Economic Theory manuscript No.
(will be inserted by the editor)

Bilateral Bargaining, Unverifiable Quality, and Options to Return*

Anke S. Kessler[1], **Christoph Lülfesmann**[2]

[1] Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, B.C., V5A 1S6, Canada, e-mail: akessler@sfu.ca
[2] Department of Economics Simon Fraser University, 8888 University Drive, Burnaby, B.C., V5A 1S6, Canada, e-mail: cluelfes@sfu.ca

Received: date / Revised version: date

Summary. The paper investigates an alternating-offers bargaining game between a buyer and a seller who face several trading opportunities. These items (goods or services) differ in their non-verifiable quality characteristics which gives rise to a moral hazard problem on the seller's part. For the special case of two goods, we completely characterize the set of subgame-perfect equilibria. We find that the seller always extends an option to return the good, while the buyer may suffer from this warranty. Also, qualitatively different types of equilibrium outcomes occur depending on the parameters of the model: (a) the seller may obtain a larger share of the surplus although the parties ex ante have symmetric bargaining positions, (b) the subgame-perfect equilibrium may entail inefficient trade, and (c) multiple equilibria may exist including equilibria with delay in negotiations. Finally, we analyze a situation where bargaining proceeds after the good was returned which is shown to reestablish uniqueness and efficiency of equilibrium.

Keywords and Phrases: Bilateral Bargaining, Unverifiable Quality Warranties
JEL Classification Numbers: C78, L14, L15, D82

* This paper has greatly benefitted from discussions with Avner Shaked and Timothy von Zandt. We also wish to thank Wolfgang Leininger, Zvika Neeman, Clemens Puppe, Wolfram Richter, Karl Schlag, Ilya Segal, and seminar participants in Dortmund, Bonn and Berkeleyfor helpful comments and discussions. Financial support by Deutsche Forschungsgemeinschaft, SFB 303 at the University of Bonn is gratefully acknowledged. Remaining errors are our own.
Correspondence to: Christoph Lülfesmann

1 Introduction

Bargaining theory is concerned with the question of how economic agents reach a mutually beneficial agreement. The basic non-cooperative framework as laid out by Rubinstein [21] posits a sequential bargaining process with alternating offers between two players over an infinite time period and discounting. In this model and the subsequent literature, it is usually assumed that the surplus over which the parties have to reach an agreement is fixed. In the context of a trade relationship between a buyer and a seller, for example, little thought has been given to the possibility that the parties face *several*, mutually exclusive trading opportunities (goods or services). If all parties are fully informed about the characteristics of these trading opportunities, one may be tempted to conclude that this possibility is irrelevant: rational parties should simply agree to trade the commodity which yields the highest net surplus. This argument, however, is only correct if the allocational problem they face (which surplus to share) is independent of the distributional problem (how to share the surplus).

In the present paper, we consider a natural framework in which this condition is not satisfied. Two parties, a buyer and a seller, bargain over the price of a commodity. The seller has various commodities at her disposal which differ in their respective characteristics. An obvious interpretation of this situation is that the commodity to be traded can be delivered in different qualities. While we assume that both parties have symmetric and complete information on the goods' characteristics, these are non-verifiable and, hence, cannot be contracted upon. This lack of contingencies implies that - as in standard bargaining models - the players are constrained to making pure price offers. In some sense, the notion of non-contractibility gives rise to an 'intermediate' information structure in which the parties are fully informed but they are unable to use their knowledge directly during the course of the bargaining game. The two features of the model, several trading opportunities and non-verifiability of quality, cause an additional conflict of interest: at any given trade price, the seller prefers to trade a good with low production costs whereas the buyer is interested only in high quality. Hence, the parties do not only disagree with respect to the distribution of payoffs for a given commodity but also on which commodity to trade, i.e., which surplus to share.

In order to facilitate the comparison of our results with the existing literature,we will stay within the alternating-offers bargaining framework as set up in Rubinstein [21], where we assume for simplicity that the two parties have symmetric bargaining positions (discount factors) and the time interval between two consecutive offers is very small. In each bargaining round, one of the parties offers a contract specifying a trade price and, possibly, a return policy. Once an agreement is reached, the seller can deliver a good which the buyer either accepts or returns to the seller. In the latter case, trade does not take place and payments are made according to the contract, i.e. the buyer may have the option to return the good without having to pay the agreed upon trade price under customer return policy. Two variants of this basic structure are analyzed. In a first version, the game ends immediately after the trade stage if trade has not been successfully completed, i.e., if the buyer has rejected delivery. This modelling may be appropriate in situations where there is a time constraint on final trade (e.g. the good is perishable). In this case, we demonstrate that (a) the seller enjoys a strategic advantage over the buyer in the sense that she captures more than half of the surplus, (b) a subgame perfect equilibrium may entail inefficient trade, and (c) multiple equilibria can exist including equilibria with delay in negotiations. Thereafter, we analyze a second scenario where the parties are allowed to renegotiate after an unsuccessful trade stage. In contrast to the first variant, we accordingly assume that bargaining may commence anew once the buyer has rejected delivery. We show that renegotiation restores uniqueness and efficiency of equilibrium, and implements equal sharing of the (largest) surplus.

While a unique subgame-perfect equilibrium prevails in the standard Rubinstein bargaining framework, our results are in line with a number of recent papers which show that multiple and inefficient equilibria can arise even in bargaining situations with complete information. Shaked [22] demonstrates that if one player is allowed to opt out each time an offer has been rejected, multiple equilibria and delay may occur. Muthoo [18, 19] studies an extension of the standard bargaining model where the proposer can retract an offer after it has been accepted by his opponent. If the players are sufficiently impatient, there exists a unique subgame-perfect equilibrium that coincides with that in the Rubinstein framework. Otherwise, however, any partition of the cake and any arbitrary delay can be supported as subgame-perfect equilibrium outcome although no offer is actually retracted. In Busch and Wen [4], the agents play a static 'disagreement game' in between bargaining rounds which endogenously determines their inside options. They show that there exist as many stationary equilibria as there are equilibria in the disagreement game. In addition, there may exist multiple non-Markov equilibria even if the disagreement game has a unique Nash equilibrium. Finally, Avery and Zemsky [3] establish multiplicity of equilibria when a player can prevent his opponent from submitting a counteroffer for some period of time (i.e., if she can 'burn money'). This paper also shows that there exist inefficient equilibria with equilibrium delay where 'money burning' in fact arises.[1]

Although we confine our attention to a specific trade relationship, it will become clear below that our results are more generally applicable to situations where allocative efficiency and distributional aspects are intertwined, as is often the case in the presence of moral hazard. Furthermore, we will see that the seller often offers a customer return policy in equilibrium. In-

[1] In models that combine a bargaining process with the structure of an infinitely repeated game, Haller and Holden [11] and Fernandez and Glazer [10] also show multiplicity of equilibria and delay. For an excellent overview of the bargaining literature, see Muthoo [19].

tuitively, offering to reimburse the buyer if he is 'dissatisfied' with the good helps to commit herself not to deliver low quality: it insures the buyer against fraudulent behavior by the seller, and relaxes the seller's incentive constraint. As our results show, a warranty provision in fact often restores efficient trade of a high-quality item thereby solving the moral hazard problem that otherwise leads to allocative inefficiencies. Not surprisingly, we find that the money back guarantee (at least weakly) raises the joint surplus in any equilibrium of the game. Also, the option to return has a positive effect on the seller's equilibrium payoff so that she always unilaterally extends it. Perhaps paradoxically and in contrast to common beliefs, however, we find that the *buyer* may suffer from such a policy.

Empirically, customer return policies and warranties are predominant in the commercial business. For example, Che [5] reports data from a recent survey on Illinois retailers. According to this study, 78 percent of those retailers give cash refunds with a receipt. Almost all electronic and catalog retailers allow return of their products within a certain period of purchase. In the insurance industry, the client is in many countries granted a legal right to step back from a signed contract within a certain period of time. In the US, consumers are covered by an implied 'warranty of merchantability' even without a written warranty. An important part of the theoretical literature on warranties and money-back guarantees emphasizes incentive effects.[2] In line with our results, these papers show that that warranties alleviate the problem of seller moral hazard (see, e.g., Cooper and Ross [6], Mann and Wissink [17].[3] However, and in contrast to the present bargaining approach, this literature presumes competitive product markets, so that allocational and distributional issues are logically separated and consumers necessarily reap all the efficiency gains of warranty provisions. While we do not doubt that a competitive framework is often appropriate, there exist important markets that are characterized by extensive bargaining between buyer and seller. In computer retailing, for example, hardware components are often assembled for a buyer's specific demand, and the realized price results from a bargaining process between the parties involved. In the insurance business, the terms of contract are usually negotiated between the contracting parties. More generally, bargaining between buyer and seller is the rule rather than the exception in situations where the trading item is tailor-made, as is often the case in business-to-business transactions.

Our approach also draws on the literature on incomplete contracts which assumes that some important variables or contingencies cannot be con-

[2] For a comprehensive survey, see Emons [8]. Other branches of this literature analyze insurance (see Emons [7]) or signalling (see Che [5]) issues.

[3] In Klein and Leffler [16], the concern for reputation can induce a seller to charge a 'quality-assuring' price and to deliver high-quality goods. Hence, their argument is similar to that made in the efficiency wage literature. Since the model is cast in a competitive equilibrium framework, however, the quality-assuring price is not generally compatible with the zero profit condition. See also Allen [2] and Shapiro [24].

tracted upon (see Hart [12] and Tirole [25] for comprehensive treatments and assessments of the incomplete contracting paradigm). Contractual incompleteness may arise for several reasons: it may be too costly to write contracts which specify all relevant contingencies,[4] or the court as the enforcing parties may lack the technical expertise to evaluate all relevant characteristics even if the parties signed a fully contingent contract. In the models analyzed in the literature, contractual incompleteness often refers to an ex-ante stage after which investments are expended and final trade is executed. Even if the object of trade cannot be identified ex ante, it is usually assumed that conditional contracts can be written at the trade stage, so that an ex-post efficient outcome prevails for any given investment level. In contrast, our paper sheds some light on the additional problem that may arise when verification problems prevent a complete description of the item at the trade stage. Interestingly, incomplete contracting can also endogenously emerge in our framework since the seller may optimally submit a non-contingent proposal. While a contingent offer would trigger trade of the efficient good at its Rubinstein price, we find that the seller's equilibrium profit is often higher when her contract offer remains incomplete. According to our results, this behavior is attractive for the seller in economic situations where the buyer himself is unable to describe all relevant aspects of the terms of trade in an offer to the seller.[5]

The paper proceeds in Section 2 with a description of the model, and an analysis of the trading stage. Section 3 completely characterizes the set of subgame perfect equilibria for the case of two goods. The subsequent subsection 4 considers the possibility of ongoing negotiations after a rejection at the trade stage. A final section 5 concludes and briefly discusses further applications.

2 The Model

Consider a risk-neutral buyer (B) who bargains with a risk-neutral seller (S) over the price of a particular item or commodity he wants to purchase. The seller has several commodities at her disposal which may differ in their

[4] In the context of a hold-up problem, Hart and Moore [13] and Aghion and Tirole [1], for example, argue that quality may be non-contractible because the exact characteristics of goods are too complicated to be fully specified in a contract. In a similar spirit, Fama [9] and Holmström [14] assume that the performance of employees is observable but non-contractible due to measurement costs.

[5] For example, consider the contracting process between an insurance provider and a customer about the terms of an insurance contract. Usually, the insurance agent advises his client and the parties agree upon some key contractual characteristics and the insurance premium (the trade price). Upon agreement, the insurance company mails a written contract to its client that details all contractual provisions. In many cases, the customer only then learns the exact 'terms of trade' while having the option to step back from the agreement within a certain period.

respective characteristics, e.g. their quality and costs of production. We assume that both parties have symmetric and perfect information concerning these attributes which are, however, not contractible. The latter assumption implies that a sales contract between the two parties must necessarily remain incomplete. In particular, a trade price which is agreed upon cannot be made contingent on the specific good to be traded. Given that the goods' characteristics are non-verifiable, a potential conflict of interests arises if the buyer prefers to trade a commodity which yields a high benefit whereas the seller prefers to trade a commodity with low production costs.

2.1 A basic framework of seller moral hazard

Suppose that there are two goods which the seller can potentially produce and deliver.[6] Let $b_i > 0$, $c_i \geq 0$, $i \in \{L, H\}$ be the buyer's gross benefit from purchasing and the seller's cost of producing good i, respectively. Without loss of generality, let $s_L \equiv b_L - c_L < b_H - c_H \equiv s_H$, i.e. good H yields a higher net surplus s and should be traded from an efficiency point of view. In order to introduce conflicting interests between the buyer and the seller we assume that $c_H > c_L$. Together, these two inequalities imply $b_H > b_L$. This is a natural specification if, for instance, good H is of superior quality (and thus has higher production costs) than good L. In the following, we will therefore frequently refer to these goods as the high-quality and low-quality good, respectively. For future reference, also define $\beta \equiv b_H - b_L > 0$ as the difference in the buyer's gross valuations of the two goods.

Trade between B and S is formalized as a two-stage game in which the seller first chooses which good to deliver and the buyer next chooses whether to accept the good or return it to the seller. Suppose first a price p has been agreed upon in the bargaining stage (see below). The buyer's payoff if trade of good i takes place is then $b_i - p$. Clearly, if p has to be paid regardless of whether the good is returned or not, it is a dominant strategy for the buyer to accept delivery (recall $b_i > 0$). But since $c_H > c_L$, the seller will then always deliver the low quality good. To avoid this inefficient outcome, the parties may include a customer return policy (money-back guarantee) m in the sales contract which specifies the refund m the buyer is entitled to if he is 'dissatisfied' with the product. Note that since quality is non-verifiable, the contract must remain silent concerning the exact cause of the customer's dissatisfaction, i.e., there are 'no questions asked' if a customer complains (we believe that this frequently observed practice can precisely be attributed to the parties' difficulties of specifying the relevant characteristics in the sales contract). Under such a return policy, the buyer

[6] The restriction to two commodities is made for simplicity. Alternatively, one could assume that there is a large number of commodities of which only two are 'relevant' in the sense that they yield a non-negative surplus. Under some additional and unrestrictive assumptions, our results also remain qualitatively unaffected for any finite number of valuable goods.

pays upon return of the item the price minus the refund and his utility is $m - p$. Hence, he will accept good i if and only if $b_i - m \geq 0$. Next, consider the seller's decision which commodity to deliver. If $m \leq b_L$, i.e. if the buyer accepts good L, it is optimal for her to trade the low-quality good due to its lower production costs. If $b_H \geq m > b_L$, however, the buyer would return a low-quality item and only accept the high-quality good. Put differently, $b_H \geq m > b_L$ is an incentive constraint which must be satisfied if the high-quality good is to be traded. Hence, for a given contract (p, m), the subgame perfect equilibrium of the trading game is

$$x_H = 1 \Rightarrow b_L \leq m \leq b_H \tag{1}$$

$$x_L = 1 \Rightarrow m \leq b_L \tag{2}$$

$$x_i = 0 \qquad \text{otherwise}$$

where x_i is a binary variable with $x_i = 1$ if good i is traded and $x_i = 0$ otherwise. Note that if m could be chosen sufficiently high, there always exists a refund policy that ensures trade of the high-quality good. In order to rule out a trivial solution to the moral hazard problem, we therefore assume that the seller is protected by limited liability, i.e., no refunds in excess of the price are possible ($m \leq p$).[7] A full money-back warranty then corresponds to setting $m = p$. We will suppose that if the buyer exercises such an option to return, both parties obtain their reservation utilities.[8]

2.2 The Bargaining Procedure

In the spirit of Rubinstein [21], the bargaining process between the buyer and the seller is modelled as a sequential game of alternating offers. The question we address is what the outcome of this game and the subsequent trading stage will be. Specifically, we ask under which circumstances the fact that the parties' information is non-verifiable imposes a constraint on the efficiency of the outcome of the overall game. Bargaining takes place in

[7] For instance, the seller may simply be wealth constrained, which is the common assumption in moral hazard frameworks with risk-neutral parties. Alternatively, one could endogenize this assumption by considering an incomplete information framework. Suppose that there exists another type of buyer with with zero valuation for either good, and the seller is uninformed on the type of buyer she faces. Since there are no gains from trade with 'fraudulent' buyers that have zero valuation, she will be unwilling to offer negative no-trade prices (or decline any offer by the buyer containing such a clause) if the fraction of those types in the market is sufficiently large.

[8] This assumption is made in order to simplify the subsequent analysis, and implies that the seller does not incur production costs before delivery (as, e.g., in case of an insurance contract). As will become clear below, however, our subsequent results in the basic framework equally apply if the seller sinks production costs: the seller's payoff does not affect the buyer's acceptance decision, and the seller will never sign a contract that does not trigger trade in continuation equilibrium.

an infinite number of discrete points in time indexed by $t = \{0, \Delta, 2\Delta, ...\}$, where $\Delta > 0$ is the time that elapses between two consecutive offers. At even times t, the buyer makes a proposal (p, m) that the seller can immediately accept (A) or reject (R). If the offer is accepted, the bargaining ends and the outcome of the trade stage above is implemented. If the seller rejects the buyer's offer, the game passes on to the next round $(t+1)\Delta$ where it is her turn to propose an agreement. For simplicity, both players are assumed to discount the future with a common discount factor $\delta = exp(-r\Delta)$, where $r > 0$ is the discount rate. The buyer's objective is to maximize his net benefit from purchasing, $u_B = \delta^t(b_i - p)$ and the seller maximizes her net profit which is given by $u_S = \delta^t(p - c_i)$. For simplicity, the respective reservation (no trade) payoffs of the two parties are normalized to zero.

3 Equilibrium Analysis

The present framework introduces two aspects which are likely to alter the bargaining equilibrium relative to the standard solution of a Rubinstein bargaining game. First, (1) and our limited liability assumption imply that there is a lower bound on the equilibrium price if the high-quality good is to be traded: for a maximal refund $m = p$, we must have $p \geq b_L$. This constitutes an upper bound on the buyer's maximum payoff from trade of good H that is equal to $b_H - b_L = \beta$. Naturally, the seller's payoff is then bounded from below in any such equilibrium. In this regard, the seller's threat of delivering the low-quality good has an effect that is qualitatively similar to that of an outside option. Second, there exists a different commodity which can be traded (good L) and each party has the option of (implicitly) proposing to trade this alternative good (in which case no refund policy is needed). If the characteristics of the two commodities were contractible, it is easy to see that this option is irrelevant. Given that the buyer's payoff from good H is bounded from above, however, it may be attractive for one of the parties to propose an agreement on good L.

[Figure 1 about here]

The bargaining set is depicted in Figure 1, which assumes a full money-back warranty $m = p$. The figure illustrates that the seller's incentive to deliver low quality gives rise to a bargaining frontier that is not closed - a prevalent feature in many bargaining situations. As already noted above, any moral hazard problem with a risk-neutral but wealth-constrained agent, for instance, will give rise to a similar situation. Thus, our results can generally be applied to situations in which the need to set incentives correctly gives rise to an a discontinuous interdependency between efficiency and distributional concerns.

The following proposition characterizes the equilibria in our framework. All derivations are relegated to the appendix.

Proposition 1 *As the time interval Δ between two consecutive offers tends to zero:*

a) If $\beta \geq \frac{1}{2}s_H$, the game has a unique SPE which involves immediate agreement on a trade price $p^* = c_H + \frac{1}{2}s_H$ and a full money-back warranty $m = p^*$; in this equilibrium, the high quality good is produced and traded.
b) If $\frac{1}{2}s_L \leq \beta < \frac{1}{2}s_H$, the game has a unique SPE with immediate agreement on a trade price $\bar{p} = b_L > p^*$ and a full money-back warranty $m = \bar{p}$; again, the high quality good is produced and traded.
c) If $\beta < \frac{1}{2}s_L$, the game has multiple equilibria in which any buyer payoff between β and $\frac{1}{2}s_L$ can be supported as an SPE. In particular, there exists an equilibrium where the low-quality good is produced and traded at a price $\tilde{p} = c_L + \frac{1}{2}s_L$, without a warranty. Furthermore, equilibria with delay in negotiations also exist.

The SPE described in part a) replicates the Rubinstein solution for a cake of size s_H and has a simple intuition: $\beta \geq \frac{1}{2}s_H \Leftrightarrow p^* \geq b_L$ implies that for a warranty agreement with a price $(p^*, m = p^*)$, the seller delivers the high-quality good which the buyer accepts (in order for S not to deliver good L at a price p_L, the equilibrium strategy for B at the trade stage must be to reject good L if he is indifferent between trade and no trade). One can thus think of the restriction on β as a condition under which the upper bound on the buyer's maximum payoff from trade of good H is not binding at the Rubinstein solution. Since good H generates a higher surplus, it is also immediate that neither of the parties can increase its payoff by making an alternative offer $p < b_L$ for which good L would be traded. Uniqueness now follows from the standard arguments.

Part b) applies to lower values of β where p^* would induce trade of the low-quality good, even with a full warranty. Hence, the seller's option of delivering good L is binding so that in any equilibrium where the high quality good is traded, the price must be at least b_L. The buyer's payoff in such an SPE equals his upper bound payoff from the high-quality good, β, and is reduced relative to a situation where the seller's option is not binding [part a)]. Conversely, the seller's payoff has increased. Uniqueness of this SPE is ensured by the condition $\beta > \frac{1}{2}s_L$, which implies that the buyer's payoff still exceeds his maximum payoff from trade of the low-quality good, which equals half of the corresponding surplus s_L.

Once β falls below $\frac{1}{2}s_L$, there are multiple equilibria. In particular, there is a SPE in which the buyer obtains $\frac{1}{2}s_L$ and the low quality good is traded at price $\tilde{p} = c_L + \frac{1}{2}s_L < b_L$. This equilibrium replicates the Rubinstein solution in a situation where the parties bargain over a cake of size s_L. Hence, everything is *as if* the high-quality good was not available. To see why this is an equilibrium, note that if the buyer sticks to his equilibrium strategy he can (in the limit) obtain $\frac{1}{2}s_L$ by rejecting all offers and proposing $\tilde{p}$ in the next period. But if S proposes an agreement on good H, the minimum price she can ask for is b_L in which case B's utility would be β. Thus, if $\beta < \frac{1}{2}s_L$, the buyer (credibly) rejects all feasible price offers by the seller that induce trade of the high-quality good. At the same time, however, an agreement on the high-quality good with a price $\bar{p} = b_L$ and a warranty $m = \bar{p}$ continues to be an equilibrium: suppose the seller always proposes a price $\bar{p} = b_L$ and grants a full refund, which gives her $s_H - \beta$ and the buyer β. Then, if B were to propose an agreement on good L, he had to offer (in the limit) at least $p = s_H - \beta$ for this offer to be accepted by S. But because $s_L - p < \beta$, this deviation is not profitable.

Intuitively, if the seller 'insists' to trade good H, the buyer can no longer guarantee himself his Rubinstein share of the surplus of good L. Conversely, if the buyer 'insists' on a relatively high equilibrium payoff from good L, the seller can do no better than likewise proposing an agreement on this good. Since the buyer is strictly better off in the latter equilibrium and the seller is strictly better off in the former equilibrium, they can be used as punishments to construct a whole range of equilibrium payoffs in this game: whenever the seller deviates from her presumed equilibrium strategy, the strategies call for both players to play thereafter according to the SPE which gives the seller his lowest payoff, $\frac{1}{2}s_L$. Similarly, any deviation of the buyer is punished by a rejection of the seller and subsequent play of the equilibrium which gives the buyer his lowest payoff, β. By the same token, one can also generate equilibria with delay in negotiations. There, the strategies call for both players to make unreasonably low (high) offers which are not accepted. If some player deviates and attempts to reach an early agreement, he is punished by having the play switch to some equilibrium which is most unfavorable for him.[9]

Figure 2 displays the set of payoffs to the buyer where an agreement is reached in the first period as a function of β, holding s_H and s_L fixed. The subgame perfect equilibrium is unique if β is sufficiently high. For low values of β, multiple equilibria exist (the dotted area). Unless $\beta \geq \frac{1}{2}s_H$, the buyer is worse off as compared to a setting in which the commodities' characteristics are contractible or, equivalently, in which good L is not available. Then, the high-quality good would always be traded and his equilibrium payoff would be determined by the Rubinstein solution, irrespective of β. A similar picture for the seller, however, would reveal that she may or may not be better off compared to a situation where no moral hazard problem exists. In particular, the seller strictly gains from her threat of delivering low-quality in an equilibrium where the high-quality good is traded at $p = b_L$, but not in an equilibrium where the low-quality good is traded.

[Figure 2 about here]

These features are qualitatively preserved when the number of trading opportunities increases: a unique and efficient equilibrium only exists if the

[9] The line of reasoning here is standard. For similar arguments, see Muthoo [18], [22], [3] and Busch and Wen [4].

incremental increase in value between two neighboring goods that are ordered according to net surplus is maximal at the efficient good (note that this condition is quite strong as it rules out monotone marginal returns in the production of quality). As the number of goods goes to infinity, an efficient equilibrium exists generically only if the net values of the highest-quality and the lowest-quality good are negative, respectively. In this case, the buyer cannot obtain a positive surplus in any SPE. However, even then the efficient SPE is usually not unique (see Kessler and Lülfesmann [15]).

Figure 2 also allows an answer to the question whether seller and buyer gain from an option to return. If the seller could not offer this option (e.g., by law), the buyer is forced to accept delivery and both parties will rationally anticipate that good L is traded, independent of the trade price p. As is easily seen, the parties then share s_L according to the Rubinstein solution in the unique equilibrium. Thus, the seller gains from granting a money-back guarantee that facilitates her commitment to deliver high quality, and therefore unilaterally extends this warranty. Perhaps surprisingly, the buyer in contrast does not necessarily gain from the possibility to (costlessly) reject delivery. Whether his payoff increases or not depends on which equilibrium is played as can be seen from Figure 2. In the parameter range where multiple equilibria exist, the buyer is always (weakly) hurt by his option to return.[10] Intuitively, if the seller can offer a customer return policy, the buyer may not capture his 'fair share' of the low-quality good, because the seller can (credibly) refuse the buyer's proposal, offer a higher price and deliver the high-quality good, which gives her a strictly larger payoff. Thus, a customer return policy can actually be *harmful* for the buyer in situation where the terms of trade are determined in bilateral negotiations.

4 Renegotiation

In this section, the original bargaining game is modified by allowing the parties to renegotiate the trade price after the seller has delivered a good which the buyer has rejected. In this renegotiation process, the sequence of offers and counteroffers is left unaltered with the exception of the identity of the first mover. Specifically, if the buyer has made the last proposal which was followed by the seller's acceptance and an unsuccessful trade stage, time moves on to the next period and it is now the seller's turn to make the first offer in renegotiations.

The important difference is the outcome of the trade stage. Formerly, we argued that the buyer agrees to trade good i if and only if his net benefit from trade exceeds his no-trade utility. Hence, under a full refund policy, the buyer accepted the delivery of good i if and only if $m = p \leq b_i$. Together

[10] For non-vanishing time intervals between offers (values of $\Delta > 0$), there exists a non-empty set of parameter values such that the buyers's payoff in the unique equilibrium with customer return policy is lower than in the equilibrium without such a warranty.

with the seller's decision which good to deliver, we thus had $p \geq b_L$ as a necessary condition on the equilibrium price if the high-quality good was to be traded and a corresponding upper bound of the buyer's utility of β. If the parties can renegotiate, however, B will in his decision whether to accept delivery compare the benefit from accepting immediately, $b_i - p$, to the discounted utility he expects to obtain in the renegotiation process. Accordingly, the seller's incentive constraint is relaxed. The implications of this effect are summarized in the proposition below.

Proposition 2 *As $\Delta \to 0$, the game with renegotiation has a unique SPE for all values of $\beta > 0$. This equilibrium involves immediate agreement on a trade price $p^* = c_H + \frac{1}{2}s_H$ with full money-back warranty $m = p^*$, and the high quality good is produced and traded.*

Hence, the possibility of renegotiation reestablishes uniqueness and efficiency of the subgame perfect equilibrium, although the contract is never actually renegotiated. To see why this must be true, suppose time is no matter of concern for the buyer ($\Delta \to 0$). Obviously, then, he never accepts the low-quality good since he can always renegotiate and propose an agreement to share the larger surplus according to the Rubinstein solution. The crucial difference to the game without renegotiation lies in the fact that the higher the buyer's equilibrium utility, the less likely it is that the seller's incentive constraint is binding. Recall from the discussion following Proposition 1 that without renegotiation, the buyer may 'insist' on a certain payoff from trade of good L that the seller cannot give him with an agreement on good H. The effect of renegotiation, in contrast, is that larger payoffs for B facilitate agreements on the high-quality good: the acceptance condition for B and the incentive compatibility condition for S can both be satisfied simultaneously. Consequently, it is now easier for S to commit herself to deliver the high-quality good. More specifically, it is easy to show that there always exists a price offer p which (a) is accepted by B if he expects to obtain a given payoff from trade of good L in the following period, (b) is incentive compatible in the sense that B would reject delivery of the low-quality good at p if he expects to obtain the same payoff in the renegotiation phase, and (c) improves upon S's share from trade of good L. As a consequence, the moral hazard problem on the seller's side effectively disappears and the unique SPE is given by the standard partition of a cake of size s_H.[11]

[11] This reasoning depends on $\Delta \to 0$. Otherwise, the buyer may again be inclined to accept the low-quality good since he prefers to trade immediately rather than waiting the length of one time interval to enter renegotiations. In this case, it can be shown that the seller's threat of delivering good L is once again binding for sufficiently low β. Since there again is an upper bound on B's equilibrium payoff of good H, the qualitative features of Proposition 1 are preserved even if renegotiation is possible (for a more detailed analyisi, see Kessler and Lülfesmann [15]).

5 Concluding Remarks

The present paper has investigated bilateral bargaining in a situation where the parties are unable to solve the allocation problem independently of the distribution problem. To this end, we have considered a buyer and a seller who face several trading opportunities whose characteristics are observable but non-contractible. We supposed that the seller can grant a money-back warranty, i.e., the buyer has the option to step back from the purchase and to return the commodity. In such a situation, the seller's threat to deliver a low-quality good imposes an upper bound on the buyer's equilibrium payoff, and therefore weakens his bargaining position vis-a-vis the seller. We have shown that if bargaining breaks down after the buyer has refused delivery, inefficient trade is not unlikely to occur even if the players are very patient. In addition, the bargaining process may be characterized by delay or multiplicity of equilibria. However, if the parties can renegotiate their initial agreement after an unsuccessful trading stage, efficiency and uniqueness of equilibrium are restored.

The qualitative findings of our model extend beyond the present trade interpretation. For example, our results on equilibrium characteristics immediately carry over to the bargaining situation that prevails in a simple moral hazard framework with a wealth-constrained agent. Furthermore, the interpretation of the trade contract as an agreement that contains a customer return policy sheds light on the incentives of producers to offer such guarantees to their customers if the exact nature of the 'dissatisfaction' with the product cannot be properly specified: for a wide range of parameter values, the seller benefits from the buyer's return option because it facilitates her commitment to deliver goods of high quality. This commitment, in turn, increases equilibrium trade prices and, ultimately, the seller's payoff from the relationship. Interestingly, the buyer does not necessarily gain from his option to return the commodity. In many equilibria where the high-quality good is traded, the buyer's payoff is lower than his payoff in the (unique) equilibrium when the seller does not grant a money back warranty. The common view of warranties as a device to insure customers against opportunistic behavior on the sellers' part is thus not necessarily justified in situations where the parties bargain over the terms of trade.

6 Appendix

6.1 Proof of Proposition 1

The proof relies on familiar arguments from bargaining theory (based on Shaked and Sutton [23]) that can be found in, e.g., the textbooks of Osborne and Rubinstein [20] and Muthoo [19]). We therefore confine ourselves to a brief sketch of the line of reasoning to show how it can be extended to the present framework (rigorous proofs of all propositions can be found in Kessler and Lülfesmann [15]).

To introduce some notation, let G^j be a subgame which begins with an offer of player $j \in \{B, S\}$ and define $M_i^j = sup\{u_j(x_i, t)$: there is an SPE of G^j with outcome $(x_i, t)\}$ as the maximum payoff of player j in a subgame in which he is the first to make a proposal and in which good i is traded in period t. Similarly, let m_i^j be the corresponding infimum. The index i is omitted if we refer to the maximum and minimum payoffs of the players in all subgames, irrespective of which good is traded. We consider each of the possible equilibria in turn. For brevity of exposition, we refrain from arguing that they are indeed subgame perfect whenever possible and concentrate instead on the issue of uniqueness.

Step a). Suppose $\beta \geq \frac{1}{1+\delta} s_H$, then the following is an SPE: the buyer always proposes $(p_B^* = c_H + \frac{1}{1+\delta} s_H, m = p_B^*)$ and accepts any offer which gives him at least $\frac{\delta}{1+\delta} s_H$ and the seller always proposes $(p_S^* = b_H - \frac{\delta}{1+\delta} s_H, m = p_S^*)$ and accepts any offer which gives her at least $\frac{\delta}{1+\delta} s_H$. An agreement is reached immediately and the high-quality good is produced and traded.

To see why this SPE must be unique , let us ignore good L for the moment so that the traditional Rubinstein model applies where $M_H^B = m_H^B = M_H^S = m_H^S = \frac{1}{1+\delta} s_H$. Using standard argumentation, it is easy to show that $M_L^j \leq \frac{1}{1+\delta} s_L$ for $j \in \{B, S\}$. Due to $s_L < s_H$, we thus have $M^j = M_H^j$. Assume $\beta \geq \frac{1}{1+\delta} s_H \Rightarrow p_j^* \geq b_L$. Then, $m^j = m_H^j$ because player j can always propose $(p_j^*, m = p_j^*)$. Furthermore, the first offer is accepted in any subgame, which establishes uniqueness. Using $\delta = \exp(-r\Delta)$ and setting $\Delta \rightarrow 0$, part a) follows.

Step b). Suppose $\frac{1}{1+\delta} s_H - \frac{1}{1-\delta^2}(s_H - s_L) < \beta < \frac{1}{1+\delta} s_H$, then the following is an SPE: B always proposes $(\bar{p}_B = b_L, m = \bar{p}_B)$ and rejects any offer which does not give him at least $\delta\beta$. S always offers $(\bar{p}_S = b_H - \delta\beta, m = \bar{b}_S)$ and rejects any offer below $\delta(s_H - \delta\beta)$. The lower bound on β eliminates the possibility that B can profitably deviate by proposing $p_B \leq b_L$ where good L would be traded. To see why, consider an offer such that $u_B > \beta$. An acceptable offer must give S at least $u_S = \delta(s_H - \delta\beta)$. Since $u_B + u_S = s_L$ if $x_L = 1$, such a deviation does not exist if $\beta \leq s_L - \delta s_H + \delta^2\beta$. Rearranging terms yields the lower bound on β. Note further that due to $\beta > 0$, this bound is not binding for $\Delta \rightarrow 0\,(\delta \rightarrow 1)$, in which case the SPE exists for all $\beta \in (0, \frac{1}{2} s_H)$.

Turning to uniqueness, suppose in addition $\beta \geq \frac{\delta}{1+\delta} s_L$. First, it can easily be verified that $m^B = \beta \ (< \frac{1}{1+\delta} s_H)$ and $M^S = s_H - \delta\beta \ (> \frac{1}{1+\delta} s_H)$. Also, because any offer that gives B at least $\delta M^B \leq \{\delta\beta, \frac{\delta}{1+\delta} s_L\}$ is accepted, S can always offer $p_S = b_H - \delta M^B \geq b_L$ if $\beta \geq \frac{\delta}{1+\delta} s_L$. Hence, $m^S = s_h - \delta M^B$. To establish $M^B = \beta$, consider a subgame G^B and suppose $M^B > \beta$, which necessarily requires $x_L = 1$. But because the buyer must offer the seller at least δm^S, $u_B \leq s_L - \delta m^S = s_L - \delta s_H + \delta^2\beta < \beta$, under the above conditions on β, a contradiction. Hence, $M^B = M_H^B = \beta$. The rest of the argument is standard and part b) follows for $\Delta \rightarrow 0$.

Step c). The proof is by construction. Consider values $\beta < \frac{\delta}{1+\delta} s_L$. Then, an SPE exists where the buyer always proposes $\tilde{p}_B = c_L + \frac{\delta}{1+\delta} s_L$ and accepts any offer which gives him at least $\frac{\delta}{1+\delta} s_L$ and the seller always proposes $\tilde{p}_S = s_L - \frac{\delta}{1+\delta} s_L$ and accepts any offer which gives her at least $\frac{\delta}{1+\delta} s_L$. An agreement is reached immediately and the low-quality good is produced and traded. It is easy to verify that under the above condition on β, the strategies are subgame perfect. In particular, there exists no profitable deviation $(p_S \geq b_L, m)$ for S which is acceptable for B and triggers trade of good H.

Now let $\Delta \to 0$ $(\delta \to 1)$. From the argument in step b), first paragraph, both this SPE and the SPE described in Proposition 1 b) co-exist for all values $\beta < \frac{1}{2} s_L$. The buyer's utility is $\frac{1}{2} s_L$ in the former and β in the latter. Using worst punishment strategies, it is straightforward to show that each point $u_B \in \{\beta, \frac{1}{2} s_L\}$ can be supported as an equilibrium payoff to B in a subgame G^B in which B makes the first offer. In particular, consider the following strategies. B proposes $p_B \in [\frac{1}{2} s_L, b_L - \beta]$ in the first period of G^B which S accepts. If S rejects, the strategies call for both parties to play the equilibrium where the low-quality good is traded. The seller's payoff in this case is $\tilde{u}_S$ which is weakly less than $p_B - c_L$ today. If B proposes less than p_B, S rejects and the equilibrium played is the equilibrium in part b). In this case, the utility of S is $\delta u_S = \delta(s_H - \delta \max\{\beta, \bar{\beta}\})$ which strictly exceeds her maximum utility along the equilibrium path, $s_L - \max\{\beta, \bar{\beta}\}$. Standard arguments can be used show that the interval covers *all* equilibrium payoffs of B in G^B. An intuitive argument for the existence of subgame perfect equilibria with delay is provided in the discussion following Proposition 2. In order to construct those equilibria, one can use the same argument as, e.g., in Muthoo [19], Chapter 7. □

6.2 Proof of Proposition 2

Suppose $\beta \geq \frac{1-\delta^2}{1+\delta} s_H$, then the following is an SPE in the game with renegotiation: B always proposes $(p_B^* = c_H + \frac{\delta}{1+\delta} s_H, m = p_B^*)$ and accepts any offer which gives him at least $\delta u_B^* = \frac{\delta}{1+\delta} s_H$ and S always proposes $(p_S^* = b_H - \frac{\delta}{1+\delta} s_H, m = p_S^*)$ and accepts any offer which gives her at least $\delta u_S^* = \frac{\delta}{1+\delta} s_H$. An agreement is reached immediately and the high-quality good is traded.

Note that in this SPE, the buyer's payoff in a subgame G^S is equal to $\frac{\delta}{1+\delta} s_H$. Hence, given that an agreement of p_B^* is reached in a subgame G^B, B will reject delivery of good L if and only if $b_L - p_B^* \leq \frac{\delta^2}{1+\delta} s_H$ (where the indifference is solved in favor of rejecting along the equilibrium path). Using the expression for p_B^*, this is equivalent to the condition $\beta \geq (1-\delta) s_H = \frac{1-\delta^2}{1+\delta} s_H$. It is now straightforward to check that the strategies are indeed subgame perfect. In order to show that the equilibrium is unique, one can apply the same argument as in Step a) of the proof of Proposition 1 (the upper bounds on M_L^B and M_L^S are still valid).□

References

1. Aghion, P. and Tirole, J.: The Management of Innovation. Quarterly Journal of Economics **109**, 1185-1209 (1994).
2. Allen, F.: Reputation and Product Quality. RAND Journal of Economics **15**, 311-327 (1984)
3. Avery, C. and Zemsky,P.: Money Burning and Multiple Equilibria in Bargaining. Games and Economic Behavior **7**, 154- 168 (1994)
4. Busch, L. and Wen, Q.: Perfect Equilibria in a Negotiation Model. Econometrica **63**, 545-565 (1995)
5. Che, Y.-K.: Customer Return Policies for Experience Goods. Journal of Industrial Economics **64**, 17-24 (1996)
6. Cooper, R. and Ross, T.W.: Product Warranties and Double Moral Hazard. RAND Journal of Economics **16**, 103-113 (1985)
7. Emons, W.: Warranties, Moral Hazard, and the Lemons Problem. Journal of Economic Theory **46**, 16-33 (1988)
8. Emons, W.: The Theory of Warranty Contracts. Journal of Economic Surveys **3**, 43-57 (1989)
9. Fama, E.F.: Agency Problems and the Theory of the Firm. Journal of Political Economy **88**, 288-307 (1980)
10. Fernandez, R. and Glazer, J. Striking for a Bargain Between Two Completely Informed Agents. American Economic Review **81**, 240 - 252 (1991)
11. Haller, H. and Holden, S.: Non-Cooperative Wage Bargaining. Journal of Economic Theory **52**, 232-236 (1990)
12. Hart, O.: Firms, Contracts, and Financial Structure. Oxford: Oxford University Press 1995
13. Hart, O. and Moore, J. Incomplete Contracts and Renegotiation. Econometrica **56**, 755-785 (1988)
14. Holmström, B.: Managerial Incentive Problems - A Dynamic Perspective. In: Essays in Economics and Management in Honor of Lars Wahlbeck, Helsinki: Swedish School of Business 1982
15. Kessler, A.S. and Lülfesmann, C.: Bilateral Bargaining under Non-Verifiable Information. SFB Discussion Paper No. A-563, University of Bonn (2000)
16. Klein, B. and Leffler, K.B.: The Role of Market Forces in Assuring Contractual Performance. Journal of Political Economy **89**, 615-641 (1981)
17. Mann, D.P. and Wissink, J.P.: Money-Back Warranties vs. Replacement Warranties: A Simple Comparison. American Economic Review **80**, 432-436 (1990)
18. Muthoo, A.: Bargaining without Commitment. Games and Economic Behavior **2**, 291-297 (1990)
19. Muthoo, A.: Bargaining Theory with Applications. Cambridge: Cambridge University Press 1999
20. Osborne, M.J. and Rubinstein, A.: Bargaining and Markets. San Diego, CA: Academic Press 1990
21. Rubinstein, A.: Perfect Equilibrium in a Bargaining Model. Econometrica **50**, 97-109 (1982)
22. Shaked, A.: Opting Out: Bazaars Versus High-Tech Markets. Investigaciones Economicas **18**, 421-432 (1994)
23. Shaked, A. and Sutton, J. Involuntary Unemployment as a Perfect Equilibrium in a Bargaining Model. Econometrica **52**, 1351-1364 (1984)

24. Premiums for High-Quality Products as Returns to Reputation. Quarterly Journal of Economics **98**, 659-679 (1983)
25. Tirole, J.: Incomplete Contracts: Where do We Stand?. Econometrica **67**, 741-781 (1999)

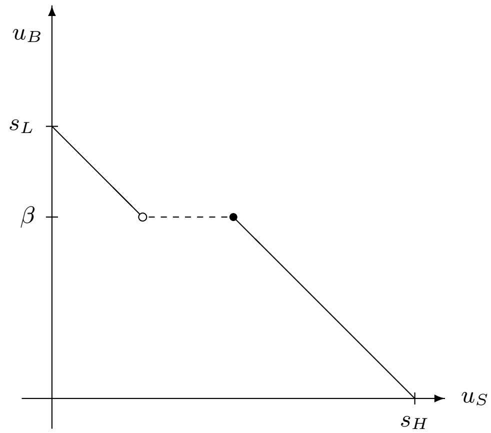

Figure 1

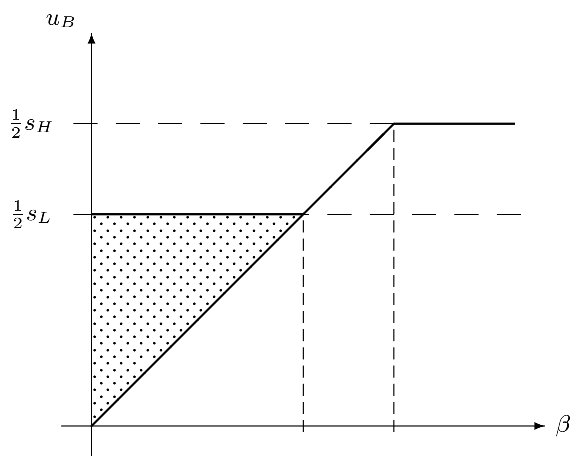

Figure 2