

Sher, Itai; Vohra, Rakesh

Article

Price discrimination through communication

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Sher, Itai; Vohra, Rakesh (2015) : Price discrimination through communication, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 10, Iss. 2, pp. 597-648,
<https://doi.org/10.3982/TE1129>

This Version is available at:

<https://hdl.handle.net/10419/150258>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/3.0/>

Price discrimination through communication

ITAI SHER

Department of Economics, University of Minnesota

RAKESH VOHRA

Department of Economics and Department of Electrical and Systems Engineering,
University of Pennsylvania

We study a seller's optimal mechanism for maximizing revenue when a buyer may present evidence relevant to her value. We show that a condition very close to transparency of buyer segments is necessary and sufficient for the optimal mechanism to be deterministic—hence, akin to classic third degree price discrimination—independently of nonevidence characteristics. We also find another sufficient condition depending on both evidence and valuations, whose content is that evidence is hierarchical. When these conditions are violated, the optimal mechanism contains a mixture of second and third degree price discrimination, where the former is implemented via sale of lotteries. We interpret such randomization in terms of the probability of negotiation breakdown in a bargaining protocol whose sequential equilibrium implements the optimal mechanism.

KEYWORDS. Price discrimination, communication, bargaining, commitment, evidence, network flows.

JEL CLASSIFICATION. C78, D82, D83.

1. INTRODUCTION

This paper examines the problem of selling a single good to a buyer whose value for the good is private information. The buyer, however, is sometimes able to support a claim about her value with evidence. Evidence can take different forms. For example, evidence may consist of an advertisement showing the price at which the consumer could buy a substitute for the seller's product elsewhere. It is not essential that a buyer present a physical document; a buyer who knows the market—and, hence, knows of

Itai Sher: itaisher@gmail.com

Rakesh Vohra: rvohra@sas.upenn.edu

A previous version of this paper circulated under the title “Optimal selling mechanisms on incentive graphs.” We are grateful to Melissa Koenig and seminar participants at The Hebrew University of Jerusalem, Tel Aviv University, the Spring 2010 Midwest Economic Theory Meetings, the 2010 Econometric Society World Congress, the 2011 Stony Brook Game Theory Festival, Microsoft Research New England, the 2012 Annual Meeting of the Allied Social Sciences Associations, the Université de Montréal, 2011 Canadian Economic Theory Conference, and the 2014 SAET Conference on Current Trends in Economics. We thank Johannes Horner and two anonymous referees for their helpful comments. Itai Sher gratefully acknowledges a Grant-in-Aid and single semester leave from the University of Minnesota. All errors are our own.

Copyright © 2015 Itai Sher and Rakesh Vohra. Licensed under the [Creative Commons Attribution-NonCommercial License 3.0](https://creativecommons.org/licenses/by-nc/3.0/). Available at <http://econtheory.org>.

DOI: 10.3982/TE1129

attractive outside opportunities—may demonstrate this knowledge through her words alone, whereas an ignorant buyer could not produce those words.

Our model is relevant whenever a monopolist would like to price discriminate on the basis of membership in different consumer segments but disclosure of membership in a segment is voluntary. This is the case with students, senior citizens, AAA members, and many other groups. Moreover, consumer segments often overlap (e.g., many AAA members are senior citizens). If the seller naively sets the optimal price within each segment without considering that consumers in the overlap will select the cheapest available price, she implements a suboptimal policy. So an optimal pricing policy must generally account for the voluntary disclosures that pricing induces.

Our model allows the monopolist not only to set prices conditional on evidence, but to sell lotteries that deliver the object with some probability. Probabilistic sale can be interpreted as delay or quality degradation.¹ Thus, our model entails a mixture of second and third degree price discrimination. Evidence and, moreover, voluntary presentation of evidence, plays a crucial role in generating the richness of the optimal mechanism. In the absence of all evidence, the optimal mechanism is a posted price. When segments are transparent to the seller, which corresponds to the case where evidence disclosure is nonvoluntary or where all consumers can prove membership *and* the lack of it for all segments, the optimal mechanism in our setting is standard third degree price discrimination. More generally, segments may not be transparent, and some consumers may not be able to prove that they do not belong to certain segments. For example, how does one prove that one is not a student? In this case, the optimal mechanism must determine prices for lotteries as a function of submitted evidence. The same lottery may sell to different types for different prices. The allocation need not be monotone in buyers' values in the sense that higher value types may receive the object with lower probability than lower value types. This can be so even when the higher value type possesses all evidence possessed by the lower value type.

We organize the analysis of the problem via the notion of an *incentive graph*. The vertices of the incentive graph are the buyer types. The graph contains a directed edge from s to t if type t can mimic type s in the sense that every message available to type s is available to type t . The optimal mechanism is the mechanism that maximizes revenue subject not to *all* incentive constraints, as in standard mechanism design, but rather only to incentive constraints corresponding to edges in the incentive graph; type t needs to be discouraged from claiming to be type s only if t can mimic s . Similar revelation principles have appeared in the literature without reference to incentive graphs (Forges and Koessler 2005, Bull and Watson 2007, Deneckere and Severinov 2008). Our innovation is to explicitly introduce the notion of an incentive graph, and to link the analysis and specific structure of the optimal mechanism to the specific structure of the incentive graph. While much of the literature deals with abstract settings, we work within the specific price discrimination application.

¹When explicit bargaining is possible, probabilistic sale can be interpreted in terms of the chance of negotiation breakdown. See Section 7.

A key result of our paper is a characterization of the incentive graphs that yield an optimal deterministic mechanism independent of the distribution over types or the assignment of valuations to types; this characterization is in terms of a property we call *essential segmentation* (Proposition 6). Essential segmentation is very close to transparency of segments. So our characterization shows that once one departs slightly from transparency, the distribution of nonevidence characteristics may be such that third degree price discrimination is no longer optimal. We also obtain a weaker sufficient condition for the optimal mechanism to be deterministic that relies on information on valuations (Proposition 4). This sufficient condition can be interpreted as saying that evidence is *hierarchical*, and it allows for solution of the model via backward induction (Proposition 15).

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 under the optimal mechanism? In this sense, our model exhibits the essential difficulty at the heart of optimal mechanism design when types are multidimensional.

Our results have both a positive and negative aspect. On the positive side, we show how to extend known results beyond the case usually studied, where types are linearly ordered, to the more general case of a tree (corresponding to hierarchical evidence). On the negative side, we establish a limit on how far the extension can go, embedding the standard revenue maximization problem in a broader framework that highlights how restrictive it is. However, even when standard results no longer apply, we develop techniques for analyzing the problem despite the ensuing complexity (Propositions 2 and 9).

In our model, randomization can be interpreted as quality degradation, but it can also be interpreted literally: We show that the optimal direct mechanism can be implemented via a bargaining protocol that exhibits some of the important features of bargaining observed in practice (Proposition 11). This model interprets random sale in terms of the probability of a negotiation breakdown. In this protocol, 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 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 more evidence as communication proceeds (Proposition 12).

The seller faces an optimal stopping problem: Should he ask for a further concession from the buyer that 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 that is optimally served with an intermediate probability; then the seller randomizes between asking for more cheap talk and proceeding to the sale. An interesting by-product 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.

The outline of the paper is as follows: [Section 2](#) presents the model. [Section 3](#) presents the benchmark of the standard monopoly problem without evidence. [Section 4](#) highlights the properties of the benchmark that may be violated in our more general model. [Section 5](#) studies the optimal mechanism. [Section 6](#) presents a revenue formula for expressing the payment made by each type in terms of the allocation. [Section 7](#) presents our bargaining protocol. The [Appendix](#) contains proofs that were omitted from the main body.

1.1 *Related literature*

This paper is a contribution to three distinct streams of work. The first, and most apparent, is the study of mechanism design with evidence. Much work in this area ([Green and Laffont 1986](#), [Singh and Wittman 2001](#), [Forges and Koessler 2005](#), [Bull and Watson 2007](#), [Ben-Porath and Lipman 2012](#), [Deneckere and Severinov 2008](#), [Kartik and Tercieux 2012](#)) examines general mechanism design environments, establishing revelation principles, and necessary and sufficient conditions for partial and full implementation. Our focus is on optimal price discrimination instead. The papers most closely related to this one are [Celik \(2006\)](#) and [Severinov and Deneckere \(2006\)](#). [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 In our setting, this would correspond to an incentive graph where directed edge (s, t) exists if and only if t has a higher value than s , and our [Proposition 4](#) applies. [Severinov and Deneckere \(2006\)](#) study a monopolist selling to buyers only some of whom are strategic. Strategic buyers can mimic any other type, whereas nonstrategic types must report their information truthfully. This setting can be seen as a special case of ours where the type of agent is a pair (S, v) that represents a strategic agent with value v or a pair (N, v) , that represents a nonstrategic agent with value v .

The second stream is third degree price discrimination. The study of third degree price discrimination has focused mainly on the impact of particular segmentations on consumer and producer surplus, output, and prices. That literature treats the segmentation of buyers as exogenous. The novelty of this paper is that segmentation is endogenous.³

The third stream is models of persuasion ([Milgrom and Roberts 1986](#), [Shin 1994](#), [Lipman and Seppi 1995](#), [Glazer and Rubinstein 2004, 2006](#), [Sher 2011, 2014](#)). These models deal with situations in which a speaker attempts to persuade a listener to take some action. Our model deals with arguments attempting to persuade the listener, i.e.,

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

³A recent paper by [Bergemann et al. \(2015\)](#) also examines endogenous segmentation. However, they assume a third party who can segment buyers by valuation. Thus, buyers are not strategic in their setting.

seller, to choose an action like lowering the price. Indeed, Glazer and Rubinstein's model can be reinterpreted as a price discrimination model where the buyer has a binary valuation for the object, assigning it either a high or low value. Our model can then be seen as a generalization from the case of binary valuations to arbitrary valuations (see Section 7.5).

A related line of work is Blumrosen et al. (2007) and Kos (2012). These papers assume that bidders can only report one of a finite number of messages. However, unlike the model we consider, all messages are available to each bidder. Hard evidence can be thought of as a special case of differentially available or differentially costly actions. One such setting is that of auctions with financially constrained bidders who cannot pay more than their budget (Che and Gale 1998, Pai and Vohra 2014). This relation potentially links our work to a broader set of concerns. In relation to our credible implementation of the optimal mechanism via a bargaining protocol (see Section 7), there is also a body of literature that studies the relation between incentive compatible mechanisms and outcomes that can be implemented in infinite horizon bargaining games with discounting (Ausubel and Deneckere 1989, Gerardi et al. 2014). 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. THE MODEL

2.1 Primitives

A seller possesses a single item he does not value. A buyer may be one of a finite number of types in a set T . The terms π_t and v_t denote, respectively, the probability of and valuation of type t . Let M be a finite set of *hard* messages, and $\sigma: T \rightrightarrows M$ is a *message correspondence* that determines the evidence $\sigma(t) \subseteq M$ available to type t . For any subset S of $\sigma(t)$, the buyer can present S . It is convenient to define: $S_t := \{m: m \in \sigma(t)\}$. Formally, S_t and $\sigma(t)$ are the same set of messages. However, we think of $\sigma(t)$ as encoding the buyer's choice set, while we think of S_t as encoding a particular choice: namely, the choice to present all messages in $\sigma(t)$. Observe that if $\sigma(t) \subseteq \sigma(s)$, then type s can also present S_t .

Assume a *zero type* $0 \in T$ with $v_0 = \pi_0 = 0$ and $\sigma(0) = \{m_0\} \subsetneq \sigma(t) \forall t \in T \setminus 0$. Thus, all types possess the single hard message available to the zero type. The zero type plays the role of the outside option. 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, which are equally available to all types, as in standard mechanism design models without evidence.

2.2 Incentive graphs

A graph $G = (V, E)$ consists of a set of vertices V and a set of directed edges E , where an edge is an ordered pair of vertices. The *incentive graph* is the graph G such that $V = T$

and E is defined by

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

So $(s, t) \in E$ means that t can *mimic* s in the sense that any evidence that s can present, t can also present. Our assumptions on the zero type imply

$$\forall t \in T \setminus 0, \quad (0, t) \in E \quad \text{and} \quad (t, 0) \notin E. \quad (2)$$

A graph $G = (V, E)$ is *transitive* if, for all types r, s , and t , $[(r, s) \in E \text{ and } (s, t) \in E] \Rightarrow (r, t) \in E$. The incentive graph is not transitive (because it is irreflexive), but (1) implies that the incentive graph satisfies a slightly weaker property we call *weak transitivity*:

$$[(r, s) \in E \text{ and } (s, t) \in E \text{ and } r \neq t] \Rightarrow (r, t) \in E \quad \forall r, s, t \in T (= V).$$

Say that an edge $(s, t) \in E$ is *good* if $v_s < v_t$ and *bad* otherwise.

2.3 Graph-theoretic terminology

Here we collect some graph-theoretic terminology used in the sequel. We suggest that the reader skip this section and return to it as needed. A *path* in $G = (V, E)$ is a sequence $P = (t_0, t_1, \dots, t_n)$ of vertices with $n \geq 1$ such that for $i = 1, \dots, n$ and $j = 0, \dots, n$, (i) $(t_{i-1}, t_i) \in E$, and (ii) $i \neq j \Rightarrow t_i \neq t_j$. If for some $i = 1, \dots, n$, $s = t_{i-1}$ and $t = t_i$, we write $(s, t) \in P$ and $t \in P$ (and also $s \in P$). Path P is an $s - t$ *path* if $t_0 = s$ and $t_n = t$, and \mathcal{P}_{s-t} is the set of all $s - t$ paths in G . The set $\mathcal{P}_t := \mathcal{P}_{0-t}$ is the set of all $0 - t$ paths and $\mathcal{P} := \bigcup_{t \in T \setminus 0} \mathcal{P}_t$ is the set of all paths originating in 0. We sometimes use the notation $P: t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_n$ for the path P . A *cycle* in G is a sequence $C = (t_0, t_1, \dots, t_n)$ of vertices such that for $i, j = 1, \dots, n$, (i) $(t_{i-1}, t_i) \in E$, (ii) $i \neq j \Rightarrow t_i \neq t_j$, and (iii) $t_0 = t_n$. Graph G is *acyclic* if G does not contain any cycles. An *undirected path* is a sequence (t_0, t_1, \dots, t_n) such that for $i = 1, \dots, n$ and $j = 0, \dots, n$, (i) either $(t_{i-1}, t_i) \in E$ or $(t_i, t_{i-1}) \in E$, and (ii) $i \neq j \Rightarrow t_i \neq t_j$. A graph $G = (V, E)$ is *strongly connected* if for all vertices $s, t \in V$, there is an $s - t$ path in G . Graph G is *weakly connected* if there is an undirected path connecting each pair of its vertices. Graph $G_0 = (V_0, E_0)$ is a *subgraph* of $G = (V, E)$ if $V_0 \subseteq V$ and $E_0 := \{(s, t) \in E : s, t \in V_0\}$. We refer to G_0 as the *subgraph of G generated by V_0* . A subgraph $G_0 = (V_0, E_0)$ of G is a *strongly* (resp., *weakly*) *connected component of G* if (i) G_0 is strongly (resp., weakly) connected, and (ii) if $V_0 \subsetneq V_1 \subseteq V$, the subgraph of G generated by V_1 is not strongly (resp., weakly) connected. Finally the *outdegree* of a vertex t is the number of vertices s such that $(t, s) \in E$.

2.4 The seller's revenue maximization problem

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

PRIMAL PROBLEM (Edges). The seller's problem is

$$\text{maximize } \sum_{t \in T} \pi_t p_t \quad (3)$$

subject to

$$\forall (s, t) \in E, \quad v_t q_t - p_t \geq v_t q_s - p_s \quad (4)$$

$$\forall t \in T, \quad 0 \leq q_t \leq 1 \quad (5)$$

$$p_0 = 0. \quad (6)$$

The seller's objective is to maximize expected revenue (3). In contrast to standard mechanism design, (3)–(6) does not require one 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 6. 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 $t \in T \setminus 0$). An allocation $q = (q_t : t \in T)$ is said to be *incentive compatible* if there exists a vector of payments $p = (p_t : t \in T)$ such that (q, p) is feasible in (3)–(6).

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 and Severinov (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 that 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 and Severinov (2008). Related arguments are presented by Bull and Watson (2007). (Note that our model satisfies Bull and Watson's normality assumption because each type t buyer can present all subsets of $\sigma(t)$.⁴)

3. THE STANDARD MONOPOLY PROBLEM

This section summarizes a special case of our problem that will serve as a benchmark: the standard monopoly problem. Call the incentive graph *complete* if for all $s, t \in T$ with $t \neq 0$, $(s, t) \in E$; with a complete incentive graph, every (nonzero) type can mimic every other type. The *standard monopoly problem* is (3)–(6) with a complete incentive

⁴Bull and Watson (2007) also explain the close relation of their normality assumption to the nested ranged condition of Green and Laffont (1986) and relate their analysis to that of the latter paper.

graph. Here, we assume without loss of generality that $T = \{0, 1, \dots, n\}$ with $0 < v_1 \leq v_2 \leq \dots \leq v_n$.

For each $t \in T$, define

$$\Pi_t = \sum_{i=t}^n \pi_i, \quad (7)$$

where Π_t , viewed as a function of t , is the *complementary cumulative distribution function* (c.c.d.f.). Define the *quasi-virtual value* of type $t \in T \setminus n$ as⁵

$$\widehat{\psi}(t) := v_t - (v_{t+1} - v_t) \frac{\Pi_{t+1}}{\pi_t}. \quad (8)$$

The quasi-virtual value of type n is simply $\widehat{\psi}(n) := v_n$. Use of the qualifier “quasi” is explained in Section 5.2. Say that quasi-virtual values are *monotone* if

$$\widehat{\psi}(t) \leq \widehat{\psi}(t+1), \quad t = 0, \dots, n-1. \quad (9)$$

The following proposition summarizes the well known properties of the standard monopoly problem.

PROPOSITION 1 (Standard monopoly benchmark). *Any instance of the standard monopoly problem satisfies the following properties.*

(i) *The allocation $q \in [0, 1]^T$ is incentive compatible exactly if q satisfies allocation monotonicity:*

$$q_t \leq q_{t+1}, \quad t = 0, 1, \dots, n-1.$$

(ii) *For any allocation q satisfying allocation monotonicity, the revenue maximizing vector of payments p such that (q, p) is feasible in (4)–(6) is given by the revenue formula*

$$p_t = v_t q_t - \sum_{i=1}^{t-1} q_i (v_{i+1} - v_i) - v_1 q_0. \quad (10)$$

(iii) *There exists an optimal mechanism satisfying the following statements:*

(a) *Deterministic allocation. Each type is allotted the object with probability 0 or 1.*

(b) *Uniform price. Each type receiving the object makes the same payment.*

(iv) *Assume quasi-virtual values are monotone. Then in any optimal mechanism, the buyer is served with probability 1 if she has a positive quasi-virtual value and with probability 0 if she has a negative quasi-virtual value.⁶ The seller’s expected revenue is equal to the expected value of the positive part of the quasi-virtual value: $\sum_{t \in T} \{\widehat{\psi}(t), 0\} \pi_t$.*

⁵Notice that because $\pi_0 = 0$, $\widehat{\psi}(0) = -\infty$.

⁶If the buyer has a zero virtual value, then for every $\alpha \in [0, 1]$, there is an optimal mechanism in which the buyer is served with probability α .

Part (i) follows from standard arguments (as in Myerson 1982), part (ii) follows from Lemmas 3 and 5 (in the Appendix), and part (iii) is an easy corollary of Proposition 4. The proof of part (iv) is given in the Appendix.

4. DEVIATIONS FROM THE BENCHMARK

Using the case of the complete incentive graph as a benchmark (as cataloged in Proposition 1), this section highlights various anomalies that arise when the incentive graph is incomplete.

OBSERVATION 1 (Nonstandard properties of feasible and optimal mechanisms).

- (i) *When the incentive graph is incomplete, all allocations may be feasible (i.e., incentive compatible). Acyclicity of the incentive graph is a necessary and sufficient condition for all allocations to be incentive compatible.*
- (ii) *When the incentive graph is incomplete, the optimal mechanism may involve any of the following properties:*
 - (a) *Price discrimination. Two types may pay different prices for the same allocation.*
 - (b) *Random allocation. Some types may be allotted the object with a probability intermediate between 0 and 1.⁷*
 - (c) *Violations of allocation monotonicity. For some good edge (s, t) , type t may be allotted the object with lower probability than s .*

REMARK 1. Whereas (ii)(a) refers to third degree price discrimination, (ii)(b) can be interpreted as a form of second degree price discrimination so that the optimal mechanism contains a mix of the two. Part (ii)(c) generalizes allocation monotonicity to arbitrary incentive graphs, and shows that allocation monotonicity, which was implied by feasibility for the complete graph (Proposition 1(i)), is not even implied by optimality for arbitrary incentive graphs.

The proof of Observation 1(i) is omitted because it is not used directly as a lemma in any of our main results. The other parts are illustrated by the following example.

EXAMPLE 1. Consider the market for a book. Some students are required to take a class for which the book is required and, consequently, they have a high value of 3 for the

⁷Strictly speaking, what differentiates this from the standard monopoly problem is that for a fixed incentive graph and assignment of values and probabilities to types, *all* optimal mechanisms may require randomization, so that randomization is essential rather than incidental. Moreover, in the standard monopoly problem (the discrete version), the set of parameter values for which there even exists an optimal mechanism involving randomization has zero measure. In contrast, when the incentive graph is incomplete, the set of parameter values inducing *all* optimal mechanisms to randomize is nondegenerate. See Remark 2.

book. Students who are not required to take the class have a low value of 1. More than half the students are required to take the class. All nonstudents have a medium value of 2.

If no type had any evidence, a posted price would be optimal. Suppose next, for illustrative purposes, all buyers have an ID card that records their student status or lack of it. Then it would be optimal to set a price of \$2 to nonstudents and \$3 to students.

Suppose more realistically that only students possess an ID identifying their student status and nonstudents have no ID. If the seller now attempted to set a price of \$2 to nonstudents and \$3 to students, no student would choose to reveal his student status. Thus, the natural form of third degree price discrimination is ruled out.⁸

In this case, the seller can benefit from a randomized mechanism. If students did not exist, the optimal mechanism would be a posted price of 2. So by a continuity argument, if the proportion of nonstudents is large enough, the optimal mechanism is such that without a student ID, a buyer faces a posted price of 2. The seller cannot charge a price higher than 2 with a student ID, since a student can receive this price when he withholds his ID. Suppose, however, that the seller can offer to sell a lower quality version of the product that yields the same payoff as receiving the object with a probability of 1/2. Let the seller offer this option only with a student ID for a price of 1/2. The low student type would be willing to select this option, while the high student type would be willing to mimic a nonstudent and obtain the high quality version for a price of 2. Indeed, it is easy to see that under our assumptions, this randomized mechanism is optimal.

Finally, let us introduce a small proportion of students who have a value $2 + \epsilon$, where ϵ is a small positive number. If these students form a sufficiently small proportion of the student population, it will still be optimal for the seller to offer a price of 2 for the high quality product without a student ID, and a price of 1/2 for the low quality product (equivalent to receiving the object with probability 1/2) with a student ID. The new medium value students will prefer the lower price of 1/2 with a student ID. However, this is a violation of allocation monotonicity, as these new students can mimic the nonstudents who have a lower value and receive the item with probability 1. ◇

REMARK 2 (Nondegeneracy). [Example 1](#) shows that random allocation is not a knife-edge phenomenon. For sufficiently small changes in the parameters—the values and probabilities of the (nonzero) types—the optimum in the last paragraph of the example remains unique, and still has the properties of random allocation and failure of allocation monotonicity. With a view to [Proposition 2](#), types with zero virtual valuation as defined by (19) (the only types eligible for random allocation at the optimum) are not a knife-edge phenomenon, but rather can be robust to small changes in the parameters.

⁸Since more than half of the students have a high value, the seller would prefer to sell the object for \$2 regardless of student status rather than to offer a discounted price of \$1 to students.

5. THE OPTIMAL MECHANISM

5.1 *Virtual values and general properties of the optimal mechanism*

Here we analyze the revenue maximization problem on arbitrary incentive graphs via a generalization of the classical analysis of optimal auctions in terms of virtual valuations. To do so, we display the dual of the seller’s problem.⁹

DUAL PROBLEM (Edges). The dual is

$$\underset{(\mu_t)_{t \in T}, (\lambda(s,t))_{(s,t) \in E}}{\text{minimize}} \sum_{t \in T} \mu_t \tag{11}$$

subject to

$$\forall t \in T \setminus 0, \quad \sum_{s: (s,t) \in E} \lambda(s, t) - \sum_{s: (t,s) \in E} \lambda(t, s) = \pi_t \tag{12}$$

$$\forall t \in T, \quad v_t \pi_t - \sum_{s: (t,s) \in E} \lambda(t, s)(v_s - v_t) \leq \mu_t \tag{13}$$

$$\forall (s, t) \in E \quad \lambda(s, t) \geq 0 \tag{14}$$

$$\forall t \in T, \quad \mu_t \geq 0. \tag{15}$$

We now use the dual to derive a generalization (19) of the classical notion of virtual value, a key to our analysis. The standard virtual value (8) employs the complementary cumulative distribution function (c.c.d.f.) Π_t (see (7)). As we now show, the dual variables $\lambda(s, t)$ can be interpreted as providing a generalization of the c.c.d.f. to arbitrary incentive graphs. However, whereas Π_t is exogenous, the quantities $\lambda(s, t)$ used to construct the analog of Π_t are endogenous. For the complete incentive graph, Π_t is the probability of all types “above” type t (including t), in the sense of having a higher value than t . On an arbitrary incentive graph, types differ not only by value (and probability), but also according to which other types they can mimic. Thus, there is no obvious way to linearly order types such that some types are above others. Nevertheless, we construct an analog of the c.c.d.f.

Next we provide insight into how the generalization is achieved. Consider first the case of the complete incentive graph (where we recall $T = \{0, 1, \dots, n\}$ and $0 = v_0 < v_1 \leq \dots \leq v_n$; see Section 3). If quasi-virtual values are monotone, then, by well known reasoning, at an optimum of the primal, the downward adjacent constraints (i.e., those

⁹The derivation of (13) requires some manipulation: When one takes the dual of (3)–(6), one initially gets the constraint

$$\sum_{s: (s,t) \in E} v_t \lambda(s, t) - \sum_{s: (t,s) \in E} v_s \lambda(t, s) - \mu_t \leq 0 \quad \forall t \in T$$

instead of (13). Using (12) to substitute $\pi_t + \sum_{s: (t,s) \in E} \lambda(t, s)$ for $\sum_{s: (s,t) \in E} \lambda(s, t)$ in the above inequality, we obtain

$$\mu_t \geq v_t \left(\pi_t + \sum_{s: (t,s) \in E} \lambda(t, s) \right) - \sum_{s: (t,s) \in E} v_s \lambda(t, s) = v_t \pi_t - \sum_{s: (t,s) \in E} \lambda(t, s)(v_s - v_t),$$

which is constraint (13).

of the form $(t, t + 1)$ bind. Moreover, we can eliminate all other incentive constraints without altering the optimal solution.¹⁰ As $\lambda(s, t)$ is the multiplier on the (s, t) incentive constraint, it follows that there is a dual optimum satisfying

$$\lambda(s, t) > 0 \quad \text{only if } t = s + 1. \quad (16)$$

So (12) simplifies to

$$\begin{aligned} \lambda(t - 1, t) - \lambda(t, t + 1) &= \pi_t, \quad t = 1, \dots, n - 1 \\ \lambda(n - 1, n) &= \pi_n. \end{aligned} \quad (17)$$

It follows that

$$\begin{aligned} \sum_{s: (t,s) \in E} \lambda(t, s) &= \lambda(t, t + 1) \\ &= \left(\sum_{i=t+1}^{n-1} [\lambda(i - 1, i) - \lambda(i, i + 1)] \right) + \lambda(n - 1, n) \\ &= \sum_{i=t+1}^n \pi_i, \end{aligned}$$

where the first equality follows from (16), the second equality is a telescoping sum, and the third follows from (17). To summarize,

$$\sum_{s: (t,s) \in E} \lambda(t, s) = \Pi_{t+1}. \quad (18)$$

Equation (18) delivers the promised relationship between the dual solution and the c.c.d.f. when the incentive graph is complete. When the incentive graph is not complete, then (18) suggests that we use $\sum_{s: (t,s) \in E} \lambda(t, s)$ instead of Π_{t+1} for the cumulative probability mass “above” t .¹¹

We are now in a position to construct an analog of the virtual value. This is done by means of constraints (13), which we call *virtual value constraints*. If we divide the

¹⁰If virtual values were not monotone, eliminating the nonadjacent constraints would require us to introduce allocation monotonicity as an additional set of constraints, which in turn would introduce new variables into the dual.

¹¹To make this vivid, imagine that all the probability (of mass 1) is concentrated on the zero type. We would like to transport this probability mass to the other types along the edges of the incentive graph so that each type t receives her allotted share π_t . The quantity $\lambda(s, t)$ is the total probability mass that travels along edge (s, t) . Then $\sum_{s: (s,t) \in E} \lambda(s, t)$ is the total probability mass flowing into type t and $\sum_{s: (t,s) \in E} \lambda(t, s)$ is the total mass flowing out of type t , in other words, the probability mass above t . Equation (12) then says that the difference between the inflow and the outflow of t is π_t , the mass that t receives. For this reason, the constraints (12) are called *flow conservation constraints*. Such constraints have been extensively studied in the literature on network flow problems. See Ahuja et al. (1993) for an extensive treatment.

constraint (13) corresponding to t by π_t , and call the resulting expression on the left-hand side $\psi(t)$, then¹²

$$\psi(t) := v_t - \frac{\sum_{s:(t,s) \in E} \lambda(t,s)(v_s - v_t)}{\pi_t}. \quad (19)$$

In the case of the complete incentive graph (assuming also (9)), using (16) and (18), at a dual optimum, (19) reduces to (8), the quasi-virtual value. This suggests that, in general, we interpret $\psi(t)$ as the *virtual valuation* of type t , an analog of the virtual valuation in traditional mechanism design. Constraints (13) and (15), and the minimization (11) establish the following relation at any dual optimum:

$$\mu_t = \max\{\psi(t), 0\} \pi_t.$$

In words, μ_t is the positive part of the virtual valuation of type t multiplied by the probability of type t . Proposition 2 below now follows from strong duality and complementary slackness. This proposition—an analog of the standard result from the theory of optimal auctions and of Proposition 1(iv)—validates our interpretation in terms of virtual values.

PROPOSITION 2. *At any optimal mechanism, a buyer type is served with probability 1 if she has a positive virtual valuation and with probability 0 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:*

$$\sum_{t \in T} \max\{\psi(t), 0\} \pi_t.$$

This result establishes one link between the standard analysis and our model. We conclude this section with additional results that a general optimal solution in our model shares with an optimum of the standard problem. These results will also be useful in the sequel. Call a feasible solution to the dual *good* if

$$\lambda(s, t) > 0 \quad \Rightarrow \quad v_s < v_t \quad \forall (s, t) \in E. \quad (20)$$

In words, a dual solution is good if the variables $\lambda(s, t)$ are only positive on good edges. For our next result, we present an elementary definition: For any feasible solution $z = (q, p)$ to the primal (3)–(6) and $(s, t) \in E$, we say that the (s, t) incentive constraint *binds* at z if the incentive constraint (4) corresponding to (s, t) holds with equality.

PROPOSITION 3. (i) *Elimination of bad edges. Eliminating incentive constraints corresponding to bad edges in the primal does not alter the optimal expected revenue. Consequently, there exists a dual optimum that is good.*

(ii) *Monotonicity along binding constraints. If $z = (q, p)$ is an optimal mechanism and the (s, t) -incentive constraint binds at z , then $q_s \leq q_t$.*

¹²Notice, in particular, that because $\pi_0 = 0$, $\psi(0) = -\infty$.

By part (i), only incentive constraints corresponding to good edges are relevant. By part (ii), we have allocation monotonicity along good edges, provided those edges correspond to binding constraints, a partial antidote to [Observation 1\(ii\)\(c\)](#).¹³ [Proposition 3\(ii\)](#) does not follow from the definition of a binding constraint alone. If we replace the word “optimal” in the proposition by “feasible,” the proposition is false. This highlights a contrast with the standard problem with complete incentive graph: Whereas in the standard problem, monotonicity follows from feasibility (i.e., incentive compatibility), with an incomplete incentive graph, monotonicity—restricting attention to binding constraints—requires optimality. Feasibility is insufficient because the standard argument requires not only the downward incentive constraint saying the higher type does not want to mimic the lower type, but also the upward constraint. In our model, we may have the downward constraint without having the upward constraint. However, for binding constraints, a higher allocation must be accompanied by a higher price, so that a violation of monotonicity would allow an increase in revenue to be achieved by allowing the higher type to receive the lower type’s higher price and allocation, which the higher type desires.

5.2 Optimality of deterministic mechanisms

This section presents an essentially complete solution for the case where the optimal mechanism is deterministic. [Proposition 4](#) presents a sufficient condition—tree structure—for the optimal mechanism to be deterministic. This sufficient condition depends on the valuations assigned to types. [Proposition 6](#) presents a characterization, a necessary and sufficient condition—essential segmentation—on the incentive graph alone for the optimal mechanism to be deterministic regardless of the assignment of values to types. [Lemma 1](#) relates the two previously mentioned results: The incentive graph satisfies our characterization (essential segmentation) if and only if our sufficient condition (tree structure) holds for all (nondegenerate) assignments of valuations to types. This justifies studying deterministic optima through the lens of our sufficient condition of tree structure. We go on to relate the optimal solution to the classic analysis of optimal auctions. [Appendix A.3](#) provides a simple algorithm to find the optimum under tree structure.

To proceed, we require some definitions. A *tree* is a graph such that for every two distinct vertices s and t , the graph contains a unique undirected path from s to t . The *monotone incentive graph* is the graph $G' = (V, E')$, where E' is the set of good edges in E (see [Section 2.2](#)). [Proposition 3\(i\)](#) suggests the relevance of G' to the seller’s problem. Observe that possibly unlike G , G' is acyclic. It follows from [Observation 1\(i\)](#) that if G' rather than G were the true incentive graph, then all allocations would be incentive compatible, yet by [Proposition 3\(i\)](#), the optimal expected revenue induced by G and G' is the same.

¹³Another partial antidote to [Observation 1\(ii\)\(c\)](#) is that $\sigma(s) = \sigma(t)$ and $v_s < v_t$ imply $q_s \leq q_t$. This condition holds at all feasible mechanisms, not just at the optimal mechanism. Note also that if $\sigma(s) = \sigma(t)$, then if s and t receive the same allocation, s and t also receive the same price.

The *transitive reduction* $G^* = (V, E^*)$ of G' is the smallest subgraph of G' with the property that G^* and G' have the same set of (directed) paths.¹⁴ So, for example, if G' contains edges (r, s) , (s, t) , and (r, t) , then in G^* , we eliminate (r, t) . Because G' is acyclic, G^* is well defined. If G^* is a tree, then we say that the environment satisfies *tree structure*. Observe that tree structure is a property not just of the incentive graph, but of the assignment of values to types because the monotone incentive graph depends on this assignment. Indeed, as the set of good edges depends on the valuation profile, G' and G^* depend on the assignment of values to types $v = (v_t : t \in T)$, and we sometimes write G'_v and G^*_v to make this dependence explicit. Unlike G and G' , the graph G^* is not (weakly) transitive.¹⁵ Finally, it is easy to verify that G^* is a tree exactly if for every $t \in T \setminus 0$, there exists a unique *directed* $0 - t$ path in G^* .

Tree structure can be interpreted as imposing a *hierarchical* structure on evidence: The condition means that if $v_r < v_s < v_t$, and type t can mimic both r and s , then type s can mimic type r , so that there is a clear hierarchy among (lower value) types that any type t can mimic. This condition holds if there is no evidence. Alternatively, suppose evidence consists of a set of provable characteristics and that the more of these characteristics one has, the more attractive the item becomes. If every pair of characteristics is either incompatible, so that no agent can have both, or clearly ranked, so that anyone who has the higher ranked characteristic also has the lower ranked characteristic, then evidence is hierarchical. While this is natural, it is also restrictive.

We now provide a sufficient condition for a deterministic optimum.

PROPOSITION 4 (Tree structure). *If G^* is a tree, then the optimal mechanism is deterministic.*

The feature of the tree structure that we exploit to prove [Proposition 4](#) is that it allows one to determine the binding incentive constraints a priori.¹⁶ These are analogous to the downward adjacent constraints that typically bind in standard problems, but now occur in the context of a tree rather than a line. This feature also accounts for the other nice properties associated with tree structure detailed below.

EXAMPLE 2. This example illustrates tree structure. Let $T = \{0, 1, \dots, 7\}$, and consider [Figure 1](#), which illustrates the graph G^* .

Edge $(s, t) \in E'$ (the edge set for the monotone incentive graph) if in [Figure 1](#) there is a directed path from s to t . For example, the arrow from 3 to 4 means that 4 can mimic 3 and the path $1 \rightarrow 3 \rightarrow 4$ means that 4 can mimic 1. Let the numbers of the types also represent their values. As usual, $\pi_0 = 0$. Suppose that $\pi_1 = \pi_2 = \pi_3 =: \pi_a$ and $\pi_4 = \pi_5 =$

¹⁴In other words, G^* is the unique subgraph of G' such that (i) G' and G^* have the same set of directed paths, and (ii) for any subgraph G'' of G' with the same set of directed paths as G' , G^* is a subgraph of G'' .

¹⁵The one exception is the (uninteresting) case where for each nonzero type t , the unique type s such that $(s, t) \in E$ is 0. In this case, G^* is weakly transitive.

¹⁶The key result in this connection is [Lemma 5](#).





FIGURE 1. A tree.

FIGURE 2. The graph G^* for the student ID example.

$\pi_6 = \pi_7 =: \pi_b$. If π_b/π_a is sufficiently small, then the unique optimal mechanism sells the object to all types except 0 and 1, sets a price of 2 for types 2, 5, and 6, and sets a price of 3 for types 3, 4, and 7. In contrast, if π_b/π_a is sufficiently large, only types 4, 5, 6, and 7 are served, and each is charged a price equal to her value. As stated by [Proposition 4](#), in both cases, the optimal mechanism is deterministic. However, the optimal mechanism does involve price discrimination as different types receive different prices. \diamond

In light of [Proposition 4](#), it is instructive to reconsider the student ID example ([Example 1](#)). When only students have an ID, the optimal mechanism was randomized. Let HS and LS stand for the high and low value students, respectively, and let NS stand for the nonstudents. Then the graph G^* is given by [Figure 2](#).

Although LS can mimic NS, there is no edge from LS to NS because nonstudents have a higher value than low value students and G^* does not contain bad edges. As there are two paths from 0 to HS, G^* is not a tree.

Tree structure is also sufficient to guarantee allocation monotonicity.

PROPOSITION 5 (Allocation monotonicity). *If G^* is a tree, then every optimal mechanism (q, p) is monotone in the sense that if (s, t) is a good edge in E , then $q_s \leq q_t$.*

This strengthens the form of allocation monotonicity of [Proposition 3\(ii\)](#), and approaches the form of allocation monotonicity present in the standard monopoly problem.¹⁷

Tree structure depends on the assignment of values to types. Next we provide a characterization of the incentive graphs that induce a deterministic optimum independently of the value assignment.

Let V° be the set of vertices with outdegree zero in G , so that V° represents the set of types that cannot be mimicked by other types. Define

$$\tilde{V} := V \setminus (V^\circ \cup 0), \quad \tilde{E} := \{(s, t) \in E : s \neq 0, t \notin V^\circ\}, \quad \tilde{G} = (\tilde{V}, \tilde{E}).$$

Let $\{G_i = (V_i, E_i) : i = 1, \dots, n\}$ be the set of weakly connected components of \tilde{G} . Say that G is *essentially segmented* if (i) G_i is strongly connected for $i = 1, \dots, n$, and (ii) for each $t \in V^\circ$, there exists i such that $\{s \in V \setminus 0 : (s, t) \in E\} \subseteq V_i$. (For the definition of weakly and strongly connected components, see [Section 2.3](#).)

PROPOSITION 6 (Essential segmentation). *Graph G induces a deterministic optimum for all assignments of values and probabilities if and only if G is essentially segmented.*

Essential segmentation is similar to, but slightly weaker than, the standard assumption associated with third degree price discrimination that the seller can distinguish between different consumer segments.¹⁸ This corresponds in our model to the situation where (nonzero) types can be partitioned into segments such that all types within a segment can mimic one another and no type in any segment can mimic any type in a different segment. This means that incentive constraints must be honored within segments but not across segments. Such an incentive graph would be induced if agents within a segment have the same evidence and the evidence of any one segment is not a subset of the evidence of any other segment. Essential segmentation is weaker than standard segmentation only insofar as it allows for types t with outdegree zero (i.e., types who cannot be mimicked by any other types), such that t can mimic types in at most one segment. When there are no such zero outdegree types, that is, when $V^\circ = \emptyset$, then essential segmentation reduces to the standard notion of segmentation. As under essential segmentation, each type can claim mimic types in only one segment, so we can interpret [Proposition 6](#) to mean that transparency of segments is very close to a necessary and sufficient condition for third degree price discrimination to be always optimal (independent of the values and probabilities of types).

Say that valuation assignment $(v_t : t \in T)$ is nondegenerate if $\forall s, t \in T, s \neq t \Rightarrow v_s \neq v_t$. [Lemma 1](#) relates tree structure to essential segmentation ([Propositions 4](#) and [6](#)).

LEMMA 1. *Graph G is essentially segmented if and only if for all nondegenerate values assignments v , G_v^* is a tree.*

¹⁷Whereas in the standard monopoly problem, allocation monotonicity follows from *feasibility*, when G^* is a tree, allocation monotonicity requires the stronger assumption of *optimality*. See the discussion following [Proposition 3](#).

¹⁸Equivalently, the seller observes a signal correlated with the consumer's value.

In the deterministic case, under an additional assumption, we can strengthen [Proposition 2](#) to arrive at a stronger characterization of the optimum in terms of virtual values. When G^* is a tree, define the *quasi-virtual value* of type t as

$$\widehat{\psi}(t) := v_t - \sum_{s: (t,s) \in E^*} (v_s - v_t) \frac{\Pi_s}{\pi_t}, \tag{21}$$

where

$$\Pi_s := \pi_s + \sum_{r: (s,r) \in E'} \pi_r. \tag{22}$$

Observe that Π_s is defined with respect to the edge set E' in the *monotone* incentive graph and $\widehat{\psi}(t)$ is defined with respect to the edge set E^* of the transitive reduction.

Whereas the virtual value as defined by (19) is endogenous, in that it depends on an optimal solution to the dual of the seller’s problem, the quasi-virtual value defined by (21) is exogenous, defined purely in terms of primitives. Terminologically, (21) is more similar to the notion of virtual value in [Myerson \(1982\)](#), whereas (19) is more similar to the notions in [Myerson \(1991\)](#) and [Myerson \(2002\)](#). Because the sign of the expression in (19) is always a reliable guide to the allotment of the object, whereas the sign of (21) is only a reliable guide to the allotment of the object under special assumptions, we call the expression in (19) the virtual value and call the expression in (21) the quasi-virtual value.

Say that quasi-virtual values are *single-crossing* if

$$(s, t) \in E' \Rightarrow (\widehat{\psi}_s \geq 0 \Rightarrow \widehat{\psi}_t \geq 0) \quad \forall s, t \in V.$$

In the standard case of the complete incentive graph, single-crossing monotone virtual values is a weakening of the common assumption of monotone quasi-virtual values (see [Section 3](#) and especially (9)). Under tree structure and single-crossing quasi-virtual values, we attain a strengthening of [Proposition 2](#), which is also an analog of the characterization of optimal auctions due to [Myerson \(1982\)](#).

PROPOSITION 7. *Assume that G^* is a tree and that quasi-virtual values are single-crossing. Then an optimal mechanism serves each type exactly if her quasi-virtual value is nonnegative.¹⁹ Formally, an optimal mechanism (q^*, p^*) is*

$$\begin{aligned} q_t^* &= \begin{cases} 1 & \text{if } \widehat{\psi}(t) \geq 0 \\ 0 & \text{otherwise} \end{cases} \\ p_t^* &= \begin{cases} \min(\{v_t\} \cup \{v_s : (s, t) \in E, \widehat{\psi}(s) \geq 0\}) & \text{if } \widehat{\psi}(t) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \end{aligned} \tag{23}$$

REMARK 3. Single-crossing virtual values are sufficient for [Proposition 7](#) because we deal with the single agent case. In the multi-agent case, we would need to assume monotone virtual values.

¹⁹One could replace the term “nonnegative” by “positive” and, correspondingly, replace the weak inequalities by strict inequalities in (23), and the proposition would continue to be true.

Assuming tree structure and single-crossing quasi-virtual values, [Proposition 2](#) shows that the optimal mechanism can be found simply by computing the quasi-virtual values. When tree structure holds, but the assumption of single-crossing quasi-virtual values fails, then it is only slightly more difficult to find the optimal mechanism. The optimal mechanism can then be found by solving a simple perfect information game between the seller and nature by backward induction. We describe this game in [Appendix A.3](#).

6. A REVENUE FORMULA

[Section 5](#) described the properties of the optimal mechanism when the incentive graph is incomplete. The analysis revealed analogies (and disanalogies) to the classical analysis. Specifically, an analog of the classical notion of virtual value plays a prominent role in our analysis. This section explores another related analogy between classical mechanism design and mechanism design with an incomplete incentive graph: revenue equivalence.

In classical mechanism design, when types are continuous, the allocation determines the revenue (up to a constant). Even in the classical case with complete incentive graph, when types are discrete as in [Section 3](#), we do not attain exact revenue equivalence, but we do attain a formula for expressing the maximal revenue consistent with an allocation in terms of that allocation, which is analogous to the formula for expressing the unique revenue in terms of the allocation in the continuous case. In this section, we derive an analog of the classical discrete revenue formula when the incentive graph is incomplete. In addition to its intrinsic interest, [Proposition 9](#) served as a useful lemma for proving several of the results in previous sections.

On the way to our result, it is useful and instructive to reformulate the seller's problem.

PRIMAL PROBLEM (Paths). The reformulation is

$$\text{maximize } \sum_{(q_t, p_t)_{t \in T}} \pi_t p_t \tag{24}$$

subject to

$$\forall (t_0, t_1, \dots, t_k) \in \mathcal{P}, \quad p_{t_k} \leq v_{t_k} q_{t_k} - \sum_{r=1}^{k-1} q_{t_r} (v_{t_{r+1}} - v_{t_r}) - v_{t_1} q_{t_0} \tag{25}$$

$$\forall t \in T, \quad 0 \leq q_t \leq 1 \tag{26}$$

$$p_0 = 0. \tag{27}$$

We call this the *path formulation*, as opposed to the *edge formulation* (3)–(6), because whereas the edge formulation contained an incentive constraint for every edge

in the incentive graph, the path formulation contains a related constraint for every path in the incentive graph.²⁰ The path formulation is a relaxation of the edge formulation, but the following proposition shows that the two are equivalent for our purposes.

PROPOSITION 8. *The path and edge formulations of the primal have the same set of optima.*

Then the next proposition follows.

PROPOSITION 9 (Discrete revenue formula). *Given any incentive compatible allocation q , the revenue maximizing prices that implement that allocation are*

$$p_t = \max_{P=(t_0, \dots, t_{n_P}) \in \mathcal{P}_t} v_{t_{n_P}} q_{t_{n_P}} - \sum_{i=1}^{n_P-1} q_{t_i} (v_{t_{i+1}} - v_{t_i}) - v_{t_1} q_{t_0}. \quad (28)$$

A path P , solving the maximization (28) will always bind at the revenue maximizing prices implementing q .

Proposition 9 is a corollary of **Proposition 8** because it follows by setting each price p_t equal to the lowest upper bound determined by (25) in the path formulation of the primal. This result generalizes the revenue formula in **Proposition 1(ii)**. That formula does not contain a maximization because we know a priori that the downward adjacent constraints are binding. Similarly, when G^* is a tree, there is a unique $0-t$ path $P = t_0 \rightarrow \dots \rightarrow t_n = t$ in G^* and for any monotone feasible allocation (i.e., an allocation satisfying $(s, t) \in E' \Rightarrow q_s \leq q_t$),²¹ **Proposition 9** simplifies to

$$p_t = v_{t_n} q_{t_n} - \sum_{i=1}^{n-1} q_{t_i} (v_{t_{i+1}} - v_{t_i}) - v_{t_1} q_{t_0}.$$

We conclude this section by displaying the dual of (24)–(27), which we call the *path formulation* of the dual; it is an alternative to the edge formulation of the dual.

The path formulation of the dual will be useful for describing the credible implementation of the optimal mechanism in **Section 7**.

²⁰The constraints (25) can be interpreted in terms of path lengths: 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. Given any path $t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_k$ in G , the induced path length is the sum of its edge lengths $\sum_{j=1}^k v_{t_j} (q_{t_j} - q_{t_{j-1}})$. Regrouping terms, (25) can then be interpreted as saying that the price p_t is bounded above by the length of any path from 0 to t .

²¹In the case of the complete incentive graph, we did not need to specify that (10) was monotone over and above being feasible (i.e., incentive compatible), because with the complete incentive graph, every incentive compatible allocation is monotone.

DUAL PROBLEM 1 (Paths). The path formulation is

$$\underset{(\mu_t)_{t \in T}, (\lambda_P)_{P \in \mathcal{P}}}{\text{minimize}} \sum_{t \in T} \mu_t \quad (29)$$

subject to

$$\forall t \in T \setminus 0, \quad \sum_{P \in \mathcal{P}_t} \lambda_P = \pi_t \quad (30)$$

$$\forall t \in T, \quad v_t \pi_t - \sum_{s: (t,s) \in E} \sum_{\{P \in \mathcal{P}: (t,s) \in P\}} \lambda_P (v_s - v_t) \leq \mu_t \quad (31)$$

$$\forall P \in \mathcal{P}, \quad \lambda_P \geq 0 \quad (32)$$

$$\forall t \in T, \quad \mu_t \geq 0. \quad (33)$$

Dual variables indexed by edges $\lambda(s, t)$ in (11)–(15) have been replaced by variables λ_P indexed by paths in (29)–(33). Any feasible solution to the path formulation $(\lambda_P : P \in \mathcal{P})$ of the dual induces a feasible solution to the edge formulation (11)–(15), for the same vector $(\mu_t : t \in T)$, via the relation²²

$$\lambda(s, t) = \sum_{P \ni (s,t)} \lambda_P \quad \forall (s, t) \in E, \quad (34)$$

where the summation on the right-hand side of (34) is over paths in \mathcal{P} containing the edge (s, t) . Say that $\lambda = (\lambda_P : P \in \mathcal{P})$ is *good* if the corresponding vector $(\lambda(s, t) : (s, t) \in E)$ defined by (34) is good (see (20)). Say that λ is to *proportional* if

$$\forall P = (t_0, \dots, t_n) \in \mathcal{P}, \quad \lambda_P = \left[\prod_{i=1}^n \frac{\sum_{P' \ni (t_{i-1}, t_i)} \lambda_{P'}}{\sum_{P' \ni t_i} \lambda_{P'}} \right] \pi_{t_n}. \quad (35)$$

Under the network flow interpretation mentioned above,²³ λ_P can be interpreted as a *flow on paths*, $\lambda(s, t)$ can be interpreted as a *flow on edges*, and the property of proportionality (35) can be interpreted as factoring the flow on paths into a (normalized) product of flows on edges. The following result will be useful for our credibility result in Section 7.

PROPOSITION 10. *Let $\mu = (\mu_t : t \in T)$ and $\lambda = (\lambda(s, t) : (s, t) \in E)$. Suppose that (λ, μ) is optimal in the edge formulation of the dual and that λ is good. (An optimal solution that is also good exists by Proposition 3(i).) Then there exists a unique proportional vector*

²²To see why $(\lambda(s, t) : (s, t) \in E)$ defined by (34) is feasible in the edge formulation of the dual (given the fixed vector μ), note first that (34) and the fact that $(\mu, (\lambda_P : P \in \mathcal{P}))$ satisfy (31) immediately imply that $(\mu, \lambda(s, t) : (s, t) \in E)$ satisfies (13). Moreover, (30) and (34) imply

$$\sum_{s: (s,t) \in E} \lambda(s, t) - \sum_{s: (t,s) \in E} \lambda(t, s) = \sum_{s: (s,t) \in E} \sum_{P \ni (s,t)} \lambda_P - \sum_{s: (t,s) \in E} \sum_{P \ni (t,s)} \lambda_P = \sum_{P \ni t} \lambda_P - \sum_{P \ni t: P \notin \mathcal{P}_t} \lambda_P = \sum_{P \in \mathcal{P}_t} \lambda_P = \pi_t.$$

So $(\lambda(s, t) : (s, t) \in E)$ satisfies (12).

²³See Ahuja et al. (1993) for a comprehensive treatment of network flows.

$\tilde{\lambda} = (\lambda_P : P \in \mathcal{P})$ related to λ by (34). Moreover, $(\tilde{\lambda}, \mu)$ is optimal in the path formulation of the dual.

In light of Proposition 10, when discussing a good optimal solution $(\mu, (\lambda(s, t) : (s, t) \in E))$, we freely interchange $(\lambda(s, t) : (s, t) \in E)$ with $(\lambda_P : P \in \mathcal{P})$, where the latter is the unique proportional vector corresponding to the former via (34).

7. A CREDIBILITY RESULT

7.1 Motivation

This section considers how the optimal mechanism might be implemented. Consider first a naive approach. The seller requests that the buyer present a cheap talk message about his type along with evidence. The seller claims, without commitment, that if the buyer submits cheap talk claim t and evidence S_t , the seller will offer to sell the object with probability q_t for an expected price of p_t .²⁴ (The precise details of whether payment is made prior to the randomization or conditional on receipt of the object can be spelled out in various ways.) A problem with this approach is that if the buyer always reported truthfully, the seller would know the buyer's value and so would often prefer to charge a higher price, and, moreover, to sell the object with probability 1 regardless of the allocation q_t .

There are many alternatives to the naive approach, whose plausibility would depend on, among other things, the seller's commitment power. Here we pursue a particularly simple and natural approach based on minimal commitment power. We introduce a simple bargaining protocol with back-and-forth communication in which the seller can commit to honoring a take-it-or-leave-it offer once she has made it. The seller cannot commit to making any specific offer in the future contingent on various events, such as presentation of evidence. She can only commit to honoring an offer for sale with probability 1 and not to any form of randomization. Our bargaining protocol implements the optimal mechanism as a sequential equilibrium.

In effect, what we have done is to reduce the entire commitment problem—the commitment to a complex randomized mechanism requiring price discrimination on the basis of evidence—to the commitment to a take-it-or-leave-it offer. Any foundation for this latter simpler commitment has a potential to serve as a foundation for commitment to our more complex mechanism as well.²⁵ Our bargaining model applies best in settings where the seller deals directly with the buyer and has the discretion to lower prices. This is the case in many firms that employ sales agents who are permitted, at

²⁴Assume that if the evidence submitted is not the maximum evidence of the type t in the cheap talk claim, the buyer neither receives the object nor makes a payment.

²⁵For example, Ausubel and Deneckere (1989) provide a foundation for seller commitment to take-it-or-leave-it offers in terms of the equilibrium of an infinite horizon bargaining game. However, further analysis would be required to integrate their approach to founding commitment to take-it-or-leave-it offers on the equilibrium of a repeated interaction with our quite different approach to founding commitment to a complex randomized price discrimination mechanisms on commitment to take-it-or-leave-it offers. This is beyond the scope of the current paper. For an analysis related to Ausubel and Deneckere (1989), see also Gerardi et al. (2014).

their discretion, to offer discounts off the list price.²⁶ It may not be appropriate for settings where interaction is more anonymous and distant.

One virtue of our bargaining protocol is that it has a closer resemblance to real-world negotiation than the direct mechanism. Specifically, a random allocation is not generated by commitment to randomize in response to certain messages and evidence, but rather arises naturally out of the randomized equilibrium strategies that can lead negotiations to break down. Another virtue of our protocol is that it provides an appealing interpretation of the dual of the seller's revenue maximization problem: At an optimal solution, the dual variables can be interpreted as encoding the buyer's reporting strategy in the bargaining protocol.

7.2 The bargaining game

We now describe our bargaining game.

- 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) or
 - (b) 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.

To summarize, the buyer opens with a cheap talk claim \hat{t} , which implicitly includes an offer to pay $v_{\hat{t}}$, the value of type \hat{t} . The seller responds either by asking for another offer or by 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. There is no discounting, so that we think of this as a fast interaction.

²⁶For example, a recent survey by the consulting firm Oliver Wyman reports that discretion is allowed in approximately 50% of the unsecured loans and more than 60% of the secured loans offered by major European banks (Efma and Oliver Wyman 2012).

7.3 Credible implementation

We now state our main result concerning the bargaining game.

PROPOSITION 11 (Credibility result). *There is a sequential equilibrium of the dynamic bargaining protocol that implements the optimal mechanism.*²⁷

We call a sequential equilibrium implementing the optimal mechanism a *credible implementation*. The remainder of this section describes the equilibrium of **Proposition 11** and highlights the main qualitative features of the equilibrium.

7.3.1 Equilibrium strategies We now describe the strategies in the equilibrium of **Proposition 11**. The description omits details of off-path play, which are to be found in the **Appendix**.

It is of interest that the seller's strategy depends on an optimal solution to the seller's revenue maximization problem (the primal (3)–(6)), and the buyer's strategy depends on an optimal solution to the dual.

The buyer's strategy. We may assume that after her type t is realized, the buyer performs a private preliminary randomization that guides her behavior throughout the game. Specifically, the buyer randomizes over paths leading to t (i.e., those in \mathcal{P}_t), selecting path P with probability

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

where λ is an optimal proportional and good solution of the path formulation of the dual. (For the meanings of “proportionality” and “good,” see (35) and the surrounding discussion in **Section 6**.) Equation (30) implies that these probabilities sum to 1. When

$$P : 0 = t_0^* \rightarrow \dots \rightarrow t_k^* \rightarrow \dots \rightarrow t_n^* = t \quad (36)$$

is the outcome of the preliminary randomization, the type t buyer reports along path (36), so that her first cheap talk claim is t_0^* ($= 0$), if she is asked to make a k th cheap talk claim, she will report t_k^* , and so on. If evidence is requested following cheap talk report t_k^* , she presents evidence $S_{t_k^*}$.²⁸ She drops out if asked for more cheap talk after t_n^* ($= t$). The buyer accepts any take-it-or-leave-it offer up to her value v_t , and rejects higher offers.

The seller's strategy. Let $(q_t : t \in T)$ be an allocation in an optimal solution to the seller's revenue maximization problem. On the equilibrium path, when the buyer has made the sequence of reports (t_0, \dots, t_k) , the seller requests evidence with probability $(q_{t_k} - q_{t_{k-1}})/(1 - q_{t_{k-1}})$ and requests another cheap talk report with the remaining probability.²⁹ (If $k = 0$, the buyer requests another cheap talk report with probability 1.) If the

²⁷If there is more than one optimal mechanism, then for every optimal mechanism, there is a sequential equilibrium of the bargaining protocol that implements it.

²⁸That $P \in \mathcal{P}_t$ implies that it is feasible for the type t buyer to present evidence $S_{t_k^*}$.

²⁹On the equilibrium path, $q_{t_k} \geq q_{t_{k-1}}$ and $1 > q_{t_{k-1}}$.

buyer presents S_{t_k} in response to the seller's evidence request, the seller makes a take-it-or-leave-it offer at price v_{t_k} . If the seller presents any other evidence S , the seller makes an offer at the price equal to the maximum value of any type that has access to S .

7.3.2 Qualitative properties of equilibrium Call any strategy profile in which the players' strategies satisfy the properties specified in [Section 7.3.1 canonical](#).³⁰ A credible implementation is canonical if the strategies it employs are canonical. The proof of [Proposition 11](#) establishes the existence of a canonical credible implementation. The following propositions provide key qualitative features of canonical equilibria.

PROPOSITION 12. *Let (t_0, t_1, \dots, t_k) be any sequence of cheap talk reports that occurs with positive probability in a canonical credible implementation. Then*

$$v_{t_0} < v_{t_1} < \dots < v_{t_n}$$

$$S_{t_0} \subseteq S_{t_2} \subseteq \dots \subseteq S_{t_n}.$$

In words, both the buyer's claimed value and the amount of claimed evidence increase as bargaining proceeds, so the buyer makes a sequence of concessions and brings up additional evidence without retracting her claim to possess previously mentioned evidence. Moreover, when the true type is t ,

$$S_{t_i} \subseteq \sigma(t) \quad \text{for } i = 0, 1, \dots, k.$$

The buyer never claims to possess evidence that she does not possess.

For a proof, see the [Appendix](#).

PROPOSITION 13. *For each natural number n , there exists an environment—a set of types, values, and probabilities, and an incentive graph—such that in any canonical credible implementation, there is a positive probability that the players communicate for at least n rounds.*

[Proposition 13](#) relates our analysis to the literature on long cheap talk ([Aumann and Hart 2003](#), [Forges and Koessler 2008](#)). The proof constructs a sequence of environments, one for each natural number n , such that in the n th environment, any canonical credible implementation involves approximately n rounds of communication. [Proposition 13](#) applies only to *canonical* credible implementations, but we believe that it is difficult to see how one could possibly bound communication in the sequence of environments constructed in the proof in some other noncanonical credible implementation. The key property generating the long communication in these environments is that the optimal mechanism contains a long chain of types linked by binding incentive constraints along which the allocation probability *strictly* increases. The length of these chains approaches infinity as n becomes large. If the reader is interested in a concrete illustration of how a credible implementation works, we recommend the proof of [Proposition 13](#) as an illustration.

³⁰For a strategy profile to be canonical, it is not required that it satisfies the off-path properties specified in the [Appendix](#).

7.3.3 *Proof outline* The proof of [Proposition 11](#) must establish four claims:

1. The equilibrium strategies described in [Section 7.3.1](#)—if followed—will implement the optimal mechanism. That is, if $(q_t, p_t : t \in T)$ is an optimal mechanism, these strategies will lead buyer type t to receive the object with probability q_t and to make an expected payment of p_t .
2. The buyer's strategy is a best reply to the seller's strategy.
3. The seller's strategy is a best reply to the buyer's strategy.
4. The off-equilibrium path requirements of sequential equilibrium are met.

The [Appendix](#) establishes the first two claims and the last. Next we informally describe the argument for the seller's best reply, which turns out to be the most interesting argument. That claim too is formally established in the [Appendix](#).

7.3.4 *Optimal stopping interpretation of the seller's equilibrium problem* Given the buyer's equilibrium strategy, the seller's problem in the dynamic bargaining protocol becomes an optimal stopping problem. We argue in the [Appendix](#) that the seller can restrict attention to *stopping strategies*, which are strategies such that

- if the buyer made reports (t_0, \dots, t_k) prior to the seller's request for evidence and presented evidence S_{t_k} , the seller offers price v_{t_k} .³¹

Because on the equilibrium path, the buyer always presents evidence S_{t_k} if asked for evidence after type reports (t_0, \dots, t_k) , the seller's problem becomes an optimal stopping problem: Following history (t_0, \dots, t_k) , if the seller stops, requesting evidence, he receives a payoff of v_{t_k} . If the seller continues, requesting more cheap talk, the buyer may drop out, yielding the seller a zero payoff. If the buyer does not drop out, presenting another cheap talk claim t_{k+1} , the seller can secure a payoff of $v_{t_{k+1}}$ by stopping in the following round. As mentioned in [Proposition 12](#), $v_{t_k} < v_{t_{k+1}}$, so, by waiting, the seller may attain a higher reward.

The stochastic process facing the seller in his optimal stopping problem is endogenous because it depends not only on the distribution of types, but also on the buyer's reporting strategy. If the buyer reported (t_1, \dots, t_k) and then the seller continued, requesting more cheap talk, the probability that the seller assigns to the event that the buyer will not drop out but instead will report $t_{k+1} = t$ is

$$\frac{\lambda(t_k, t)}{\sum_{s: (s, t_k) \in E} \lambda(s, t_k)}, \quad (37)$$

where the variables λ come from an optimal solution to the dual of the seller's revenue maximization problem (11)–(15). Equation (12) implies that the remaining probability, which is the probability that the buyer drops out, is

$$\frac{\pi_{t_k}}{\sum_{s: (s, t_k) \in E} \lambda(s, t_k)}. \quad (38)$$

³¹Ultimately, in constructing the equilibrium, we will modify the seller's optimal stopping strategy off the equilibrium path to insure buyer optimization.

The probabilities (37) and (38) depend only on t_k , and not the entire history (t_0, \dots, t_k) . So the stochastic process facing the seller is a Markov chain with $n + 1$ states, where n is the number of types (including the zero type). The first n states correspond to the types and state t is interpreted to occur when the buyer claims to be of type t . The last state d —the *drop-out state*—corresponds to the event where the buyer drops out. For two types, r and t , the transition probability between r and t is given by (37) when r is substituted for t_k .³² The transition probability from type t to d is (38) when t is substituted for t_k . State d is the unique absorbing state. The process begins in state 0. It follows from the fact that on the equilibrium path, reported values are strictly increasing (see Proposition 12), that no state other than d is accessible from itself, and so the process reaches the absorbing state d in at most $n + 1$ steps.

Say that a sequence of type reports $P = (t_0, \dots, t_k)$ is *buyer on-path* if³³

$$\begin{aligned} &\text{in the preliminary randomization, some type } t \text{ selects a path } P' \\ &\text{with initial subsequence } P \text{ with positive probability.} \end{aligned} \tag{39}$$

The use of the term “buyer” in buyer on-path is motivated by the fact that whether a sequence of reports occurs on the equilibrium path depends not only on the buyer’s strategy, but also on the seller’s strategy. A *buyer on-path* sequence of reports P is one such that there exists some seller strategy that, when coupled with the buyer’s equilibrium strategy, leads P to occur with positive probability.

The seller’s equilibrium strategy is such that on a sequence (t_0, \dots, t_k) , the seller stops with positive probability if $q_{t_k} > 0$ and continues with positive probability if $q_{t_k} < 1$. The key to establishing that the seller’s equilibrium strategy is a best reply to the buyer’s equilibrium strategy is, therefore, the following lemma, which is proven via the complementary slackness conditions relating optimal solutions of the seller’s revenue maximization problem to its dual.

LEMMA 2. *Let $(q_t : t \in T)$ be the allocation in some revenue maximizing mechanism and let $P = (t_0, \dots, t_k)$ be a buyer on-path report sequence. If the buyer uses her equilibrium strategy and the seller is restricted to stopping strategies, the following statements hold:*

- (i) *If $q_{t_k} > 0$, then, conditional on P , the seller is weakly better off stopping immediately than continuing for one more step and then stopping.*
- (ii) *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.*

³²The transition probability is 0 if $(r, t) \notin E$.

³³Interestingly, (39) holds exactly if, in the preliminary randomization, type t_k selects path P with positive probability.

PROOF. Using the probabilities (37), the seller's preference between stopping now and stopping in one step is determined according to the resolution of the inequality

$$\underbrace{v_{t_k}}_{\text{stopping now}} \underset{\geq}{\leq} \underbrace{\sum_{t:(t_k,t) \in E} v_t \frac{\lambda(t_k,t)}{\sum_{s:(s,t_k) \in E} \lambda(s,t_k)}}_{\text{stopping in one step}}, \quad (40)$$

where the variables λ come from a good optimal dual solution. Using (13) to substitute $\pi_{t_k} - \sum_{s:(t_k,s) \in E} \lambda(t_k,s)$ for the denominator in the fraction on the right-hand side of (40) and rearranging terms, the resolution of (40) is equivalent to the resolution of

$$\psi(t_k) \underset{\geq}{\leq} 0.$$

So stopping now is preferable to stopping in exactly one step exactly if the virtual value of type t_k is nonnegative. Using the same logic as in the proof of Proposition 2, if $q_t < 1$, then by complementary slackness, $\mu_t = 0$, implying via (14) and (19) that $\underset{\geq}{\leq}$ becomes \leq , establishing part (ii) of the lemma. Alternatively, if $q_t > 0$, then by complementary slackness, (13) holds with equality, which implies via (15) and (19) that $\underset{\leq}{\geq}$ becomes \geq , establishing part (i) of the lemma. \square

To complete the proof, we now use Lemma 2 to argue by backward induction that a stopping strategy that coincides with the seller's equilibrium strategy following all histories containing a buyer on-path sequence of type reports is optimal among stopping strategies. The remaining details of the proof are given in the Appendix.

7.4 The deterministic case

We now discuss the credible implementation in the deterministic case. Section 5.2 above characterized conditions under which the optimal mechanism is deterministic.

First, consider the standard monopoly problem studied in Section 3. Then Proposition 6 implies an optimal mechanism that is deterministic. Indeed, the optimal mechanism is simply a posted price. We could describe the canonical credible equilibrium, but in this case, there is a simpler credible implementation. Let p^* be the optimal posted price. Then $p^* = v_{t^*}$ for some type t^* . There is an equilibrium in which all types t^* with $v_t < p^*$ drop out immediately. All remaining types claim to be of type t^* . In the latter case, the seller immediately requests evidence. As with the complete graph, all types have the same evidence; this common evidence is presented. The seller makes a take-it-or-leave-it offer of p^* that is accepted.³⁴ This sort of equilibrium can be generalized to the case where G^* is a tree.

PROPOSITION 14. *If G^* is a tree—so that by Proposition 4, there is an optimal mechanism that is deterministic—then there is a pure strategy credible implementation with only one round of cheap talk communication.*

³⁴The equilibrium can be supported off the equilibrium path in several ways.

Let (q^*, p^*) be an optimal deterministic mechanism. In the equilibrium described by [Proposition 14](#), all types with $q_t^* = 0$ drop out immediately. All types with $q_t^* = 1$ claim to be the first type r on the unique $0 - t$ path in G^* with $q_r^* = 1$. Notice that here the argument already relies not only on the fact that the optimal mechanism is deterministic, but also on the assumption that G^* is a tree. In response, the seller requests evidence, and if the buyer presents S_r , the seller offers price p_s^* . If the seller had an incentive to offer a higher or lower price, then similarly (q^*, p^*) would not be optimal. For example, if the seller had an incentive to offer a higher price, then the seller could do better with a mechanism (q, p) with $q_r = p_r = 0$. If any buyer were to deviate from the prescribed behavior, then after the buyer presents evidence S , the seller offers the price $\max\{v_t : S \subseteq \sigma(t)\}$.

7.5 Relation to [Glazer and Rubinstein \(2004, 2006\)](#)

This section discusses the connection—centering on our credibility result—between the monopolist’s optimal price discrimination problem and the literature on optimal persuasion. [Glazer and Rubinstein \(2004, 2006\)](#) studied a persuasion game with a binary decision in which a privately informed speaker attempts to persuade a listener to accept a request by presenting evidence.³⁵ They showed that the optimal persuasion mechanism—optimal for the listener—when the listener moves first—committing to a response rule—is an equilibrium of the game when the speaker moves first and the listener moves second. Their credibility results are closely related to our [Proposition 11](#). Indeed, [Proposition 11](#) may be viewed as a “partial” generalization of their results. This is surprising as, at first glance, the monopolist’s problem appears quite different from the persuasion problem studied by [Glazer and Rubinstein](#).

We now briefly describe a few more details of the model of [Glazer and Rubinstein \(2006\)](#).³⁶ The speaker knows the state in a set $x \in X$ and the listener does not. In state x , the speaker has a set of hard messages $\sigma(x)$. The listener would like to accept the speaker’s request if the state belongs to A (the *accept states*) and reject on the complementary set R (the *reject states*). If uncertain, the listener aims to minimize the probability of error (i.e., accepting in R or rejecting in A). [Glazer and Rubinstein \(2006\)](#) assumed that the speaker can only present one message in $\sigma(x)$, whereas we assume that the buyer can present any subset of $\sigma(t)$. However, one can encode the possibility where an agent can present any subset of his messages as a special case within a model where the speaker can present only one message. Specifically, suppose that there is some subset B of the set of messages, which we can call *base messages*, and for every subset S of B , there exists a message m_S such that $m_S \in \sigma(t)$ exactly if $S \subseteq \sigma(t)$. Call this case the case of *unrestricted evidence*. Henceforth, we restrict attention to the special case of the

³⁵Strictly speaking, in [Glazer and Rubinstein \(2004\)](#), the speaker makes a cheap talk claim, and the listener partially verifies the speaker’s claim. This can equivalently be translated into a model where the speaker presents evidence.

³⁶[Sher \(2014\)](#) discusses the relation between the models of [Glazer and Rubinstein \(2004, 2006\)](#), and presents a more general model embedding both.

Glazer–Rubinstein model with unrestricted evidence. Our result is a *partial* generalization of the Glazer–Rubinstein result insofar as it generalizes their result in this special case.

A persuasion rule maps every message in the range of σ to a probability of acceptance. The listener commits to a persuasion rule. Knowing this rule, the speaker presents a message and the rule determines the outcome. The *optimal persuasion rule* is thus the persuasion rule that minimizes the probability of listener error. Glazer and Rubinstein (2006) showed that deterministic persuasion rules are optimal.³⁷ In the *game without commitment*, the speaker presents a message and the listener responds by either accepting or rejecting the speaker's request, without having committed to a response. The Glazer and Rubinstein (2006) credibility result is that there exists a sequential equilibrium of the game without commitment that implements the optimal persuasion rule. We refer to this equilibrium as the *credible implementation* of the optimal rule.

With unrestricted evidence, the Glazer–Rubinstein model may be viewed as a special case of our model. Specifically, suppose that there are two values, a low value v_L and a high value v_H with $0 < v_L < v_H$. All types other than the zero type value the object at either v_L or v_H . We refer to types with the low value as *low types* and types with a high value as *high types*. In general, there are many low types and many high types that differ according to the evidence that they possess. We refer to this special case as the *binary values model*. The correspondence between the binary values model and the persuasion model is as follows. The states in the persuasion model correspond to the nonzero types in the binary values model, the accept states corresponding to the low types, and the reject states corresponding to the high types. The probability of a state is the probability of the corresponding type. The buyer corresponds to the speaker and the seller corresponds to the listener. The buyer is attempting to persuade the seller to offer a low price. The seller accepts the request by offering a low price and rejects the request by offering a high price. Glazer and Rubinstein (2006) remark that their model is essentially unchanged if instead of choosing the persuasion rule to minimize the expected error, the listener chooses the persuasion rule to minimize the expected cost of error, where the cost of error depends on the state. The cost of error at an accept states is v_L , because if the seller offers a high price when the buyer is of a low type, he loses a potential payoff of v_L . The cost of error at a reject state is $v_H - v_L$, because if the seller offers a low price when the seller is of a high type, he loses the additional increment $v_H - v_L$ that he could

³⁷As explored by Sher (2014), the optimal mechanism is not necessarily a persuasion rule in the above sense. In general, the listener could do better with one round of back-and-forth cheap talk communication between the speaker and listener (where the listener randomizes her communication) followed by evidence presentation. The listener's communication can be interpreted as a request for specific evidence. However, it follows from the analysis of Sher (2014) that in the case of unrestricted evidence, persuasion rules are indeed optimal, and, hence, not only the optimal persuasion rule, but also the optimal mechanism, is deterministic. Sher (2014) refers to the persuasion rules of Glazer and Rubinstein (2006) as *static* persuasion rules, and refers to the mechanisms with back-and-forth cheap talk communication preceding evidence presentation as *dynamic* persuasion rules. The persuasion rules studied by Glazer and Rubinstein (2004) are dynamic in this sense. Related analyses occur in Forges and Koessler (2005), Bull and Watson (2007), and Deneckere and Severinov (2008).

have received if he had offered a high price. When $v_H = 2v_L$, minimizing the expected cost of error coincides with minimizing the error probability.

In the binary values model, the revenue maximizing mechanism is deterministic and corresponds to the optimal persuasion rule in the persuasion model, where a type is accepted in the persuasion model exactly if he receives a low price of v_L in the revenue maximizing mechanism. All other types are rejected by the optimal persuasion rule; in the revenue maximization problem, these rejected types pay a high price of v_H if they are high types and are not served if they are low types (since they are unwilling to pay the high price).

Although the optimal mechanism is deterministic, the credible implementation of the optimal mechanism may involve randomization. In the binary values model, if an optimal solution is good, then there are two types of paths P for which $\lambda_P > 0$: either P is of the form (i) $0 \rightarrow r$ or of the form (ii) $0 \rightarrow s \rightarrow t$, where s is a low type and t is a high type. This implies, in particular, that if s is a low type and $P = (0, s)$, then $\lambda_P = \pi_s$. So in the credible implementation of the revenue maximizing mechanism, all low types initially tell the truth (after claiming to be the zero type) and high types initially randomize over lies, pretending to be one of the low types they can mimic, before admitting to being high types if asked for a second cheap talk report (unless the high type receives a zero payoff, in which case she might admit to being a high type without claiming to be a low type first). Recall that when a type claims to be low, the seller may not ask for a second cheap talk report for fear that a truly low type would drop out.

The credible implementation of the optimal persuasion rule is similar, and can be derived from the revenue maximization problem and its dual. For expositional ease, assume that the optimal good dual solution is such that for all high types t , the path $P' = (0, t)$ is such that $\lambda_{P'} = 0$. Since we are considering a good dual solution, for any path $P = (0, s, t)$, where s is a low type and t is a high type, $\lambda_P = \lambda(s, t)$, where the latter comes from the corresponding optimal solution of the edge formulation of the dual. Then at the credible implementation of the optimal persuasion rule, at each accept state t , the speaker presents evidence S_t . At each reject state s , the speaker presents evidence S_t with probability $\lambda(t, s)/\pi_s$ if $(t, s) \in E$ and presents any evidence not of this form with probability 0. (Since λ is good, the speaker presents evidence S_t with positive probability only for accept states t .) The listener accepts the speaker's request if the evidence is S_t for some t such that $p_t = v_L$ at the revenue maximizing mechanism and rejects otherwise.

A few points are worth emphasizing. First, here we have identified an alternative sufficient condition for the optimal mechanism to be deterministic: that the environment be an instance of the binary values model. Unlike the sufficient conditions of [Section 5.2](#), this new sufficient condition is completely independent of the incentive graph, and depends only on valuations. Second, in contrast to [Proposition 13](#), which shows that, in general, outside the binary values model, the credible implementation of the seller's revenue maximizing mechanism may require many rounds of communication, the credible implementation of the optimal persuasion rule requires only one round of communication. Why is this? Our analysis sheds light on this question. In our model, the maximal number of rounds of cheap talk communication in the canonical credible implementation corresponds roughly to the length of the longest chain of binding

incentive constraints along which the allocation strictly increases. **Proposition 3**(i) implies that in the binary values model, the only relevant incentive constraints correspond to edges of the form $(0, \ell)$ and (ℓ, h) , where ℓ is a low type and h is a high type. So we must only consider chains with two edges. Since, in the canonical credible implementation, (i) the initial claim to be the zero type is for simplicity of description of the mechanism and not strictly necessary, and (ii) all high types have the same value, so if a high type's claim to be a low type is rejected, it is not necessary for the high type to specify which high type she is, it is indeed possible to reduce communication from three rounds to one.

APPENDIX

A.1 Preliminaries

This section collects results and notation useful for proofs of the main results. For any path $P: t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_n$, define the *path length* of P to be n the number of edges (t_{i-1}, t_i) in P . (This differs from the notion of path length discussed in [footnote 20](#).) Also for any path P in G , write $P = (t_0^P, \dots, t_{n_P}^P)$; this notation is sometimes convenient.

Let $G = (V, E)$ be the incentive graph. Let E' be the set of good edges in E and let $G' = (V, E')$. The *monotone primal* is the program that performs the maximization in (3) subject to (5) and (6), and

$$\forall (s, t) \in E', \quad v_t q_t - p_t \geq v_t q_s - p_s. \quad (41)$$

The *original primal* is (3)–(6).

Let $z = (q, p)$ be an optimal solution in the monotone primal and let $B_z = \{(s, t) \in E' : v_t q_t - p_t = v_t q_s - p_s\}$. So, B_z is the set of binding constraints at z . Define $G'_z := (V, B_z)$ and let $\tilde{G}_z^* = (V, B_z^*)$ be the transitive reduction of G'_z .

LEMMA 3. *Let $z = (q, p)$ be an optimum of the monotone primal and let $P: 0 = t_0 \rightarrow \dots \rightarrow t_n$ be a path in G'_z . Then*

$$p_{t_n} = \sum_{i=1}^n v_{t_i} (q_{t_i} - q_{t_{i-1}}). \quad (42)$$

Moreover, for each $t \in V \setminus 0$, there exists a $0 - t$ path in G'_z .

PROOF. We must have $q_0 = 0$. Otherwise redefine $q_0 := 0$ and add the same sufficiently small increment ε to the payment of all types t other than 0. Feasibility of the original mechanism implies feasibility of the new mechanism and revenue increases, contradicting optimality of the original mechanism. Now (42) follows from an easy induction on n . Finally, let S be the set of nonzero types t such that there is no $0 - t$ path in G'_z . If $S \neq \emptyset$, increase the payment of all types in S by sufficiently small ε , arriving at a new feasible mechanism with higher revenue, a contradiction. \square

LEMMA 4. *The monotone primal and original primal have the same set of optima.*

PROOF. Since the monotone primal is a relaxation of the original primal, it is sufficient to show that any optimum of the monotone primal is feasible in the original primal. Let $z = (q, p)$ be optimal in the monotone primal and choose $(t, s) \in E$. By Lemma 3, there exists a path $P: 0 = t_0 \rightarrow \dots \rightarrow t_n = t$ in G'_z . Let j be the largest index i such that $v_{t_i} < v_s$; j exists because $v_{t_0} = v_0 < v_s$.³⁸ Since E is weakly transitive, $(t_j, s) \in E'$. So

$$\begin{aligned} v_s q_s - p_s &\geq v_s q_{t_j} - p_{t_j} = v_s q_{t_j} - \sum_{i=1}^n v_{t_i} (q_{t_i} - q_{t_{i-1}}) \\ &\geq v_s q_{t_j} - \sum_{i=1}^n v_{t_i} (q_{t_i} - q_{t_{i-1}}) + \sum_{i=j+1}^n (v_s - v_{t_i}) (q_{t_i} - q_{t_{i-1}}) \\ &= v_s q_t - p_t, \end{aligned}$$

where the first inequality follows from feasibility of z in the monotone primal, the first equality follows from Lemma 3, and the second inequality follows from the fact that $v_s < v_{t_i}$ for $i = j + 1, \dots, n$ by choice of j and Proposition 3(ii) applied to the monotone primal. \square

LEMMA 5. *If G^* is a tree, then for any optimum $z = (q, p)$ of the monotone primal, $\bar{G}_z^* = G^*$.*

PROOF. If $E^* \subseteq B_z$, then for all $(s, t) \in E^*$, \bar{G}_z^* contains an $s - t$ path. But (s, t) is the unique $s - t$ path in G' and $E' \supseteq \bar{B}_z^*$. So $(s, t) \in \bar{B}_z^*$. So $E^* \subseteq \bar{B}_z^*$. But since for every $(s, t) \in B_z$, there is an $s - t$ path in G^* , $E^* = \bar{B}_z^*$. So assuming for contradiction that $\bar{G}_z^* \neq G^*$, it follows that $E^* \not\subseteq B_z$. So G^* contains a path $P: 0 = t_0 \rightarrow \dots \rightarrow t_n$ such that $(t_{n-1}, t_n) \notin B_z$. Let P be the shortest path with this property. It follows that

$$v_{t_n} q_{t_n} - p_{t_n} > v_{t_n} q_{t_{n-1}} - p_{t_{n-1}}. \quad (43)$$

By Lemma 3, there exists a $0 - t_n$ path $Q: 0 = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_m = t_n$ in G'_z . Since G^* is a tree, P is the unique $0 - t_n$ path in G^* . Since G^* is the transitive reduction of G' and $(s_{m-1}, t_n) = (s_{m-1}, s_m) \in B_z \subseteq E'$, there exists $i \in \{0, \dots, m-2\}$ such that $t_i = s_{m-1}$ ³⁹ and

$$v_{t_n} q_{t_n} - p_{t_n} = v_{t_i} q_{t_i} - p_{t_i}. \quad (44)$$

So

$$p_{t_{n-1}} - p_{t_i} > v_{t_n} (q_{t_{n-1}} - q_{t_i}) \geq v_{t_{n-1}} (q_{t_{n-1}} - q_{t_i}), \quad (45)$$

where the strict inequality follows by subtracting (44) from (43), and the weak inequality follows from the fact that edge (t_{n-1}, t_n) is in E' , and so is good, and that by choice of P and Proposition 3(ii) (applied to when the incentive graph is G'), $q_j \geq q_{j-1}$ for

³⁸Condition (2) implies $s \neq 0$.

³⁹We have $i \neq m - 1$ because $(t_{n-1}, t_n) \notin B_z$. This also implies that $n \geq 2$.

$j = i, \dots, n - 1$. Conclude by observing that (45) contradicts the incentive constraint $(t_i, t_{n-1}) \in E'$.⁴⁰ \square

A.2 Proofs

PROOF OF PROPOSITION 1 (iv). In view of Proposition 2, it is sufficient to show that in the standard monopoly problem, $\widehat{\psi}(t) = \psi(t) \forall t \in T$. To this end, define $\lambda(s, t) = \Pi_t$ if $t = s + 1$ and define $\lambda(s, t) = 0$ otherwise. If virtual values (19) are defined with respect to the above-defined λ 's, then virtual values and quasi-virtual values coincide. However, as virtual values are defined with respect to an optimal dual solution, it remains to show that the above λ 's are part of an optimal dual solution. Define $\mu_t = \max\{\widehat{\psi}(t), 0\}\pi_t$. Next define $q_t = 1$ if $\widehat{\psi}(t) > 0$ and $q_t = 0$ otherwise. Let s be the smallest type with $q_s = 1$, and define $p_t = v_s$ if $q_t = 1$ and $p_t = 0$ otherwise. Using the monotonicity of quasi-virtual values and the complementary slackness conditions, it is straightforward to verify that (q, p) and (λ, μ) are feasible, and, moreover, optimal, in the primal and dual, respectively. \square

PROOF OF PROPOSITION 3. (i) To establish that eliminating bad edges does not alter optimal revenue, it is sufficient to show that any optimal solution to the monotone primal is feasible in the original primal. Let $z = (q, p)$ be optimal in the monotone primal and choose $(t, s) \in E$. By Lemma 3, there exists a path $P: 0 = t_0 \rightarrow \dots \rightarrow t_n = t$ in G'_z . Let j be the largest index i such that $v_{t_i} < v_s$; j exists because $v_{t_0} < v_s$. Since E is transitive, $(t_j, s) \in E'$. So

$$\begin{aligned} v_s q_s - p_s &\geq v_s q_{t_j} - p_{t_j} = v_s q_{t_j} - \sum_{i=1}^j v_{t_i} (q_{t_i} - q_{t_{i-1}}) \\ &\geq v_s q_{t_j} - \sum_{i=1}^j v_{t_i} (q_{t_i} - q_{t_{i-1}}) + \sum_{i=j+1}^n (v_s - v_{t_i}) (q_{t_i} - q_{t_{i-1}}) \\ &= v_s q_t - p_t, \end{aligned} \tag{46}$$

where the first inequality follows from (41), both equalities follow from Lemma 3, and the second inequality follows from (i) $v_s < v_{t_i}$ for $i = j + 1, \dots, n$ by choice of j , and (ii) $q_{t_i} \geq q_{t_{i-1}}$ for $i = j + 1, \dots, n$ by Proposition 3(ii) applied to incentive graph G' . Condition (46) implies that z is feasible in the original primal.

(ii) Since (s, t) binds, $v_t(q_t - q_s) = p_t - p_s$. 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. \square

PROOF OF PROPOSITION 4. As G^* is a tree, denote the unique $0 - t$ path in G^* as $P^t: 0 = s_0^t \rightarrow s_1^t \rightarrow \dots \rightarrow s_{n_t}^t = t$.

⁴⁰The inclusion $(t_i, t_{n-1}) \in E'$ follows by weak transitivity of E' .

LEMMA 6. Let $\hat{q} = (\hat{q}_t : t \in V)$ be an optimum of

$$\text{maximize } \sum_{t \in V} q_t \hat{\psi}(t) \pi_t \tag{47}$$

subject to

$$q_s \leq q_t \quad \forall (s, t) \in E^* \tag{48}$$

$$0 \leq q_t \leq 1 \quad \forall t \in V \setminus 0 \tag{49}$$

$$q_0 = 0. \tag{50}$$

Let⁴¹

$$\hat{p}_t = \sum_{i=1}^{n_t} v_{s_i^t} (\hat{q}_{s_i^t} - \hat{q}_{s_{i-1}^t}) \quad \forall t \in T. \tag{51}$$

If G^* is a tree, then (\hat{q}, \hat{p}) is optimal in the primal (3)–(6). Moreover, $\sum_{t \in V} \hat{p}_t \pi_t = \sum_{t \in V} \hat{q}_t \hat{\psi}(t) \pi_t$.

PROOF. Any optimum of the monotone primal satisfies (50). Proposition 5 implies that any optimum of the monotone primal satisfies (48). Since G^* is a tree, Lemmas 3 and 5 imply that any optimum (q, p) of the monotone primal satisfies

$$p_t = \sum_{i=1}^{n_t} v_{s_i^t} (q_{s_i^t} - q_{s_{i-1}^t}) \quad \forall t \in T. \tag{52}$$

So adding the constraints (50) and (52) to the monotone primal will not alter the set of optima. Using the fact that the monotone primal only contains good edges, (52) and (48) imply all of the constraints in the monotone primal, so the incentive constraints may now be removed. Using (50) and (52),

$$\begin{aligned} \sum_{t \in V} p_t \pi_t &= \sum_{t \in V} \pi_t \sum_{i=1}^{n_t} v_{s_i^t} (q_{s_i^t} - q_{s_{i-1}^t}) \\ &= \sum_{t \in V} v_t q_t \pi_t - \sum_{t \in V} \sum_{i=1}^{n^t-1} (v_{s_{i+1}^t} - v_{s_i^t}) q_{s_i^t} \pi_t \\ &= \sum_{t \in V} v_t q_t \pi_t - \sum_{(t,s) \in E^*} (v_s - v_t) q_t \pi_s = \sum_{t \in V} q_t \hat{\psi}(t) \pi_t. \end{aligned} \tag{53}$$

Plugging (53) into the monotone primal objective, it follows that (\hat{q}, \hat{p}) is optimal in the monotone primal and, hence, by Lemma 4, is also in the original primal. \triangleleft

As the constraint matrix corresponding to (48) and (49) is totally unimodular (see Proposition 2.6 in Part III.1 of Nemhauser and Wolsey 1999), it follows that (47)–(50) has an integer optimal solution (see Proposition 2.3 in Part III.1 of Nemhauser and Wolsey

⁴¹When $t = 0$, (51) reduces to $\hat{p}_0 = 0$.

1999), and so by Lemma 6 that the original primal has an integer optimal solution, corresponding to a deterministic optimum. \square

PROOF OF PROPOSITION 5. Proposition 5 follows from Lemmas 4 and 5 and Proposition 3(ii). \square

PROOF OF LEMMA 1 AND PROPOSITION 6.

LEMMA 7. *If G is essentially segmented, then G_v^* is a tree for all nondegenerate valuation assignments v .*

PROOF. For each $t \in T$, let $\ell(t)$ be the maximum length of a $0 - t$ path in G_v^* . (For the definition of path length, see Appendix A.1.) If G_v^* is not a tree, then for some $t \in V_0$, G_v^* contains two distinct $0 - t$ paths P and Q . Choose such a t to minimize $\ell(t)$.

Since G_v^* is the transitive reduction of an acyclic graph, there exist vertices r, s such that $r \in P \setminus Q$ and $s \in Q \setminus P$. Note that $r, s \in \tilde{V}$. Since v is nondegenerate, assume without loss of generality, $v_r < v_s$. If there were an $r - s$ path in G , then G_v^* would contain two distinct $0 - s$ paths, P' and Q' , where $r \in P'$ and Q' coincides with Q up to s . But since $\ell(s) < \ell(t)$, this contradicts the choice of t . So there is no $r - s$ path in G . Since both r and s are on paths to the common vertex t , it now follows that G cannot be essentially segmented. \triangleleft

LEMMA 8. *If G is essentially segmented, then G induces a deterministic optimum for all assignments of values and probabilities.*

PROOF. When v is nondegenerate, the result follows from Lemma 7 and Proposition 4. The result is extended to degenerate v via the maximum theorem.⁴² \triangleleft

Define a *bad configuration* to be a triple of nonzero types (r, s, t) such that $r \neq s$, $(r, t) \in E$, $(s, t) \in E$, and $(r, s) \notin E$.

LEMMA 9. *If G contains a bad configuration, then there exists an assignment of values and probabilities such that no optimal mechanism is deterministic.*

PROOF. Let (r, s, t) be a bad configuration. Then suppose that $v_r = 1$, $v_s = 2$, and $v_t = 3$. Define

$$\delta = \min \left(\left\{ \frac{\pi_s}{\pi_t}, \frac{\pi_t}{\pi_r} \right\} \cup \left\{ \frac{\pi_r}{\pi_u} : u \in T \setminus \{0, r, s, t\} \right\} \right).$$

⁴²In view of Proposition 8, we may apply the maximum theorem to the path formulation of the primal rather than to the edge formulation. In particular, the path formulation makes it easy to establish that the feasible region is lower hemicontinuous in the valuation profile. Note that for the purpose of the proof, we may assume that valuations are all below a certain bound. Then for sufficiently large M , we may append the constants $-M \leq p_t \leq M \forall t \in T$ without changing the set of optima in the path formulation. This ensures that the feasible region is compact.

If probabilities are chosen so that δ is sufficiently large, then all optimal mechanisms are such that

$$q_s = q_t = 1, \quad p_s = p_t = 2, \tag{54}$$

where we have used the assumption that $(s, t) \in E$. If δ is sufficiently large, then since $(r, s) \notin E$ and $(r, t) \in E$, at the optimal mechanism, p_r will be chosen as large as possible subject to the (r, t) and $(0, r)$ incentive constraints and (54), so that $q_r = p_r = 1/2$. \triangleleft

LEMMA 10. *If G is not essentially segmented, then G contains a bad configuration.*

PROOF. If condition (ii) of essential segmentation fails, then there exist $i, j \in \{1, \dots, n\}$ with $i \neq j$ and $r \in V_i, s \in V_j, t \in V^\circ$ such that $(r, t) \in E$ and $(s, t) \in E$. Since V_i and V_j are distinct weakly connected components of \tilde{G} , $(r, s) \notin E$. So (r, s, t) is a bad configuration.

Next assume (i) fails. Then there exists $i \in \{1, \dots, n\}$ and $r, s \in V_i$ such that G_i contains no directed $r - s$ path but G_i contains an undirected $r - s$ path $P = (r = t_0, \dots, t_n = s)$. Because G , and, hence, also \tilde{G} , is weakly transitive, there exists $k \in \{1, \dots, n\}$ such that $(t_k, t_{k-1}) \in E$ and $(t_{k-1}, t_k) \notin E$. Since $t_{k-1} \in V_i$ and, hence, $t_{k-1} \notin V^\circ$, there exists t such that $(t_{k-1}, t) \in E$, and by weak transitivity of G , $(t_k, t) \in E$. It follows that (t_{k-1}, t_k, t) is a bad configuration. \triangleleft

LEMMA 11. *If G contains a bad configuration, then there exists a nondegenerate valuation assignment v such that G_v^* is not a tree.*

PROOF. Let (r, s, t) be a bad configuration. Set $v_r < v_s < v_t$ and assign the remaining values so as to be nondegenerate. The graph G^* must contain $0 - t$ paths P and Q such that $r \in P$ and $s \in Q$. Because G' contains neither an $r - s$ path (because $(r, s) \notin E$ and G' is weakly transitive) nor an $s - r$ path (because $v_r < v_s$), and every path in G^* is a path in G' , $r \notin Q$. So G^* is not a tree. \triangleleft

Lemmas 8–10 imply Proposition 6. Lemmas 7, 10, and 11 imply Lemma 1. \square

PROOF OF PROPOSITION 7. Consider the relaxed program that results from removing the monotonicity constraints (48) from (47)–(50). The allocation q^* defined in (23) is clearly an optimum of the relaxed program. But by single-crossing quasi-virtual values, q^* satisfies the monotonicity constraints, and so is optimal in (47)–(50). Moreover, the assumption of single-crossing quasi-virtual values implies that (q^*, p^*) from (23) satisfies (51). So by Lemma 6, (q^*, p^*) is optimal in the primal (3)–(6). \square

PROOF OF PROPOSITION 8. Let \mathcal{P}' be the set of paths originating in 0 in the monotone incentive graph G' . Note that \mathcal{P}' does not contain any cycles. (Terminology aside, G' is acyclic.) Define the *monotone path primal* to be the program that performs the maximization in (24) subject to

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

and (27). To avoid ambiguity, we rename the monotone primal defined in Appendix A.1 the *monotone edge primal*. In contrast, (3)–(6) and (24)–(27) will be referred to, respectively, as the *original edge primal* and the *original path primal*.

LEMMA 12. *If (q, p) is optimal in the monotone path primal, then (q, p) is feasible in the monotone edge primal.*

PROOF. Let (q, p) be optimal in the monotone path primal. Assume for contradiction that for some $(s, t) \in E'$, (q, p) violates the (s, t) incentive constraint in the monotone edge primal so that

$$p_t > v_t(q_t - q_s) + p_s. \quad (55)$$

Optimality of (q, p) in the monotone path primal implies that there exists a path $P: 0 = t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_n = s$ such that

$$p_s = v_{t_n} q_{t_n} - \sum_{i=1}^{n-1} q_{t_i} (v_{t_{i+1}} - v_{t_i}) - v_{t_1} q_{t_0}. \quad (56)$$

Setting $t_{n+1} = t$ and plugging (56) into (55), we obtain

$$p_t > v_{t_{n+1}} q_{t_{n+1}} - \sum_{i=1}^n q_{t_i} (v_{t_{i+1}} - v_{t_i}) - v_{t_1} q_{t_0}. \quad (57)$$

Because $(t_n, t_{n+1}) = (s, t) \in E'$, $P': t_0 \rightarrow \dots \rightarrow t_n \rightarrow t_{n+1}$ belongs to \mathcal{P}' . (P' cannot be a cycle as G' is acyclic.) But then (57) contradicts the feasibility of (q, p) in the monotone path primal. \triangleleft

LEMMA 13. (i) *If (q, p) is feasible in the original edge primal, then (q, p) is feasible in the original path primal.*

(ii) *The monotone path primal, the original path primal, the monotone edge primal, and the original edge primal all have the same value.*

PROOF. Choose $P: t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_k$ in \mathcal{P} . The instance of (25) in the original path primal follows from the incentive constraints (4) corresponding to edges (t_{i-1}, t_i) for $i = 1, \dots, k$ in the original edge primal. This establishes part (i). The monotone path primal is a relaxation of the original path primal. So part (ii) follows from part (i) and Lemmas 4 and 12. \triangleleft

We are now in a position to prove Proposition 8. If (q, p) is optimal in the original edge primal, by Lemma 13, (q, p) is optimal in the original path primal. Now suppose (q, p) is optimal in the original path primal. Then (q, p) is feasible in the monotone path primal, and so by Lemma 13(ii), is optimal in the monotone path primal. So by Lemma 12 and Lemma 13(ii), (q, p) is optimal in the monotone edge primal. So by Lemma 4, (q, p) is optimal in the original edge primal. \square

PROOF OF PROPOSITION 10. Let $\mu = (\mu_t: t \in T)$, $\lambda = (\lambda(s, t): (s, t) \in E)$, and $\tilde{\lambda} = (\lambda_P: P \in \mathcal{P})$. Let (λ, μ) be a good optimal solution to the edge formulation of the

dual. For any $(s, t) \in E$ and path P in G , define

$$\phi(s, t) := \frac{\lambda(s, t)}{\sum_{r: (r, t) \in E} \lambda(r, t)} \quad \text{and} \quad \Phi_P := \left[\prod_{i=1}^{n_P} \phi(t_{i-1}^P, t_i^P) \right] \pi_{t_{n_P}^P}. \quad (58)$$

(For notation t_i^P and $t_{n_P}^P$, see Appendix A.1.) The term Φ_P is defined for both paths P originating in 0 and paths P not originating in 0. (In contrast, λ_P is only defined for paths P originating in 0.)

In view of Proposition 8, strong duality, the fact that (24)–(27) is the dual of (29)–(33), and that the objective function value for both the edge and path formulations of the dual is determined by μ , if $(\tilde{\lambda}, \mu)$ is feasible in the path formulation of the dual, then $(\tilde{\lambda}, \mu)$ is optimal.

Suppose that $\tilde{\lambda}$ satisfies (34) and (35). Then (34) implies that

$$\sum_{P \ni t} \lambda_P = \sum_{s: (s, t) \in E} \sum_{P \ni (s, t)} \lambda_P = \sum_{s: (s, t) \in E} \lambda(s, t) \quad \forall t \in T \setminus 0. \quad (59)$$

So (34) and (35) imply that

$$\lambda_P = \Phi_P \quad \forall P \in \mathcal{P}. \quad (60)$$

So the proportional vector $\tilde{\lambda}$ corresponding to λ is unique if it exists.

To complete the proof, it is sufficient to argue that if $\tilde{\lambda}$ is defined from λ via (60), then $(\tilde{\lambda}, \mu)$ is feasible in the path formulation of the dual, and satisfies (34) and (35).

Let $\ell(t)$ be the maximum length of a $0 - t$ path in the monotone incentive graph G' . (For the definition of path length, see Appendix A.1.) First, we argue by induction on $\ell(t)$ that $\tilde{\lambda}$ satisfies (30). For the base case, when $\ell(t) = 1$, $(s, t) \in E'$ exactly if $s = 0$. So the only $0 - t$ path not containing a bad edge is $P^* = (0, t)$. So $\sum_{P \in \mathcal{P}_t} \lambda_P = \lambda_{P^*} = \Phi_{P^*} = \pi_t$, where the first equality follows from the assumption that λ is good and the last equality follows from (58) given that $P^* = (0, t)$. For the inductive step,

$$\sum_{P \in \mathcal{P}_t} \lambda_P = \pi_t \sum_{s: (s, t) \in E} \phi(s, t) \frac{\sum_{P \in \mathcal{P}_s} \lambda_P}{\pi_s} = \pi_t \sum_{s: (s, t) \in E} \phi(s, t) = \pi_t.$$

The first two equalities follow from (58), (60), and the fact that if $P \ni t$, then either (s, t) is a bad edge, in which case $\phi(s, t) = 0$, or P contains a bad edge, in which case $\Phi_P = 0$. The third equality follows from the fact that if $(s, t) \in E$ is such that $\phi(s, t) > 0$, then $(s, t) \in E'$ so that $\ell(s) < \ell(t)$ and we can apply the inductive hypothesis. The fourth equality follows from (58). So $(\tilde{\lambda}, \mu)$ satisfies (30). Further below, we show that $(\tilde{\lambda}, \mu)$ also satisfies (31).

For any $t \in T$, let $\mathcal{P}_{t \rightarrow}$ be the set of paths in G originating in t (where a path contains at least one edge so that (t) is not a path). Next, we argue that $\tilde{\lambda}$ satisfies (34) and (35). To do so, we argue by backward induction on $\ell(t)$ that

$$\sum_{s: (s, t) \in E} \lambda(s, t) = \pi_t + \sum_{P \in \mathcal{P}_{t \rightarrow}} \Phi_P \quad \forall t \in T. \quad (61)$$

For the base case, when $\ell(t)$ is maximal, then for all $s \in T$, if $(s, t) \in E$, then (s, t) is a bad edge and if $P \in \mathcal{P}_{t \rightarrow}$, then P contains a bad edge. So both (12) and (61) reduce to

$\sum_{s:(s,t) \in E} \lambda(s, t) = \pi_t$. Since λ satisfies (12), (61) holds in this case. For the inductive step,

$$\begin{aligned} \sum_{s:(s,t) \in E} \lambda(s, t) &= \pi_t + \sum_{s:(t,s) \in E} \lambda(t, s) = \pi_t + \sum_{s:(t,s) \in E} \left\{ \phi(t, s) \sum_{r:(s,r) \in E} \lambda(s, r) \right\} \\ &= \pi_t + \sum_{s:(t,s) \in E} \left\{ \phi(t, s) \left[\pi_s + \sum_{P \in \mathcal{P}_{s \rightarrow}} \Phi_P \right] \right\} = \pi_t + \sum_{P \in \mathcal{P}_{t \rightarrow}} \Phi_P, \end{aligned}$$

where the first equality follows from (12), the second equality follows from (58), the third equality follows from the inductive hypothesis and the fact for all $(t, s) \in E$, if $\phi(t, s) > 0$, then $(s, t) \in E'$, and hence, $\ell(s) > \ell(t)$, and the fourth equality follows from (58) and the fact that if $(t, s) \in E$ and $P \ni t$, then either (s, t) is a bad edge, in which case $\phi(s, t) = 0$, or P contains a bad edge, in which case $\Phi_P = 0$. We have now established (61).

Next we have

$$\begin{aligned} \lambda(s, t) &= \phi(s, t) \sum_{r:(r,t) \in E} \lambda(r, t) = \phi(s, t) \left\{ \pi_t + \sum_{P \in \mathcal{P}_{t \rightarrow}} \Phi_P \right\} \\ &= \left\{ \frac{\sum_{P \in \mathcal{P}_s} \Phi_P}{\pi_s} \right\} \times \phi(s, t) \times \left\{ \pi_t + \sum_{P \in \mathcal{P}_{t \rightarrow}} \Phi_P \right\} = \sum_{P \ni (s,t)} \lambda_P, \end{aligned}$$

where the first equality follows from (58), the second from (61), the third from the fact, established above, that $\tilde{\lambda}$ satisfies (30), and the fourth from (58) and the fact that for all $(s, t) \in E$, $P \in \mathcal{P}_s$, and $P' \in \mathcal{P}_{t \rightarrow}$, then if $t_0^P \rightarrow \dots \rightarrow t_{n_P}^P = s \rightarrow t = t_0^{P'} \rightarrow \dots \rightarrow t_{n_{P'}}^{P'}$ contains a cycle, then since λ is good, $\Phi_P \times \phi(s, t) \times \Phi_{P'} = 0$. So $\tilde{\lambda}$ satisfies (34). Equation (34) implies (59), which, together with (58), (60), and another application of (34), implies that $\tilde{\lambda}$ satisfies (35). Finally, (34) and (13) imply that $(\tilde{\lambda}, \mu)$ satisfies (31) and, hence, is feasible in the path formulation of the dual. \square

REMARK 4. Let (λ, μ) be a good optimum of the edge formulation of the dual and let $(\tilde{\lambda}, \mu)$ be the corresponding proportional optimum of the path formulation, where $\tilde{\lambda}$ is related to λ by (34). Then (61) and the fact that $\tilde{\lambda}$ satisfies (60) imply that

$$\sum_{P' \in \mathcal{P}: P \subseteq P'} \lambda_{P'} = \frac{\lambda_P}{\pi_t} \sum_{s:(s,t) \in E} \lambda(s, t) \quad \forall t \in T, \forall P \in \mathcal{P}_t.$$

PROOF OF PROPOSITION 11. The proof is implemented in two steps. First we present a Bayesian Nash equilibrium implementing the optimal mechanism. Then we explain how this can be strengthened to a sequential equilibrium.

Step 1: Bayesian Nash equilibrium.

Strategies. The players' on-path play was described in Section 7.3.1. Here we specify off-path play. This is important to ensure that the strategies are mutual best replies. When the preliminary randomization yields path P in (36), then even off the equilibrium path, when asked for a cheap talk report, the buyer reports along path P as described in Section 7.3.1 unless the buyer herself previously deviated from reporting along path P ,

in which case she drops out. Similarly, when asked for evidence, the buyer's strategy is as described in Section 7.3.1 unless the buyer previously deviated from reporting along P , in which case she drops out without presenting evidence. To specify the seller's off-equilibrium play, we assume that at Step 3 of the dynamic bargaining protocol, if the seller has observed a sequence of reports that is not buyer on-path (see (39)) or if the last report \hat{t} that the seller observed is such that $q_{\hat{t}} = 1$, the seller requests evidence with probability 1. At Step 3, if the buyer has presented evidence S , and either $S \neq S_{t_k}$, where t_k was the last type report before evidence was requested, or the sequence of type reports presented was not buyer on-path, then the seller makes an offer at a price equal to the maximum value of any type that had access to S (that is, $\max\{v_r : S \subseteq \sigma(r)\}$). Note that because the buyer's and seller's randomizations are determined by dual and primal optima, complimentary slackness and Proposition 3(ii) imply that if t_{k-1} and t_k are two consecutive cheap talk reports that were made by the buyer on the equilibrium path, then $q_{t_{k-1}} \leq q_{t_k}$. So the seller's strategy as described here and in Section 7.3.1 is well defined. In what follows, we refer to the buyer and seller strategies defined here and in Section 7.3.1 as ζ^* and ξ^* , respectively.

The strategies implement the optimal mechanism. We show that the strategies ζ^* and ξ^* , if followed, implement the same outcome as the optimal mechanism. For any path $P: t_0 \rightarrow t_1 \rightarrow \dots \rightarrow t_n$ in \mathcal{P}_t define k_P to be the smallest index i such that $q_{t_i} = 1$ if such an i exists and $k_P = n$ otherwise. Complementary slackness and Proposition 3(ii) imply that if $P \in \mathcal{P}_t$ and $\lambda_P > 0$, then $q_{t_{k_P}} = q_t$. Using notation t_i^P and $t_{n^P}^P$ defined in Appendix A.1, and adopting the convention $q_{t_{-1}^P} := 0$, the strategy profile (ζ^*, ξ^*) induces a probability of sale for type t buyer of

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

where we have used the fact that by (30), $\sum_{P \in \mathcal{P}_t} \lambda_P / \pi_t = 1$. Similarly, the expected payment of the type t buyer induced by (ζ^*, ξ^*) is

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

where the second equality uses the fact that if $\lambda_P > 0$, then $q_{t_k^P} \geq q_{t_{k-1}^P}$ by complementary slackness in the edge formulation, (34), and Proposition 3(ii). The last equality uses complementary slackness. It follows that we implement the optimal mechanism.

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 $P = (t_0, \dots, t_k)$ that 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.⁴³ However, it now follows from the arguments like those establishing that ζ^* and ξ^* , if followed, implement the optimal mechanism, that the buyer's payoff from this deviation would be $v_t q_s - p_s$. Incentive compatibility ((4) in the primal) implies this deviation would yield a payoff inferior to $v_t q_t - p_t$, which, as we argued above, is the payoff that the type t buyer would attain if she used ζ^* .

Seller optimization. We now fill in some missing details from the argument of Section 7.3.4 that ξ^* is a best reply to ζ^* . We start with two lemmas establishing facts that were mentioned in Section 7.3.4.

LEMMA 14. *There exists a seller best reply to ζ^* that 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 nonterminal history h satisfying (i) following h , it is the seller's turn to make an offer (Step 5), and (ii) the sequence of cheap talk reports $P = (t_0, \dots, t_k)$ in h is buyer on-path (see (39)). Statements (i) and (ii) imply that the buyer presented evidence S_{t_k} (since she uses ζ^*). 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 t with $v_t \geq p$, at which point ξ' requests evidence and then behaves as in a stopping strategy, so that if the buyer then presents S_t , ξ' makes an offer of v_s . 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 that is a stopping strategy and also a best reply to ζ^* .⁴⁴ \triangleleft

LEMMA 15. *If, using ζ^* , the buyer previously reported $P = (t_1, \dots, t_k)$ and then the seller continued, requesting more cheap talk, the probability that the seller assigns to the event that the buyer will not drop out but instead will report $t_{k+1} = t$ is given by (37).*

⁴³Observe in particular that if $P = (t_0, \dots, t_k)$, $\lambda_P > 0$, 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 that ends in the last type s in P for which $\sigma(s) \subseteq \sigma(t)$, and so drop out after s is reached.

⁴⁴For ξ^0 to be a stopping strategy, we may have to make some additional modifications conditional on histories that occur with zero probability and, hence, do not affect the seller's payoff.

PROOF. By Remark 4, the probability that in the preliminary randomization, the buyer selected a sequence $P'' \supseteq P$ is

$$\sum_{t \in T} \sum_{\{P'' \in \mathcal{P}_t : P \subseteq P''\}} \pi_t \frac{\lambda_{P''}}{\pi_t} = \sum_{P'' \in \mathcal{P} : P \subseteq P''} \lambda_{P''} = \frac{\lambda_P}{\pi_{t_k}} \sum_{s : (t_k, s) \in E} \lambda(t_k, s). \tag{62}$$

Similarly, the probability that the buyer selected $P'' \supseteq P' = (t_0, \dots, t_k, t)'$ is

$$\frac{\lambda_{P'}}{\pi_t} \sum_{s : (t, s) \in E} \lambda(t, s). \tag{63}$$

Equations (34) and (35) imply that

$$\frac{\lambda_{P'} / \pi_t}{\lambda_P / \pi_{t_k}} = \frac{\lambda(t_k, t)}{\sum_{s : (s, t) \in E} \lambda(s, t)}.$$

It follows that dividing (63) by (62) yields (37). ◁

The definition of ξ^* implies the existence of a stopping strategy $\tilde{\xi}^*$ that agrees with ξ^* on all histories that contain a buyer on-path sequence of type reports (see (39)). Once Lemma 2 is established and in view of Lemma 14, to complete the argument that the ξ^* is a best reply to ζ^* , it is sufficient to show that $\tilde{\xi}^*$ is optimal among stopping strategies. We argue by backward induction. As we go, note that (35) implies that a sequence of type reports P is buyer on-path if and only if $\lambda_P > 0$. Consider a history $P = (t_0, \dots, t_k)$ with $\lambda_P > 0$. First let P be such a history of maximal length.⁴⁵ (Recall that the length of P is the number of edges in P .) In this case, we must have $q_{t_k} = 1$,⁴⁶ and clearly it is optimal to stop as required by $\tilde{\xi}^*$. 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(ii) that for all $P' = (t_0, \dots, t_k, t_{k+1})$ 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 2 now implies that following P , stopping immediately would be optimal as required by $\tilde{\xi}^*$. Next suppose that $q_{t_k} < 1$. Then by Lemma 2, the seller would be weakly better off continuing one step and then stopping than stopping immediately, and so continuing and then following $\tilde{\xi}^*$ (which by backward induction, is optimal) would be even better, again as required by $\tilde{\xi}^*$.

Step 2: Sequential equilibrium. The above argument presented a strategy profile and established that it is a Bayesian Nash equilibrium implementing the optimal mechanism. We now strengthen this to a sequential equilibrium. First, we modify the strategies ζ^* and ξ^* to form strategies ζ^{**} and ξ^{**} , respectively. Strategy ξ^{**} agrees with ξ^* at any

⁴⁵Such a history exists because T is finite and λ has no bad edges; in other words, a sequence of cheap talk reports cannot form a cycle.

⁴⁶Suppose 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 (31) imply that there must exist $t_{k+1} \in T$ with $v_{t_k} < v_{t_{k+1}}$ and $\sum_{P' \ni (t_k, t_{k+1})} \lambda_{P'} > 0$. Since λ is good, $t_{k+1} \notin (t_0, \dots, t_k)$. So consider the path $P'' = (t_0, \dots, t_k, t_{k+1})$. So (35) implies that $\lambda_{P''} = \lambda_P (\sum_{P' \ni (t_k, t_{k+1})} \lambda_{P'} / \sum_{P' \ni t_{k+1}} \lambda_{P'}) (\pi_{t_{k+1}} / \pi_{t_k}) > 0$, contradicting the assumption that P was of maximal length.

seller information set that occurs with positive probability given (ζ^*, ξ^*) , as well as any seller information set at which evidence has not yet been presented. At any seller information set that 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 $\max\{v_r : S \subseteq \sigma(r)\}$. Similarly, ζ^* agrees with ζ^{**} at any buyer information set that 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. At any type t buyer information set I that occurs with zero probability given (ζ^*, ξ^*) only because the type t buyer has taken a sequence of actions that would have been taken with positive probability by some other buyer type according to ζ^* (and the seller has taken actions consistent with ξ^*)—call such information sets *undetected*—the buyer continues by following some type t best reply to ξ^{**} conditional on I . At any other information set, the buyer drops out.

LEMMA 16. (i) *The strategy profile (ζ^{**}, ξ^{**}) is a Bayesian Nash equilibrium of the dynamic communication protocol.*

(ii) *The strategy profiles (ζ^{**}, ξ^{**}) and (ζ^*, ξ^*) induce the same probability distribution over terminal histories.*

PROOF. Part (ii) follows by construction. So consider part (i). If ξ is a seller strategy profile such that (ζ^*, ξ) and (ζ^{**}, ξ) induce the same probability distribution over terminal histories, then part (ii) of the lemma and the fact that (ζ^*, ξ^*) is a Bayesian Nash equilibrium implies that ξ is not a profitable seller deviation at (ζ^{**}, ξ^{**}) . So consider a ξ such that (ζ^*, ξ) and (ζ^{**}, ξ) induce different probability distributions over terminal histories. The definition of ζ^{**} means that (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 the 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 (i) of the lemma. \square

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.⁴⁷ For each $\epsilon > 0$, we construct a totally mixed buyer strategy ζ^ϵ such that $\zeta^\epsilon \rightarrow \zeta^{**}$ as $\epsilon \rightarrow 0$. Enumerate the

⁴⁷There is no corresponding issue for the buyer's beliefs because the seller has no private information.

types t'_1, \dots, t'_n in T so that $i < j \Rightarrow v_{t'_i} \geq v_{t'_j}$. The strategy 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 strategy⁴⁸). 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 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; 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 that 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. Lemma 16(ii) implies 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 S . In either event, the seller will use the off-equilibrium beliefs derived above and infer that the buyer is of the type 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 is undetected, 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. \square

PROOF OF PROPOSITION 12. This follows from the equilibrium strategies and the fact that at an optimal good dual solution, $\sum_{P \in \mathcal{P} : (s,t) \in P} \lambda_P > 0$ implies that both $v_s < v_t$ and $S_s \subseteq S_t$. The former inequality depends on the fact that λ is good, while the latter inclusion depends only on (1). \square

PROOF OF PROPOSITION 13. To prove the proposition, we present a series of environments, parameterized by n , such that the n th environment requires $n + 1$ rounds of communication in the equilibrium constructed in Section 7.3.1. Some assertions made

⁴⁸The resulting beliefs do not depend on which totally mixed seller strategy is used.




FIGURE 3. The incentive graph.

below are proved in Sher and Vohra (2011), a previous working paper version of this paper.

Let $T = X \cup Y \cup \{0\}$, where $X = \{x_1, \dots, x_n\}$ and $Y = \{y_0, y_1, \dots, y_n\}$. 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 to the types in Y as y -types. As usual, we assume that $v_0 = \pi_0 = 0$. Moreover, we assume that

$$v_{x_1} < v_{x_2} < \dots < v_{x_n} < v_{y_0} < v_{y_1} < v_{y_2} < \dots < v_{y_n}. \tag{64}$$

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

$$E = \{(0, t) : t \in T \setminus \{0\}\} \cup \{(x_i, x_j) \in X \times X : i < j\} \cup \{(x_i, y_j) \in X \times Y : i \leq j\} \cup \{(y_i, y_j) \in Y \times Y : i < j\}. \tag{65}$$

This can be represented pictorially as in Figure 3.

Each directed path in Figure 3 corresponds to an edge in the incentive graph (65). 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

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

$$\sigma(0) = \{m_0\},$$

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

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

$$\delta_n := 0$$

$$\delta_{i-1} := \delta_i + \frac{v_{x_i} \pi_{x_i} - (v_{x_{i+1}} - v_{x_i})[\delta_i + \sum_{j=i+1}^n \pi_{x_j}]}{v_{y_i} - v_{x_i}} \quad \forall i = 1, \dots, n.$$

Note that when $i = n$, we define $(v_{x_{i+1}} - v_{x_i})[\delta_i + \sum_{j=i+1}^n \pi_{x_j}] := 0$, and similarly in (68) below, when $i = n$, we define $(v_{x_{i+1}} - v_{x_i})[(\sum_{j=i+1}^n \pi_{x_j} + \sum_{j=i+1}^n \pi_{y_j})/\pi_{x_i}] := 0$. We also

assume that

$$v_{y_i} - (v_{y_{i+1}} - v_{y_i}) \frac{\sum_{j=i+1}^n \pi_{y_j}}{\pi_{y_i}} > 0 \quad \forall i = 0, 1, \dots, n-1 \quad (66)$$

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i}) \frac{\sum_{j=i+1}^n \pi_{x_j}}{\pi_{x_i}} > (v_{x_{i+1}} - v_{x_i}) \frac{\delta_i}{\pi_{x_i}} \quad \forall i = 1, \dots, n-1 \quad (67)$$

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i}) \left[\frac{\sum_{j=i+1}^n \pi_{x_j} + \sum_{j=i+1}^n \pi_{y_j}}{\pi_{x_i}} \right] - (v_{y_i} - v_{x_i}) \frac{\pi_{y_i}}{\pi_{x_i}} < 0 \quad \forall i = 1, \dots, n. \quad (68)$$

For any profile of valuations satisfying (64), there are many probability distributions $(\pi_t : t \in T)$ such that (66)–(68) are satisfied. Inequality (66) 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 renormalized to sum to 1), then the optimal allocation would allocate the object to each type with probability 1.

Similarly, (67) implies that if there were only x -types, then it would be optimal to allocate the object to all types. But the assumption (67) for x -types is a stronger assumption than the corresponding assumption (66) for y -types. Indeed, a simple induction using (67) implies

$$\delta_i > 0 \quad \forall i = 0, 1, \dots, n-1,$$

and, hence, (67) also implies

$$v_{x_i} - (v_{x_{i+1}} - v_{x_i}) \frac{\sum_{j=i+1}^n \pi_{x_j}}{\pi_{x_i}} > 0 \quad \forall i = 1, \dots, n-1.$$

Inequality (68) says that if the set of types were of the form $T_i := \{x_i, x_{i+1}, \dots, x_n\} \cup \{y_i, y_{i+1}, \dots, y_n\} \cup \{0\}$ for $i = 1, \dots, n$, 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

$$q_{y_i} := 1 \quad \forall i = 0, 1, \dots, n \quad (69)$$

$$p_{y_i} := v_{y_0} \quad \forall i = 0, 1, \dots, n. \quad (70)$$

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

$$q_{x_1} := \frac{v_{y_1} - v_{y_0}}{v_{y_1} - v_{x_1}} \quad (71)$$

$$p_{x_1} := v_{x_1} q_{x_1} \quad (72)$$

$$q_{x_i} := \frac{(v_{y_i} - v_{y_0}) - (v_{x_i} q_{x_{i-1}} - p_{x_{i-1}})}{v_{y_i} - v_{x_i}} \quad \forall i = 2, \dots, n \quad (73)$$

$$p_{x_i} := v_{x_i} (q_{x_i} - q_{x_{i-1}}) + p_{x_{i-1}} \quad \forall i = 2, \dots, n. \quad (74)$$

It is straightforward to verify that

$$0 < q_{x_1} < q_{x_2} < \dots < q_{x_n} < 1.$$

Moreover, the mechanism given by (69)–(74) is the unique optimal mechanism. It is also the case that at every dual optimal solution, we have that

$$\begin{aligned} \lambda(x_{i-1}, x_i) &> 0 \quad \forall i = 2, \dots, n \\ \lambda(0, x_1) &> 0 \\ \lambda(x_h, x_i) &= 0 \quad \forall i = 2, \dots, n, \forall h < i - 1 \\ \lambda(0, x_i) &= 0 \quad \forall i = 2, \dots, n. \end{aligned}$$

It follows that the unique path P from 0 to x_n with $\lambda_P > 0$ is $P = (0, x_1, x_2, \dots, x_n)$, so type x_n is asked to submit a cheap talk report $n + 1$ times in the equilibrium constructed in Section 7.3.1. To describe the equilibrium in more detail, each x -type x_i uses the sequence of reports $(0, x_1, x_2, \dots, x_i)$, 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: $(0, x_1, x_2, \dots, x_i, y_i)$ and $(0, y_0, y_1, \dots, y_i)$. 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 requesting evidence. \square

A.3 Algorithm for optimal mechanism under tree structure

Here we show how under tree structure, the optimal mechanism can then be found by solving a simple perfect information game between the seller and nature by backward induction. So assume that G^* is a tree. Then for each type t , let D_t be the set of types s such that the unique path in G^* from 0 to s passes through t ; t itself belongs to D_t . A *descendant* of t is a type in D_t other than t . So D_t consists of t and all of the descendants of t . For example in Figure 1, $D_3 = \{3, 4, 7\}$. Formally, $D_t := \{t\} \cup \{s : (t, s) \in E'\}$, where E' is the edge set of the monotone incentive graph. A *child* of type t is type s such that $(t, s) \in E^*$. Let C_t be the set of children of t . So in Figure 1, $C_1 = \{2, 3\}$; in contrast, $D_1 = \{1, 2, 3, 4, 5, 6, 7\}$.

The ex ante probability that the buyer's type belongs to D_t is Π_t (see (22)). If s is a child of t , then using Bayes' rule, the ex ante probability that the buyer's type belongs to D_s given that the buyer's type belongs to D_t is $\Pr(D_s|D_t) := \Pi_s/\Pi_t$. The ex ante probability that the buyer's type is t given that the buyer's type belongs to D_t is $\Pr(t|D_t) := \pi_t/\Pi_t$. Observe that

$$\Pr(t|D_t) + \sum_{s \in C_t} \Pr(D_s|D_t) = 1.$$

GAME AGAINST NATURE. 1. Start with $t := 0$.

2. (a) With probability $\Pr(t|D_t)$, the game ends and the seller receives a payoff of 0.

(b) For each child s of t , with probability $\Pr(D_s|D_t)$, nature selects type s . Go to Step 3.

3. The seller observes s and selects STOP or CONTINUE.

(a) If the seller selects STOP, the game ends and the seller receives a payoff of v_s .

(b) If the seller selects CONTINUE, redefine $t := s$ and go to Step 3.

Observe that if the game reaches a history where nature has selected a type s that has no children, then $\Pr(s|D_s) = 1$, so even if the seller were to select CONTINUE (which would not be a wise choice), the game would end in one step. It follows that the game terminates after nature has selected a type at most n times, where n is the number of types.

PROPOSITION 15. *Assume that G^* is a tree. The optimal mechanism can be found by solving the game against nature by backward induction: For each type s on which the seller stops, the seller sells the object for a price of v_s and to s all descendants of s . The seller does not sell the object to any other type.*

PROOF. Let \mathcal{S} and \mathcal{C} be shorthand, respectively, for STOP and CONTINUE. Let \mathcal{P}^* be the set of all paths in G^* that begin at 0 (including the degenerate path (0)). A strategy in the extensive form game can be represented as a function $\eta: \mathcal{P}^* \rightarrow \{\mathcal{S}, \mathcal{C}\}$. Because G^* is a tree, for each $t \in T \setminus 0$, there is a unique $0 - t$ path $P^t: 0 = s_0^t \rightarrow s_1^t \rightarrow \dots \rightarrow s_{n_t}^t = t$ in G^* . So we can rewrite η as a function $\eta: T \rightarrow \{\mathcal{S}, \mathcal{C}\}$, where $\eta(t) = \eta(P^t)$. Define

$$T^\eta = \{t \in T: [\forall s \in T, (s, t) \in E' \Rightarrow \eta(s) = \mathcal{C}] \text{ and } [\eta(t) = \mathcal{S}]\}.$$

Let E_η be the seller's expected payoff to strategy η in the game against nature. Then, using the fact that $\Pi_0 = \Pi_{s_0^t} = 1$,

$$E^\eta = \sum_{t \in T^\eta} \left(\prod_{i=1}^{n_t} \frac{\Pi_{s_i^t}}{\Pi_{s_{i-1}^t}} \right) v_t = \sum_{t \in T^\eta} v_t \Pi_t.$$

For any strategy η , define $q_t^\eta = 1$ if either $t \in T^\eta$ or there exists $s \in T^\eta$ with $(s, t) \in E'$. Otherwise, define $q_t^\eta = 0$. Then q^η is feasible in (47)–(50). For all $(t, s) \in E'$, let $P^{t,s}: t = r_0^{t,s} \rightarrow \dots \rightarrow r_{n_t,s}^{t,s} = s$ be the unique $s - t$ path in G^* . Then

$$\begin{aligned} \sum_{t \in V} q_t^\eta \hat{\psi}(t) \pi_t &= \sum_{t \in T^\eta} \sum_{s \in D_t} \hat{\psi}(s) \pi_s = \sum_{t \in T^\eta} \sum_{s \in D_t} \left(v_s - \sum_{r: (s,r) \in E^*} (v_r - v_s) \frac{\Pi_r}{\pi_s} \right) \pi_s \\ &= \sum_{t \in T^\eta} \left[v_t \pi_t + \sum_{s \in D_t \setminus t} \left(\left[v_{r_{n_t}^{t,s}} + \sum_{i=1}^{n_t,s} (v_{r_i^{t,s}} - v_{r_{i-1}^{t,s}}) \right] \pi_s \right) \right] = \sum_{t \in T^\eta} v_t \Pi_t = E^\eta, \end{aligned}$$

where the third equality follows by separating and regrouping the terms multiplying each probability π_s , and the last equality follows from the fact that $v_{r_{n_t}^{t,s}} +$

$\sum_{i=1}^{n_{t,s}} (v_{r_i^{t,s}} - v_{r_{i-1}^{t,s}}) = v_s$. So q^η achieves the objective function value E^η . Next, choose any deterministic feasible solution q to (47)–(50) (i.e., a feasible solution such that $q_t \in \{0, 1\}$ for all t); then the strategy η defined by $\eta(t) = \mathcal{S}$ if $q_t = 1$ and by $\eta(t) = \mathcal{C}$ otherwise is such that $q = q^\eta$. That this is so depends on the monotonicity constraints (48). Lemma 6 and the fact—explained in the proof of Proposition 4—that the program (47)–(50) admits a deterministic optimum now imply that the backward induction solution to the game achieves the seller's optimal revenue in the primal (3)–(6). The remaining claims in the proposition are straightforward to verify. \square

REFERENCES

- Ahuja, Ravindra K., Thomas L. Magnanti, and James B. Orlin (1993), *Network Flows*. Prentice Hall, Englewood Cliffs, New Jersey. [608, 617]
- Aumann, Robert J. and Sergiu Hart (2003), “Long cheap talk.” *Econometrica*, 71, 1619–1660. [621]
- Ausubel, Lawrence M. and Raymond J. Deneckere (1989), “A direct mechanism characterization of sequential bargaining with one-sided incomplete information.” *Journal of Economic Theory*, 48, 18–46. [601, 618]
- Ben-Porath, Elchanan and Barton L. Lipman (2012), “Implementation with partial provability.” *Journal of Economic Theory*, 147, 1689–1724. [600]
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015), “The limits of price discrimination.” *American Economic Review*, 105, 921–957. [600]
- Blumrosen, Liad, Noam Nisan, and Ilya Segal (2007), “Auctions with severely bounded communication.” *Journal of Artificial Intelligence Research*, 28, 233–266. [601]
- Bull, Jesse and Joel Watson (2007), “Hard evidence and mechanism design.” *Games and Economic Behavior*, 58, 75–93. [598, 600, 603, 626]
- Celik, Gorkem (2006), “Mechanism design with weaker incentive compatibility constraints.” *Games and Economic Behavior*, 56, 37–44. [600]
- Che, Yeon-Koo and Ian Gale (1998), “Standard auctions with financially constrained bidders.” *Review of Economic Studies*, 65, 1–21. [601]
- Deneckere, Raymond and Sergei Severinov (2008), “Mechanism design with partial state verifiability.” *Games and Economic Behavior*, 64, 487–513. [598, 600, 603, 626]
- Efma and Oliver Wyman (2012), “Pricing at European retail banks.” http://www.efma.com/ressources/studies/2012/1-iwv8j_e_study.pdf. [619]
- Forges, Françoise and Frédéric Koessler (2005), “Communication equilibria with partially verifiable types.” *Journal of Mathematical Economics*, 41, 793–811. [598, 600, 626]
- Forges, Françoise and Frédéric Koessler (2008), “Long persuasion games.” *Journal of Economic Theory*, 143, 1–35. [621]

- Gerardi, Dino, Johannes Hörner, and Lucas Maestri (2014), “The role of commitment in bilateral trade.” *Journal of Economic Theory*, 154, 578–603. [601, 618]
- Glazer, Jacob and Ariel Rubinstein (2004), “On optimal rules of persuasion.” *Econometrica*, 72, 1715–1736. [600, 625, 626]
- Glazer, Jacob and Ariel Rubinstein (2006), “A study in the pragmatics of persuasion: A game theoretical approach.” *Theoretical Economics*, 1, 395–410. [600, 625, 626]
- Green, Jerry R. and Jean-Jacques Laffont (1986), “Partially verifiable information and mechanism design.” *Review of Economic Studies*, 53, 447–456. [600, 603]
- Kartik, Navin and Olivier Tercieux (2012), “Implementation with evidence.” *Theoretical Economics*, 7, 323–355. [600]
- Kos, Nenad (2012), “Communication and efficiency in auctions.” *Games and Economic Behavior*, 75, 233–249. [601]
- Lipman, Barton L. and Duane J. Seppi (1995), “Robust inference in communication games with partial provability.” *Journal of Economic Theory*, 66, 370–405. [600]
- Milgrom, Paul and John Roberts (1986), “Relying on information of interested parties.” *RAND Journal of Economics*, 17, 18–32. [600]
- Moore, John (1984), “Global incentive constraints in auction design.” *Econometrica*, 52, 1523–1535. [600]
- Myerson, Roger B. (1982), “Optimal coordination mechanisms in generalized principal–agent problems.” *Journal of Mathematical Economics*, 10, 67–81. [605, 614]
- Myerson, Roger B. (1991), *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts. [614]
- Myerson, Roger B. (2002), “Incentive-dual methods in game theory: Virtual utility and dual reduction revisited.” Working paper, Chicago University. [614]
- Nemhauser, George L. and Laurence A. Wolsey (1999), *Integer and Combinatorial Optimization*. Wiley, New York. [631, 632]
- Pai, Mallesh M. and Rakesh Vohra (2014), “Optimal auctions with financially constrained buyers.” *Journal of Economic Theory*, 150, 383–425. [601]
- Severinov, Sergei and Raymond Deneckere (2006), “Screening when some agents are nonstrategic: Does a monopoly need to exclude?” *RAND Journal of Economics*, 37, 816–840. [600]
- Sher, Itai (2011), “Credibility and determinism in a game of persuasion.” *Games and Economic Behavior*, 71, 409–419. [600]
- Sher, Itai (2014), “Persuasion and dynamic communication.” *Theoretical Economics*, 9, 99–136. [600, 625, 626]
- Sher, Itai and Rakesh Vohra (2011), “Price discrimination through communication.” CMS-EMS discussion paper, Northwestern University. [642]

Shin, Hyun Song (1994), "The burden of proof in a game of persuasion." *Journal of Economic Theory*, 64, 253–264. [600]

Singh, Nirvikar and Donald Wittman (2001), "Implementation with partial verification." *Review of Economic Design*, 6, 63–84. [600]

Vohra, Rakesh V. (2011), *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, Cambridge. [601]

Submitted 2011-12-4. Final version accepted 2014-5-26. Available online 2014-5-28.