CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1536

"Price Discrimination Through Communication"

Itai Sher
University of Minnesota

Rakesh Vohra
Northwestern University

June 2011

JEL Classification: C78, D82, D83.

Keywords: price discrimination, communication, bargaining, commitment, evidence, network flows.

Price Discrimination Through Communication

Itai Sher* Rakesh Vohra†

June 2011

Abstract

We study a seller's optimal mechanism for maximizing revenue when the buyer may present evidence relevant to the buyer's value, or when different types of buyer have a differential ability to communicate. We introduce a dynamic bargaining protocol in which the buyer first makes a sequence of concessions in a cheap talk phase, and then at a time determined by the seller, the buyer presents evidence to support his previous assertions, and then the seller makes a take-it-or-leave-it offer. Our main result is that the optimal mechanism can be implemented as a sequential equilibrium of our dynamic bargaining protocol. Unlike the optimal mechanism to which the seller can commit, the equilibrium of the bargaining protocol also provides incentives for the seller to behave as required. We thereby provide a natural procedure whereby the seller can optimally price discriminate on the basis of the buyer's evidence.

JEL Classification: C78, D82, D83.

Keywords: price discrimination, communication, bargaining, commitment, evidence, network flows.

1 Introduction

In economics it is common to model communication as cheap talk, but cheap talk seems useless for some fundamental economic interactions. Consider a buyer and seller negotiating over price. If some buyer type could persuade the seller to lower the price via some cheap talk message, then all buyer types could achieve the discount in the same way. Yet, many buyer seller transactions involve communication prior to price setting. To understand why, we allow for differential communication ability among buyers. We model this by giving different types of buyers access to different sets of messages. These messages can be interpreted as

*Economics Department, University of Minnesota. Email: isher@umn.edu.

†Kellogg School of Management, MEDS Department, Northwestern University. Email: r-vohra@kellogg.northwestern.edu.

hard evidence. We also allow for cheap talk, which can serve a valuable role in combination with evidence.

Evidence can take many different forms. An example is an advertisement that shows the price at which the consumer could buy a substitute for the seller's product elsewhere. However, the buyer need not present a physical document; a buyer who knows the market may demonstrate this through her words alone, whereas an ignorant buyer could not produce those words. As another example, when purchasing a house, the buyer may claim that a loan with favorable terms for which she qualifies has a cap below the asking price. The seller may verify this, or alternatively, he may believe that if the buyer did not know of such a loan, she would not have thought of mentioning it. The seller could also take control of the process. In the early days of the internet it may have made sense for car dealers to ask potential buyers for their email address. Having an email address is a signal the buyer is more likely to be surfing the web doing price comparisons, so the dealer would have an incentive to offer a lower price.[1]

In this setting we study the optimal direct mechanism that maximizes the seller's expected revenue. Say that type t *can mimic* type s if every message available to s is also available to t. With evidence, we need only impose a subset of the incentive constraints which we would have to consider if there were only cheap talk. In particular, we only impose an incentive constraint that discourages type t from claiming to be type s if t can mimic s.

In contrast to Myerson (1981), for example, the optimal mechanism in our setting will involve both *price discrimination* and *randomization*. Different buyer types will receive different prices and receive the object with different probabilities depending on the evidence that they can present.

When *all* incentive constraints must be respected, we know that the downward adjacent constraints will bind. In our setting, the absence of some incentive constraints makes it difficult to say *a priori* which of them will bind at optimality; if type t can mimic both lower value types s and r, but s and r cannot mimic each other, which type will t want to mimic at the optimal mechanism? This makes the optimal direct mechanism difficult to interpret. To remedy this, we show that the optimal direct mechanism can be implemented via a natural bargaining protocol in which the buyer and seller engage in several rounds of cheap talk communication followed by the presentation of evidence by the buyer and then a take-it-or-leave-it offer by the seller. This implementation also suggests that in addition to the usual determinants of bargaining (patience, outside option, risk aversion, commitment) the persuasiveness of arguments is also relevant.

Communication in the sequential equilibrium of our bargaining protocol is monotone in two senses: the buyer makes a sequence of concessions in which she claims to have successively higher valuations and at the same time the buyer admits to having more and

[1] We thank Simon Board for this example.

more evidence as communication proceeds. To see how the two are related, imagine that the buyer has evidence suggesting she is of an intermediate value. When the buyer and seller are arguing over whether the price should be low or intermediate, the buyer would like to withhold this evidence, but once the buyer admits to an intermediate value, she would like to present this evidence to prevent a high price.

Throughout, the buyer's communication is disciplined by the need to present the supporting evidence at the end. The seller decides when to exit the cheap talk phase and enter the evidence presentation phase. The seller faces an optimal stopping problem: should he ask for a further concession from the buyer which would yield additional information about the buyer's type but risk the possibility that the buyer will be unwilling to make an additional concession and thus drop out?

The seller's optimal stopping strategy is determined by the optimal mechanism. The seller asks for another cheap talk message when the buyer claims to be of a type that is not optimally served and requests supporting evidence in preparation for an offer and sale when the buyer claims to be of a type that is served. Most interesting is when the buyer claims to be of a type which is optimally served with an intermediate probability; then the seller randomizes between asking for more cheap talk and proceeding to the sale.

The buyer's strategy is determined by an optimal solution to the dual of the seller's optimal mechanism problem. In particular, an optimal dual solution determines a probability distribution of paths through types connected by binding incentive constraints. The buyer randomizes according to this distribution and her reporting strategy–her sequence of concessions–follows such a path. The fact that the reports are concessions–the buyer admits to successively higher values–follows from a nontrivial lemma that binding incentive constraints point from higher to lower value types. That the buyer claims to have successively more and more evidence follows from the fact that we need only consider incentive constraints in the direction of increasing evidence.

An interesting byproduct of the analysis is that the optimal mechanism can be implemented with no more commitment than the ability to make a take-it-or-leave-it offer.

When the optimal mechanism is deterministic, we show that the back-and-forth cheap talk communication in the bargaining protocol collapses to a single stage. Nevertheless, randomization is still required on the buyer's part. A much stronger assumption is required to eliminate all randomization from the bargaining protocol. Contrariwise, when the optimal mechanism requires randomization, the bargaining protocol requires several rounds of cheap talk. In this case, sequential communication is required to ease the seller's commitment requirements. We present a family of examples which contains arbitrarily many rounds of communication. Finally, we show that with binary values our model has a close connection to the Glazer & Rubinstein (2004) model of optimal persuasion.

The outline of the paper is as follows: In section 2, we present the model. In section

3, we study the optimal mechanism. In sections 4 and 5, we present our dynamic bargaining protocol and prove that the optimal mechanism can be implemented as a Bayesian Nash equilibrium of our bargaining protocol. Section 6 strengthens the solution concept to sequential equilibrium. Section 7 presents the family of examples showing that communication may contain arbitrarily many rounds. Section 8 examines several special cases of the model which have some additional structure. Section 9 concludes. An appendix contains proofs which were omitted from the main body.

1.1 Prior Literature

Our work is closely related to the models of persuasion (Milgrom & Roberts 1986, Shin 1994, Lipman & Seppi 1995, Glazer & Rubinstein 2004, Glazer & Rubinstein 2006, Sher 2011, Sher 2010). These models deal with situations in which a speaker attempts to persuade a listener to take some action. Our model deals in particular with arguments attempting to persuade a seller to lower his price. Our result has an interesting relation to the credibility result of Glazer & Rubinstein (2004); that paper studied persuasion with respect to a binary decision involving no exchange of money. A detailed discussion of the relationship is presented in Section 8.3.

This paper is also a contribution to the body of research on mechanism design with evidence (Green & Laffont 1986, Singh & Wittman 2001, Forges & Koessler 2005, Bull & Watson 2007, Ben-Porath & Lipman 2008, Deneckere & Severenov 2008, Kartik & Tercieux 2009). These papers study general mechanism design environments, establishing revelation principles and necessary and sufficient conditions for partial and full implementation. In contrast, our focus is on optimal price discrimination. A related application has been investigated by Severenov & Deneckere (2006) in which some agents are strategic and may mimic any other type whereas others are nonstrategic, and the latter must report their information truthfully. Celik (2006) studies an adverse selection problem in which higher types can pretend to be lower types but not vice versa, and shows that the weakening of incentive constraints does not alter the optimal mechanism.[2]

A related line of work is Blumrosen, Nisan & Segal (2007) and Kos (2011) which assumes that bidders can only report one of a finite number of messages. However, unlike the models we consider, all messages are available to each bidder. There is a also a body of literature that studies the relation between between incentive compatible mechanisms and outcomes that can be implemented in infinite horizon bargaining games with discounting (Ausubel & Deneckere 1989, Gerardi, Horner & Maestri 2010). This literature does not study the role of evidence, which is our main focus. Moreover our results are quite different both in substance and technique. Finally, our work contributes to the linear programming approach to mechanism design (Vohra 2011).

[2]Technically, a closely related analysis is that of Moore (1984).

2 The Model

Suppose a seller possesses a single item he does not value which he would like to sell to a buyer. Let T be a finite set of buyer types. π_t and v_t are respectively the probability of and valuation of type t. There is also a finite set M of *hard* messages. For any type $t \in T$ there is a finite set $\sigma(t) \subseteq M$ of messages which are available to type t. σ is the **message correspondence**. We may interpret the message correspondence in terms of evidence. We assume that for any subset of S of $\sigma(t)$, the buyer can present S. In particular, the buyer can present all evidence in $\sigma(t)$. It is convenient to define: $S_t := \{m : m \in \sigma(t)\}$. Of course formally, S_t and $\sigma(t)$ are the same set of messages. However, if we think of $\sigma(t)$ as encoding the buyer's choice set, we think of S_t as encoding a particular choice: namely the choice to present all messages in $\sigma(t)$. Note that a type $s \neq t$ may also be able to present S_t if $\sigma(t) \subseteq \sigma(s)$.

Our assumption that the seller can present any subset of messages is technically stronger than the assumption of **normality** of Bull & Watson (2007). However, for our purposes, the two are equivalent.[3][4]

We assume that there is a **zero type** $0 \in T$ with $v_t = \pi_t = 0$ and $\sigma(t) = \{m_0\}$ where $m_0 \in \sigma(t), \forall t \in T$. It is also convenient to assume that for all $t \in T \setminus 0$, $\sigma(t) \neq \sigma(0)$. The zero type plays the role of the outside option. We assume that for all $t \in T \setminus 0, v_t > 0$ and $\pi_t > 0$.

In addition to the hard messages M, we assume that the buyer has access to an unlimited supply of cheap talk messages. These cheap talk messages are available to all buyer types. In the bargaining protocol described in Section 4 we restrict the set of cheap talk messages to correspond with the set T of types (but allow many messages to be sent). Nothing would be gained if we allowed the buyer access to a larger set of cheap talk messages in the bargaining protocol.

2.1 Incentive Graphs

It is useful for the analysis to define a directed graph. The set of vertices in this graph is the set T of types, and the set of edges $E \subseteq T \times T$, where:

$$(s, t) \in E \Leftrightarrow [\sigma(s) \subseteq \sigma(t) \text{ and } s \neq t] \tag{1}$$

Notice that our assumptions on the zero type are such that:

$$\forall t \in T \setminus 0, (0, t) \in E \tag{2}$$

[3] In particular, for any normal message structure, we can construct a message correspondence satisfying our assumption which leads to the same optimal mechanism.

[4] Another related assumption from the literature, which is also essentially equivalent for our purposes is the **nested range condition** of Green & Laffont (1986).

It is also true (but less important) that for all $t \in T, (t, 0) \notin E$. We refer to a graph as just described as an **incentive graph**. Note that E is transitive, except that self edges of the form (t, t) are excluded. The term "transitivity" is to be understood with this qualification below.

3 The Optimal Mechanism

In this section, we study the optimal mechanism. Section 3.1 formulates the problem and studies its properties. Section 3.2 provides a useful reformulation of the problem.

3.1 Properties of the Optimum

We consider an optimal mechanism design problem that is formulated below. q_t is the probability that type t receives the object and p_t is the expected payment of type t.

Primal Problem (Edges)

$$
\begin{aligned}
\text{maximize} \quad & \sum_{t \in T} \pi_t p_t && (3)\\
\text{subject to} \quad & && \\
\forall (s,t) \in E, \quad & v_t q_t - p_t \geq v_t q_s - p_s && (4)\\
\forall t \in T, \quad & 0 \leq q_t \leq 1 && (5)\\
& p_0 = 0 && (6)
\end{aligned}
$$

The seller's objective is to maximize expected revenue (3). The problem (3-6) resembles a standard mechanism design problem with the exception that the optimal mechanism does not have to honor all incentive constraints, but only incentive constraints for pairs of types (s, t) with $(s, t) \in E$. Indeed the label "edges" refers to the fact that there is an incentive constraint for each edge of the incentive graph, and is to be contrasted with the formulation in terms of paths to be presented in section 3.2. The interpretation is that we only impose an incentive constraint saying that t should not want to claim to be s if t **can mimic** s in the sense that any evidence that s can present can also be presented by t. The individual rationality constraint is encoded by (6) and the instances of (4) with $s = 0$ (recall that $(0, t) \in E$ for all $t \in T \setminus 0$).

Although they did not explicitly study the notion of an incentive graph, the fact that in searching for the optimal mechanism we only need to consider the incentive constraints in (4) follows from Corollary 1 of Deneckere & Severenov (2008), which may be viewed as a version of the revelation principle for general mechanism design problems with evidence. More specifically, given a social choice function f mapping types into outcomes, these authors show that when agents can reveal all subsets of their evidence, there exists a (possibly dynamic) mechanism Γ which respects the right of agents to decide which of their own evidence to present and is such that Γ implements f if and only if f satisfies all (s,t)-incentive constraints for which $\sigma(s) \subseteq \sigma(t)$. This justifies the program (3-6) for our problem. For further details, the reader is referred to Deneckere & Severenov (2008). Related arguments are presented by Bull & Watson (2007). (Note that our model satisfies their normality assumption because each the type t buyer can present all subsets of $\sigma(t)$).[5]

In our analysis, the dual of (3-6) will play an important role. In particular, the dual will allow us to identify the buyer's strategy in our dynamic bargaining protocol.

Dual Problem (Edges)

$$\text{minimize} \quad \sum_{t\in T} \mu_t \tag{7}$$

subject to

$$\forall t \in T \setminus 0, \quad \sum_{s:(s,t)\in E} \lambda(s,t) - \sum_{s:(t,s)\in E} \lambda(t,s) = \pi_t \tag{8}$$

$$\forall t \in T, \quad v_t\pi_t - \sum_{s:(t,s)\in E} \lambda(t,s)(v_s - v_t) \leq \mu_t \tag{9}$$

$$\forall (s,t) \in E \quad \lambda(s,t) \geq 0, \tag{10}$$

$$\forall t \in T, \quad \mu_t \geq 0 \tag{11}$$

The dual has a network interpretation. The multipliers $\lambda(s,t)$ on the incentive constraints can be interpreted as a flow on the edge (s,t) of the incentive graph. Each nonzero type t is a demand vertex with demand equal to the probability of t, π_t. Constraint (8) is a flow conservation constraint saying that the net flow of vertex t (the inflow minus the outflow) is equal to the demand of vertex t. So that supply equals demand, we view vertex 0 (with $(0,t) \in E, \forall t \in T \setminus 0$) as a supply vertex with supply $\sum_{t\in T\setminus 0} \pi_t = 1$.

Next we interpret constraint (9). It is convenient to introduce the notation:[6]

$$\psi_t := v_t - \frac{\sum_{s:(t,s)\in E} \lambda(t,s)(v_s - v_t)}{\pi_t} \tag{12}$$

Evaluated at a dual optimum, we may interpret ψ_t as the **virtual valuation** of type t. ψ_t is analogous to the virtual valuation in traditional mechanism design (i.e., when we impose *all* incentive constraints, not just those in the incentive graph). Constraint (9) together with the minimization (7) serve to establish the following relation, which must hold at the

[5]Bull & Watson (2007) also explain the close relation of their normality assumption to the nested ranged condition of Green & Laffont (1986) and relate their analysis to that of the latter paper.

[6]Notice in particular that because $\pi_0 = 0$, $\psi_0 = -\infty$.

optimum:

$$\mu_t = \max\{\psi_t, 0\}\pi_t \tag{13}$$

In other words μ_t is the positive part of the virtual valuation of type t multiplied by the probability of type t. The following proposition now follows from strong duality and complementary slackness:

PROPOSITION 3.1 *At any optimal mechanism a buyer type is served with probability one if she has a positive virtual valuation and with probability zero if she has a negative virtual valuation. Types with zero virtual valuation are served with some (possibly zero) probability. The seller's revenue is equal to the expected value of the positive part of the virtual valuation:*

$$\sum_{t\in T} \max\{\psi_t, 0\}\pi_t$$

Let us compare Proposition 3.1 to the standard mechanism design problem in which we impose all incentive constraints. In this case, assume wlog that the set of types is $\{0, 1, \ldots, n\}$ and $i < j \Rightarrow v_i < v_j$. In that problem (with monotone virtual valuations), we know that the downward adjacent constraints bind (even without imposing a monotonicity constraint) and moreover at the optimum, we would have:[7]

$$\sum_s \lambda(i, s) = \lambda(i, i+1) = \sum_{j=i+1}^{n} \pi_j \tag{14}$$

so that the virtual value can be written:

$$\psi_i = v_i - (v_{i+1} - v_i)\frac{\sum_{j=i+1}^{n} \pi_j}{\pi_i}$$

Once we know which incentive constraints are binding, it is easy to solve for the exact values of the multipliers $\lambda(s, t)$ and hence to determine the virtual values. In contrast, in our case with only a subset of incentive constraints, we do not know *a priori* which constraints will bind. For this reason, we do not know in which "direction" the cumulative distribution function which typically features in the expression for the virtual valuation should point. In (12), the flow λ emerging from an optimal dual solution gives that direction (or rather those directions). The flow conservation constraints (8) relate the flow emanating from t to the cumulative probability mass of types "above" t which can reach t by passing through a sequence of binding incentive constraints.[8] Such a conceptualization should be useful more generally for multi-dimensional mechanism design problems with or without evidence.

[7] To be precise, in this case, (14) always holds at *some* optimal solution.

[8] The possibility of flow on cycles may initially appear to interfere with this interpretation. However, as we explain below, it is always possible to find an optimal dual solution without any flow on cycles.

Despite the differences between our problem and the standard problem, once the virtual values are found, Proposition 3.1 shows that the solution to our problem is similar to the solution to the standard problem in the sense that seller serves only types with non-negative virtual value (with probability one if the virtual value is positive) and earns the expected positive part of the virtual value.

We now present some examples which highlight some differences between our problem and the standard problem:

Example 1 *Let $T = \{0, 1, \ldots, 7\}$, and consider the following diagram, illustrating the incentive graph:*

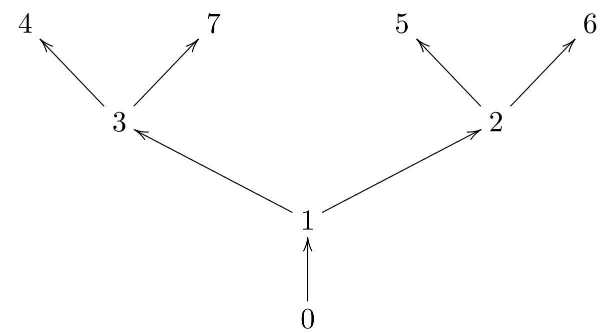

Figure 1: An Incentive Graph

Suppose the edge $(s,t) \in E$ if in Figure 1 there is a directed path from s to t. For example, $(1,7) \in E$ even though in Figure 1 an edge from 1 to 7 is absent. Such an incentive graph can be induced by a message correspondence in which each type t has message m_t and in addition for each s such that $(s,t) \in E$, t has message m_s (where $s \neq t \Rightarrow m_s \neq m_t$).

Suppose, moreover that the numbers of the types also represent their values for the object so that for $t = 0, 1, \ldots, 7$, $v_t = t$. Suppose moreover that $\pi_0 = 0$, $\pi_1 = \pi_2 = \pi_3 =: \pi_a$ and $\pi_4 = \pi_5 = \pi_6 = \pi_7 =: \pi_b$, and define:

$$K := \frac{\pi_b}{\pi_a}$$

If K is sufficiently small, then the unique optimal mechanism is given by the following table:

t	q_t	p_t
7	1	3
6	1	2
5	1	2
4	1	3
3	1	3
2	1	2
1	0	0
0	0	0

In particular, type 2 receives the object for a price of 2, type 3 receives the object for a price of 3, and type 1 is not served. Types 4 and 7 mimic type 3, and types 5 and 6 mimic type 2. None of the types receiving the higher price of 3 can mimic any of the types receiving the lower price of 2. This example illustrates that, in contrast to the case where all incentive constraints are imposed, the optimal solution may satisfy:

Price Discrimination *Different types pay different prices.*

Next observe that if K is sufficiently large, then the optimal mechanism becomes:

t	q_t	p_t
7	1	7
6	1	6
5	1	5
4	1	4
3	0	0
2	0	0
1	0	0
0	0	0

In this case, types 2 and 3 are no longer served, and the seller achieves prefect price discrimination for types 4, 5, 6, and 7. This illustrates that the optimal mechanism involves endogenous segmentation. Buyer types are segmented into different classes with different prices, but ex ante, we do not know how the types will be grouped into which classes.

Example 2 *Let $T = \{0, 1, 2, 3, 4\}$. Consider the incentive graph in Figure 2:*

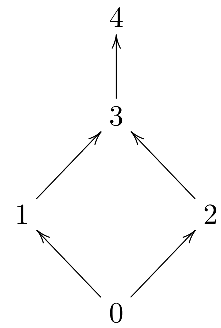

Figure 2: An Incentive Graph

As in Example 1, edge $(s, t) \in E$ if in the above diagram there is a directed path from s to t.

Suppose again that types correspond to values so that $v_t = t$ for all types t. Finally the probabilities of the types satisfy the following relations:

$$\pi_2 > K\pi_4 > K^2\pi_1 > K^3\pi_3 > \pi_0 = 0 \tag{15}$$

where K is some positive number. If (15) holds for K sufficiently large, then the unique optimal mechanism is given by the following table:

t	q_t	p_t
4	1	2
3	2/3	2/3
2	1	2
1	2/3	2/3
0	0	0

To see this, observe that if K is sufficiently large, then the optimal mechanism must extract the full surplus from type 2. The next priority will be to extract as much surplus as possible from 4 given that she can mimic 2, which determines 4's payment and allocation. Following this, we would like to extract as much surplus from 1 as possible subject to the incentive compatibility of the previously determined allocations and payments for 2 and 4. Since 1 can only mimic the zero type we can set $p_1 = q_1$, so the question becomes: how high can we set q_1? We can only set $q_1 = 2/3$ because that is the point at which 4 becomes indifferent between mimicking 1 and 2. For any higher value of q_1, 4 would strictly prefer to mimic 1 than to mimic 2, and the lost revenue from 4 would not be compensated by the increased revenue from 1. Finally, in the case of 3 we have little leeway. Of the types that 3 can mimic (1 and 2), 3 prefers the payment and allocation of 1. If we attempted to set $q_3 > q_1$

at a price increment that 3 would find attractive, 4 (who can mimic 3) would also find this attractive, and the seller would lose too much money on 4 for this to be worthwhile.

This example illustrates two features that an optimal mechanism may possess.

Random Allocation *Some types receive the item with a probability strictly between zero and one.*

Failure of Allocation Monotonicity *A higher value type t can mimic a lower value type s, and nevertheless, t receives the item with lower probability than s.*

In our example, the higher value type t is 3, and the lower value type s is 2. Note that random allocation introduces a second form of price discrimination which is distinct from that found in Example 1, and more akin to second-degree price discrimination.

It is important to note that this example is not knife-edge. Indeed in this example, it is easy to see that for sufficiently small changes in the parameters $(v_t, \pi_t : t \in T \setminus 0)$, the optimum will remain unique and will still have the properties of random allocation and allocation monotonicity. With a view to Proposition 3.1, types with zero virtual valuation (the only types eligible for random allocation at the optimum) are not a knife-edge phenomenon, but rather are robust to small changes in the parameters.

In light of the above examples, it is useful to present the following proposition which states some important properties of optimal solutions.

PROPOSITION 3.2 *There exists an optimal solution to the dual satisfying:*

$$\lambda(s,t) > 0 \Rightarrow v_s < v_t \qquad \forall (s,t) \in E \tag{16}$$

All optimal solutions to the primal and dual satisfy:

$$\lambda(s,t) > 0 \Rightarrow q_s \leq q_t \qquad \forall (s,t) \in E \tag{17}$$

Proof: In Appendix. ■

(16) says roughly that the binding incentive constraints point from higher value types to strictly lower value types. Referring to edges $(s,t) \in E$ with $v_s < v_t$ as **good edges** and with $v_s \geq v_t$ as **bad edges**, a flow λ satisfying (16) is said to **avoid bad edges.** In the Examples 1 and 2 we considered incentive graphs without bad edges, although our theory allows for bad edges.

(17) says that the allocation is increasing along the binding incentive constraints. This can be thought of as a weaker form of the allocation monotonicity property discussed in Example 2. In particular, (17) says that insofar as allocation monotonicity is violated, it must be violated only along non-binding incentive constraints.

3.2 A Reformulation in Terms of Paths

There is a natural reformulation of problem (3-6) in terms of paths, which will be essential for our bargaining protocol. Given an allocation $q = (q_t : t \in T)$, for each edge $(s,t) \in E$, interpret $v_t(q_t - q_s)$ as the "length" of the edge. A path is a sequence of vertices $t_0 \to t_1 \to \cdots \to t_k$ with $k \geq 1$ and $(t_i, t_{i+1}) \in E$ for all $i = 1, \ldots, k$. The length of such a path is $\sum_{j=1}^{k} v_{t_j}(q_{t_j} - q_{t_{j-1}})$. In this paper, a path will always assumed to be *simple*, i.e., paths containing cycles are excluded. Let $\mathcal{P}$ be the set of all paths beginning in 0. For any paths P and P', write $P' \subseteq P$ if P' is an initial subsequence of P and $t \in P$ if t is a vertex in P.

Notice that if we add the IC constraints (4) corresponding to each edge on a path beginning at $t_0 = 0$, and use $p_0 = 0$, we obtain that

$$p_{t_k} \leq \sum_{j=1}^{k} v_{t_j}(q_{t_j} - q_{t_{j-1}}) = v_{t_k} q_{t_k} - \sum_{r=1}^{k-1} q_{t_r}(v_{t_{r+1}} - v_{t_r}) - v_{t_1} q_{t_0}. \tag{18}$$

This observation leads to the following relaxed formulation of (3-6).

Primal Problem (Paths)

$$\text{maximize} \quad \sum_{t \in T} \pi_t p_t \tag{19}$$

$$\text{subject to}$$

$$\forall (t_0, t_1, \ldots, t_k) \in \mathcal{P}, \quad p_{t_k} \leq v_{t_k} q_{t_k} - \sum_{r=1}^{k-1} q_{t_r}(v_{t_{r+1}} - v_{t_r}) - v_{t_1} q_{t_0} \tag{20}$$

$$\forall t \in T, \quad 0 \leq q_t \leq 1 \tag{21}$$

$$p_0 = 0 \tag{22}$$

(18) says that the price p_t is bounded above by the length of any path from 0 to t. This formulation is relaxed because while the constraints (4) imply the constraints (20), the converse is not true. Nevertheless, we establish the relevance of this program below.

To write down the dual to this problem, denote by $\mathcal{P}_t$ the set of paths that begin with 0 and terminate with t (where $t \in T \setminus 0$) and by $\mathcal{P}_{t,s}$ the set of paths that contain the edge $(t,s) \in E$. If $P \in \mathcal{P}_{t,s}$, we also write $(t,s) \in P$. The dual to the path formulation is:

Dual Problem (Paths)

$$
\begin{aligned}
&\text{minimize} && \sum_{t \in T} \mu_t && (23)\\
&\text{subject to} && && \\
&\forall t \in T \setminus 0, && \sum_{P \in \mathcal{P}_t} \lambda_P = \pi_t && (24)\\
&\forall t \in T, && v_t \pi_t - \sum_{s:(t,s) \in E} \sum_{P \in \mathcal{P}_{t,s}} \lambda_P (v_s - v_t) \leq \mu_t && (25)\\
&\forall P \in \mathcal{P}, && \lambda_P \geq 0, && (26)\\
&\forall t \in T, && \mu_t \geq 0 && (27)
\end{aligned}
$$

PROPOSITION 3.3 *Any optimal solution* $(\lambda_P : P \in \mathcal{P})$ *to (23-27) induces an optimal solution* $(\lambda(s,t) : (s,t) \in E)$ *to (7-11) via:*

$$
\lambda(s,t) = \sum_{P \in \mathcal{P}_{s,t}} \lambda_P \qquad \forall (s,t) \in E \qquad (28)
$$

Similarly, any optimal solution to (3-6) is an optimal solution to (19-22).

Proof: In Appendix. ■

The edge and path formulations of our problem are not equivalent in terms of the set of feasible solutions; however, Proposition 3.3 shows that the two formulations have a common optimum; this holds for both the primal and the dual. Henceforth, whenever we refer to an optimal dual solution, we mean an optimal solution which is common to both the path and edge formulations. A similar comment applies to the primal. Notice finally that in the above theorem when discussing optimal dual solutions, we did not explicitly mention the vector $\mu = (\mu_t : t \in T)$. This is because the optimal μ is induced from the other optimal dual variables via (12-13). Similarly we will often omit mention of μ below with this understanding in mind.

The near equivalence of the edge and path formulations of the dual is closely related to a well known path decomposition of network flows. Whereas the edge formulation specified a flow $\lambda(s,t)$ on edges (s,t), the path formulation specifies a flow λ_P on paths P. Indeed, parallel to the discussion of the edge formulation, (24) is a flow conservation constraint and (25) is related to the virtual valuation. The path decomposition mentioned above tells us that any flow on edges can be decomposed as a flow on paths and cycles. The decomposition (28) of Proposition 3.3 only decomposes the optimal flow on edges as a flow on paths, excluding cycles, and indeed (23-27) does not contain any variables indexed by cycles. Cycles can be excluded at the optimum in standard network flow problems such as the minimum cost flow problem, but as our problem differs somewhat,[9] to exclude cycles, we must appeal to (16) of Proposition 3.2, which tells us that we can always find an optimal dual solution avoiding bad edges. Any such optimum cannot have any cycles in its decomposition. Using the duality theorem, this also allows us to eliminate constraints corresponding to cycles from the path formulation of the primal.

Next observe that $\sum_{P:t\in P} \lambda_P$ is the total amount of flow that goes through t. This includes flow that terminates in t ($\sum_{P:\in\mathcal{P}_t} \lambda_P$) as well as flow that passes through t. Given an optimal dual solution λ denote by $\phi(s,t|\lambda)$ the fraction of all flow that either terminates or passes through t which came via s. Notice that

$$\phi(s,t|\lambda) = \frac{\lambda(s,t)}{\sum_{r:(r,t)\in E} \lambda(r,t)} = \frac{\sum_{P\in\mathcal{P}_{s,t}} \lambda_P}{\sum_{P:t\in P} \lambda_P} \tag{29}$$

We shall refer to $\phi(s,t|\lambda)$ as the **normalized flow** on (s,t).

For any path $P = (t_0, \ldots, t_k)$ (where t_0 is not necessarily 0), define:

$$\Phi(P|\lambda) := \prod_{i=1}^{k} \phi(t_{i-1}, t_i|\lambda) \tag{30}$$

Furthermore, let $\tau(P)$ be the terminal vertex of path P.

Lemma 3.4 *There exists an optimal dual solution λ satisfying (16) such that*

$$\lambda_P = \Phi(P|\lambda)\pi_{\tau(P)} \qquad \forall P \in \mathcal{P} \tag{31}$$

$$\sum_{P:(t_0,\ldots,t_k)\subseteq P} \lambda_P = \prod_{i=1}^{k} \phi(t_{i-1}, t_i|\lambda) \sum_{P:t_k\in P} \lambda_P \qquad \forall (t_1, \ldots, t_k) \in \mathcal{P} \tag{32}$$

Proof: See Appendix. ■

In general, any flow on edges has many path decompositions, all of which lead to the same objective function value. Property (31) of Lemma 3.4 says that we may always choose a particular path decomposition which has a certain special relation to the flow on edges. In particular, we may choose the path decomposition so that the flow on any path P is equal to the the probability of the terminal vertex of P multiplied by the product of normalized flows on edges in P. Any flow satisfying (31) also satisfies (32). We will call a flow on paths satisfying (31-32) a **proportional flow**.

[9] If instead of (23-27), we were dealing the the closely related minimum cost flow problem, we could argue that cycles could be excluded at the optimum because the extreme points of that problem do not contain cycles. However our problem is not quite identical to the minimum cost flow problem because it contains the additional constraint (9), which prevents us from immediately appealing to the standard argument.

4 The Bargaining Game

We show that the optimal mechanism can be implemented as a sequential equilibrium of a dynamic bargaining protocol in which the seller does not commit to his strategy ahead of time. The dynamic bargaining protocol is as follows:

Dynamic Bargaining Protocol

1. Nature selects a type $t \in T$ for the buyer with probability π_t.
2. The buyer either:
   (a) drops out and the interaction ends, or
   (b) makes a cheap talk report of $\hat{t}$ (where $\hat{t}$ is a type in T).
3. The seller either:
   (a) requests another cheap talk message, in which case we return to step 2 (this occurs at most $|T|$ times),
   (b) or requests evidence.
4. The buyer can
   (a) drop out and the interaction ends, or
   (b) present evidence $S \subseteq \sigma(t)$.
5. The seller makes a take-it-or-leave-it-offer.

Note At step 3, when the seller requests a cheap talk message or evidence, the seller does not specify which cheap talk message or which evidence is to be furnished.

The protocol is a model of bargaining between seller and buyer. As our main goal is to interpret the optimal mechanism, there is no discounting so that we think of this as a fast interaction. The buyer opens first with a claim/offer about the most she can pay. The seller can respond either by asking for another offer or demanding proof in return for sale at an announced price. Note that the buyer's cheap talk claims contain information about the evidence that she possesses as well as her value.

5 Equilibrium

We now describe an equilibrium of the bargaining protocol which implements the optimal mechanism. In this section, for economy of exposition, we employ a relatively weak solution concept, namely *Bayesian Nash equilibrium*. This requires only that the strategies of the players are mutual best replies. In section 6 we show how to strengthen our results using the stronger solution concept of *sequential equilibrium*, which requires sequential rationality off the equilibrium path with respect to beliefs that are consistent with the structure of the game.

Our plan is as follows. In section 5.1, we present the equilibrium strategies of the two players. In section 5.2, we verify that the strategies of section 5.1–if followed–implement the optimal mechanism. Section 5.3 establishes that the buyer's strategy is a best reply to the seller's strategy. Section 5.4 establishes that the seller's strategy is a best reply to the buyer's strategy. We show that the seller's problem may be interpreted as an optimal stopping problem. In section 5.5 we state a theorem bringing together the various arguments presented in this section. We also highlight some interesting qualitative properties of the equilibrium.

5.1 Equilibrium Strategies

Here we exhibit the equilibrium strategies that implement the optimal mechanism in the dynamic bargaining protocol. The seller's strategy depends on an optimal solution to the primal and the buyer's strategy depends on an optimal solution to the dual (these problems have been defined in section 3).

We first describe the buyer's strategy. Throughout the description we use t to denote the type chosen by nature. We may assume that after her type t is realized, the buyer performs a private preliminary randomization which guides her behavior throughout the course of the game. In particular the buyer randomizes over paths in $\mathcal{P}_t$ selecting path P with probability:

$$\frac{\lambda_P}{\pi_t}$$

where λ is an optimal dual solution avoiding bad edges and satisfying (31-32). Observe that (24) implies that these probabilities sum to one. Throughout the description of the buyer's strategy, $(t_0^*, \ldots, t_n^*)$ denotes the outcome of the preliminary randomization. The type t buyer reports along path $(t_0^*, \ldots, t_k^*, \ldots, t_n^*)$. If evidence is requested following cheap talk report t_k^*, she presents evidence $S_{t_k^*}$. She drops out if asked for more cheap talk after $t_n^*(= t)$.

We now present this buyer strategy a little more formally. The description is conditional on the realization of the buyer's type and the outcome of the preliminary randomization. In this case, we have three parts: first, what cheap talk reports to make; second, what evidence to offer when requested to do so; third, what offers to accept.

Buyer's Equilibrium Strategy Part 1

When the buyer is asked for the $(k+1)$th report:

- if the previous cheap talk reports were $(t_1^*, \ldots, t_k^*)$ and $k < n$, the buyer makes cheap talk claim t_{k+1}^*.
- Otherwise, the buyer drops out.

Buyer's Equilibrium Strategy Part 2

- If the buyer made reports $(t_0^*, \ldots, t_k^*)$ (for some $k \leq n$) prior to the seller's request for evidence, then following this request, the buyer presents evidence $S_{t_k^*}$.
- If the buyer made reports which do not correspond to an initial subsequence of the outcome of her preliminary randomization, then following the evidence request, she drops out.

Buyer's Equilibrium Strategy Part 3

If the buyer has a strict preference concerning the seller's take-it-or-leave-it offer, she follows her preference, and if indifferent, she accepts.

Next we present a description of the seller's equilibrium strategy in two parts. In the first part, we specify whether the seller asks for a cheap talk message or for evidence as a function of the history of cheap talk messages (i.e. sequence of types) sent by the buyer. In the second part, we describe how the seller responds when the buyer offers evidence in response to an evidence request. The seller's strategy depends on an optimal allocation $(q_t : t \in T)$ in the primal problem.[10]

In what follows it is useful to define $\lambda_{(t_0)} := 1$. (Since $\mathcal{P}$ was defined so as to exclude (t_0), this has not been previously defined).[11] In interpreting the seller's strategy, it is useful to keep in mind that if the buyer uses the strategy defined above and $\lambda_P = 0$, then the probability that the seller will see the sequence of reports P is zero; this follows from the fact that λ has been chosen to satisfy (31).[12]

[10]An optimal allocation $(q_t : t \in T)$ is an allocation for which there exists $(p_t : t \in T)$ such that $(q_t, p_t : t \in T)$ is optimal in the primal.

[11]$\mathcal{P}$ does not include (t_0) because we defined a path so that it must contain at least two vertices.

[12]In particular, (31) implies that $\sum_{P':P \subseteq P'} \lambda_{P'} > 0 \Leftrightarrow \lambda_P > 0$; this also implicitly relies on the fact that for all $t \neq 0$, $\pi_t > 0$.

Seller's Equilibrium Strategy Part 1

If the buyer made reports $P = (t_0, \ldots, t_k)$, then

- if $\lambda_P > 0$, then:
  - if $q_{t_{k-1}} = 1$, the seller requests evidence t_k.
  - if $q_{t_{k-1}} < 1$:
    - with probability $\frac{1-q_{t_k}}{1-q_{t_{k-1}}}$, the seller requests another cheap talk message, and
    - with probability $\frac{q_{t_k}-q_{t_{k-1}}}{1-q_{t_{k-1}}}$, the seller requests evidence.

    (Here we define $q_{t_{-1}} := 0$.)[13]
- if $\lambda_P = 0$, the seller requests evidence.

Seller's Equilibrium Strategy Part 2

If the buyer made reports $P = (t_0, \ldots, t_k)$ prior to the seller's request for evidence, and presented evidence S, then

- if $\lambda_P > 0$ and $S = S_{t_k}$, then the seller makes an offer at price v_{t_k}.
- Otherwise, the seller makes an offer at price:

$$\max\{v_r : S \subseteq \sigma(r)\}. \tag{33}$$

We emphasize again–as argued below–the above strategies constitute a Bayesian Nash equilibrium (i.e., they are mutual best replies), but not a sequential equilibrium. For sequential equilibrium, see section 6.

In what follows we refer to the buyer and seller strategies defined in this section as ζ^* and ξ^* respectively.

5.2 The Strategies Implement the Optimal Mechanism

We show that the strategies ζ^* and ξ^*, if followed, implement the same outcome as the optimal mechanism.

[13](17) of Lemma 3.2 and (28) imply that if $\lambda_P > 0$, then $q_{t_{k-1}} \leq q_{t_k}$.

For any path $P \in \mathcal{P}_t$, let $n_P + 1$ be the length of P (i.e., the number of vertices in P) and t_i^P be the ith vertex in P so that we may write $P = (t_0^P, \ldots, t_{n_P}^P)$, and moreover, let:

$$k_P := \begin{cases} \min\left\{k : q_{t_k^P} = 1\right\} & \text{if } q_{t_{n_P}^P} = 1; \\ n_P, & \text{otherwise.} \end{cases} \tag{34}$$

(17) of Lemma 3.2 and (28) imply that $q_{t_{k_P}^P} = q_{t_{n_P}^P} = q_t$ whenever $\lambda_P > 0$. Also recall the convention (from Section 5.1) that $q_{t_{-1}^P} = 0$. The strategy profile (ζ^*, ξ^*) induces a probability of sale for type t buyer of:

$$\begin{aligned} &\sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \sum_{k=0}^{k_P} \left[\prod_{i=1}^{k-1} \frac{1-q_{t_i^P}}{1-q_{t_{i-1}^P}}\right] \frac{q_{t_k^P} - q_{t_{k-1}^P}}{1-q_{t_{k-1}^P}} \\ &= \sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \sum_{k=0}^{k_P} \left(q_{t_k^P} - q_{t_{k-1}^P}\right) = \sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} q_{t_{k_P}} = q_t, \end{aligned}$$

where we have used the fact that by (24), $\sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} = 1$. Similarly, the expected payment of the type t buyer induced by (ζ^*, ξ^*) is:

$$\begin{aligned} &\sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \sum_{k=0}^{k_P} \left[\prod_{i=1}^{k-1} \frac{1-q_{t_i^P}}{1-q_{t_{i-1}^P}}\right] \frac{q_{t_k^P} - q_{t_{k-1}^P}}{1-q_{t_{k-1}^P}} v_{t_k^P} \\ &= \sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \sum_{k=0}^{k_P} \left(q_{t_k^P} - q_{t_{k-1}^P}\right) v_{t_k^P} = \sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \sum_{k=0}^{n_P} \left(q_{t_k^P} - q_{t_{k-1}^P}\right) v_{t_k^P} \\ &= \sum_{P\in\mathcal{P}_t} \frac{\lambda_P}{\pi_t} \left[v_{t_{n_P}^P} q_{t_{n_P}^P} - \sum_{k=1}^{n_P-1} q_{t_k^P} \left(v_{t_{k+1}^P} - v_{t_k^P}\right) - v_{t_1^P} q_{t_0^P} \right] = p_t, \end{aligned}$$

where the second equality uses Theorem 3.2 and (34) and the last equality uses complementary slackness. It follows that we implement the optimal mechanism.

5.3 Buyer Optimization

Here we prove that ζ^* is a buyer best reply to ξ^*. If the type t buyer had a profitable deviation she would have a profitable pure strategy deviation including some sequence of reports $P = (t_0, \ldots, t_k)$ which she would make before dropping out. We may assume that $P \in \mathcal{P}_s$ for some $s \in T$ with $\sigma(s) \subseteq \sigma(t)$ and $\lambda_P > 0$ because at any moment that it becomes evident to the seller that one of these conditions is violated, the buyer can no longer attain a positive utility given the seller's strategy and so the buyer may as well drop out.[14] However,

[14] Observe in particular that if $P = (t_0, \ldots, t_k)$, $\lambda_P > 0$ and $\sigma(t_i) \not\subseteq \sigma(t)$, then $\sigma(t_{i+1}) \not\subseteq \sigma(t)$, and so once t_i is reached any seller offer will be weakly above v_t. So the type t buyer may as well select the truncation of P which ends in the last type s in P for which $\sigma(s) \subseteq \sigma(t)$, and so drop out after s is reached.

it now follows from the arguments like those of section 5.2 that the buyer's payoff from this deviation would be $v_t q_s - p_s$. Incentive compatibility ((4) in the primal problem) implies this deviation would yield a payoff inferior to $v_t q_t - p_t$, which by the argument of section 5.2, is the payoff that the type t buyer would attain if she used ζ^*.

5.4 Seller Optimization: An Optimal Stopping Problem

In this section, we argue that ξ^* is a best reply to ζ^*.

Before proceeding it is useful to consider a few facts. Consider a seller strategy $\underline{\xi}$ which always requests another cheap talk message. One can show that $\lambda_P > 0$ exactly if P is a path that would be observed with positive probability if the seller used $\underline{\xi}$ against ζ^* (see footnote 12). So $\lambda_P > 0$ implies that P is a path, or sequence of reports, that the seller would observe with positive probability if he did not bring the cheap talk communication phase to an end by requesting evidence. Next, observe that given the buyer's strategy ζ^*, whenever the seller requests evidence following a sequence of reports $P = (t_0, \ldots, t_k)$ (with $\lambda_P > 0$), the buyer will present evidence S_{t_k}

Call a seller strategy a **stopping strategy** if it agrees with part 2 of the definition of the seller's strategy ξ^* (see section 5.1). If the seller uses a stopping strategy against ζ^*, then following any sequence of reports $P = (t_0, \ldots, t_k)$, if the seller requests evidence, the buyer will present evidence S_{t_k}, and the seller will make an offer at price v_{t_k}.

Lemma 5.1 *There exists a seller best reply to ζ^* which is a stopping strategy.*

Proof: Let ξ be a best reply to ζ^*. There exists a deterministic best reply to any buyer strategy, so for simplicity assume that ξ is deterministic. Consider a non-terminal history h satisfying (i) following h, it is the seller's turn to make an offer (step 5), and (ii) h occurs with positive probability if the players use strategy profile (ζ^*, ξ). Let $P = (t_0, \ldots, t_k)$ be the sequence of cheap talk reports which were made in h. (i-ii) imply that the buyer presented evidence S_{t_k}. Suppose that conditional on h, ξ offers a price p different than v_{t_k}. Then we may assume that $v_{t_k} < p$ because given ζ^*, all buyer types consistent with h have value at least equal to v_{t_k}. Now consider a seller strategy ξ' that agrees with ξ except on histories following the sequence of cheap talk reports P. Following P, ξ' continues to request cheap talk reports until the buyer presents a cheap talk report s with $v_s \geq p$, at which point ξ' requests evidence, and then behaves as in a stopping strategy, making an offer of v_s if the appropriate evidence is presented and offering (33) otherwise. Then notice that conditional on the initial sequence of reports P, ξ and ξ' will lead to the same collection of buyer types being served, but each such buyer type will pay a weakly higher price under ξ' than under ξ. Since ξ was a best reply, it follows that ξ' is also a best reply. By a sequence of such modifications we can turn the strategy ξ into a seller strategy ξ^0 which is a stopping strategy and also a best reply to ζ^*.[15] ■

Lemma 5.1 implies that in searching for a best reply to ζ^*, we can restrict attention to stopping strategies. Since ξ^* is a stopping strategy, it suffices to show that ξ^* is better than all other stopping strategies. This allows us to think of the seller's problem as an optimal stopping problem. *Stopping* corresponds to requesting evidence, and *continuing* corresponds to requesting another cheap talk message. Conditional on stopping, there is no further decision for the seller to make because we restrict attention to strategies where the seller offers a price v_{t_k} where t_k was the last cheap talk claim made by the buyer,[16] and all buyer types which have not dropped out by this point will accept so that the seller's payoff will be v_t. If the seller continues, with some probability the buyer drops out, giving the seller payoff of zero, and with some probability the buyer makes another report t_{k+1}. The stochastic process which the seller is facing is endogenous because the distribution of reports t_{k+1} is determined by an optimal dual solution λ. Note that stopping strategies allow the seller to randomize the decision of whether to stop.

Next we characterize the beliefs that the seller has as bargaining progresses.

Lemma 5.2 *For any seller strategy ξ, if $P = (t_0, \ldots, t_k)$ is a sequence of cheap talk reports that the seller observes with positive probability given (ζ^*, ξ), from the seller's perspective, the conditional probability that the buyer would–if given the opportunity–present another cheap talk message (rather than dropping out) and moreover would present cheap talk message t_{k+1} given that she has already presented P is:*

$$\frac{\sum_{P:(t_k,t_{k+1})\in P} \lambda_P}{\sum_{P:t_k\in P} \lambda_P}$$

Proof: The probability that in the preliminary randomization, the buyer selected a sequence P with $(t_0, \ldots, t_k) \subseteq P$ is:

$$\sum_{t\in T} \sum_{P\in\mathcal{P}_t:(t_0,\ldots,t_k)\subseteq P} \pi_t \frac{\lambda_P}{\pi_t} = \sum_{P\in\mathcal{P}:(t_0,\ldots,t_k)\subseteq P} \lambda_P = \prod_{i=1}^{k} \phi(t_{i-1}, t_i|\lambda) \sum_{P:t_k\in P} \lambda_P \qquad (35)$$

where the last equality follows from our assumption that λ satisfies (31-32). Similarly, the probability that the buyer selected P with $(t_0, \ldots, t_k, t_{k+1}) \subseteq P$ is

$$\prod_{i=1}^{k} \phi(t_{i-1}, t_i|\lambda) \sum_{P:t_{k+1}\in P} \lambda_P \qquad (36)$$

[15] In order for ξ^0 to be a stopping strategy, we may have to make some additional modifications conditional on histories which occur with zero probability, and hence do not affect the seller's payoff.

[16] Conditional on stopping, ζ^* is such that (on the equilibrium path) the buyer will always present evidence S_{t_k}.

Dividing (36) by (35) and using (29), the desired conditional probability is:

$$\frac{\phi(t_k, t_{k+1}|\lambda) \sum_{P:t_{k+1}\in P} \lambda_P}{\sum_{P:t_k\in P} \lambda_P} = \frac{\sum_{P:(t_k,t_{k+1})\in P} \lambda_P}{\sum_{P:t_k\in P} \lambda_P}$$

■

Lemma 5.3 *Suppose the buyer uses ζ^* and let $P = (t_0, \ldots, t_k)$ and $\lambda_P > 0$. Restricting attention to stopping strategies that induce history P with positive probability:*

1. *If $q_{t_k} > 0$, then conditional on P, the seller is weakly better off stopping immediately then continuing for one more step and then stopping.*
2. *If $q_{t_k} < 1$, then conditional on P, the seller is weakly better off continuing for one more step and then stopping than stopping immediately.*

Proof: Using Lemma 5.2, the seller's preference between stopping now and stopping in one step is determined according to the resolution of the following inequality:

$$\underbrace{v_{t_k}}_{\text{stopping now}} \gtreqless \underbrace{\sum_{t_{k+1}\in T} v_{t_{k+1}} \frac{\sum_{P:(t_k,t_{k+1})\in P} \lambda_P}{\sum_{P:t_k\in P} \lambda_P}}_{\text{stopping in one step}}$$

This is equivalent to:

$$v_{t_k} \sum_{\{P:t_k\in P\}} \lambda_P \gtreqless \sum_{t_{k+1}\in T_{k+1}} v_{t_{k+1}} \sum_{\{P:(t_k,t_{k+1})\in P\}} \lambda_P \tag{37}$$

Using (24), the LHS of (37) can be rewritten:

$$v_{t_k} \sum_{P\in\mathcal{P}_{t_k}} \lambda_P + v_{t_k} \sum_{t_{k+1}} \sum_{\{P:(t_k,t_{k+1})\in P\}} \lambda_P = v_{t_k}\pi_{t_k} + v_{t_k} \sum_{t_{k+1}} \sum_{\{P:(t_k,t_{k+1})\in P\}} \lambda_P.$$

Substituting into (37) and rearranging gives

$$v_{t_k}\pi_{t_k} \gtreqless \sum_{t_{k+1}\in T} (v_{t_{k+1}} - v_{t_k}) \sum_{P\in\mathcal{P}_{t_k,t_{k+1}}} \lambda_P \tag{38}$$

To analyze (38) we invoke complementary slackness. If $q_t < 1$, then $\mu_t = 0$, which implies via (25) that $\gtreqless$ becomes $\leq$ establishing part 2 of the lemma. On the other hand if $q_t > 0$, then (25) holds with equality, which implies that $\gtreqless$ becomes $\geq$ establishing part 1 of the lemma. ■

We now argue by backward induction that ξ^* is optimal among all stopping strategies. Consider a history $P = (t_0, \ldots, t_k)$ with $\lambda_P > 0$. First let P be such a history of maximal length.[17] (Here the length of P is the number of vertices in P). In this case, we must have $q_{t_k} = 1$,[18] and clearly it is optimal to stop as required by ξ^*. Now suppose we have established the result for all histories P' (with $\lambda_{P'} > 0$) that are longer than P. First suppose that $q_{t_k} > 0$. It follows from Proposition 3.2 that for all $P' = (t_0, \ldots, t_k, t_{k+1})$ with $\lambda_{P'} > 0$, $q_{t_{k+1}} > 0$. It follows from the inductive hypothesis that conditional on any such P', it would be optimal for the seller to stop. Lemma 5.3 now implies that following P, stopping immediately would be optimal as required by ξ^*. Next suppose that $q_{t_k} < 1$. Then by Lemma 5.3, the seller would be weakly better off continuing one step and then stopping than stopping immediately, and so continuing and then following ξ^* (which by backwards induction, is optimal) would be even better, again as required by ξ^*.

It now follows from Lemma 5.1 that ξ^* is a best reply to ζ^*.

5.5 Summary of the Argument

We summarize the argument given above.

Theorem 5.4 (ζ^*, ξ^*) *is a Bayesian Nash equilibrium of the dynamic bargaining protocol which implements the optimal mechanism.*

The following proposition gives some of the qualitative properties of the equilibrium.

PROPOSITION 5.5 *Let $(t_0, t_1, \ldots, t_k)$ be any sequence of cheap talk reports that occur with positive probability in the equilibrium described above. Then:*

$$v_{t_0} < v_{t_1} < \cdots < v_{t_n} \tag{39}$$

$$S_{t_0} \subseteq S_{t_2} \subseteq \cdots \subseteq S_{t_n} \tag{40}$$

Proof: This follows from the fact that $\sum_{P \in \mathcal{P}_{s,t}} \lambda_P > 0$ implies that both $v_s < v_t$ and $S_s \subseteq S_t$. The former inequality uses (28) and the fact we have chosen an optimal λ to avoid bad edges in accordance with Lemma 3.2, while the latter inclusion does not even depend on the optimality of λ but merely invokes (1). ■

In each round of the bargaining protocol the seller can ask the buyer for another cheap talk message or for the presentation of evidence supporting the buyer's current cheap talk

[17] Such a history exists because T is finite λ has no bad edges; in other words, a sequence of cheap talk reports cannot form a cycle.

[18] Suppose that $q_{t_k} < 1$. Then by complementary slackness $\mu_{t_k} = 0$. But then the fact that $v_{t_k}\pi_{t_k} > 0$ and (25) imply that there must exist $t_{k+1} \in T$ with $v_{t_k} < v_{t_{k+1}}$ and $P' \in \mathcal{P}_{t_k, t_{k+1}}$ such that $\lambda_{P'} > 0$. It follows that $\phi(t_k, t_{k+1}|\lambda) > 0$. Since λ avoids bad edges, $t_{k+1} \notin (t_0, \ldots, t_k)$. So consider the path $P'' = (t_0, \ldots, t_k, t_{k+1})$. By Lemma 3.4, $\lambda_{P''} = \lambda_P \phi(t_k, t_{k+1}|\lambda) \frac{\pi_{t_{k+1}}}{\pi_{t_k}} > 0$, contradicting the assumption that P was of maximal length.

communication, ending the cheap talk phase. While not required by the protocol, in equilibrium, the buyer will make a sequence of concessions, claiming to have successively higher valuations or else drop out if the seller asks him to make concessions too many times ((39) of Proposition 5.5). The cheap talk phase also contains cheap talk claims about the evidence the buyer could present if called upon to do so. In equilibrium, these claims are also increasing, in the sense that the buyer claims to have more and more evidence as the protocol progresses ((40) of Proposition 5.5). The motive for withholding evidence is that the evidence may suggest that the buyer has a higher value. As the buyer admits to having successively higher values, she admits to possessing successively more evidence which she previously withheld.

All along, the buyer's communication is constrained by the need to present supporting evidence at the end. When the seller finally requests evidence, the buyer presents the evidence she claimed to have at her last cheap talk claim–which due to the increasing nature of the evidential claims is the cumulative evidence that she has claimed to have during the procedure. When the buyer presents the evidence she claimed to have in his most recent cheap talk claim, the seller makes an offer to sell the object at the value the buyer claimed to have during her most recent cheap talk claim. If the buyer were to fail to present this evidence, the seller would offer to sell the object only at a high price. The seller faces an optimal stopping problem: should he ask for a further concession from the buyer, risking the possibility that the buyer will be unwilling to make one and thus drop out? Updating on the equilibrium reporting strategy, each further cheap talk report gives the seller more information about the buyer's type, but at some point it is no longer worthwhile for the seller to acquire additional information for fear of driving the buyer away.

6 Sequential Rationality off Equilibrium

In the previous section we exhibited a Bayesian Nash equilibrium of the dynamic bargaining protocol which implements the optimal mechanism. In this section, we establish that this can be extended to sequential equilibrium. Ordinarily, the definition of sequential equilibrium is problematic when players have a continuum of pure strategies. However, in our game only the seller has a continuum of pure strategies (the buyer has finitely many strategies) and the seller has no private information. Thus, the issue of disciplining the buyer's beliefs about the seller's type–which would be the source of the problem–is moot.[19]

Theorem 6.1 *There exists a sequential equilibrium of the dynamic bargaining protocol which induces the same probability distribution over terminal histories as* (ζ^*, ξ^*) *and thereby implements the optimal mechanism.*

[19] It is also inessential that the seller has a continuum of strategies as we could restrict the seller to make offers from the set $\{v_t : t \in T\}$ without materially changing the game.

The proof is in the appendix. The strategies associated with this equilibrium are somewhat more complicated than those presented in Section 5.1. All the added complexity involves only behavior off the equilibrium path. This result assures us that the equilibrium studied in the previous section does not rely on non-credible threats.

7 Many Rounds of Communication

In this section, we present an example that illustrates various of the ideas discussed above. In particular, the example is of interest because when the equilibrium constructed in Section 5.1 is applied to this example, it involves many rounds of cheap talk communication. Indeed, the equilibrium involves arbitrarily many rounds depending on the choice of the number n below. The appendix contains proofs of various claims we make below about this example.

Let $T = X \cup Y \cup \{0\}$, where $X = \{x_1, \ldots, x_n\}, Y = \{y_0, y_1, \ldots, y_n\}$. So we partition the set of types (other than the zero type) into two sets X and Y. Let us refer to the types in X as x-types, and the types in Y as y-types. As usual, we assume that $v_0 = \pi_0 = 0$. Moreover we assume that:

$$v_{x_1} < v_{x_2} < \cdots < v_{x_n} < v_{y_0} < v_{y_1} < v_{y_2} < \cdots < v_{y_n} \tag{41}$$

This means that within the set of x-types and within the set of y-types, valuations are strictly increasing in the indices of the type. However, all y-types have higher valuations than all x-types. The incentive graph is given by:

$$\begin{aligned} E = \{(0,t) : t \in T \setminus 0\} \cup \{(x_i, x_j) \in X \times X : i < j\} \cup \{(x_i, y_j) \in X \times Y : i \leq j\} \\ \cup \{(y_i, y_j) \in Y \times Y : i < j\} \end{aligned} \tag{42}$$

This can be represented pictorially:

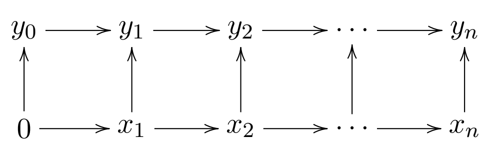

Figure 3: The Incentive Graph

Each directed path in Figure 3 corresponds to an edge in the incentive graph (42). So all types can mimic the zero type, all x-types can mimic lower index x-types, all y-types can mimic lower index y-types, and y-types can also mimic x-types with a weakly lower index.

The above incentive graph can be induced by the message structure

$$\sigma(t) = \{m_s : (s,t) \in E\} \cup \{m_t\} \qquad \forall t \in T \setminus 0$$
$$\sigma(0) = \{m_0\}$$

Where we assume that if $s \neq t$, then m_s and m_t are distinct messages.

We now make some assumptions which allow us to explicitly solve for the optimal mechanism. First, a definition is useful. Working backwards from n, we recursively define:

$$\delta_n := 0 \tag{43}$$
$$\delta_{i-1} := \delta_i + \frac{v_{x_i}\pi_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\delta_i + \sum_{j=i+1}^{n} \pi_{x_j}\right]}{v_{y_i} - v_{x_i}} \qquad \forall i = 1, \ldots, n \tag{44}$$

Note that when $i = n$, we define $(v_{x_{i+1}} - v_{x_i})\left[\delta_i + \sum_{j=i+1}^{n} \pi_{x_j}\right] := 0$, and similarly in (47) below, when $i = n$, we define $(v_{x_{i+1}} - v_{x_i})\left[\frac{\sum_{j=i+1}^{n} \pi_{x_j} + \sum_{j=i+1}^{n} \pi_{y_j}}{\pi_{x_i}}\right] := 0$. We also assume that:

$$v_{y_i} - (v_{y_{i+1}} - v_{y_i})\frac{\sum_{j=i+1}^{n} \pi_{y_j}}{\pi_{y_i}} > 0 \quad \forall i = 0, 1, \ldots, n-1 \tag{45}$$

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i})\frac{\sum_{j=i+1}^{n} \pi_{x_j}}{\pi_{x_i}} > (v_{x_{i+1}} - v_{x_i})\frac{\delta_i}{\pi_{x_i}} \qquad \forall i = 1, \ldots, n-1 \tag{46}$$

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\frac{\sum_{j=i+1}^{n} \pi_{x_j} + \sum_{j=i+1}^{n} \pi_{y_j}}{\pi_{x_i}}\right] - (v_{y_i} - v_{x_i})\frac{\pi_{y_i}}{\pi_{x_i}} < 0 \qquad \forall i = 1, \ldots, n \tag{47}$$

For any profile of valuations satisfying (41), there are many probability distributions $(\pi_t : t \in T)$ such that (45-47) are satisfied. (45) implies that if (aside from the zero type) there were only y-types (where we take the restriction of the incentive graph to these types and the probabilities re-normalized to sum to one), then the optimal allocation would allocate the object to each type with probability 1.

Similarly, (46) implies that if there were only x-types, then it would be optimal to allocate the object to all types. But the assumption (46) for x-types is a stronger assumption than the corresponding assumption (45) for y-types. Indeed, a simple induction using (46) implies:

$$\delta_i > 0 \qquad \forall i = 0, 1, \ldots, n-1, \tag{48}$$

and hence (46) also implies:

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i})\frac{\sum_{j=i+1}^{n} \pi_{x_j}}{\pi_{x_i}} > 0 \qquad \forall i = 1, \ldots, n-1 \tag{49}$$

(47) says that if the set of types were of the form $T_i := \{x_i, x_{i+1}, \ldots, x_n\} \cup \{y_i, y_{i+1}, \ldots, y_n\} \cup \{0\}$ for $i = 1, \ldots, n$, and the incentive graph were $\{(s,t) \in (T_i \setminus \{y_i\}) \times (T_i \setminus \{y_i, 0\}) : s \neq t\} \cup \{(x_i, y_i)\}$, then it would be optimal not to allocate the object to x_i. Notice that T_1 differs from T because T contains y_0 whereas T_1 does not; we do not define a set T_0 because there is no type x_0.

Given the above assumptions, the optimal prices and allocation for the y-types are given by:

$$q_{y_i} := 1 \qquad \forall i = 0, 1, \ldots, n \tag{50}$$

$$p_{y_i} := v_{y_0} \qquad \forall i = 0, 1, \ldots, n \tag{51}$$

For the x-types the optimal allocation and prices can be defined recursively as follows:

$$q_{x_1} := \frac{v_{y_1} - v_{y_0}}{v_{y_1} - v_{x_1}} \tag{52}$$

$$p_{x_1} := v_{x_1} q_{x_1} \tag{53}$$

$$q_{x_i} := \frac{(v_{y_i} - v_{y_0}) - (v_{x_i} q_{x_{i-1}} - p_{x_{i-1}})}{v_{y_i} - v_{x_i}} \qquad \forall i = 2, \ldots, n \tag{54}$$

$$p_{x_i} := v_{x_i}(q_{x_i} - q_{x_{i-1}}) + p_{x_{i-1}} \qquad \forall i = 2, \ldots, n \tag{55}$$

It is straightforward to verify that:

$$0 < q_{x_1} < q_{x_2} < \cdots < q_{x_n} < 1 \tag{56}$$

Moreover, the mechanism given by (50-55) is the unique optimal mechanism. It is also the case that at every dual optimal solution we have that:

$$\lambda(x_{i-1}, x_i) > 0 \qquad \forall i = 2, \ldots, n \tag{57}$$

$$\lambda(0, x_1) > 0 \tag{58}$$

$$\lambda(x_h, x_i) = 0 \qquad \forall i = 2, \ldots, n, \forall h < i-1 \tag{59}$$

$$\lambda(0, x_i) = 0 \qquad \forall i = 2, \ldots, n \tag{60}$$

It follows that the unique path P from 0 to x_n with $\lambda_P > 0$ is $P = (0, x_1, x_2, \ldots, x_n)$, so type x_n is asked to submit a cheap talk report $n+1$ times in the equilibrium constructed in Section 5.1. To describe the equilibrium in more detail, each x-type x_i, uses the sequence of reports

$(0, x_1, x_2, \ldots, x_i)$, dropping out if the seller requests another message after x_i. Each y-type y_i randomizes over two sequences of reports at the preliminary phase: $(0, x_1, x_2, \ldots, x_i, y_i)$ and $(0, y_0, y_1, \ldots, y_i)$. If the seller receives the report y_0, he requests evidence and then given that evidence $S_{y_0} = \{m_0, m_{y_0}\}$ is presented, he makes a take-it-or-leave-it offer at price v_{y_0}. If the seller receives the report x_i, he randomizes between asking for another cheap talk report and and requesting evidence.

8 Special Cases

In this section, we discuss some special cases of our model under which the optimal mechanism or dynamic bargaining protocol simplify.

8.1 Deterministic Optima

We have already seen in Example 2 and Section 7 that in general there may not exist an optimal mechanism which is deterministic. Our first result concerns the case where there exists an optimal mechanism which is deterministic.

PROPOSITION 8.1 *If there exists an optimal mechanism which is deterministic (i.e., with $q_t \in \{0, 1\}, \forall t \in T$), then there exists a sequential equilibrium of the dynamic bargaining protocol implementing the optimal mechanism in which the buyer always makes only one cheap talk report before the seller requests evidence.*

The proof of this theorem depends on a modification of the arguments supporting the results in sections 5-6. Due to the overlap, we only sketch the proof here.

First, we present the strategies. As in section 5.1, the seller's strategy depends on the allocation $q = (q_t : t \in T)$ at the optimal mechanism.

Seller's Equilibrium Strategy

1. After the buyer's first report t, the seller requests evidence.
2. If the buyer presented report t and then evidence S, then
   - if $q_t = 1$ and $S = S_t$, the seller makes an offer at price v_t.
   - otherwise, the seller makes an offer at price:

$$\max\{v_r : S \in \sigma(r)\}.$$

To describe the buyer's strategy, we require an optimal dual solution λ avoiding bad edges and satisfying (31-32). If the optimal allocation is deterministic, then (17) of Lemma 3.2 and (28) imply that for all $t \in T$ with $q_t = 1$ and all $P = (t_0, \ldots, t_k) \in \mathcal{P}_t$ with $\lambda_P > 0$, there exists a unique vertex t_i such that $q_{t_j} = 1$ for all $j \geq i$ and $q_{t_j} = 0$ for all $j < i$. Call this vertex $\theta(P)$. The following specifies the buyer's behavior when she is of type t:

Buyer's Equilibrium Strategy

1. At the cheap talk stage,

   - if $q_t = 1$, the buyer presents cheap talk message r with probability:

$$\frac{\sum_{P \in \mathcal{P}_t : \theta(P) = r} \lambda_P}{\pi_t}. \tag{61}$$

   - otherwise, the buyer drops out.

2. At the evidence presentation stage, if $q_t = 1$ and the buyer presented cheap talk message r with $S_r \subseteq \sigma(t)$, then the buyer presents evidence S_r.

As in section 5.1, we assume that the buyer will accept a seller offer unless she strictly prefers to reject. The above descriptions of the strategies are incomplete insofar as they do not specify behavior in all counterfactual histories. Following the arguments similar to those in section 5 and the proof of Theorem 6.1 in the appendix, one can extend these strategies to all counterfactual histories in such a way that they form a sequential equilibrium implementing the optimal mechanism. The key observation in modifying the above-mentioned arguments is that a deterministic mechanism induces a deterministic seller strategy in section 5. Moreover, this deterministic strategy decides whether to request evidence and what price to offer (given that appropriate evidence is presented) based only on the last cheap talk report presented. Once randomization has been eliminated from the seller's strategy, there is no point in having the buyer present a sequence of reports. If the buyer is of a type that will ultimately be served, she may as well immediately present a message which triggers an evidence request (under the seller's strategy in section 5).[20] If the buyer will not be served, she may as well drop out immediately. We modify the seller's strategy accordingly to be based only on this single cheap talk report. The seller experiences some loss of information relative to the old equilibrium, but this does not force him to reconsider his strategy because in the old equilibrium his strategy was in all essential respects measurable with respect to the information that he now receives.

One can go further and eliminate the seller's evidence request altogether. Since the seller now immediately requests evidence conditional on any report, we may as well consider a game in which the buyer presents his evidence at the same time that she presents her cheap talk message, and the seller simply makes a take-it-or-leave-it offer. If $q_t = 1$, the type t

[20] We do not mean that holding fixed ξ^*, it is a best reply for the buyer to claim to be the type that will ultimately be served. Rather we may alter ξ^* such that it is a best reply for the buyer to claim to be the type that will ultimately be served, and this alteration will not affect anything of substance.

buyer presents the report-evidence pair (r, S_r) with probability (61). However we emphasize that such a simplification is only possible when the optimal mechanism is deterministic, because when the optimal mechanism is random, the buyer generally cannot predict with certainty whether the seller will request evidence and proceed to the take-it-or-leave-it offer prior to making her claim.

We conclude this section by providing an example which shows that despite implementing a deterministic mechanism the equilibria of the bargaining protocol described in this section still involve randomization on the part of the buyer. Such randomization is generally unavoidable when implementing the optimal deterministic mechanisms with our dynamic bargaining protocol.

Example 3 *Suppose that* $T = \{0, x_1, x_2, y, z\}$ *where* $v_0 = \pi_0 = 0$*,* $v_{x_i} = 1$ *and* $\pi_{x_i} = 1/5$ *for* $i = 1, 2$*,* $v_y = v_z = 2$*, and* $\pi_y = \pi_z = 3/10$*. Suppose further that* $\sigma(0) = \{m_0\}, \sigma(x_1) = \{m_0, m_1\}, \sigma(x_2) = \{m_0, m_2\}, \sigma(y) = \{m_0, m_1, m_2\}, \sigma(z) = \{m_0, m_3\}$*. Then the optimal mechanism is such that* $q_0 = p_0 = 0$*,* $q_t = p_t = 1$ *for* $t \in \{x_1, x_2, y\}$ *and* $q_z = 1$ *and* $p_z = 2$*. On the other hand, in any sequential equilibrium of the dynamic bargaining protocol in which the seller uses a pure strategy, the buyer must use a mixed strategy. One class of equilibria implementing the optimum are such that if the seller sees a cheap talk claim* $t \in \{x_1, x_2, z\}$*, he requests evidence, and if the corresponding evidence* S_t *is presented, the seller makes a take-it-or-leave-it offer at price* v_t*; on the other side, the type* 0 *buyer drops out immediately, types* x_1, x_2*, and* z *report their types truthfully at the first stage of the cheap talk phase, and type* y *claims to be type* x_1 *with probability* α *and type* x_2 *with probability* $1 - \alpha$ *where* $1/3 \leq \alpha \leq 2/3$*. If* α *were greater than* $2/3$*, then conditional on seeing a cheap talk report* x_1 *and evidence* $S_{x_1} = \{m_0, m_1\}$*, the seller would prefer to make a take-it-or-leave-it offer at price* 2 *(or even* $2 - \epsilon$ *to guarantee acceptance) than to make an offer at price* $v_{x_1} = 1$*, as required. Similarly, the equilibrium would unravel if* $\alpha < 1/3$*. Certainly, within this class of equilibria, the type* y *buyer cannot play a pure strategy. Indeed, there does not exist any pure strategy sequential equilibrium implementing the optimum.*

8.2 Deterministic Equilibria

Motivated by Example 3, in this section we present a condition under which there exist pure strategy equilibria of the dynamic bargaining protocol implementing the optimal mechanism (which of course also implies that the optimal mechanism is deterministic). In this case, we also present an explicit formula for the virtual value and an explicit solution for the optimal mechanism. The case considered here encompasses the standard case in which all types can mimic all other types (the complete graph) and the virtual values are monotone, but is considerably more general. Our explicit solutions for the virtual valuations and the optimal mechanism generalize the well-known solutions for the standard case.

Define a **(directed) tree** to be a graph such that for some vertex 0 (called the **root**) and any other vertex t, there is a unique path from 0 to t. A graph $G = (V, E)$ is **acyclic** if it does not contain a cycle (that is, there does not exist a sequence $(t_0, t_1, \dots, t_n)$ of vertices with $(t_{i-1}, t_i) \in E$ for $i = 1, \dots, n$ and $t_0 = t_n$. For any graph G the **transitive reduction of** G is the smallest subgraph of G whose transitive closure is equal to the transitive closure of G.[21] For any acyclic graph, the transitive reduction exists and is unique.

Given an incentive graph (T, E), define

$$E^+ := \{(s,t) \in E : v_s < v_t\}$$

So E^+ is the set of good edges. It is easy to see that the graph $G^+ := (T, E^+)$ is acyclic. If the transitive reduction of G^+ is a tree, then we say that the incentive graph has **tree structure**. It is easy to see that if G^+ is a tree, it has root 0. The standard case of the complete graph has tree structure. In that case, we enumerate the types $\{0, 1, \dots, n\}$ so that if $i < j$ $v_i < v_j$, the transitive reduction of G^+ contains precisely the edges of the form $(i, i+1)$; so there will be edges only among adjacent types pointing from a lower type to a higher type. More generally, the transitive reduction will have a similar character, with edges connecting adjacent types, but the types will not necessarily form a linear order but rather will allow for a more general tree structure. Example 1 falls into this more general case of tree structure. Figure 1 actually displays not the entire graph G^+ (or G)[22] but rather only the edges in the transitive reduction of G^+.

If the incentive graph has tree structure, then for any $t \in T$, let P^t refer to the unique path from 0 to t in the transitive reduction of G^+. Let $\varphi(t)$ be the unique vertex preceding t in P^t. $\varphi(t)$ is the **predecessor** of t.

Under tree structure, for each $t \in T$, define:

$$\widehat{\psi}_t := v_t - \sum_{s:t=\varphi(s)} (v_s - v_t) \frac{\pi_s + \sum_{r:(s,r)\in E^+} \pi_r}{\pi_t}$$

We refer to $\widehat{\psi}_t$ as type t's **quasi-virtual valuation on a tree**. Such quasi-virtual valuations have expressions analogous to the virtual valuation in Myerson (1981) but adapted to the tree structure of the incentive graph. However, if the notion of virtual value is defined in terms of the Lagrangian for the mechanism designer's problem as in section 10.5 of Myerson (1991), or equivalently if it is defined as in (12), the virtual value need not coincide with the quasi-virtual value.

[21] For a more formal definition and treatment see Aho, Garey & Ullman (1972).
[22] In Example 1, it so happens that $G^+ = G$, but this is inessential.

We say that the incentive graph has **single-crossing quasi-virtual valuations** if:

$$(s,t) \in E^+ \Rightarrow \left(\widehat{\psi}_s \geq 0 \Rightarrow \widehat{\psi}_t \geq 0\right)$$

A stronger condition implies single-crossing quasi-virtual valuations is **monotone quasi-virtual valuations**:

$$(s,t) \in E^+ \Rightarrow \widehat{\psi}_s \leq \widehat{\psi}_t$$

The advantage of this latter stronger condition is that it would be more promising if we were to attempt to extend Proposition 8.2 to a setting with multiple buyers.

PROPOSITION 8.2 *Assume that the incentive graph has tree structure and single-crossing quasi-virtual valuations. Then there exists an optimal dual solution such that:*

$$\widehat{\psi}_t = \psi_t, \qquad \forall t \in T \tag{62}$$

An optimal mechanism is given by:

$$q_t = \begin{cases} 1, & \text{if } \widehat{\psi}_t \geq 0; \\ 0, & \text{otherwise.} \end{cases} \qquad p_t = \begin{cases} \min\left(\{v_t\} \cup \{v_s : (s,t) \in E, \widehat{\psi}_s \geq 0\}\right), & \text{if } \widehat{\psi}_t \geq 0; \\ 0, & \text{otherwise.} \end{cases} \tag{63}$$

Moreover, there exists a deterministic equilibrium involving only one round of cheap talk in the bargaining protocol implementing the optimal mechanism. In this equilibrium, a type t buyer claims to be the first type r in P^t with $\widehat{\psi}_r \geq 0$ if $\widehat{\psi}_t \geq 0$ and immediately drops out if $\widehat{\psi}_t < 0$. The seller offers price v_r following a cheap talk claim r if r is the unique type in P^r with non-negative virtual valuation and the appropriate verifying evidence S_r is presented.

Proof: In Appendix. ■

In addition to tree structure, Example 1 satisfies monotone quasi-virtual values (hence also single-crossing quasi-virtual values), and so illustrates the proposition. In that example, the types with non-negative virtual values are the types who are served, and the price paid by any such type t is the value of the first type s with non-negative virtual value on the unique path from 0 to t in the transitive reduction of G^+. The difference between the solutions for K small and K large in that example is explained by the fact that when K is small, all types except 0 and 1 have non-negative quasi virtual values, whereas when K is large only types $4, 5, 6$, and 7 have non-negative virtual values, so that increasing K causes virtual values to turn positive further up on the tree.

8.3 Binary Values

In this section, we consider the special case in which all types (except type 0) are of one of two values v_L and v_H with $0 < v_L < v_H$. Notice however that different types with the same value generally have different messages. This case is closely related to the persuasion models presented in Glazer & Rubinstein (2004), Glazer & Rubinstein (2006), and Sher (2010). That model involved a speaker and listener. The speaker knows the state in a set $x \in X$ and the listener does not. In each state x the speaker has a set of messages $\sigma(x)$. The listener must make a binary decision, to either accept or reject the speaker's request. The listener would like to accept the speaker's request if $x \in A$ (the set of **accept states**), and reject the speaker's request if x belongs to the complementary set R (the set of **reject states**.)

One difference is that in the model of persuasion, it is assumed that the speaker can only present one message in $\sigma(x)$ whereas in our model, the buyer can produce any subset of $\sigma(t)$. However define:

$$\sigma^*(t) := \{S : S \subseteq \sigma(t)\} \tag{64}$$

This is a message correspondence in which the individual messages are in fact sets of messages. Assuming the buyer can produce one message in $\sigma^*(t)$ is equivalent to assuming that the buyer may present any number of messages from $\sigma(t)$. Message structures of the form (64) are strictly less general than general message structures when one assumes that the speaker/buyer can only present one hard message. Such message structures correspond roughly to the class of normal message structures (see Bull & Watson (2007), Sher (2010)) and for the purpose of the persuasion model, there is no additional generality in allowing for message correspondences that are normal but not of the form (64).

Glazer & Rubinstein (2004) proved a credibility result showing that there is no value to commitment for the listener in this problem, and Sher (2010) generalized this result to allow for a broader class of message correspondences.[23] The dynamic game implementing the optimum involves the speaker presenting for a cheap talk claim, followed by an evidence request by the listener and then a possibly random decision (accept or reject) contingent on the cheap talk claim and evidence presented.

[23]In the context of a different version of the model which does not allow evidence presentation to be preceded by back-and-forth cheap talk communication, Glazer & Rubinstein (2006) proved a distinct credibility result.

PROPOSITION 8.3 *With binary values, the model presented here is equivalent to the optimal persuasion model in Sher (2010) (a generalization of Glazer & Rubinstein (2004)) when the message structure is given by (64), the set of states is the set of types $T \setminus 0$, and:*

$$A := \{t : v_t = v_L\} \tag{65}$$

$$R := \{t : v_t = v_H\} \tag{66}$$

Sher (2010) is strictly more general than the binary value version of the model presented here because it does not require the message correspondence to be of the form (64).

The translation between the price discrimination and persuasion models established by Proposition 8.3 is as follows: The speaker is a buyer who would like to persuade the listener, a seller, to charge a low price v_L. The seller would like to accept the request and charge a low price if the buyer has a low value and reject the request charging a high price if the buyer has a high value.

With binary values our main theorem corresponds to the credibility result for the optimal persuasion model.

Sher (2010) showed that when the message structure is normal, then there exists an optimal dynamic persuasion rule which is deterministic.[24] An optimal dynamic persuasion rule corresponds to an optimal mechanism in our setting.[25] It then follows from Propositions 8.1 and 8.3 that:

PROPOSITION 8.4 *With binary values there always exists an optimal mechanism which is deterministic. There is a sequential equilibrium implementing the optimal mechanism in the bargaining protocol which contains only one round of cheap talk communication.*

Sher (2010) also showed that under normality, the dynamic persuasion model could be solved by solving a maximum flow problem. This also yields a method of solution for the optimal mechanism in our model with binary values, as well as the equilibrium of the bargaining protocol supporting this optimal mechanism. In particular, define A and R as in (65-66). Then form a graph with with vertices $\{x, y\} \cup (T \setminus 0)$. Here x is a source and y is a sink. The graph has an edge (x, s) for all $s \in A$, and edge (t, y) for all $t \in R$ and an edge (s, t) whenever $\sigma(s) \subseteq \sigma(t)$, $s \in A$, and $t \in R$. Each edge of the form (x, s) has capacity π_s, each edge of the form (t, y) has capacity π_t, and each edge of the form (s, t) has infinite

[24]Glazer & Rubinstein (2006) showed that in general the optimal *static* persuasion rule is deterministic regardless of normality; see Sher (2010) for the distinction between static and dynamic persuasion rules.

[25]The reader may wonder about the qualifier "dynamic" on persuasion rules given that in our setting, the optimal mechanism is static. The reason for this is that Sher (2010) considered message structures that are not normal, for which the optimal mechanism must generally be dynamic even when the listener (corresponding to our seller) can commit. In the special case of normal message structures (such as those satisfying (64) and hence such as the ones studied here), it is possible to prove that one can restrict attention to static mechanisms in the persuasion problems.

capacity. A minimum cut in this graph corresponds to an optimal mechanism.[26] All types on the source side of the cut are served and receive price v_L. All types t on the sink side of the cut are either not served if $t \in A$ or served at the high price v_H if $t \in R$. The buyer's reporting strategy corresponds to a maximum flow in this graph. In the equilibrium of the bargaining protocol implementing the optimal mechanism, each type in A tells the truth (or drops out if not served), whereas the types t in R on the source side of the cut claim to be type $s \in A$ with probability proportional to the flow on (s, t), and types $t \in R$ on the sink side of the cut truthfully reveal their type. In this way, the equilibrium supporting the optimal mechanism will generally involve randomization. For further details on this construction, see Sher (2010).

9 Conclusion

This paper examined the seller's optimal mechanism for maximizing revenue when different types of the buyer have a differential ability to communicate. The main result showed that the optimal mechanism could be implemented as a sequential equilibrium of a natural bargaining protocol yielding a natural procedure whereby the seller could optimally price discriminate on the basis of the buyer's evidence.

Our problem shares with the problem of multi-dimensional mechanism design the quality that the binding incentive constraints are not known *a priori*. Hence, a succinct description of the optimal mechanism in these cases is not possible. However, the techniques we develop here suggest that it may still be possible to provide an interpretation of the optimal mechanism in terms of a dynamic bargaining protocol.

10 Appendix: Proofs

Notation

Here we collect some useful notation. $\mathcal{P}_{[s-t]}$ is the set of all $s - t$ paths (i.e., paths from s to t). So recalling that $\mathcal{P}_t$ was defined to be the set of $0 - t$ paths, we have $\mathcal{P}_t = \mathcal{P}_{[0-t]}$, and similarly $\mathcal{P} = \bigcup_{t \in T} \mathcal{P}_t$. Finally, we define $\mathcal{P}_{t\rightarrow} := \bigcup_{s \neq t} \mathcal{P}_{[t-s]}$.

For $P = (t_0, \ldots, t_n) \in \mathcal{P}$ and $0 \leq k \leq n$, we write:

$$(t_0, \ldots, t_k) \subseteq P \tag{67}$$

Note that for (67), $(t_0, \ldots, t_k)$ must an *initial* subsequence of P. If $P = (t_0, \ldots, t_n)$ if $i = 0, \ldots, n$, we write: $t_i \in P$ and $(t_{i-1}, t_i) \in P$.

[26]For a definition of network theoretic notions referred to here but not defined in the text, the reader is referred to Ahuja, Magnanti & Orlin (1993).

Results

Proof of Proposition 3.2

The proposition is established by means of two lemmas.

Lemma 10.1 *All optimal solutions to the primal and dual satisfy (17).*

Proof: Let (q, p) and λ be, respectively, optimal solutions to the primal (3-6) and the dual (7-11). If $\lambda(s, t) > 0$, then complementary slackness implies:

$$p_t = v_t(q_t - q_s) + p_s.$$

So if $q_s > q_t$, then $p_s > p_t$. But then redefining $q_t := q_s$ and $p_t := p_s$ yields another feasible primal solution. In particular, t's incentive constraints are still satisfied, as are those of all types who can mimic t by transitivity of E and this leads to higher revenue, a contradiction. ■

To prove that we can find an optimal dual solution λ satisfying (16) it is convenient to consider the path formulation of the dual (7-11). A **cycle** is a sequence of vertices $(t_0, t_1, \ldots, t_n)$ with $(t_{i-1}, t_i) \in E$ for $i = 1, \ldots, n$, $t_0 = t_n$, and all vertices $t_1, \ldots, t_n$ distinct. Let $\mathcal{C}$ be the set of cycles. Note that $\mathcal{P} \cap \mathcal{C} = \emptyset$. Define $\mathcal{P}^* := \mathcal{P} \cup \mathcal{C}$. Then by a standard result we can find $(\lambda_P : P \in \mathcal{P}^*)$ (known as a **path decomposition**) such that:[27]

$$\lambda(s, t) = \sum_{P \in \mathcal{P}^*} \lambda_P \tag{68}$$

[27] For example, in our problem, starting with $(\lambda(s, t) : (s, t) \in E)$, we can construct $(\lambda_P : P \in \mathcal{P}^*)$ as follows.

1. Define $\lambda_P := 0, \forall P \in \mathcal{P}^*$, $\widehat{\pi}_t := \pi_t, \forall t \in T \setminus 0$.
2. Find $P := (t_0, \ldots, t_n) \in \mathcal{P}^*$ with $\lambda_P = 0$, $\lambda(t_{i-1}, t_i) > 0, \forall i = 1, \ldots, n$ and such that if $P \in \mathcal{P}$, then $\widehat{\pi}_{t_n} > 0$. If no such P exists, $(\lambda_P : P \in \mathcal{P}^*)$ is the desired path flow. Otherwise, if such a P exists, go to step 3.
3. If $P \in \mathcal{P}$, let $\gamma := \min(\{\widehat{\pi}_{t_n}\} \cup \{\lambda(t_{i-1}, t_i) : i = 1, \ldots, n\})$, and let $\lambda_P := \gamma$, $\lambda(t_{i-1}, t_i) := \lambda(t_{i-1}, t_i) - \gamma, \widehat{\pi}_{t_n} := \widehat{\pi}_{t_n} - \gamma$. If $P \in \mathcal{C}$, let $\gamma := \min\{\lambda(t_{i-1}, t_i) : i = 1, \ldots, n\}$, and let $\lambda_P := \gamma$, $\lambda(t_{i-1}, t_i) := \lambda(t_{i-1}, t_i) - \gamma$. Go to step 2.

The validity of the algorithm is established by arguing inductively that as the algorithm progresses, (i) the value of the sum $\lambda(s, t) + \sum_{P \in \mathcal{P}^* : (s,t) \in P} \lambda_P$ remains constant, (ii) $(\lambda(s, t) : (s, t) \in E)$ satisfies (8) with $\widehat{\pi}_t$ playing the role of π_t, and $(\lambda_P : P \in \mathcal{P}^*)$ satisfies (70) with $\pi_t - \widehat{\pi}_t$ playing the role of π_t, and (iii) as long as $(\lambda(s, t) : (s, t) \in E)$ is not uniformly zero, we can find an appropriate P to continue the algorithm in step 2.

Let $\mathcal{P}^*_{s,t}$ be the set of paths (or cycles) in $\mathcal{P}^*$ such that $(s,t) \in P$. Given this translation, we can rewrite the dual (7-11) as follows:

$$\begin{aligned}
&\text{minimize} && \sum_{t \in T} \mu_t && (69)\\
&\text{subject to} && && \\
&\forall t \in T \setminus 0, && \sum_{P \in \mathcal{P}_t} \lambda_P = \pi_t && (70)\\
&\forall t \in T, && v_t \pi_t - \sum_{s:(t,s)\in E} \sum_{P \in \mathcal{P}^*_{t,s}} \lambda_P (v_s - v_t) \le \mu_t && (71)\\
&\forall P \in \mathcal{P}^*, && \lambda_P \ge 0, && (72)\\
&\forall t \in T, && \mu_t \ge 0 && (73)
\end{aligned}$$

Observe that constraint (70) involves only a summation of variables indexed by simple paths leading to t and no cycles. (69-73) differs from (23-27) in that in (71) of $\mathcal{P}^*_{s,t}$ plays the role of $\mathcal{P}_{s,t}$ in (25), and in (72) we quantify over $\mathcal{P}^*$ whereas in (26) we quantify over $\mathcal{P}$.

Lemma 10.2 *There is an optimal solution to the dual (7-11) with no bad edges.*

Proof: Suppose not. Amongst all optimal dual solutions choose one with the smallest discrepancy, where we measure the discrepancy by:

$$\sum_{(s,t)\in E: v_s \ge v_t} \sum_{P \in \mathcal{P}^*_{s,t}} \lambda_P \left[1 + (v_s - v_t)\right].$$

Define:

$$\mathcal{P}^+ := \{(t_0, \dots, t_n) : t_0 = 0, (t_{i-1}, t_i) \in E \text{ and } \lambda(t_{i-1}, t_i) > 0 \text{ for } i = 1, \dots, n\}$$

There must be some $P \in \mathcal{P}^+$ with a bad edge. In particular, if (s,t) is the bad edge with $\lambda(s,t) > 0$, then by (70), there must be $P' \in \mathcal{P}_s$ with $\lambda_P > 0$. To arrive at P, simply append the edge (s,t) to P'. We may choose $P = (t'_0, \dots, t'_n) \in \mathcal{P}^+$ of minimal length containing a bad edge (where "length" refers to the number of vertices in P). We must have $n \ge 2$. Moreover because P is of minimal length, (t'_{n-2}, t'_{n-1}) is a good edge and (t'_{n-1}, t'_n) is a bad edge. If $P \notin \mathcal{P}$, then there exists integer k with $0 < k < n$ such that $(t'_k, \dots, t'_n) \in \mathcal{C}$ (so that $t'_k = t'_n$). Define:

$$Q := \begin{cases} (t'_0, \dots, t'_n), & \text{if } P \in \mathcal{P};\\ (t'_k, \dots, t'_n), & \text{if } P \notin \mathcal{P}. \end{cases}$$

It is possible to choose the path decomposition $(\lambda(s,t) : (s,t) \in E)$ with $\lambda_Q > 0$.[28] Observe that the discrepancy does not depend on the choice of path decomposition.

[28] In particular, in using the algorithm in footnote 27, the first time that we arrive at step 2, choose $P = Q$.

Case 1 $\mu_{t'_{n-1}} > 0$ or $v_{t'_{n-1}} = v_{t'_n}$.

If $t'_{n-2} \neq t'_n$, let Q' be obtained from Q by replacing (t'_{n-2}, t'_{n-1}) and (t'_{n-1}, t'_n) with (t'_{n-2}, t'_n).[29] Construct a new solution to the dual by decreasing λ_Q by ϵ and increasing $\lambda_{Q'}$ by ϵ. All constraints (70) are still satisfied. The same is true for constraint (71) for any vertex in $T \setminus \{t'_{n-2}, t'_{n-1}\}$ because for any edge (s,t) not equal to (t'_{n-2}, t'_{n-1}) or (t'_{n-1}, t'_n), $\sum_{P \in \mathcal{P}^*_{s,t}} \lambda_P$ remains unchanged.

For vertex t'_{n-1}, the LHS of constraint (71) changes by $\epsilon(v_{t'_n} - v_{t'_{n-1}}) \leq 0$ so that the constraint is still satisfied when we add $\epsilon(v_{t'_n} - v_{t'_{n-1}})$ to $\mu_{t'_{n-1}}$ (for ϵ sufficiently small, $\mu_{t'_{n-1}}$ remains nonnegative by the assumptions of Case 1).

For vertex t'_{n-2}, the LHS of (71) changes by $\epsilon(v_{t'_{n-1}} - v_{t'_{n-2}}) - \epsilon(v_{t'_n} - v_{t'_{n-2}}) = -\epsilon(v_{t'_n} - v_{t'_{n-1}}) \geq 0$. So (71) remains feasible when we add $-\epsilon(v_{t'_n} - v_{t'_{n-1}})$ to $\mu_{t'_{n-2}}$.

The total change in the dual objective function value is zero. Because $(v_{t'_{n-2}} - v_{t'_n}) < (v_{t'_{n-1}} - v_{t'_n})$, we have a new optimal dual solution with smaller discrepancy.

If $t'_{n-2} = t'_n$, then $Q = (t'_{n-2}, t'_{n-1}, t'_n)$, and arguments similar to the above show that reducing λ_Q by ϵ, and modifying $\mu_{t'_{n-1}}$ and $\mu_{t'_{n-2}}$ accordingly will lead to another feasible solution with the same objective function value and smaller discrepancy.

Case 2 $\mu_{t'_{n-1}} = 0$ and $v_{t'_{n-1}} > v_{t'_n}$.

In this case, $v_{t'_{n-1}} \pi_{t'_{n-1}} > 0$ and (71) imply there exists vertex u with $v_u > v_{t'_{n-1}}$ and $(t'_{n-1}, u) \in E$ and moreover $\lambda(t'_{n-1}, u) > 0$. Defining $K = (t'_0, \ldots, t'_{n-1}, u)$, we must have that $K \in \mathcal{P}$ because by the way that K was constructed, every edge of K is good. Moreover, it is possible to choose the path decomposition so that $\lambda_K > 0$ (in addition to $\lambda_Q > 0$).[30] We assume that $t'_{n-2} \neq t'_n$ and explain how to modify the proof if $t'_{n-2} = t'_n$ in footnote 31. Let Q' be obtained from Q by replacing (t'_{n-2}, t'_{n-1}) and (t'_{n-1}, t'_n) with (t'_{n-2}, t'_n). Let K' be obtained from K by replacing (t'_{n-2}, t'_{n-1}) and (t'_{n-1}, u) with (t'_{n-2}, u). Choose positive ϵ and δ satisfying:

$$\delta(v_u - v_{t'_{n-1}}) = \epsilon(v_{t'_{n-1}} - v_{t'_n}). \tag{74}$$

Decrease the flow on Q by ϵ, decrease the flow on K by δ, increase the flow on Q' by ϵ, increase the flow on K' by δ.

Observe that (70) still holds for all vertices. (71) still holds for all vertices as well; we explain why for vertices t'_{n-1} and t'_{n-2}. For t'_{n-1}, the LHS of (71) changes by $\epsilon(v_{t'_n} - v_{t'_{n-1}})$ due to the decrease in λ_Q and changes by $\delta(v_u - v_{t'_{n-1}})$ due to the decrease in λ_K. (74) implies that these changes cancel out. For t'_{n-2}, the LHS of (71) changes by $+\epsilon(v_{t'_{n-1}} - v_{t'_{n-2}})$

[29] This is possible because of transitivity of E. Similar comments apply to the construction in Case 2.

[30] In particular, in the algorithm in footnote 27, the first time that we arrive at step 2, choose $P = Q$. However in step 3, if $\gamma = \lambda(s,t)$ for some $(s,t) \in K \setminus \{(t_{n-1}, u)\}$, set $\lambda_Q = \gamma - \epsilon$ for some small $\epsilon > 0$. This will allow us to choose $P = K$, the second time that we arrive as step 2, and then proceed with the algorithm in the ordinary way.

due to the decrease in λ_Q, changes by $-\epsilon(v_{t'_n} - v_{t'_{n-2}})$ due to the increase in $\lambda_{Q'}$, changes by $+\delta(v_{t'_{n-1}} - v_{t'_{n-2}})$ due to the decrease in λ_K, and changes by $-\delta(v_u - v_{t'_{n-2}})$ due to the increase in $\lambda_{K'}$. Again, (74) implies that these changes cancel out.

The dual objective value does not change. The change in the discrepancy is $-\epsilon(1 + (v_{t'_{n-1}} - v_{t'_n})) < 0$ if $v_{t'_n} > v_{t'_{n-2}}$ and $-\epsilon(1+(v_{t'_{n-1}} - v_{t'_n})) + \epsilon(1+(v_{t'_{n-2}} - v_{t'_n})) = -\epsilon(v_{t'_{n-1}} - v_{t'_{n-2}}) < 0$ if $v_{t'_n} \leq v_{t'_{n-2}}$. Thus, in both cases the discrepancy declines, a contradiction.[31] ■

Proof of Proposition 3.3

As explained in the proof of Proposition 3.2, (69-73) is equivalent to (7-11). (69-73) differs from (23-27) only insofar as the former contains variables λ_P for $P \in \mathcal{C}$. Note however that (16) of Proposition 3.2 and (68) that there exists an optimal solution of (69-73) satisfying:

$$\lambda_P = 0, \qquad\qquad \forall P \in \mathcal{C} \tag{75}$$

But (23-27) is equivalent to (69-73) with the additional constraint (75). Note also that (68) reduces to (28) under (75). This establishes the part of the proposition pertaining to the dual.

With regard to the primal, (19-22) is the dual of (23-27) which implies via the first part of the theorem that (19-22) and (3-6) have the same value, and because (19-22) is a relaxation of (3-6), every optimal solution of (3-6) is an optimal solution of (19-22). ■

Proof of Lemma 3.4

Start with an optimal dual solution, λ' that avoids bad edges. Throughout write $\Phi(P)$ (resp. $\phi(s,t)$) for $\Phi(P|\lambda')$ (resp. $\phi(s,t|\lambda')$). We show that $\lambda_P := \Phi(P)\pi_{\tau(P)}$ is the desired dual solution. Note that λ also avoids bad edges.

Define $\mathcal{P}_t^\ell$ be the set of all paths P ending in t such that (i) P contains at most ℓ edges, and (ii) if P has fewer than ℓ edges, then P begins in 0. For vertex h, define

[31] Here we explain how to modify the proof if $t'_{n-2} = t'_n$. In this case, decrease λ_Q by ϵ, decrease λ_K by δ, and increase $\lambda_{K'}$ by δ; we do not introduce Q' because it would correspond to a self-loop at t'_{n-2} and we have excluded such edges. All constraints of the form (70) are still satisfied because Q is a cycle. The argument that (71) still holds for t'_{n-1} is the same as the one given above. For t'_{n-2}, the LHS of (71) changes by $\epsilon(v_{t'_{n-1}} - v_{t'_{n-2}}) + \delta(v_{t'_{n-1}} - v_{t'_{n-2}}) - \delta(v_u - v_{t'_{n-2}}) = \epsilon(v_{t'_{n-1}} - v_{t'_n}) - \delta(v_{t'_{n-1}} - v_u) + \epsilon(v_{t'_n} - v_{t'_{n-2}}) = \epsilon(v_{t'_n} - v_{t'_{n-2}}) = 0$, where the second to last equality follows from (74). The discrepancy changes by $-\epsilon(1 + (v_{t'_{n-1}} - v_{t'_n})) < 0$.

$\mathcal{P}^\ell_{[h-t]} := \mathcal{P}^\ell \cap \mathcal{P}_{[h-t]}$. We argue that:

$$\sum_{P\in\mathcal{P}^\ell_t} \Phi(P) = 1 \tag{76}$$

The base case $\ell = 1$ is immediate. Next observe that:

$$\begin{aligned}\sum_{P\in\mathcal{P}^\ell_t} \Phi(P) &= \sum_{P\in\mathcal{P}^{\ell-1}_{[0-t]}} \Phi(P) + \sum_{h\in T\setminus 0}\sum_{(j,h)\in E} \phi(j,h) \sum_{\left\{P\in\mathcal{P}^{\ell-1}_{[h-t]}\right\}} \Phi(P) \\ &= \sum_{P\in\mathcal{P}^{\ell-1}_{[0-t]}} \Phi(P) + \sum_{h\in T\setminus 0} \sum_{\left\{P\in\mathcal{P}^{\ell-1}_{[h-t]}\right\}} \Phi(P) = \sum_{P\in\mathcal{P}^{\ell-1}_t} \Phi(P) = 1\end{aligned}$$

The first equality relies on the fact that if $(t_0,\ldots,t_k)$ is a cycle, $\prod_{i=1}^k \phi(t_{i-1},t_i) = 0$ (because λ' has no bad edges). The last equality follows from the inductive hypothesis. Noting that for sufficiently large ℓ, $\mathcal{P}^\ell_t = \mathcal{P}_t$, (76) implies that λ satisfies (24).

Let $\mathcal{P}_{t\to} := \bigcup_{s\neq t} \mathcal{P}_{[t-s]}$ and $|P|$ denote the number of edges in the path P. Then,

$$\begin{aligned}\pi_t + \sum_{P\in\mathcal{P}_{t\to}} \Phi(P)\pi_{\tau(P)} &= \pi_t + \sum_{s:(t,s)\in E} \phi(t,s)\left[\pi_s + \sum_{P\in\mathcal{P}_{s\to}} \Phi(P)\pi_{\tau(P)}\right] \\ = \pi_t + \sum_{s:(t,s)\in E}\phi(t,s)\left[\sum_{P:s\in P}\lambda'_P\right] &= \sum_{P\in\mathcal{P}_t}\lambda'_P + \sum_{s\in T\setminus t}\sum_{P\in\mathcal{P}_{t,s}}\lambda'_P = \sum_{P:t\in P}\lambda'_P\end{aligned}$$

where the second equality follows inductively, assuming the derivation for s with $\max\{|P| : \Phi(P) > 0, P\in\mathcal{P}_{s\to}\} < \max\{|P| : \Phi(P) > 0, P\in\mathcal{P}_{t\to}\}$.

It now follows (using 76) that:

$$\begin{aligned}\sum_{P\in\mathcal{P}_{t,j}} \lambda_P &= \left[\sum_{P\in\mathcal{P}_t}\Phi(P)\right] \times \left(\phi(t,j)\left[\pi_j + \sum_{P'\in\mathcal{P}_{j\to}} \Phi(P')\pi_{\tau(P')}\right]\right) \\ = \phi(t,j)\left[\pi_j + \sum_{P'\in\mathcal{P}_{j\to}}\Phi(P')\pi_{\tau(P')}\right] &= \phi(t,j)\left[\sum_{P:j\in P}\lambda'_P\right] = \sum_{P\in\mathcal{P}_{t,j}}\lambda'_P\end{aligned} \tag{77}$$

This implies that λ satisfies (25) and has the same objective function value as λ'. Using (76), it is routine to verify that λ satisfies (32). Finally (77) and the definition of λ imply that λ satisfies (31). ■

Proof of Theorem 6.1

The first step is to modify the equilibrium strategies. Recall from section 5.1 that ζ^* and ξ^* refer to the buyer and seller strategies respectively. ζ^{**} and ξ^{**} will denote respectively the modified buyer and seller strategies which will form the required sequential equilibrium. We now formally define these modified strategies.

Seller's Modified Strategy

1. ξ^{**} agrees with ξ^* at any seller information set which occurs with positive probability given (ζ^*, ξ^*), as well as any seller information set at which evidence has not yet been presented.

2. At any seller information set which occurs with zero probability given (ζ^*, ξ^*) at which the buyer previously presented evidence S, ζ^{**} requires the seller to make a take-it-or-leave-it offer at price:

$$\max\{v_r : S \subseteq \sigma(r)\}$$

Buyer's Modified Strategy

1. ζ^* agrees with ζ^{**} at any buyer information set which occurs with positive probability given (ζ^*, ξ^*), as well as any information set where the seller decides whether to accept a take-it-or-leave-it-offer.

2. At any type t buyer information set I which occurs with zero probability given (ζ^*, ξ^*) only because the type t seller has taken a sequence actions which would have been taken with positive probability by some other buyer type according to ζ^*, (and the seller has taken actions consistent with ξ^*), the buyer continues by following some type t best reply to ξ^{**} conditional on I.

3. At any other information set, the buyer drops out.

Lemma 10.3 *1. (ζ^{**}, ξ^{**}) is a Bayes-Nash equilibrium of the dynamic communication protocol.*

2. (ζ^{**}, ξ^{**}) and (ζ^*, ξ^*) induce the same probability distribution over terminal histories.

Proof: Consider part 1. If ξ is a seller strategy profile such that (ζ^*, ξ) and (ζ^{**}, ξ) induce the same probability distribution over terminal histories, then part 1 of the lemma and the fact that (ζ^*, ξ^*) is a Bayes-Nash equilibrium implies that ξ is not a profitable seller deviation at (ζ^{**}, ξ^{**}). So, consider an ξ such that (ζ^*, ξ) and (ζ^{**}, ξ) induce different probability distributions over terminal histories. The definition of ζ^{**} means (ζ^*, ξ) differs from (ζ^{**}, ξ)

only insofar as sometimes the buyer drops out in the latter when he would not have done so in the former. This implies that the seller's payoff under (ζ^*, ξ) is weakly higher than the seller's payoff under (ζ^{**}, ξ), which, in turn, implies that ξ is not a profitable seller deviation at (ζ^{**}, ξ^{**}). Using the same argument as for the seller, if the buyer has a profitable deviation ζ at (ζ^{**}, ξ^{**}), then (ζ, ξ^*) and (ζ, ξ^{**}) must induce a different probability distribution over terminal histories. But this means that (ζ, ξ^*) and (ζ, ξ^{**}) differ only in that in the latter, following certain histories the seller makes the offer $\max\{v_r : S \subseteq \sigma(r)\}$, where S is the evidence that has been presented by the buyer, whereas in the former, the seller would have made a different offer. Notice that if the buyer has presented S, she must have been of a type t such that $S \subseteq \sigma(t)$. But this implies that $v_t \leq \max\{v_r : S \subseteq \sigma(r)\}$, which in turn implies that buyer's payoff is weakly higher under (ζ, ξ^*) than under (ζ, ξ^{**}), so that ζ is not a profitable buyer deviation at (ζ^{**}, ξ^{**}). This establishes part 2 of the lemma. ■

To complete the proof, we show that the players' strategies are sequentially rational off the equilibrium path, where the seller's off equilibrium beliefs are consistent with the structure of the game as required by sequential equilibrium.[32] For each $\epsilon > 0$, we construct a totally mixed buyer strategy ζ^ϵ such that $\zeta^\epsilon \to \zeta^{**}$ as $\epsilon \to 0$. Enumerate the types $t'_1, \ldots, t'_n$ in T so that $i < j \Rightarrow v_{t'_i} \geq v_{t'_j}$. ζ^ϵ is the buyer strategy in which with probability $1 - \epsilon^i$, the type t_i buyer plays (her part of) ζ^{**} and with probability ϵ^i, she randomizes uniformly over all type t pure strategies. So a type with a higher index (and hence a lower value) trembles with a probability that approaches zero faster than a type with a lower index. Off the equilibrium path, the seller's beliefs about the buyer's type are the limiting beliefs derived via Bayes' rule using ζ^ϵ (and *any* totally mixed seller strategy[33]). It follows that in any off equilibrium path history, if the seller can infer that the buyer has deviated from ζ^{**}, the seller will infer that that the buyer is the highest value type that could have performed the actions consistent with that history; so if no evidence has been presented, the seller will infer that the buyer is of a highest value type, and if evidence has been presented, the seller will infer that the buyer has the highest value among those types who could have presented the evidence.

First we establish that given any seller information set I which occurs with zero probability under (ζ^{**}, ξ^{**}), ξ^{**} is a seller's best reply to ζ^{**} given the seller's off equilibrium beliefs derived above. Part 1 of Lemma 10.3 imply that I also occurs with zero probability under (ζ^*, ξ^*). First suppose that at I, the buyer has not yet presented evidence. Then no matter what the seller does, the buyer will drop out at the next opportunity, so the seller is best replying. Next consider I at which the buyer has presented evidence S. Because I has zero probability under (ζ^*, ξ^*), according to ζ^{**}, either the buyer should have dropped out prior to presenting evidence or the buyer should have presented evidence different from

[32]There is no corresponding issue for the buyer's beliefs because the seller has no private information.
[33]The resulting beliefs do not depend on which totally mixed seller strategy is used.

S. In either event, the seller will use the off equilibrium beliefs derived above and infer that the buyer is of the t such that $v_t = \max\{v_r : S \subseteq \sigma(t)\}$, and so it will be optimal to offer the maximal price that the type t buyer will accept, namely, $\max\{v_r : S \subseteq \sigma(t)\}$, as required by ξ^{**}.

Finally, we establish that given any buyer information set I that occurs with zero probability in equilibrium, ζ^{**} is a buyer best reply to ξ^{**}. Again, I has zero probability under (ζ^*, ξ^*). If at I, the seller has made a take-it-or-leave-it-offer, or if I falls under part 2 of the definition of the buyer's modified strategy ζ^{**}, then the result is immediate from the definitions of ζ^* and ζ^{**}. In any other case, the buyer cannot possibly attain a positive utility, and by dropping out as required by ζ^{**}, she attains a utility of zero.

Proofs of Claims from Section 7

In this section, we prove various claims made in the course of the discussion of the example in Section 7. First we provide a simple proof that for any profile of valuations satisfying (41), there are many probability distributions $(\pi_t : t \in T)$ such that (45-47) are satisfied. In particular, it is straightforward to verify that there will always exists a number K sufficiently large (where sufficiently large depends on the profile of valuations) such that any probability distribution satisfying:

$$\pi_{y_0} > K\pi_{y_1} > K^2\pi_{y_2} > \cdots > K^n\pi_{y_n} > K^{n+1}\pi_{x_1} > K^{n+2}\pi_{x_2} > \cdots > K^{2n}\pi_{x_n}$$

will satisfy (45-47).

Claim 10.4 *The mechanism defined by (50-55) satisfies (56).*

Proof: (41) implies that $0 < q_{x_1} < 1$. Let $2 \leq i \leq n$. We argue inductively that

$$q_{x_{i-1}} < q_{x_i} < 1.$$

(52-55) implies that:

$$v_{y_{i-1}} - v_{y_0} = v_{y_{i-1}} q_{x_{i-1}} - p_{x_{i-1}} \tag{78}$$

$$v_{y_i} - v_{y_0} = (v_{y_i} - v_{x_i})(q_{x_i} - q_{x_{i-1}}) + v_{y_i} q_{x_{i-1}} - p_{x_{i-1}} \tag{79}$$

By the inductive hypothesis, $q_{x_{i-1}} < 1$. (78) and (41) then imply that:

$$v_{y_i} - v_{y_0} > v_{y_i} q_{x_{i-1}} - p_{x_{i-1}},$$

which together with (79) and (41) implies that:

$$q_{x_{i-1}} < q_{x_i}$$

(78), (41) and $q_{x_{i-1}} < 1$ imply that:

$$v_{x_i} - v_{y_0} < v_{x_i} q_{x_{i-1}} - p_{x_{i-1}} \tag{80}$$

On the other hand, if $q_{x_i} \geq 1$, the numerator is weakly larger than the denominator in (54), which implies that:

$$v_{x_i} - v_{y_0} \geq v_{x_i} q_{x_{i-1}} - p_{x_{i-1}},$$

which contradicts (80). We have now established (56). ■

Claim 10.5 *The mechanism defined by (50-55) is feasible in the primal (3-6). Moreover the incentive constraints corresponding to pairs of the form* $(x_i, x_{i+1}), (x_i, y_i), (0, x_1), (0, y_0)$ *and* (y_j, y_k) *with* $j < k$ *hold with equality whereas all other incentive constraints hold with strict inequality.*

Proof: For $i < j$, we have:

$$\begin{aligned} v_{x_j}(q_{x_j} - q_{x_i}) &= \sum_{k=i+1}^{j} v_{x_j}(q_{x_k} - q_{x_{k-1}}) \\ &\geq \sum_{k=i+1}^{j} v_{x_k}(q_{x_k} - q_{x_{k-1}}) = \sum_{k=i+1}^{j} (p_{x_k} - p_{x_{k-1}}) = p_{x_j} - p_{x_i}, \end{aligned}$$

where we have used Claim 10.4 and (54). Moreover, the inequality is strict exactly when $i < j - 1$. This establishes the (x_i, x_j)-incentive constraints when $i < j$. Next choose y_j and $i \leq j$. To establish the (x_i, y_j) incentive constraint, we observe:

$$\begin{aligned} v_{y_j}(q_{y_j} - q_{x_i}) &= v_{y_j}(q_{y_j} - q_{x_j}) + \sum_{k=i+1}^{j} v_{y_j}(q_{x_k} - q_{x_{k-1}}) \\ &\geq v_{y_j}(q_{y_j} - q_{x_j}) + \sum_{k=i+1}^{j} v_{x_k}(q_{x_k} - q_{x_{k-1}}) = p_{y_j} - p_{x_j} + \sum_{k=i+1}^{j} (p_{x_k} - p_{x_{k-1}}) = p_{y_j} - p_{x_i} \end{aligned}$$

where, among other things, we have used $q_{y_j} = 1$ and (52-55). The inequality is strict exactly when $i < j$. The $(0, x_i)$-incentive constraints follow from a simple induction, which can also be used to show that all inequalities other than the $(0, x_1)$ inequality are strict. The (y_i, y_j)-constraints for $i < j$ follow from the fact that all y-types receive the same allocation, and the $(0, y_i)$-incentive constraints follow from (41). ■

We now define a solution to the dual.

$$\begin{aligned}
\lambda(x_i, x_{i+1}) &= \sum_{j=i+1}^{n} \pi_{x_j} + \delta_i & \forall i = 1, \ldots, n-1 \\
\lambda(0, x_1) &= \sum_{i=1}^{n} \pi_{x_i} + \delta_0 \\
\lambda(x_i, y_i) &= \delta_{i-1} - \delta_i & \forall i = 1, \ldots, n \\
\lambda(y_i, y_{i+1}) &= \sum_{j=i+1}^{n} \pi_{y_j} - \delta_i & \forall i = 0, 1, \ldots, n-1 \\
\lambda(0, y_0) &= \sum_{i=0}^{n} \pi_{y_i} - \delta_0 \\
\lambda(s, t) &= 0 \text{ if not otherwise specified above.}
\end{aligned} \tag{81}$$

Claim 10.6 *(81) together with (12-13) define an optimal dual solution.*

Proof: First we argue that the flow conservation constraints (8) are satisfied. First pick $x_i \in X \setminus \{x_1, x_n\}$. We have:

$$\begin{aligned}
\sum_{s:(s,x_i)\in E} \lambda(s, x_i) - \sum_{s:(x_i,s)\in E} \lambda(x_i, s) &= \lambda(x_{i-1}, x_i) - \lambda(x_i, x_{i+1}) - \lambda(x_i, y_i) \\
&= \sum_{j=i}^{n} \pi_{x_j} + \delta_{i-1} - \left(\sum_{j=i+1}^{n} \pi_{x_j} + \delta_i \right) - (\delta_{i-1} - \delta_i) = \pi_{x_i}
\end{aligned}$$

The cases of x_1 and x_n are similar, the latter of which uses $\delta_n = 0$ (see (43)). Next consider $y_i \in Y \setminus \{y_0, y_n\}$

$$\begin{aligned}
\sum_{s:(s,y_i)\in E} \lambda(s, y_i) - \sum_{s:(y_i,s)\in E} \lambda(y_i, s) &= \lambda(x_i, y_i) + \lambda(y_{i-1}, y_i) - \lambda(y_i, y_{i+1}) \\
&= (\delta_{i-1} - \delta_i) + \left(\sum_{j=i}^{n} \pi_{y_j} - \delta_{i-1} \right) - \left(\sum_{j=i+1}^{n} \pi_{y_j} - \delta_i \right) = \pi_{y_i}
\end{aligned}$$

The cases of y_0 and y_n are similar. (9) and (11) are automatically satisfied by (12-13). An easy induction using (46) establishes that $\lambda(x_i, y_i) > 0$ for $i = 1, \ldots, n$. Next observe that

$$\delta_{n-1} = \frac{v_{x_n} \pi_{x_n}}{v_{y_n} - v_{x_n}} < \pi_{y_n} = \sum_{j=(n-1)+1}^{n} \pi_{y_j},$$

where the first equality follows from (44) for $i = n-1$ and the inequality follows from (47) for $i = n$. We proceed inductively, assuming that $\delta_i < \sum_{j=i+1}^{n} \pi_{y_j}$. We have:

$$
\begin{aligned}
\delta_{i-1} &= \frac{v_{x_i}\pi_{x_i} - (v_{x_{i+1}} - v_{x_i})\sum_{j=i+1}^{n} \pi_{x_j} + (v_{y_i} - v_{x_{i+1}})\delta_i}{v_{y_i} - v_{x_i}} \\
&< \frac{v_{x_i}\pi_{x_i} - (v_{x_{i+1}} - v_{x_i})\sum_{j=i+1}^{n} \pi_{x_j} + (v_{y_i} - v_{x_{i+1}})\sum_{j=i+1}^{n} \pi_{y_j}}{v_{y_i} - v_{x_i}} \\
&= \sum_{j=i+1}^{n} \pi_{y_j} + \frac{v_{x_i}\pi_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\sum_{j=i+1}^{n} \pi_{x_j} + \sum_{j=i+1}^{n} \pi_{y_j}\right]}{v_{y_i} - v_{x_i}} \\
&< \sum_{j=i}^{n} \pi_{y_j},
\end{aligned}
$$

where the first equality follows from (44), the first inequality follows from (41) and the inductive hypothesis, and the second inequality follows from (47). It follows that $\lambda(y_i, y_{i+1}) > 0$ for $i = 0, \ldots, n-1$, which in turn implies that $\lambda(0, y_0) > 0$. That $\lambda(x_i, x_{i+1}) > 0$ for $i = 1, \ldots, n-1$ and $\lambda(0, x_1) > 0$ follows from (48). For all other edges (s,t) not mentioned above, we have $\lambda(s,t) = 0$. We have now established that the potential solution defined by (81) and (12-13) satisfies (10), and moreover that this solution is dual feasible. We now argue for optimality.

Using Claim 10.5, it is straightforward to verify that for all $(s,t) \in E$, if $\lambda(s,t) > 0$, the (s,t)-incentive constraint holds with equality. For all $x_i \in X \setminus \{x_n\}$:

$$
\begin{aligned}
\psi_{x_i} =& v_{x_i} - (v_{x_{i+1}} - v_{x_i})\frac{\lambda(x_i, x_{i+1})}{\pi_{x_i}} - (v_{y_i} - v_{x_i})\frac{\lambda(x_i, y_i)}{\pi_{x_i}} \\
=& v_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\frac{\sum_{j=i+1}^{n} \pi_{x_j} + \delta_i}{\pi_{x_i}}\right] - (v_{y_i} - v_{x_i})\left[\frac{\delta_{i-1} - \delta_i}{\pi_{x_i}}\right] \\
=& v_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\frac{\sum_{j=i+1}^{n} \pi_{x_j} + \delta_i}{\pi_{x_i}}\right] \\
& - (v_{y_i} - v_{x_i})\left[\frac{\delta_i}{\pi_{x_i}} + \frac{v_{x_i}\pi_{x_i} - (v_{x_{i+1}} - v_{x_i})\left[\delta_i + \sum_{j=i+1}^{n} \pi_{x_j}\right]}{(v_{y_i} - v_{x_i})\pi_{x_i}} - \frac{\delta_i}{\pi_{x_i}}\right] \\
=& 0,
\end{aligned}
$$

where the second to last inequality uses (44). A similar argument shows that $\psi_{x_n} = 0$. For all $y_i \in Y \setminus \{y_n\}$:

$$
\psi_{y_i} = v_{y_i} - (v_{y_{i+1}} - v_{y_i})\frac{\lambda(y_i, y_{i+1})}{\pi_{y_i}} = v_{y_i} - (v_{y_{i+1}} - v_{y_i})\left[\frac{\sum_{j=i+1}^{n} \pi_{y_j} - \delta_i}{\pi_{y_i}}\right] > 0,
$$

where the inequality follows from (45) and (48). We also have $\psi_{y_n} = v_{y_n} > 0$. It now follows from (50) that for all types t, if $\mu_t > 0$ (which is equivalent to $\psi_t > 0$ by (13)), then $q_t = 1$. That any type with $\psi_t < 0$ receives an allocation of $q_t = 0$ holds vacuously because there are no (nonzero) types with $\psi_t < 0$. We have now established that the complementary slackness conditions hold, and hence we obtain the desired result. ■

Claim 10.7 *The optimal solution defined by (50-55) is the unique optimal solution in the primal. Moreover, every optimal dual solution satisfies (57-60).*

Proof: First suppose that for some $h < i$, $\lambda(x_h, x_{i+1}) > 0$ at an optimal dual solution. The fact that by Claim 10.5, the primal optimal solution defined by (50-55) is such that the (x_h, x_{i+1}) incentive constraint holds with strict inequality would then contradict complementary slackness. Similarly we must have $\lambda(0, x_{i+1}) = 0$ for all $i \geq 0$. Next suppose that $\lambda(x_i, x_{i+1}) = 0$. Then by the structure of the incentive graph and the flow conservation constraint (8), we must have $\lambda(x_h, x_{i+1}) > 0$ for some $h < i$ or $\lambda(0, x_{i+1}) > 0$, which we have just seen is impossible. The part of the claim about the dual solution follows.

As established in the proof of Claim 10.6, in the optimal dual solution given by (81) and (12-13), $\mu_{y_i} > 0$ for all $y_i \in Y$. It follows from complementary slackness that $q_{y_i} = 1$, or in other words, that (50) holds. Moreover, at any optimal solution we must have $p_0 = q_0 = 0$. The rest of the mechanism (50-55) is then determined by solving the equations:

$$v_t q_t - p_t = v_t q_s - p_s \qquad \forall (s,t) \in E \text{ with } \lambda(s,t) > 0,$$

where λ is given by (81). ■

Proof of Proposition 8.2

Tree structure implies:

$$(t, r) \in E^+ \Leftrightarrow [t = \varphi(r) \text{ or } \exists s, (t = \varphi(s) \text{ and } (s, r) \in E^+)] \tag{82}$$

$$\Leftrightarrow [t = \varphi(r) \text{ or } \{t \neq \varphi(r) \text{ and } \exists! s, (t = \varphi(s) \text{ and } (s, r) \in E^+)\}] \tag{83}$$

Define:

$$\lambda(s,t) = \begin{cases} \pi_t + \sum_{r:(t,r)\in E^+} \pi_r, & \text{if } s = \varphi(t); \\ 0, & \text{otherwise.} \end{cases} \tag{84}$$

We have:

$$\sum_{s:(s,t)\in E} \lambda(s,t) = \lambda(\varphi(t),t) = \pi_t + \sum_{r:(t,r)\in E^+} \lambda(t,r) = \pi_t + \sum_{s:t=\varphi(s)} \left[\pi_s + \sum_{r:(s,r)\in E^+} \pi_r\right]$$
$$= \pi_t + \sum_{s:t=\varphi(s)} \lambda(t,s) = \pi_t + \sum_{s:(t,s)\in E} \lambda(t,s),$$

where we used (82-83). It follows that λ satisfies the flow conservation constraints (8).

Defining μ_t via (12-13), the optimal solution satisfies the virtual valuation constraints, and we have a feasible solution to the dual, for which by construction we have $\psi_t = \widehat{\psi}_t$.

Next I establish that (63) is feasible in the primal, or in other words, incentive compatible. Assume for contradiction that for some $(s,t) \in E$,

$$v_t q_t - p_t < v_t q_s - p_s, \tag{85}$$

Then because $v_t q_t - p_t \geq 0$, we must have $\widehat{\psi}_s \geq 0$ and $p_s < v_t$. However (1) and (63) imply that for some r with $(r,t) \in E$ with $\widehat{\psi}_r \geq 0$, $p_s = v_r$. But then single crossing quasi-virtual valuations imply that $\widehat{\psi}_t \geq 0$, and in turn, (63) implies that $p_t \leq p_s$. Since $q_s = q_t = 1$, this contradicts (85). So (63) is primal feasible.

Next suppose that $\lambda(s,t) > 0$. Then $(s,t) \in E^+$, and moreover $s = \varphi(t)$. If $\widehat{\psi}_s < 0$, then by single-crossing virtual valuations, the fact that $s = \varphi(t)$, and the transitivity of E^+, for all $(r,t) \in E^+$, $\widehat{\psi}_r < 0$. So either $q_t = p_t = 0$ or $q_t = 1$ and $p_t = v_t$. In either case, $v_t q_t - p_t = 0 = v_t \times 0 - 0 = v_t q_s - p_s$. If $\widehat{\psi}_s \geq 0$, then single-crossing virtual valuations imply that $\widehat{\psi}_t \geq 0$. It then follows using an argument similar to the one in the previous paragraph that $p_s = p_t$ and $q_s = q_t = 1$, so that again $v_t q_t - p_t = v_t q_s - p_s$. To summarize, we have shown that:

$$\lambda(s,t) > 0 \Rightarrow v_t q_t - p_t = v_t q_s - p_s, \qquad \forall (s,t) \in E \tag{86}$$

Next suppose that $\mu_t > 0$. Then $\psi_t = \widehat{\psi}_t > 0$, which implies that $q_t = 1$, so that:

$$\mu_t > 0 \Rightarrow q_t = 1 \qquad \forall (s,t) \in E \tag{87}$$

Next suppose that $\psi_t < 0$. Then $\widehat{\psi}_t = \psi_t < 0$, so that $q_t = 0$. If follows that:

$$\psi_t < 0 \Rightarrow q_t = 0 \qquad \forall t \in T \tag{88}$$

(86-88) are the complementary slackness optimality conditions for the primal and dual. It follows that (63) is an optimal mechanism.

The part of the proposition concerning the equilibrium of the bargaining protocol now follows from Proposition 8.1 and the way the players' strategies are constructed from the primal and dual solutions (using the modification of the strategies from section 5.1 given in section 8.1). In particular, notice that in the path decomposition of λ, P^t is the unique path P in $\mathcal{P}_t$ with $\lambda_P > 0$.

References

Aho, A., Garey, M. & Ullman, J. (1972). The Transitive Reduction of a Directed Graph, *SIAM Journal of Computing* **1**: 131–137.

Ahuja, R., Magnanti, T. & Orlin, J. (1993). *Network Flows*, Prentice Hall.

Ausubel, L. & Deneckere, R. (1989). A Direct Mechanism Characterization of Sequential Bargaining with One-Sided Incomplete Information, *Journal of Economic Theory* **48**: 18–46.

Ben-Porath, E. & Lipman, B. (2008). Implementation and Partial Provability. Unpublished Manuscript.

Blumrosen, L., Nisan, N. & Segal, I. (2007). Auctions with Severely Bounded Communication, *Journal of Artificial Intelligence Research* **28**: 233–266.

Bull, J. & Watson, J. (2007). Hard Evidence and Mechanism Design, *Games and Economic Behavior* **58**: 75–93.

Celik, G. (2006). Mechanism design with Weaker Incentive Compatibility Constraints, *Games and Economic Behavior* **56**: 37–44.

Deneckere, R. & Severenov, S. (2008). Mechanism Design with Partial State Verifiability, *Games and Economic Behavior* **64**: 487–513.

Forges, F. & Koessler, F. (2005). Communication Equilibria with Partially Verifiable Types, *Journal of Mathematical Economics* **41**: 793–811.

Gerardi, D., Horner, J. & Maestri, L. (2010). The Role of Commitment in Bilateral Trade. Cowles Foundation Discussion Paper No. 1760.

Glazer, J. & Rubinstein, A. (2004). On Optimal Rules of Persuasion, *Econometrica* **72**: 1715–1736.

Glazer, J. & Rubinstein, A. (2006). A Study in the Pragmatics of Persuasion: A Game Theoretical Approach, *Theoretical Economics* **1**: 395–410.

Green, J. & Laffont, J. (1986). Partially Verifiable Information and Mechanism Design, *Review of Economic Studies* **53**: 447–456.

Kartik, N. & Tercieux, O. (2009). Implementation with Evidence: Complete Information. Unpublished Manuscript.

Kos, N. (2011). Communication and Efficiency in Auctions. Unpublished Manuscript.

Lipman, B. & Seppi, D. (1995). Robust Inference in Communication Games with Partial Provability, *Journal of Economic Theory* **66**: 370–405.

Milgrom, P. & Roberts, J. (1986). Relying on Information of Interested Parties, *Rand Journal of Economics* **17**: 18–32.

Moore, J. (1984). Global Incentive Constraints in Auction Design, *Econometrica* **52**: 1523–1535.

Myerson, R. (1981). Optimal Auction Design, *Mathematics of Operations Research* **6**: 58–73.

Myerson, R. (1991). *Game Theory: Analysis of Conflict*, Harvard University Press.

Severenov, S. & Deneckere, R. (2006). Screening when Some Agents are Nonstrategic: Does a Monopoly Need to Exclude?, *RAND Journal of Economics* **37**: 816–849.

Sher, I. (2010). Persuasion and Dynamic Communication. Unpublished Manuscript.

Sher, I. (2011). Credibility and Determinism in a Game of Persuasion, *Games and Economic Behavior* **71**: 409–419.

Shin, H. (1994). The Burden of Proof in a Game of Persuasion, *Journal of Economic Theory* **64**: 253–264.

Singh, N. & Wittman, D. (2001). Implementation with Partial Verification, *Review of Economic Design* **6**: 63–84.

Vohra, R. (2011). *Mechanism Design: A Linear Programming Approach*, Cambridge University Press.