed that decision makers are loss averse, i.e. that “losses loom larger than gains” (Kahneman and Tversky, 1979, p.279) and thus the utility function is steeper in the loss domain (see section 4.3 for a formal characterization of this property).

The simulation is based on the pointwise generation of utility functions of these general types, which are then multiplied at the supporting points to form $n(x)$. By comparing each such point of the product to neighboring points, we can identify local maxima and study their number and properties.

4.2 Generation of utility functions

The main goal of the simulation is to study the effect of different shapes of the two players’ utility functions on the occurrence of multiple local optima of $n(x)$. Since we want to approximate continuous functions, we need to generate a large number of points on the utility function in the interval under consideration. As already stated, we assume without loss of generality that payoffs as well as utilities are scaled to the $[0, 1]$ -interval.

For simplicity, we present the algorithm we are using to generate arbitrary utility functions of a given shape for the case of convex functions. An s-shaped utility function can be obtained by generating a convex function for one part of the domain, and a concave function for the remainder. The concave part can be generated by a straightforward adaptation of the algorithm, or by generating a convex function and then taking its mirror image.

For this exposition, we use the following notation: We represent the utility function by $m + 2$ points, which are equally spaced on the x axis representing the price. We denote the values on the x -axis by $X = (x_0 = 0, x_1, \dots, x_m, x_{m+1} = 1)$, and the corresponding utility values by $U = (u(x_i))$, where $u_0 = 0$ and $u_{m+1} = 1$. Thus, m additional utility values u_1, \dots, u_m have to be randomly generated.

Since we want to study arbitrary utility functions of the specific shapes, it is important that the generated functions cover the set of possible functions as completely as possible. Any point below the diagonal of the unit square (i.e., any point for which $u(x) \leq x$) can be a point on a convex utility function. Therefore, the utility functions we generate should cover the entire lower triangle of the two dimensional unit square.

This requirement rules out the straightforward approach to generate random convex functions, which would proceed as follows: Generate m random values on the zero-one interval and sort them. Denote the smallest value by r_1 , and the largest by r_m . Then assign

$$u_i = \sum_{j=1}^i r_j / \sum_{j=1}^m r_j \quad (19)$$

While this approach obviously generates a convex function, all functions generated by this approach will be very similar. The value assigned to $u(x_1)$ is always the

smallest of m randomly generated values, and thus likely to be very similar across runs, and this similarity continues along the entire function.

We therefore use a different approach, which is similar to the bisection method for arbitrary utility functions described in Dias and Vetschera (2018). In contrast to Dias and Vetschera (2018), we here generate utility values for a large number of equidistant points, and not for a small number of predefined attribute values, and we only generate functions of a specified shape.

In step k of the algorithm, the interval $[0, 1]$ is divided into 2^k subintervals of equal length. The function is evaluated at the endpoints of these subintervals.

The algorithm proceeds as follows. In the first step $k = 1$, $m = 1$, $X^1 = (0, 1/2, 1)$ and $U^1 = (0, r, 1)$, where r is an arbitrary utility value $r < 1/2$ (to keep a convex shape). In subsequent steps $k \geq 2$ of the algorithm, 2^{k-1} new points are generated so that at the end of step k , the utility function is represented by $2^k + 1$ points (including the end points $u(0) = 0$ and $u(1) = 1$). At the beginning of each step, new vectors X^k and U^k of length $2^k + 1$ are created and existing points are copied to the even-numbered positions in these vectors, i.e. $x_{2i}^k = x_i^{k-1}$ and $u_{2i}^k = u_i^{k-1}$, $\forall i \in \{1, \dots, 2^{k-1} - 1\}$. For each of these existing points, a slope s_{2i}^k between the slopes of the neighboring intervals is randomly generated:

$$s_{2i}^k = w_{2i}^k \cdot (u_{2i}^k - u_{2i-2}^k) + (1 - w_{2i}^k) \cdot (u_{2i+2}^k - u_{2i}^k) \quad \forall i \in \{1, \dots, 2^{k-1} - 1\} \quad (20)$$

where w_{2i}^k is a random number uniformly distributed between zero and one. Since all intervals are of equal length, we can represent the slope by the difference of utility values and s_{2i}^k is $x_{2i}^k - x_{2j-2}^k$ times the actual slope.

New points are placed in the middle of the intervals between existing points and receive odd numbers:

$$x_i^k = (x_{i-1}^k + x_{i+1}^k)/2 \quad \forall i \in \{1, 3, \dots, 2^k - 1\}. \quad (21)$$

For each of these points, a random utility value is generated between the lower bound

$$l_i^k = \max\left(u_{i-1}^k + s_{i-1}^k/2; u_{i+1}^k - s_{i+1}^k/2\right) \quad \forall i \in \{1, 3, \dots, 2^k - 1\} \quad (22)$$

and the upper bound

$$b_i^k = (u_{i-1}^k + u_{i+1}^k)/2 \quad \forall i \in \{1, 3, \dots, 2^k - 1\} \quad (23)$$

as

$$u_i^k = q_i^k \cdot l_i^k + (1 - q_i^k) b_i^k \quad \forall i \in \{1, 3, \dots, 2^k - 1\}. \quad (24)$$

In (24), q_i^k is a random value drawn from a uniform distribution in $[0, 1]$. The utility value u_i^k thus is constrained from above by the line connecting the neighboring points and from below by the two slopes at the neighboring points. This ensures that the slopes are increasing and the utility function has a convex shape. Figure 2 illustrates this approach graphically.

To avoid that the resulting function becomes too close to linear, the random values used in (24) can be drawn from a subset of the zero-one interval. In the simulations described below, we used the interval $[0.25, 0.75]$ for variable w_j^k in equation (20) and the interval $[0.8, 1]$ for q_i^k in equation (24). The high value of q_i^k in equation (24) guarantees a markedly convex shape of the function. Allowing u_i^k to be close to the straight line connecting u_{i-1}^k and u_{i+1}^k would force the entire function to be close to linear in that interval.




Fig. 2 Random generation of a convex function

4.3 Experimental setup

As explained in the introduction to this section, we want to study three types of utility functions:

1. Convex utility functions;
2. s-shaped utility functions based on the assumption of loss aversion, which use a concave function $u(x - r)$ for $x \geq r$ and $-\lambda u(r - x)$ for $x < r$, where $\lambda > 1$ is a parameter measuring loss aversion;
3. General s-shaped utility functions, which are composed of a concave function $u_1(x)$ for values $x > r$ above a reference level r and a different convex function $u_2(x)$ for $x < r$.

For all functions, we used 8,196 ($= 2^{13}$) intervals. We used the following approaches to generate the involved concave or convex functions:

1. A power function $u(x) = x^\alpha$;
2. A piecewise linear function generated by the algorithm described in 4.2 for $k = 2$ (i.e. 4 segments), using linear interpolation within these segments;
3. A general convex or concave function generated by the above algorithm for $k = 13$, i.e. each point was individually generated.

Combining these three function types with the three shapes gives 9 types of utility functions. Note that for the case of a convex power function, we have already shown in section 3 that $n(x)$ has only one (global) maximum. This already known result helps to verify correctness of the simulation implementation.

For each such setting, we generated 100,000 problems consisting of one monotonically decreasing utility function for the Buyer and one monotonically increasing utility function for the Seller. For the two types of s-shaped functions, we randomly generated reference points for buyers and sellers in the zero-one interval, which mark the boundary between the convex and the concave parts. All reference points

were rounded to the nearest subinterval boundary (multiple of $1/8196$) so that each s-shaped utility function was evaluated also exactly at the corresponding reference point to study how often maxima coincide with the reference point of one party. For buyers, values $x < r$ are considered as gains and therefore formed the concave part, and values $x > r$ (losses) are the convex part. For sellers, gains and losses are reversed.

4.4 Results

Table 1 provides a first overview of the average number and distribution of maxima for the different utility functions. As we have already shown analytically, the convex power function generates a unique maximum. However, this is not true for convex functions in general, in particular the piecewise linear function with just four segments has the highest average number of maxima of all functions studied. According to a Wilcoxon test, all differences between different shapes (within the same type) and different types (within the same shape) are significant.

Table 1 Distribution (in %) and average number of maxima for different utility functions

Shape	Type	N. Maxima				Mean
		1	2	3	> 3	
Convex	Piecewise	25.0	43.6	12.2	19.3	2.258
	General	81.8	10.3	4.1	3.9	1.346
	Power	100.0	0.0	0.0	0.0	1.000
Loss aversion	Piecewise	84.3	12.4	3.1	0.2	1.193
	General	78.0	12.0	4.8	5.3	1.434
	Power	79.7	18.9	1.5	0.0	1.218
S-Shaped	Piecewise	72.8	21.1	5.7	0.4	1.337
	General	63.7	22.6	6.9	6.8	1.640
	Power	63.5	32.0	4.4	0.0	1.409

Reference points are a characteristic feature of s-shaped utility functions. The question whether the location of the Buyer's and the Seller's reference points has an impact on the number of maxima is therefore of particular interest when studying these utility functions. Figures 3 and 4 show the number of maxima depending on the locations of both reference points. The two axes correspond to the location of the reference point for the Seller on the horizontal axis and the reference point of the Buyer on the vertical axis. In the area above the main diagonal of this square, the reference point for the Buyer is above that of the Seller, this means there exists an interval of prices which fulfill both expectations (since the Buyer expects to pay more than the Seller expects to receive).

Both figures show that multiple maxima are much more likely to occur when $r_S > r_B$, i.e., when the expectations of both sides are not compatible. In the converse case, when expectations are compatible, unique maxima are much more common. For the power function and the piecewise linear function exhibiting loss aversion, our simulation results contain no case with more than two maxima in which expectations are compatible, and the cases with two maxima happen only in extreme situations in which both parties have very high or both have very low




Fig. 3 Location of reference points and number of maxima (green=1, yellow=2, red ≥ 3), utility function with loss aversion




Fig. 4 Location of reference points and number of maxima (green=1, yellow=2, red ≥ 3), general s-shaped utility function

reference points (e.g., the situation depicted on the right part of Figure 1). These are also the regions in which the general function exhibits some cases of three or more maxima even for compatible expectations.

For general s-shaped functions, the situation is somewhat different. There are in general more cases with multiple maxima. For the power function, cases with three or more maxima occur at medium reference levels, when expectations are compatible. However, in these cases, the Buyer's reference point is only slightly above the Seller's reference point, so the range of prices compatible with both parties' expectations is rather narrow. A similar picture can be observed for the piecewise linear function. In contrast, for the general function, multiple optima occur mostly when both aspiration levels are extreme, as was the case for utility functions with loss aversion.

Compatibility of expectations thus seems to be a major issue determining whether multiple maxima of the Nash objective function exist. Table 2 analyzes this relationship in more detail.

The distributions of the number of maxima differ drastically when expectations are compatible and when they are not. For utility functions exhibiting loss aversion, the fraction of problems with a unique maximum increases from around 60% to well over 90%, and an even more dramatic change can be observed for general s-shaped utility functions. A χ^2 -test shows that these differences are all highly significant. Still, it has to be noted that compatibility of expectation does not

Table 2 Distribution of the number of maxima (in %) according to compatibility of expectations

Shape	Type	Compatible	Number of maxima				χ^2
			1	2	3	>3	
Loss aversion	Piecewise	No	68.7	24.6	6.2	0.5	*** 18368.9
		Yes	99.9	0.1	0.0	0.0	
	General	No	59.6	21.7	8.8	10.0	*** 19551.5
		Yes	96.2	2.4	0.8	0.6	
	Power	No	60.2	36.9	3.0	0.0	*** 23545.3
		Yes	99.2	0.8	0.0	0.0	
S-Shaped	Piecewise	No	53.3	35.1	10.9	0.7	*** 19694.9
		Yes	92.5	7.0	0.5	0.0	
	General	No	38.3	37.2	12.0	12.5	*** 28164.4
		Yes	89.1	8.0	1.8	1.1	
	Power	No	34.4	57.0	8.6	0.0	*** 36809.3
		Yes	92.8	7.0	0.2	0.0	

guarantee uniqueness of the maximum. There is still a number of cases with two, and for the general s-shaped utility function even three or more, maxima.

Compatibility of expectations not only changes the number of maxima, but also where the maxima are located. We distinguish between maxima that are located in the interval between the two reference points, which in the case of compatible expectations is also the range of prices that fulfills both expectations. In the case of incompatible expectations, this is the range of prices in which the expectations of neither party are fulfilled. The second group of possible locations are those prices that are above or below both reference points, i.e., outside that middle region. Here expectations of one party are fulfilled and expectations of the other party are not fulfilled. Finally, we consider maxima that are located exactly at one of the reference points.

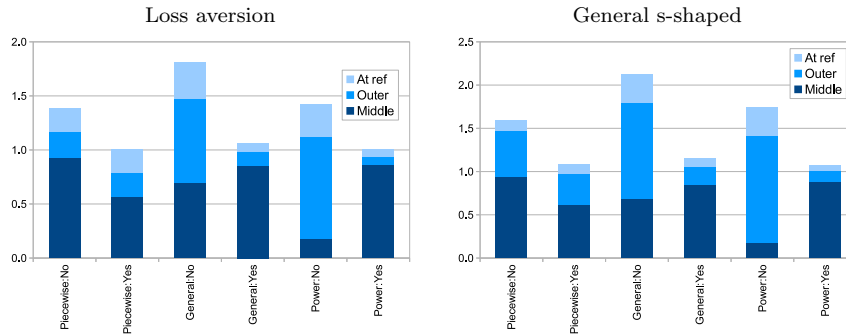
**Fig. 5** Distribution of locations of maxima according to compatibility of expectations (indicated by “:Yes” or “:No” after the name of the utility function)

Figure 5 shows how the distribution of maxima into these classes changes when expectations are compatible compared to when they are incompatible, considering loss-aversion (left) and general s-shaped (right) utility functions. Compatibility of expectations in all cases drastically reduces the number of maxima that occur

outside the interval between reference points. For the general and the power function, the number of maxima in between the reference points increases, partially offsetting the large decrease in maxima outside that interval. For the piecewise linear functions, it decreases, but the decrease in the number of maxima outside the interval is smaller, so the total decrease is about the same for all utility functions. In case of compatibility, there are also fewer maxima located exactly at a reference point.

5 Conclusions

The Zeuthen-Hicks bargaining model is a strategic bargaining model that describes a negotiation process between two parties leading to an axiomatic solution via successive concessions. Since each step in the process increases the product of utilities, the usual argument in literature is that it will end at the Nash bargaining solution, even if the parties have only local information about the utility functions. However, depending on the shape of the utility functions, the outcome may be a local optimum, rather than the true Nash solution (the global maximum).

In the present paper, we address this potential problem of convergence to a local optimum and thus missing the true Nash solution. This problem might not only affect the Zeuthen-Hicks bargaining model, but also any other strategic model that utilizes only local information and claims that the process will converge to the Nash bargaining solution. Our work therefore emphasizes a potential problem in linking strategic and axiomatic approach to bargaining, that so far has received little attention in literature.

In this work, we were able to characterize some types of utility functions for which the product of utilities has a single local maximum, which is therefore the global maximum: if the utility functions are both concave, or if they are both exponential functions or both power functions (even if convex), then there is a single optimum and the Nash solution will be reached. We also characterized the conditions for this to occur if the functions are piece-wise linear.

However, there are other, reasonable types of utility functions for which the possibility of multiple local optima exists. We therefore applied Monte-Carlo simulation to study the cases of general convex function and s-shaped functions and explore under which conditions multiple local optima are most likely to occur for such functions. For this purpose, we devised a method that generates a wide range of different functions with a given curvature (concave, convex, or s-shaped).

The simulations showed that the occurrence of 2, 3 or even more local optima is not uncommon, particularly for the 4-segment piece-wise linear and the general utility functions. Concerning shape, s-shaped functions have multiple local optima more frequently than convex functions. Considering the reference points in s-shaped functions, which reflect a priori expectations of each party, the results indicate that the relative location of these reference points has a significant impact on the number and location of local optima. When the expectations are compatible, unique optima are much more common. Nevertheless, the compatibility of the expectations does not guarantee the uniqueness of the optimum. When expectations are compatible, the optima tend to be located in-between the two reference points, but sometimes they can be located at one of them, or even outside the

interval defined by the reference points. When the expectations are incompatible, the number of maxima located outside that interval increases.

Our results provide some insights pertaining to Zeuthen-Hicks and similar bargaining processes, assuming the parties (or a mediator) have only local information about the utility function. For risk neutral or risk averse parties, the concavity of their utility functions ensures that any local maximum of the product of utilities will be the global maximum. The same happens if both parties have an exponential or a power utility function, even if one or both are convex. In these cases, a concession aiming at increasing $n(\cdot)$ will surely lead the parties closer to the Nash bargaining solution. If utilities are not concave, then a party might move farther away from the Nash solution (the global optimum) even when it makes a concession increasing $n(\cdot)$. This can happen when the offer goes in the direction of a local optimum inferior to the global maximum, as illustrated in Figure 1. Even if the parties perceive how the utility varies for both of them in the neighborhood of the offers on the table, they have no way of knowing they are converging towards a sub-optimal solution. If the parties (or a mediator) know their utility functions are s-shaped, and in particular if they have incompatible reference points, then they should be aware that multiple local optima are very likely to exist.

There are several possibilities how parties can reduce the danger of being trapped in a local optimum. One possibility is to make ambitious initial offers, surely below (for the Buyer) and surely above (for the Seller) the Nash global maximum. Subsequently, parties should make concessions in relatively small steps, in order to have a low risk of jumping from an offer x such that $n'(x) > 0$ (considering the Buyer's perspective) to an offer $x + \delta$ that also has $n'(x + \delta) > 0$, without realizing there was a maximum (possibly the global maximum) in-between x and $x + \delta$. Another possibility is to include randomization, to make sure that a large fraction of the search space is sampled and the process "jumps out" of a possible local optimum. However, such a strategy would need careful explanation and understanding by both parties, otherwise unmotivated and surprising changes in offers could be considered as a potentially unfair bargaining tactic aiming at confusing the opponent.

Our research therefore provides a first step towards understanding the potential problems that arise out of local optima in axiomatic bargaining solutions, and the implications such local optima have for the processes in strategic bargaining approaches. Of course, this wider topic can be generalized to other axiomatic solutions, and other process models. Even in the context of the Nash bargaining solution and the Zeuthen Hicks model, there is room for additional research considering e.g. other utility functions, or additional factors that could influence the number and occurrence of local optima. For instance, it would be interesting to study the effects of different concession strategies (not necessarily equal for the two parties) on the likelihood of ending the negotiation in a sub-optimal solution, in situations such that the shape of their utility functions does not guarantee the existence of a unique local maximum.

Acknowledgement

We thank three anonymous reviewers for their helpful comments that contributed to improving the paper.

References

- Bishop RL (1964) A Zeuthen-Hicks theory of bargaining. *Econometrica* 32:410–417
- Chatterjee K (2010) Non-cooperative bargaining theory. In: Kilgour DM, Eden C (eds) *Handbook of Group Decision and Negotiation*, Springer, Dordrecht, book section 8, pp 141–150
- Coddington A (2010) *Theories of the Bargaining Process*. Routledge Library Editions. Economics. Keynesian & Post-Keynesian Economics, Routledge, London
- Crawford VP (1980) A note on the Zeuthen-Harsanyi theory of bargaining. *The Journal of Conflict Resolution* 24(3):525–535
- Dias LC, Vetschera R (2018) On generating value functions in stochastic multicriteria acceptability analysis. forthcoming
- Fandel G (1985) On the applicability of group decision-making concepts to wage bargaining. In: Haimes YY, Chankong V (eds) *Decision Making with Multiple Objectives*, Springer, Berlin et al., pp 532–548.
- Filzmoser M, Hippmann P, Vetschera R (2016) Analyzing the multiple dimensions of negotiation processes. *Group Decision and Negotiation* 25:1169–1188
- Harsanyi JC (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
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291
- Kersten GE, Lai H (2010) Electronic negotiations: Foundations, systems, and processes. In: Eden C, Kilgour DM (eds) *Handbook of Group Decision and Negotiation*, Springer, pp 361–392
- Kíbris O (2010) Cooperative game theory approaches to negotiation. In: Kilgour DM, Eden C (eds) *Handbook of Group Decision and Negotiation*, Springer, Dordrecht, book section 9, pp 151–166
- Kilgour DM, Eden C (2010) Introduction to the handbook of group decision and negotiation. In: Kilgour DM, Eden C (eds) *Handbook of Group Decision and Negotiation*, Springer, Dordrecht, pp 1–7
- Nash JF (1950) The bargaining problem. *Econometrica* 18(2):155–162
- Pennings JME, Smidts A (2003) The shape of utility functions and organizational behavior. *Management Science* 49(9):1251–1263
- Rieger MO, Wang M, Hens T (2015) Risk preferences around the world. *Management Science* 61(3):637–648
- Rubinstein A (1982) Perfect equilibrium in a bargaining model. *Econometrica* 50(1):97–109
- Saraydar E (1971) Uncertainty, the bargaining problem, and the Nash-Zeuthen solution. *Theory and Decision* 1(3):309–319
- Schoemaker PJ (1990) Are risk-attitudes related across domains and response modes. *Management Science* 36(12):1451–1463
- Sutton J (1986) Non-cooperative bargaining theory: An introduction. *The Review of Economic Studies* 53(5):709–724
- Thompson W (1994) Cooperative models of bargaining. In: Aumann RJ, Hart S (eds) *Handbook of Game Theory with Economic Applications*, vol II, Elsevier, Amsterdam, pp 1238–1284
- Vetschera R (2006) Preference structures of negotiators and negotiation outcomes. *Group Decision and Negotiation* 15:111–125

- Vetschera R (2007) Preference structures and negotiator behavior in electronic negotiations. *Decision Support Systems* 44(1):135–146
- Vetschera R (2013) Negotiation processes: an integrated perspective. *Euro Journal on Decision Processes* 1(1-2):135–164
- Wagner RH (1979) On the unification of two-person bargaining theory. *The Journal of Conflict Resolution* 23(1):71–101