Auctions to Learn Consumer Demand for a Product with a Short Selling Horizon

Matthew L. Potoff • Damian R. Beil

Stephen M. Ross School of Business, University of Michigan
701 Tappan St., Ann Arbor, Michigan 48109–1234, USA
potoff@umich.edu • dbeil@umich.edu

September, 2007

A firm plans to sell a new product to the mass market at a fixed price. In order to learn about consumers' willingness to pay for the new product, online auctions are used as a test market, allowing the firm to simultaneously learn demand information for various price points. The firm faces a stopping time problem for when it stops holding the auctions and switches to mass-market selling via a fixed price. Additional auctions increase expected profits by improving the firm's pricing decision, but also incur an opportunity cost of lost time to sell to the mass market before the end of a short selling horizon. Structural properties of the optimal stopping policy are presented, and the illustrative case in which mass-market sales follow a Bass diffusion is discussed. To help the firm infer willingness to pay from auction bids, it is shown that a second price sealed-bid auction with a rebate against the future fixed price induces bidders to truthfully reveal their valuations.

1. Introduction

For a firm selling a new, innovative product, pricing is the crucial step to harvest the value created in developing and producing the new product. Yet pricing is particularly challenging for an innovative product, whose "newness" can naturally produce uncertainty about its exact demand distribution. Complicating matters further, obsolescence or competitive entry threats can shorten sales horizons during which the firm can capitalize on its product's market potential.

As a motivating example, consider SawStop, a small entrepreneurial startup, which recently developed a table saw with an innovative safety device that stops the saw blade the instant it contacts human skin. After showcasing its invention at tradeshows, initial orders were satisfied from SawStop at a price of approximately $2,850 (Mehler 2005), whereas existing industry lines without the feature sold for between $1000 and $2,500. Eventually the firm began distributing to specialty woodworking stores, selling the saw at an MSRP of $3,270 (Johnson 2006). While SawStop is the first saw with such safety features, an industry trade group, representing large tool manufacturers (Black & Decker, Bosch, Dewalt, Ryobi, and others), announced that it expects to have even better guarding mechanisms ready by sometime in 2007 (Skrzycki 2006). With such an innovative product, SawStop had to estimate consumers' willingness to pay for this breakthrough safety feature, but do so quickly given the looming threat of competitive entry.

In this paper we develop a framework and analysis of how a firm, with a product developed and ready to sell but unsure of consumer demand (willingness to pay), can strike a balance between two competing desires. On the one hand, the firm wishes to spend time gathering information on consumer demand with which to determine a fixed price, and on the other hand, the firm wishes to start mass market sales quickly to take advantage of the short sales horizon. Because the product has already been developed and is ready for sale, the firm is able to use a test market approach to gathering data. Unlike a traditional test market that requires inventory positioned in test stores and tightly controlled localized conditions (such as prices), we propose an approach that utilizes online auctions as the test market channel. Online consumer auctions are popular and growing – eBay had 222 million registered users in 2006 (eBay.com 2007), and by 2010 the online auction industry is expected to reach $65 billion in sales (Johnson and Tesch 2005). In addition to being relatively easy to set up and run, the key benefit of an online auction test market is the ability to observe consumer bids instead of just purchase/no purchase decisions.

In our model we divide the selling horizon into two phases: an initial phase in which the product is sold exclusively by auction and demand information is gathered; and a secondary stage in which a fixed price is set and the mass market is reached via posted price retailers (e.g., nation-wide rollout to retailers). To our knowledge, this study is the first to analytically model auctions to inform a mass-market, fixed-price selling phase. Balancing the firm's two competing desires, between information on the one hand and more time to exploit the market on the other, is the crux of this research problem, boiling down to a stopping time decision on when to abandon the auction phase and begin the mass market sales phase.

In making this stopping time decision, we find that the firm must pay close attention to several key factors. First is its point estimates of purchase probabilities at various prices, which are driven by observations of consumer bids in the auctions. Second is the anticipated shape of the mass-market sales trajectory (which reflects factors such as how the product will diffuse through the marketplace due to innovative and imitative purchases, Bass 1969). Third is the amount of time remaining in the selling horizon, which, depending on the shape of the sales trajectory, might encourage or discourage the use of auctions for market research. Furthermore, in using auctions for market research in this setting, the firm can use a proactive auction design that eases solution of the inverse problem between bids and willingness to pay, facilitating purchase probability estimations. Below we flesh out more context for this work, beginning with background in auctions for market research and then moving to aspects of dynamic learning and stopping time literatures.

Auctions for market research. A great deal of experimental and empirical work has been conducted using auctions as a means for market research into willingness to pay, mostly in laboratory settings. A seminal laboratory study by Hoffman et al. (1993) using products in the packaged beef industry has been followed by other studies examining willingness to pay for new products, e.g., pesticide-free fruit (Roosen et al. 1998) and milk produced with the aid of bovine growth hormone (Fox et al. 1994). Experimental economists have widely tested how subjects behave in various forms of laboratory auctions (see the text by Kagel and Roth 1995), as well as how various auctions perform at eliciting "homegrown" (not induced by the experimenter) willingness-to-pay information (Ruström 1998, Noussair et al. 2004). Econometric analyses of empirical data have taken bid data from auctions in the field and reverse engineered it to reveal information on participants' willingness to pay; see, for example, the recent text by Paarsch and Hong (2006).

Inspired by this lineage of experimental and empirical auction work to estimate willingness to pay, our paper develops an analytical (as opposed to experimental or empirical) model of how online auctions might be used for demand learning during the beginning of a short selling horizon. In an empirical study, Paarsch (1997) used historical auction field data to suggest an optimal reserve price for government timber auctions. While similar in spirit to our work, we treat the amount of data gathered as a decision variable controlled by the firm, who trades off more data (and a better informed fixed price) against more rapid entry to a mass market having a short selling horizon. Our use of a segmented explore then exploit approach also distinguishes our work from two other streams of analytical research on using auctions for demand learning. In operations management, Pinker et al. (2007) model how observations of previous auctions can be used to inform future auctions' lot sizes. In computer science, approximation algorithms have been developed to decide prices in so-called on-line auctions whereby bidders arrive sequentially and upon bidding the firm must immediately decide, based on past bid data, whether to reject or accept the bidder's offer (Bar-Yossef et al. 2002). At the other extreme, for the case in which the firm holds a single auction in which all potential consumers in the market participate, but the firm does not know the consumer willingness-to-pay distribution, Segal (2003) creates a mechanism whereby the price for each bidder is set based on a demand distribution inferred from other bidders' bids.

In our study, the firm uses auctions for the express purpose of eliciting willingness-to-pay information from consumers. Using selling mechanisms in such marketing research experiments has the advantage that actual purchasing decisions can be tested (as opposed to hypothetical surveys). Auctions, in particular, endogenously reveal the demand level at various price levels, as opposed to simulated store experiments, which only reveal demands at whatever fixed price the experimenter has set. But, in general, studies of auction experiments (in the lab or field) to elicit consumers' willingness to pay, of particular interest to marketers, have not paid attention to issues such as data censoring due to perceived outside options, affiliated beliefs about the value of such outside options, or affiliated beliefs about the quality of the tested product itself (Harrison et al. 2004). Such issues, if not controlled with the test market design, should be addressed when reverse engineering demand information from collected data. However, from a practical point of view it is likely easier to attempt to control for these factors via a proactive approach to test market design, and this motivates the second price sealed-bid auction with rebate introduced in §3.

The second price sealed-bid auction with rebate is a first step at mechanism design for the express purpose of demand curve learning in field auctions to inform a fixed price that will be available to consumers later. To our knowledge it is the first auction mechanism for market research designed to take into account a future fixed price purchasing opportunity. That said, our results characterizing how the firm balances the tradeoffs between more data and a shorter sales horizon, i.e., the stopping time decision, are robust to whatever format of auction used, provided that the inverse problem mapping bids to willingness to pay is somehow solved. These results are related to explore and exploit stopping time literature that we discuss next.

Explore and exploit. Once a product is ready for sale, demand information can be gathered by observing actual sales under different conditions, for example, by varying posted prices (see Balvers and Cosimano 1990 and references to the early literature therein, Aviv and Pazgal 2005, Lin 2006), inventory levels (e.g., Scarf 1959, Lovejoy 1990, Lariviere and Porteus 1999, Chen and Plambeck 2005), or product assortment (Caro and Gallien 2007). Our approach differs from these dynamic learning papers. First, we use auctions as a sales channel meant to gather demand information to set a fixed price. Consumer bids in auctions offer more detail on valuations than purchase/no purchase responses to a fixed price. Furthermore, auction observations are not censored by inventory stockouts.

Second, in the demand learning literature cited above, the firm controlling pricing (or inventory) faces a tradeoff between immediate gains by myopic control and future gains by control which aids learning. For instance, higher posted prices allow the firm to test the high end of the willingness-to-pay distribution at high per-unit profits, but censor data from the low end. In our model the firm's tradeoff is similar in spirit but quite different in structure from those above. In particular, our firm faces a tradeoff between short-term losses of product launch delay and long-term gains from demand learning, where after making the decision to stop auctioning (the control decision in our model), the firm sells via fixed price and learning ceases. Thus, rather than continual learning over the short horizon by adjusting prices (or inventory) upwards and downwards in response to demand history, we adopt a segmented approach more akin to test market research methods whereby a pure explore phase (auctions) is followed by a pure exploit phase (fixed price).

In introducing a stopping time problem framework for the use of auctions for market research, our work is related to sequential learning models in economics and management science. In the economics literature, for example, studies have analyzed how the timing of investments by individual firms facing uncertainty affect aggregate innovation diffusions (Jensen 1982) and economic investment cycles (Bernanke 1983). In the management science literature, McCardle (1985) studies the adoption of a new technology option whose profitability can be learned, with the firm sequentially deciding whether to adopt, reject with a fixed payoff, or continue learning at a cost. In our setting, the firm chooses among N "options" (prices) which have uncertain but correlated value. Lippman and McCardle (1991) study a firm facing an infinite number of options of uncertain but independent value; the firm learns about each option by sequentially collecting costly information, but each option is examined one at a time and either accepted or permanently discarded. In contrast, a main benefit of using auctions is that our firm can simultaneously gather information on all N prices by examining bids. Kornish and Keeney (2005) analyze simultaneous data collection on two, independent, options, modelling the choice of which of two influenza vaccines to produce before a flu season. Our model's short selling season context implies a finite horizon model, as does Kornish and Keeney's production planning setting.

As is common in the sequential learning literature, these papers, and ours, specify threshold results of the following type: "adopt" an option if its expected value based on current information is high enough. We model the cost of information gathering as the implicit cost of delaying product launch, which in turn depends upon the evolution of demand from the moment of launch up to end of the selling horizon; accordingly, we examine how threshold results depend on demand evolution aspects, such as product diffusion (Bass 1969).

The remainder of the paper is organized as follows. The model and its assumptions are described in §2. The second price sealed-bid auction with rebate is described in §3. The main results of the paper, characterizing the firm's dynamic decision policy on when to continue the auction phase and when to commence the mass-market phase, are contained in §4. Subsection 4.2 discusses the special case when fixed-price sales follow a diffusion process per the Bass model. Section 5 concludes. Proofs of results are included in the Online Appendix.

2. The Model

We model a risk-neutral, revenue-maximizing firm selling a product over a finite horizon of length T. In this stylized model, T can be thought of as the demand window for the new product (Dolan 1993, Cohen et al. 1996), and could model the time to product obsolescence, the length of the selling season for a fashion product, or the time until low-cost imitator entry dissolves the firm's ability to charge a profitable margin for an innovative product. The firm is uncertain about consumers' willingness-to-pay, or valuation distribution, for the product. More precisely, for N possible prices ordered $p_1 < p_2 < \cdots < p_N$, the firm is uncertain of w_i, the true fraction of customers whose valuation is at least p_i. As is standard in sequential learning problems (McCardle 1985, Kornish and Keeney 2005), we assume that the true w_i's are static over the horizon, although the firm's beliefs about the w_i's may change based on its observations. Each individual consumer's valuation for the product is assumed to be private, independent, and identically distributed.

The selling horizon is divided into an auction phase and a fixed price phase. During the auction phase, the firm updates its beliefs about consumer demand based on data gathered from the auctions. The firm chooses the time at which to abandon the auction learning phase and switch to selling to the mass market via a fixed price. We begin by describing the fixed-price phase.

Fixed-price phase. Once a fixed price is set, say at p_i, at some time, $t \in [0, T]$, the mass market sales process begins. We envision this process as capturing sales in whatever fixed price retail channels the firm chooses to use (e.g., brick and mortar stores, catalogs, internet sales). We define m as the *total market potential*, that is, the size of the market segment for which the product is targeted. (For example, for SawStop's innovative safety saw, the target market could be a segment from the population of private, commercial, and educator woodworkers.) Of course, not every consumer in the targeted market would purchase the product; only those consumers whose valuation for the product equals or exceeds the price, p_i, would make a purchase. If w_i is the expected fraction of consumers willing to purchase the product at price p_i, then we say the *market potential* for p_i is $w_i m$.

When the firm sets the price at p_i, it cannot necessarily expect to sell $w_i m$ units before the end of the time horizon. The *rate of market saturation* that occurs z time units after market entry is denoted by $a(z)$, where $\int_{z=0}^{T-t} a(z)dz \leq 1$ represents the total fraction of the market that can be served over horizon $[t, T]$, for t the market entry time and T the end of the horizon. For example, $a(z) = \lambda \leq 1/T$ for $z \in [0, T]$ corresponds to constant market saturation rate of λ at the fixed price channel(s) – stores, web sites, etc. – yielding an expected sales rate of $w_i m \lambda$. While $a(z)$ could describe a constant expected rate of sales, $a(z)$ could also be a more general function of time; one such model, empirically shown to have descriptive and predictive power for the growth and decline of new product market saturation rates over time, is the Bass diffusion model (Bass 1969, Mahajan et al. 2000). For generality, we will leave the form of $a(z)$ unspecified. One exception is §4.2, an illustrative section which applies our results to the case of Bass diffusion. The total number of sales expected over the horizon $[t, T]$ is $w_i m \int_{z=0}^{T-t} a(z)dz$. However, a discount rate of r is applied to future revenue. For convenience we define the *market size*

$$M(t) \triangleq m \int_{z=0}^{T-t} a(z)e^{-rz}dz,$$

which can be thought of as the discounted revenue the firm would gain from time t until the end of the horizon if every arriving customer purchased the product for \$1. Given p_i and expected purchase fraction w_i, the expected discounted future revenue accruing to the firm upon stopping the auction phase and setting a fixed price at time t is equal to $p_i w_i M(t)$.

Auction phase. The auction phase starts at the beginning of the horizon, i.e., at time $t = 0$. During this phase, the firm only sells the product via fixed-duration auctions. We normalize T such that an auction lasts 1 period. The demand data gathered during an auction which begins at time $t-1$ and ends at time t is summarized by a signal, s_t, and the information captured by auctions up to time t is expressed by a sufficient statistic, $\mathcal{B}_t$, which is a function of $(s_1, s_2, \ldots, s_t)$ (see DeGroot 1970 for a discussion on sufficient statistics). The stochastic processes which generate such data are left general, with exceptions made explicit when needed (Corollaries 1 and 3, and Proposition 3). The expected fraction of customers who will purchase at prices less than or equal to p_i, conditional on sufficient statistic $\mathcal{B}_t$, is expressed by $w_i(\mathcal{B}_t) = E[w_i|\mathcal{B}_t]$. The firm's initial demand information at the beginning of the horizon is described by $\mathcal{B}_0$. For generality, we do not specify the updating scheme used by the firm, but do assume that point estimates of purchase probabilities are unbiased and take all information available at time t into account:

$$E[w_i(\mathcal{B}_{t+1})|\mathcal{B}_t] = w_i(\mathcal{B}_t) \quad \text{for all } i = 1, \ldots, N. \tag{1}$$

That is, point estimates are martingale. Note that all Bayesian updating schemes satisfy the martingale assumption, but this need not be true in general. For example, exponential smoothing, in which more recent data are given larger weights in determining beliefs, does not satisfy the martingale assumption. Section 4 provides an example using Bayesian updating and a multidimensional Beta-multinomial information structure.

We have not assumed a particular auction format or set of bidder beliefs and behaviors for the auction phase; all we assume is that, following each auction, the firm is able to update its purchase probabilities for each price based on the data signals it gathers during the auction. In §3 we propose a possible auction format, specify accompanying assumptions on bidder behavior, and discuss how the updating problem – essentially an inverse problem mapping bids to willingness to pay – could be solved. Other auction formats, appropriate to other assumptions on bidder behavior, could exist – for example, while §3 assumes rational behavior of bidders subject to that section's assumptions, other authors have examined computing willingness to pay from auction data without assuming much in the way of structural bidder rationality (e.g., Chan et al. 2007).

Firm's stopping time problem. Since the time horizon is limited, and any amount of time spent holding auctions subtracts away from the time to sell to the mass market, it is necessary to stop holding auctions at some point. The firm must choose when to stop gathering demand information via auctions and commit to a fixed price for mass-market sales. Because each auction lasts one period, we can cast the firm's problem as a discrete time stopping problem with information updating. At stage t, the firm decides between $N+1$ alternatives: either commit to one of N fixed prices, $\{p_i\}_{i=1}^N$, which has an expected discounted payoff of $p_i w_i(\mathcal{B}_t)M(t)$ over the interval $[t,T]$, or continue auctioning for at least one more period to gather additional demand information. $\delta \triangleq e^{-1r}$ is the discount applied to a payoff delayed one period into the future. Accordingly, the firm solves the following dynamic program.

$$J_t(\mathcal{B}_t) = \max\Big\{ \underbrace{\max_j\{p_j w_j(\mathcal{B}_t)M(t)\}}_{\text{set a fixed price and enter mass market}} \quad , \quad \underbrace{\delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]}_{\text{continue auction phase}}\Big\}, \tag{2}$$

with the boundary condition $J_T(\mathcal{B}_T) = 0 \ \forall \ \mathcal{B}_T$.

Several modeling assumptions tacit in the above dynamic program warrant discussion before moving on. First, the variable cost to run an auction (running the auction website) is zero. We relax this assumption in §4.3 and examine the case with nonzero auctioning costs. Second, the model assumes the product's marginal cost is zero; a positive, constant marginal cost could easily be incorporated, with the expected fixed-price profit of price p_i becoming $(p_i - [\text{marginal cost}])w_i(\mathcal{B}_t)M(t)$.

Third, we do not include inventory holding costs, and it is assumed that the firm has sufficient capacity or inventory to meet its demand. This is for simplicity, in order to focus on demand learning for a new product, rather than focusing on issues such as managing inventory before and after a new product launch. Two papers, Kumar and Swaminathan (2003) and Ho et al. (2002), have focused on the latter, with a Bass model backdrop. Both introduce optimal (or near-optimal) policies whereby a capacitated firm delays launching the product to build an initial inventory stockpile that mitigates shortages during the selling horizon. The agreement between this type of policy on the one hand, and our framework of launch delay for learning on the other, suggests that omitting capacity and inventory decisions from our analysis would not greatly jeopardize our qualitative insights about demand learning; in fact, concerns such as capacity limitations may prompt the firm to delay mass-market entry for reasons beyond demand learning alone, allowing it to gather even more demand information than it otherwise would.

Fourth, we assume that the mere presence of a test market does not appreciably affect the underlying demand landscape for the product. That is, demand information can be gathered from the test-market auctions without influencing the characteristics of the demand, namely the size $M(t)$ of the potential mass market that can be reached (i.e., given the opportunity to purchase) during $[t,T]$, and the unknown, underlying distribution of such consumers' willingness to pay. Intuitively, any test-market sales (via auctions or any other means) may drive up demand by helping to spread positive word of mouth for a product, or have an opposite effect if the product is not well received. We leave such issues to future work, and focus here on learning a latent demand landscape that is unaffected by the fact that we examine the demand with a test market. We observe that while demand-influencing effects of test markets may tend to be encouraged by a large number of test-market sales, auctions tend in the opposite direction because they sell relatively few items. Note that a single-item auction may have many bidders but still *sell* only one item.

Finally, we do not include revenue from the test market within the dynamic program. Effectively, this assumes that the test market auction sales will be dwarfed by mass-market, posted-price sales. To put this into perspective, if a firm sold 40,000 units in a year via the mass market (e.g., through a nation-wide rollout to retailers), such mass-market sales would occur at an average rate of over 100 units per day. Intuitively, such a firm is unlikely to indulge in continuing the auction phase simply to capture additional auction revenue, which comes in at a much slower rate (auctions held by eBay and other retailers can last over a day, see Lucking-Reiley 2000). Despite revenue from the auctions not being a key concern, they do offer pricing information for the many units that will be sold during the mass-market phase. In summary, we model auctions as market research opportunities prior to mass-market sales of the product. In this spirit, the auction format for market research introduced in the next section seeks to maximize information gleaned from the auctions, rather than maximizing per-item auction revenue. (In contrast, the latter is traditionally the objective of a profit-maximizing auctioneer in the auctions and operations literature.)

3. Bid Data Inverse Problem

In the above modeling section we did not specify a particular auction format nor detail assumptions on bidder behavior. Instead, the main model, and all the main analyses which follow this section (i.e., §4), are general in that they only require that after each auction, the firm is able to use data signals from the auction to update its purchase probability estimates for the N prices. The present section provides an example of how the inverse problem between bid data and willingness to pay estimates might be solved, and highlights auction design as an approach to making the inversion easier.

An updating scheme must account for the type of signals which it processes. The signals provided by the auction will depend on the auction format, bidder beliefs, and bidder behavior. In the recent auction marketing research literature, Harrison et al. (2004) pointed out the need to account for data censoring due to issues such as bidder beliefs about outside purchasing options (e.g., bidders knowing the item will eventually be sold via fixed price) or affiliated bidder beliefs about the outside option or product quality (which can influence the bidding in an open-bid, e.g., English format). One could attempt to address these issues with an elaborate updating scheme that carefully reverse engineers consumers' true willingness to pay from their observed bids, and indeed sophisticated reverse engineering has been done in studies on field data from auctions featuring confounding factors such as publicly announced reserve prices that censor the available observations (Paarsch 1997). However, another option, available if the researcher controls the auction design, is to design an auction format which makes the willingness-to-pay distribution more easily recoverable from the observed bids. This motivates the auction format discussed next, which is designed to (at least theoretically) elicit truthful bidding from consumers. Truthful bidding simplifies the firm's belief updating process by eliminating the need for it to reverse engineer bid data to recover valuations. The auction format introduced below is a first step at auction design for the express purpose of learning demand to inform a fixed price that will later be available to consumers. To our knowledge it is the first auction mechanism for market research designed to take into account a future fixed price purchasing opportunity.

Sealed second-price auction with rebate. We propose the following auction format. The auctioneer sets but does not announce a reserve price for the item. Each bidder arriving during the auction submits a single sealed bid to the auctioneer. For an auction which began at time $t-1$, the firm collects all bids which arrive between time $t-1$ and t; of these bids, the highest bidder is awarded the item and pays the second highest bid or the reserve price, whichever is higher. Once a fixed price, p, is set, the auction winner is given a rebate equal to the amount of her auction payment in excess of p. If the auction payment was less than p, no rebate is given. The intuition behind this format is as follows. The rebate against the fixed price addresses the issue of bid adjustment due to the existence of an outside option (future fixed price), and the sealed-bid format and private reserve price avoids upper and lower truncation of the received bids. As we explain below, the second price sealed-bid auction with rebate ensures truthful bidding in equilibrium, even if the bidders anticipate a future outside option of purchasing at a fixed price. Before presenting this equilibrium result, we state behavioral assumptions for this equilibrium analysis.

Assumptions for Section 3. We apply the following behavioral assumptions for the equilibrium analysis of this section. These assumptions, which are not used in §4, help formalize our example of how the bid-to-valuation inversion might be aided by auction design. We assume that individual consumers bid in (at most) one auction, they are risk-neutral, have use for a single unit, and if they fail to purchase via auction they will purchase at the fixed price provided it does not exceed their valuation. Consumers are patient, that is, indifferent between receiving the good at auction or purchasing it later at the fixed price provided the payment is the same. We envision the auction phase as likely to last only a few weeks at most, which naturally would make the time between purchase and rebate relatively short. Consumers who arrive to the auction account for the fact that the item will eventually be available to them at a fixed price. However, they view the future fixed price as exogenous, an assumption best made when many bidders participate in the auction phase, thereby drowning out any one particular bid's influence on the firm's pricing decision. Finally, we assume that the underlying willingness-to-pay distribution is the same for all consumers. This assumption means that the firm can recover mass-market willingness to pay directly from observed bids. (In practice, consumers participating in the auction phase may value the item more and wish to acquire the item as quickly as possible, or may value the item less and use the auction as a chance to own the item at a price below what they anticipate the later fixed price will be.) Under these assumptions, we have the following proposition.

Proposition 1. (Truthful bidding.) *Truthful bidding is a dominant strategy in the second price sealed-bid auction with rebate described above.*

The second price sealed-bid auction with rebate is "detail free" in the sense that bidders can compute a bidding strategy without having to estimate the underlying consumer valuation distribution, number of bidders, distribution of the fixed price, etc. This is particularly important for a new product whose consumer valuation distribution is quite possibly unknown to individual consumers. The proposed rebate is to our knowledge novel in auction design, but from an implementation standpoint is akin to standard, post-purchase price matching guarantees common in many consumer markets, including electronics, office products, appliances, and books (Jain and Srivastava 2000, Srivastava and Lurie 2001). Note that the rebate is necessary; if it were eliminated, bidders would account for the future fixed price by bidding below their true valuation (see Lemma 1 in the Online Appendix).

While the above game-theoretical analysis can be useful for a firm wishing to solve the bid-to-willingness-to-pay inverse problem, the question of how well a given auction format actually induces truthful bidding in practice is an empirical one. While consensus has not been reached in the literature on this issue (e.g., Noussair et al. 2004), one prominent alternative, truthful auction-like mechanism used widely by experimental economists eliminates the connection between payments and other bidders' bids. Called the Becker-DeGroot-Marschak (BDM) mechanism (dating back to the seminal paper in the psychology literature by these authors, Becker et al. 1964), this "auction" accepts sealed bids and generates a random selling price at which all bids exceeding it transact. In the absence of a future fixed price option, the BDM mechanism is incentive compatible, or truthful (e.g., Kagel and Roth 1995, p79). In the proof of Proposition 1, we show that, in the presence of a fixed price option (and under the assumptions of §3), the BDM mechanism remains incentive compatible if a rebate is used. Thus, the rebate approach proposed in this section preserves incentive compatibility for both second price sealed-bid (Vickrey) and BDM mechanisms, which are the two most widely studied demand-revelation mechanisms in experimental economics (Noussair et al. 2004).

4. Analysis of Stopping Time Problem

A firm armed with a method of inverting bids to willingness to pay is in a position to collect data and update its purchase probability estimates for the N prices. In general, the auction format or assumptions could require non-identity inversion mappings when recovering willingness to pay from observed auction data. However, for convenience we will refer to bids and willingness to pay interchangeably when discussing data gained during the auction phase (this interchange will only be precise in some settings, e.g., the auction format and assumptions described in §3).

Examples. We begin by providing simple, illustrative examples of updating and the stopping problem. We first illustrate how Bayesian updating converts bids to purchase probability estimates when the prior distribution of purchase probabilities comes from the multidimensional Beta family. Then, we provide a simple example in which the firm chooses to continue holding auctions rather than stopping and committing to a fixed price.

Multidimensional Beta-multinomial Bayesian updating. For a set of non-negative parameters $\gamma_0, \gamma_1, \ldots, \gamma_N$, the multidimensional Beta$(\gamma_0, \ldots, \gamma_N)$ distribution over non-negative random variables $x_0, \ldots, x_N$, $\sum_i x_i = 1$, is defined as

$$f(x_0, x_1, \ldots, x_N) = x_0^{\gamma_0 - 1} x_1^{\gamma_1 - 1} \cdots x_N^{\gamma_N - 1} \frac{\Gamma(\gamma_0 + \gamma_1 + \cdots + \gamma_N)}{\Gamma(\gamma_0)\Gamma(\gamma_1)\cdots\Gamma(\gamma_N)} \qquad \text{where} \quad \Gamma(k) = \int_0^\infty x^{k-1} e^{-x} dx.$$

For a positive integer argument k, $\Gamma(k) = (k-1)!$. The multidimensional Beta distribution can take on a vast variety of shapes, depending on the parameter values.

The multidimensional Beta distribution describes the joint probability of $N+1$ mutually exclusive events, where x_i is interpreted as the probability of the i^{th} event. For purchase probabilities $w_1 \geq w_2 \geq \cdots \geq w_N$, the value $w_i - w_{i+1}$ is the probability of the event that a random customer's valuation lies in interval $[p_i - p_{i+1})$. Suppose the initial joint prior distribution of the event probabilities for the $N+1$ intervals $[0, p_1), [p_1, p_2), \ldots, [p_N, \infty)$ follows Beta$(\gamma_{0,0}, \gamma_{1,0}, \ldots, \gamma_{N,0})$,

$$f(w_1, w_2, \ldots, w_N) = \\ (1 - w_1)^{\gamma_{0,0} - 1} (w_1 - w_2)^{\gamma_{1,0} - 1} \cdots (w_{N-1} - w_N)^{\gamma_{N-1,0} - 1} w_N^{\gamma_{N,0} - 1} \frac{\Gamma(\gamma_{0,0} + \gamma_{1,0} + \cdots + \gamma_{N,0})}{\Gamma(\gamma_{0,0})\Gamma(\gamma_{1,0})\cdots\Gamma(\gamma_{0,N})}.$$

The parameters of this distribution comprise the initial sufficient statistic, in other words, $\mathcal{B}_0 = (\gamma_{0,0}, \gamma_{1,0}, \ldots, \gamma_{N,0})$. The expected value of w_i based on this initial prior is $w_i(\mathcal{B}_0) = \sum_{j \geq i} \gamma_{j,0} / \sum_j \gamma_{j,0}$. We next discuss how the prior is updated based upon the number of bids observed in each of the $N+1$ price intervals. (Bayesian updating with a multidimensional Beta prior is also discussed in Silver 1965.)

Let b_t be the number of unique bids (observations) received during the t^{th} auction, lasting from $t-1$ to t. For $i = 1, \ldots, N-1$ let $b_{i,t}$ be the number of such bids that are in interval $[p_i, p_{i+1})$, with $b_{0,t}$ the number of bids in $[0, p_1)$ and $b_{N,t}$ the number in $[p_N, \infty)$. The true probability that a bid will lie in $[p_i, p_{i+1})$ is $w_i - w_{i+1}$. In other words, $(b_{0,t}, \ldots, b_{N,t})$ follows a multinomial distribution with $N+1$ different event possibilities and b_t trials.

After the t^{th} auction, the parameters of the prior distribution are updated. For each i, parameter $\gamma_{i,t} \triangleq \gamma_{i,t-1} + b_{i,t}$. That is, the number of observations made within the i^{th} price interval is added to the i^{th} interval's parameter. The new sufficient statistic is $\mathcal{B}_t = (\gamma_{0,t}, \gamma_{1,t}, \ldots, \gamma_{N,t})$. At stage t, the expected value of w_i, based on the current prior, is $w_i(\mathcal{B}_t) = \sum_{j \geq i} \gamma_{j,t} / \sum_j \gamma_{j,t}$. This estimate for the fraction of customers who will purchase at price p_i can essentially be thought of as the ratio of the number of "successes" (valuations observed in interval $[p_i, \infty)$), to the total number of observations, starting from an initial set of parameters $\mathcal{B}_0$.

Simple stopping problem. Assume that priors are updated according to the multidimensional Beta-multinomial information structure. Suppose $T = 2$, with $M(0) = 1000$, $M(1) = 970$, and $M(2) = 0$. The firm has an initial sufficient statistic $\mathcal{B}_0 = (0.5, 3, 1.5)$ with the following prices and priors: $p_1 = 14$, $p_2 = 32$, $w_1(\mathcal{B}_0) = 0.6$, $w_2(\mathcal{B}_0) = 0.3$. If the number of arrivals per auction is deterministic and set to 1, and $\delta = 1$, then at time $t = 0$, the expected value of continuing for one more auction is

$$\begin{aligned}
&E\left[\max\left\{p_1 w_1(\mathcal{B}_1) M(1),\ p_2 w_2(\mathcal{B}_1) M(1)\right\} \middle| \mathcal{B}_0\right] \\
&= \ M(1)\Big((1 - w_1(\mathcal{B}_0)) \max[7.00, 8.00] + (w_1(\mathcal{B}_0) - w_2(\mathcal{B}_0)) \max[9.33, 8.00] \\
&\quad + w_2(\mathcal{B}_0) \max[9.33, 13.33]\Big) \ = \ 9,700.
\end{aligned}$$

We have used the fact that $M(2) = 0$ and it is always optimal to stop if $t = 1$. On the other hand, if the firm stops at $t = 0$, the expected profit is $\max[p_1 w_1(\mathcal{B}_0) M(0),\ p_2 w_2(\mathcal{B}_0) M(0)] = 9,600$. In this case, at $t = 0$, it is optimal to continue the auction for one more period. A necessary, but not sufficient, condition for continuing the auction phase is that additional information can alter the optimal price choice. In this example, the valuation of the single additional observed bid is either below p_1, between p_1 and p_2, or above p_2, and the corresponding optimal price for each possibility is p_2, p_1, and p_2, respectively. In general, the incremental benefit of choosing a better price must be tempered with the loss in market size (e.g., $M(0)$ vs. $M(1)$).

4.1 Structure of the Stopping Problem

We now shed light on when is it optimal to stop the auction phase and turn to a fixed price. We begin by proving structural properties of the stopping time problem, and then characterize how the dynamics of mass-market sales affect the firm's decision.

Given the unrestricted richness of the sufficient statistic, $\mathcal{B}_t$, it would be nice to find a simple 1-dimensional statistic of $\mathcal{B}_t$ which admits a threshold optimal stopping policy. For example, the expected fraction of customers willing to purchase at price p_i, $w_i(\mathcal{B}_t)$, is a 1-dimensional statistic. However, given the sheer volume of potential histories, it is possible that vastly different sufficient statistics share the same purchase probabilities but have different stopping policies. For example, it could be that $w_i(\widehat{\mathcal{B}}_t) = w_i(\mathcal{B}_t)$, but $\widehat{\mathcal{B}}_t$ describes a history with 10 data points for which gathering more data is optimal, while $\mathcal{B}_t$ is describes a history with 10,000 data points for which new data is likely to have little impact on priors and stopping is optimal. Thus some common ground between histories, or their sufficient statistics, must be imposed before a 1-dimensional statistic can admit a stopping threshold. To this end, define a set of sufficient statistics at time t by Ω_t, which, for example, could be the set of sufficient statistics having 10 data points. If certain behavior over the sufficient statistics in set Ω_t exist, a threshold result can obtain. Below we establish conditions that ensure a stopping policy threshold exists over the $w_i(\mathcal{B}_t)$'s. In the proposition, $\mathcal{B}_{t+1}$ and $\widehat{\mathcal{B}}_{t+1}$ are the stage $t+1$ updates of stage t sufficient statistics $\mathcal{B}_t$ and $\widehat{\mathcal{B}}_t$, respectively. The term $\mathbf{w}_{-i}(\mathcal{B}_t) = (w_1(\mathcal{B}_t), \ldots, w_{i-1}(\mathcal{B}_t), w_{i+1}(\mathcal{B}_t), \ldots, w_N(\mathcal{B}_t))$ is the vector of point estimates excluding that of price i, and $\mathbf{p} = (p_1, p_2, \ldots, p_N)$ denotes the vector of prices.

Proposition 2. (Effect of point estimates.) *Let Ω_t be a set of sufficient statistics such that, for any $\mathcal{B}_t$, $\widehat{\mathcal{B}}_t \in \Omega_t$, if $w_i(\widehat{\mathcal{B}}_t) \geq w_i(\mathcal{B}_t)$ and $\mathbf{w}_{-i}(\widehat{\mathcal{B}}_t) = \mathbf{w}_{-i}(\mathcal{B}_t)$ then*

$$E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t] \geq E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t], \tag{3}$$

for any bounded, nondecreasing function ϕ. Then for any $\mathcal{B}_t \in \Omega_t$ there exist thresholds $h_{it}(\mathbf{w}_{-i}(\mathcal{B}_t), \Omega_t, \mathbf{p})$, $i = 1, \ldots, N$, such that if $w_i(\mathcal{B}_t) \geq h_{it}(\mathbf{w}_{-i}(\mathcal{B}_t), \Omega_t, \mathbf{p})$ it is optimal to stop the auction and commit to price p_i, and if $w_i(\mathcal{B}_t) < h_{it}(\mathbf{w}_{-i}(\mathcal{B}_t), \Omega_t, \mathbf{p})$ for all i, it is optimal to continue the auction phase. The thresholds, $h_{it}(\mathbf{w}_{-i}(\mathcal{B}_t), \Omega_t, \mathbf{p})$, are nondecreasing in p_j and $w_j(\mathcal{B}_t)$, $j \neq i$.

To interpret Proposition 2 it is helpful to consider the following example, stated as a corollary. Let b_{t+1} denote the number of bidders in the $t+1^{\text{st}}$ auction.

Corollary 1. *The results of Proposition 2 apply to the case where priors are updated using the multidimensional Beta-multinomial information structure and Ω_t is any set of sufficient statistics containing the same number of observations, if $[b_{t+1}|\mathcal{B}_t]$ and $[b_{t+1}|\widehat{\mathcal{B}}_t]$ have the same distribution for all $\mathcal{B}_t, \widehat{\mathcal{B}}_t \in \Omega_t$.*

By restricting to sufficient statistics having the same number of observations, the effect of the purchase probability point estimates on the stopping decision can be characterized. Corollary 1 states that, for the multidimensional Beta-multinomial information structure, if the firm finds it optimal to stop and choose price p_i at a certain history of S observations, then it is also optimal to stop and choose price p_i at any other S-observation history whose purchase probability point estimates are at least as large for price p_i but no larger for the other prices. To keep the comparisons between different histories apples-to-apples, the corollary requires that the number of bids received in the $t+1^{\text{st}}$ auction should not depend on the particular S-observation history at time t. (As just one example, this would hold for Poisson bidder arrivals even with an arrival rate that depends on t and S.) The upshot of Corollary 1 is that, while it is still possible that the firm chooses "incorrectly" (the optimal decision under limited information may not be that taken under full information), as it gathers more demand information the firm will eventually make a choice once a particular pricing option looks particularly dominant.

This threshold structure of Proposition 2 is illustrated in Figure 1 for a case in which the firm is deciding between just two prices. For readability, where the meaning is clear, we suppress the dependence of $w_i(\mathcal{B}_t)$ on $\mathcal{B}_t$. For values of (w_1, w_2) near the origin, the firm's optimal decision is to continue the auction phase. Holding w_2 fixed, when w_1 increases (moves to the right), the firm continues to prefer perpetuating the auction phase until w_1 becomes so large that it reaches the threshold $h_{1t}(w_2, \Omega_t, \mathbf{p})$, which is the boundary separating the 'continue the auction phase' region from the 'stop and choose p_1' region. At this point the firm prefers stopping and selecting price p_1, and the firm continues to prefer this for any further increase in w_1. This occurs even though continuing the auction phase and gathering more bid data would allow the firm to update its priors on both prices simultaneously. Finally, note that Proposition 2 does not guarantee uniqueness of an optimal price once the decision to stop has been made. The choice of price will be such that the value of $p_i w_i M(t)$ is maximized; it is conceivable that many choices of price may give the same maximal value, as represented by the diagonal line $w_1 p_1 = w_2 p_2$ dividing the stopping regions in Figure 1.

At its core, Proposition 2 examines the tradeoff between more information and faster entry to market. Condition (3) in the proposition is used to show that the more optimistic the firm is today about the market, the less the firm expects to regret entering. The underlying martingale assumption made on the updating process (equation (1)) also plays a role in Proposition 2, as it helps to ensure that it is better to stop today and set the price at p_i rather than wait until tomorrow and set the price at the *same* p_i. Proposition 2's result is similar in spirit to threshold results in other sequential learning studies (e.g., McCardle 1985, Kornish and Keeney 2005), although Proposition 2 treats an N-dimensional case, all N payoff options are updated in every period, and the options' payoffs are correlated rather than assumed to be independent.

At this point it is worth fleshing out the reason why thresholds in Proposition 2 depend on the set of histories, Ω_t. For this purpose, we use multidimensional Beta-multinomial updating as a backdrop. Consider a setting in which the firm is deliberating between two prices, as depicted in

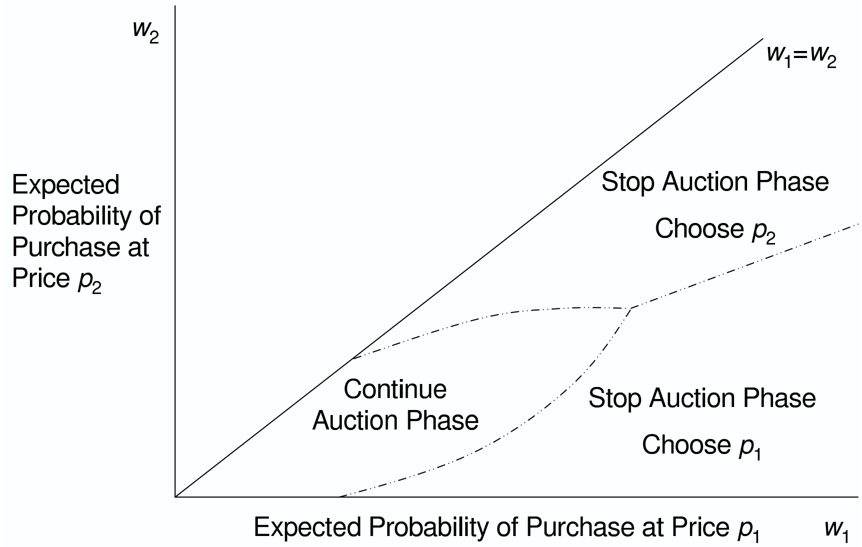

Figure 1: Structure of the stopping problem when there are two prices ($N = 2$).

Figure 1. To gain intuition as to why the stopping decision depends on the set of sufficient statistics in addition to the purchase probability point estimates, consider two extreme cases, where Ω_t and Ω'_t are sets of sufficient statistics with 3 and 90 observations, respectively. Suppose $w_1(\mathcal{B}_t) = 2/3$ and $w_2(\mathcal{B}_t) = 1/3$. One might ask, for the stopping time decision, is it important whether or not $\mathcal{B}_t \in \Omega_t$ or Ω'_t? Consider what happens to the priors after one additional bid is received. If $\mathcal{B}_t \in \Omega_t$, the updated point estimates will be $(w_1(\mathcal{B}_{t+1}), w_2(\mathcal{B}_{t+1})) = (\frac{2}{4}, \frac{1}{4})$, or $(\frac{3}{4}, \frac{1}{4})$, or $(\frac{3}{4}, \frac{2}{4})$, each with probability 1/3, depending upon which interval, $[0, p_1)$, or $[p_1, p_2)$, or $[p_2, \infty)$, contains the new bid. On the other hand, if $\mathcal{B}_t \in \Omega'_t$, the updated priors will be $(w_1(\mathcal{B}_{t+1}), w_2(\mathcal{B}_{t+1})) = (\frac{60}{91}, \frac{30}{91})$, or $(\frac{61}{91}, \frac{30}{91})$, or $(\frac{31}{91}, \frac{61}{91})$, each with probability 1/3. Clearly, the number of preexisting observations (3 versus 90) has a dramatic effect on the expected change in priors caused by continuing the auction phase. In particular, it can be shown that the expected benefit of gathering exactly one additional auction's worth of data is always decreasing in the number of previous observations, for the multidimensional Beta-multinomial structure (see Proposition 8 in the Online Appendix).

Next, it is shown how the thresholds described in Proposition 2 are affected by the shape of the market size function, $M(t)$, as time elapses.

Proposition 3. *Suppose the stochastic processes generating auction data are stationary, i.e.,* $[\mathcal{B}_{t+k}|\mathcal{B}_t = \mathcal{B}]$ *has the same distribution as* $[\mathcal{B}_{s+k}|\mathcal{B}_s = \mathcal{B}]$ *for all* s, t, k.

1. **(Longer remaining horizon can encourage auctioning.)** *If the market size function,* $M(t)$*, is log-concave over* $[t-1, T-1]$*, then if the optimal decision at time* t *is to continue auctioning, the optimal decision under the same priors (sufficient statistic) at time* $t-1$ *is also to continue auctioning.*

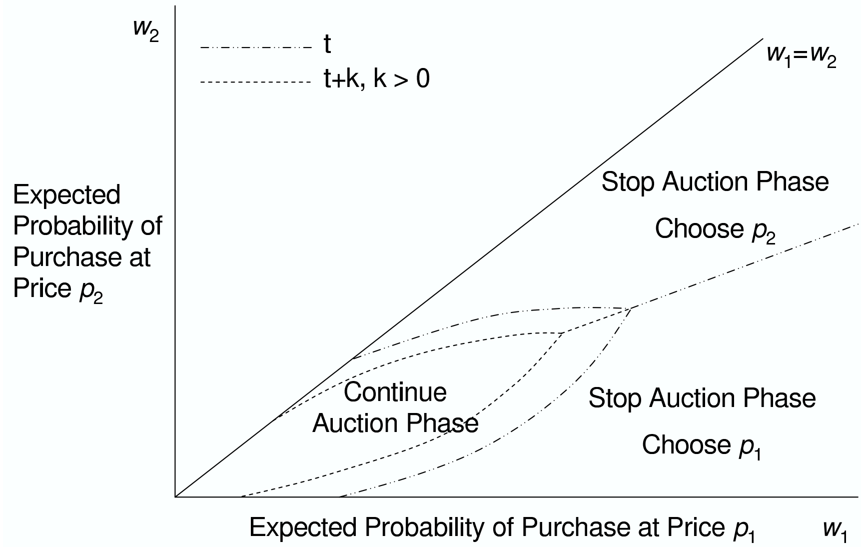

Figure 2: For a given set of sufficient statistics Ω, the continuation region shrinks as time elapses if $M(t)$ is log-concave. Below the diagonal, the $(- \cdot -)$ lines are $h_{it}(w_{-i}(\mathcal{B}), \Omega, \mathbf{p})$, and the $(\cdots)$ lines are $h_{i(t+k)}(w_{-i}(\mathcal{B}), \Omega, \mathbf{p})$, for $i = 1, 2$. For readability, dependence on sufficient statistic $\mathcal{B} \in \Omega$ is suppressed in the figure labels.

2. **(Longer remaining horizon can discourage auctioning.)** *If the market size function, $M(t)$, is log-convex over $[t-1, T-1]$, then if the optimal decision at time $t-1$ is to continue auctioning and with probability one the firm will enter before time $T-1$, the optimal decision under the same priors (sufficient statistic) at time t is also to continue auctioning.*

Proposition 3 shows that, in the firm's stopping time decision, the amount of time remaining in the horizon plays a key role. Part 1 says that, with a log-concave market size function, the more time remaining in the horizon, the more the firm is willing to delay the mass market entry in order to gather data. If a firm prefers to continue auctions with a given set of priors (or sufficient statistic) and remaining horizon, the firm would also prefer to continue auctioning with the same priors and more time remaining. As Figure 2 illustrates, the stopping regions become larger as the remaining time horizon shrinks. Part 1 requires log-concavity of the market size function $M(t)$. This assumption holds if $M(t)$ is concave, which in turn holds if the rate of market saturation $a(z)$ is (weakly) decreasing in z, the time since mass market launch.

In contrast, if the rate of market saturation is highly convex in the time from mass market launch, early market entry is extremely valuable to the firm, in order to give the saturation curve time to pick up speed and make large gains late in the horizon. This is the intuition behind part 2 of Proposition 3, which says that with a log-convex market size function, the decision to continue auctioning with a long time horizon would, under the same priors, continue to be optimal for a shorter horizon. That is, under a highly convex market size function, the stopping regions can shrink even as the horizon becomes shorter — late in the horizon there is so little time to ramp up sales that the firm has less to lose by delaying entry, and accordingly is more willing to continue auctioning.

Proposition 3 makes use of assumptions on the auction data process and the firm's entry decisions. First, the stochastic processes generating auction data are time insensitive, meaning that auctions are assumed to generate the same types and amounts of data, regardless of when they take place. This would not be the case, for example, if the firm knew bidding traffic would be very slow in early periods but increase dramatically for later periods. If auctioning offsets the cost of waiting only if traffic is high (i.e., auctioning can be beneficial only when the horizon is short but not when it is long), this would violate part 1 of the proposition if the market size is log-concave. Conversely, if traffic drops off in later periods, the stopping region might grow over time even for a log-convex market size function. Finally, part 2 requires that the firm wishing to continue auctioning at time $t-1$ enters the market by time $T-2$. While technical, this assumption is not unreasonable; if $M(t)$ is log-convex, the firm's market size decreases convexly as the firm delays its entry time towards T. To see why the entry assumption is important, suppose the firm continues at time $t-1$ due primarily to the profitable prospect of gathering exactly $T-t$ periods of data before entering. For a firm with the same priors but facing a horizon one period shorter, gathering $T-t$ periods of data would take it to the end of the horizon, at which point no sales can be made, lessening the appeal of continuing to auction at time t.

In summary, a longer time horizon can encourage or discourage furthering the auctioning phase, depending on the shape of the market size function. A log-concave market size function implies that as the remaining horizon shrinks, *more* sales would be lost by postponing entry to the market, and the impetus to enter the market becomes larger. In contrast, delay becomes *less* costly as the remaining horizon shrinks under a log-convex market size function, as the number of sales that delaying market entry sacrifices decreases as time elapses. The managerial insight is that, while a longer time horizon might naturally be seen to invite more data gathering under an optimal auction market research strategy, this intuition is sensitive to the trajectory of sales in the mass market.

Up to this point we have left the form of the market size function, $M(t)$, unspecified. However, this function could depend on any number of parameters. For example, the following subsection discusses the case in which mass market sales follow a Bass diffusion, where the relevant parameters capture the effects of two customer groups: innovators, who purchase based solely on external influences (e.g., mass market advertising), and imitators, who purchase based solely on internal factors (e.g., recommendations of past purchasers). Other parameters could model influences such as the number of stores willing to stock the product, or the geographic footprint of the firm's marketing campaign. The next result describes how the firm's stopping time decision can be sensitive to any such parameter, which we label as x.

Proposition 4. (Effect of market size parameters.) *Let x be some parameter of the market size function, $M(t)$, and let $I \subseteq \mathbb{R}$ be a subset of the real line. If $\frac{\partial}{\partial x}\left(\frac{M(s)}{M(t)}\right)$ is continuous and $\frac{\partial}{\partial x}\left(\frac{M(s)}{M(t)}\right)|_x \geq 0$ (≤ 0) $(= 0)$ for all $s \geq t$ and $x \in I$, the stopping region shrinks (grows) (does not change) as the market size parameter, x, increases within set I.*

Proposition 4 describes how the firm's market entry decision depends on parameters that influence the market size. The proposition implies that, if $M(s)/M(t)$ is monotonic in x for all $s \geq t$, the firm's stopping time decision is of a threshold type in the market parameter, x. For example, if $M(s)/M(t)$ is monotonically increasing in x, then holding t, $\mathcal{B}_t$, and all other parameter values fixed, once the firm prefers market entry with parameter value x_0 it continues to prefer market entry for any parameter value larger than x_0.

The derivative condition describes how changing x affects the percent of market size lost by delaying entry. Intuitively, a parameter's impact on the stopping decision is affected by both today's market size *and* future market sizes. A main insight of the proposition is that ratios of the market sizes can be used to describe the effect on the stopping decision. To put the proposition in these terms, for any given parameter x that affects the market size, a larger value of x encourages market entry if increasing x always boosts the percentage of market size that would be lost by postponing entry. The opposite conclusion holds if larger x instead always reduces the percent of market size sacrificed by delaying entry, in which case increasing x favors further auctioning.

Proposition 4 may appear to require rather restrictive assumptions on the behavior of $M(t)$. However, the proposition's conditions can be established for natural market models and parameters. The results in the next section show precisely this for a market diffusion model, for which the parameters in question are coefficients of innovation and imitation. But first, a very simple application of Proposition 4 directly utilizes the fact that $M(t)$ is proportional to the total market potential, m.

Corollary 2. (Effect of total market potential.) *The firm's stopping decision does not depend on the particular size of the total market potential, m.*

In other words, Corollary 2 says that without loss of generality, the firm can ignore the particular size of m in its stopping time analysis since m is a scaler in the market size function, $M(t)$. This follows from Proposition 4, which says that the market size parameters affect the firm's stopping time behavior based only the ratio of current and future market sizes. While revenues depend strongly on the total possible market size m, the stopping decision depends on factors such as point estimates, time, and the shape of the sales trajectory (market size function shape) over time. Of course, this could change if fixed costs are associated with holding auctions; see §4.3 where we show that, under nonzero auctioning costs, the stopping regions would shrink as m grows.

4.2 Diffusion Sales Process

A widely applied model of new product adoption is the diffusion model first pioneered by Bass (1969) and since applied in many studies of new product introduction, forecasting, and sales management (Mahajan et al. 1990). The Bass model presents new product adoption as a diffusion process, akin to that of a contagious agent. In the model, there are two types of customers: innovators, who purchase based solely on external influences (e.g., mass market advertising), and imitators, who purchase based solely on internal factors (e.g., recommendations of past purchasers). In this subsection we will assume that the rate of market saturation can be described by the following differential equation, where A is the antiderivative of a:

$$a(s) = \frac{dA(s)}{ds} = \alpha(1 - A(s)) + \beta A(s)(1 - A(s)). \tag{4}$$

Here α is the coefficient of innovation and β is the coefficient of imitation as described by Bass. To solve (4), we set a boundary condition $A(0) = 0$. This accounts for the assumption (discussed on page 9) that the number of sales during the auction phase is negligible, and is consistent with imitation effects being insensitive to unsuccessful purchase attempts (Kumar and Swaminathan 2003). With this boundary condition, the diffusion equation becomes $A(s) = \left(1 - e^{-(\alpha+\beta)s}\right) / \left(1 + \beta e^{-(\alpha+\beta)s}/\alpha\right)$. Product adoptions, which follow trajectory $A(s)$, are taken to occur within the population of potential purchasers (e.g., Mahajan and Peterson 1978, Kalish 1985, Kalish and Lilien 1986). If the size of this population for fixed price p_i is $w_i m$, and mass-market sales begin at time t, then $w_i m A(T - t)$ is the expected number of sales for the horizon $[t, T]$. For the sake of tractability, we will ignore discounting, and set $M(t) = mA(T - t)$ for $t \leq T$. That is,

$$M(t) = \begin{cases} \frac{m\left(1-e^{-(\alpha+\beta)(T-t)}\right)}{1+\frac{\beta}{\alpha}e^{-(\alpha+\beta)(T-t)}} & \text{for } t \in [0, T] \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

During a diffusion process, cumulative sales increase in the time from market launch. An example of the diffusion process is shown in Figure 3. How the shape of the sales trajectory depends on α and β can be understood by concentrating on each type of sales individually. Innovative sales increase at a decreasing rate (concavely) in the time since market launch, since their growth rate is proportional to the size of the untapped market (the first term on the righthand side of (4)). In contrast, the imitative sales rate is proportional to cumulative sales *and* the size of the untapped market, as captured by the second term on the righthand side of (4). When the size of the untapped market is still large, imitative sales grow at an increasing rate (convexly) as more cumulative sales spark more imitative purchases, analogous to the initial phase of an epidemiological outbreak. The convex imitative sales growth is only temporary, however, as a dwindling untapped market size inevitably drags the growth rates of both imitative and innovative sales to zero.

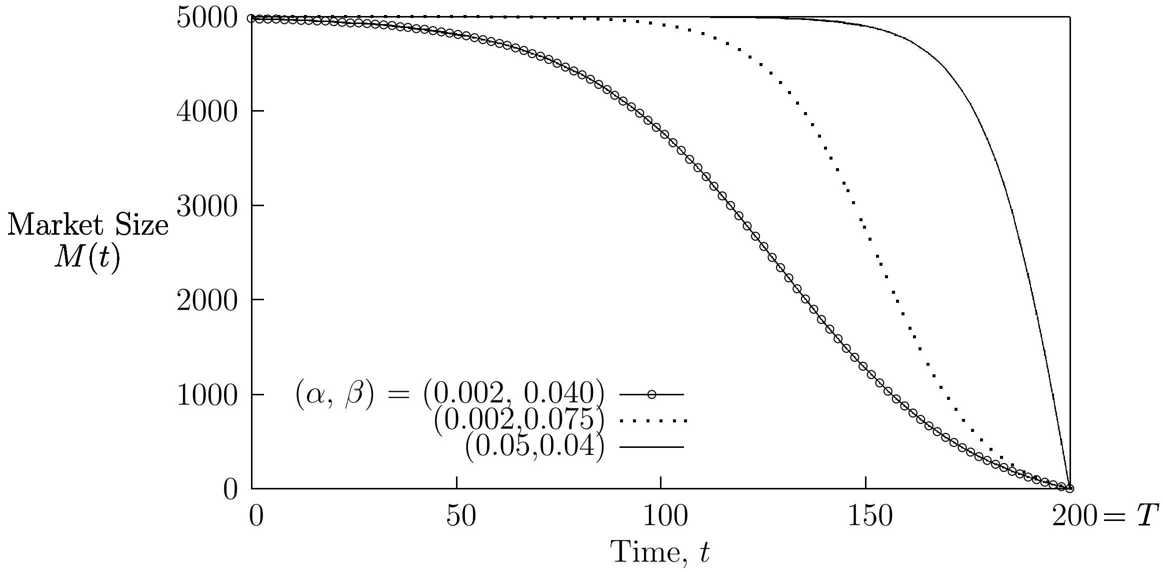

Figure 3: Sensitivity of sales trajectory to market parameters. A larger coefficient of imitation, β, enhances convexity in the sales trajectory due to "outbreak" effects. If the coefficient of innovation, α, outweighs the coefficient of imitation ($\alpha > \beta$), the sales trajectory is everywhere concave.

We next see how Propositions 2, 3, and 4 play out when sales follow a Bass diffusion. Since Proposition 2 does not depend on the market size function, its result clearly applies to the Bass diffusion setting. That is, when sales follow a Bass diffusion, stopping decision thresholds over purchase probability point estimates exist for updating structures satisfying the conditions of Proposition 2 (e.g., multidimensional Beta-multinomial updating, per Corollary 1). Our next result applies Proposition 3 to characterize how these thresholds change with time.

Corollary 3. (Longer remaining horizon encourages auctioning for Bass sales diffusion.) *Let the market size be given by* (5)*, and suppose the stochastic processes generating auction data are stationary. If the optimal decision at time t is to continue auctioning, then the optimal decision at $t-1$ under the same sufficient statistic is also to continue auctioning.*

Corollary 3 says that a longer time horizon always favors further auctioning when sales follow a Bass diffusion and auction data processes are stationary. While the imitative effect can initially cause pronounced convexity in the sales curve (see Figure 3), no matter how large β is, the market size curve is still log-concave. That is, it never becomes "too convex," and a longer horizon always makes auctioning more desirable. Next, we further explore how the shape of the market size function impacts the firm's stopping decision. The following results apply Proposition 4 to describe how the stopping regions are affected by changes in the coefficients of innovation and imitation. First, we see that increasing the innovation coefficient always encourages delaying market entry.

Corollary 4. (Innovator effect encourages auctioning.) *Suppose the market size, $M(t)$, is given by* (5)*. The stopping region always shrinks as the coefficient of innovation, α, increases.*

Corollary 4 further characterizes how the shape of the sales curve affects the firm's stopping time (market entry) decision. It says that for any coefficients of imitation and innovation, and for any time t and sufficient statistic, $\mathcal{B}_t$, the stopping region of Proposition 2 shrinks with the coefficient of imitation, α. Interestingly, increasing the coefficient of imitation can have the exact opposite effect, as the next result shows.

Corollary 5. (Imitator effect can discourage auctioning.) *Suppose the market size, $M(t)$, is given by* (5)*, and let coefficient of imitation $\beta = \beta_0$. There exists a $t_0 \leq T$ such that for $t \geq t_0$, there exists a $\delta > 0$ such that the stopping region grows in β for $\beta \in (\beta_0 - \delta, \beta_0 + \delta)$. Furthermore, there exists α_0 such that $\alpha < \alpha_0$ implies $t_0 = 0$.*

The first part of Corollary 5 states that when we are sufficiently close to the end of the time horizon ($t \geq t_0$), the stopping region grows with the coefficient of imitation, as long as the coefficient of imitation is near the original value, β_0. The second part of the corollary states that the above property holds for all $t \in [0, T]$ when the coefficient of innovation is sufficiently small, relative to β_0. The key to Corollaries 4-5 lies in how strengthening innovation or imitation rates change the relative market loss caused by delaying entry, compared to immediate entry. For 'innovative' consumers, the innovative sales rate is proportional to only the size of the untapped market, which decreases as sales accumulate. Thus, the more sales can exhaust the potential market before reaching the end of the horizon (the higher the coefficient of innovation), the smaller the relative sales loss caused by delaying entry.

In contrast, the rate of imitative sales is proportional to cumulative sales as well as the size of the untapped market. The "contagion" effect initially causes the imitative sales rate to increase rapidly. During this phase, delaying entry reduces the snowball effect of sales. If sales do not have enough time to saturate the market before reaching the end of the selling horizon, under stronger contagion effects (higher coefficient of imitation), delaying market entry results in a larger loss relative to entering immediately. This is the case if the coefficient of innovation is very small (making the initial ramp-up very slow), or the time horizon is short ($t \geq t_0$). Like a sudden epidemic outbreak that wanes quickly due to a lack of susceptible individuals, with a long enough selling horizon or strong enough imitative or innovative effects, the market eventually becomes saturated before the end of the horizon. In this case, the more sales that accumulate before reaching the end of the horizon (the higher the coefficient of imitation), the smaller the relative sales loss caused by delaying entry. This explains why Corollary 5 only applies locally near β_0.

4.3 Nonzero Auctioning Costs

In this subsection we explore the case in which the variable cost to run auctions and gather bid data during a single period is c, paid at the end of each auction period. The dynamic program formulation (2) changes in the natural way, adding $-\delta c$ to the last term in the maximization (equation (A.21) in the Online Appendix). We have the following results.

Proposition 5. (Effects of point estimates and remaining time under nonzero auctioning costs.) *For the firm's stopping time decision under nonzero auctioning costs, Propositions 2 (thresholds in purchase probability point estimates) and part 1 of Proposition 3 (stopping regions grow with remaining time for log-concave market size function) hold as before.*

However, Proposition 4 (the effect of market size parameters) must be slightly changed to accommodate the effect of nonzero auctioning costs.

Proposition 6. (Effect of market size parameters under nonzero auctioning costs.) *Let x be some parameter of the market size function $M(t)$, and $I \subseteq \mathbb{R}$. If $\frac{\partial}{\partial x}\left(\frac{M(s)}{M(t)}\right)$ and $\frac{\partial}{\partial x}M(t)$ are continuous, and $\frac{\partial}{\partial x}\left(\frac{M(s)}{M(t)}\right)\big|_x$ and $\frac{\partial}{\partial x}M(t)\big|_x$ are both ≥ 0 (≤ 0) ($= 0$) for all $s \geq t$ and $x \in I$, then the stopping region shrinks (grows) (does not change) as the market size parameter, x, increases within set I.*

Proposition 6 immediately implies that stopping regions shrink with the total potential market size m, which scales $M(t)$. The managerial insight here is that, when auctions are costly to operate, both shape *and* magnitude of the market size function are important. Intuitively, the magnitude is important to help offset the variable cost of running auctions, and shape is important to ensure that not too much of the market is sacrificed by delay. The importance of magnitude is what prevents part 2 of Proposition 3 from holding as before; we can have cases where, even with a log-convex market size function, the stopping region does not shrink with time. If the market size dips far below the auctioning cost, say $p_N M(t+1) < c$, then no matter what the shape of $M(t)$, continuing the auction phase at $t-1$ will not imply that under the same sufficient statistic continuing at t would be optimal.

Finally, as auctions become more costly, the firm naturally becomes more inclined to enter the mass market sooner:

Proposition 7. (Costlier auctioning encourages mass-market entry.) *The stopping region grows as c, the variable cost to run auctions and gather bid data during a single period, increases.*

5. Conclusions

This paper presents a framework and analysis for deploying online auctions as a demand-learning tool. By using online auctions for initial sales of a new product for which demand is uncertain, information on consumer willingness to pay can be gathered before a fixed price is set for the mass market (e.g., posted-price retailers). We study how long a firm should delay market entry in order to learn demand information: On one hand, delaying mass-market entry allows demand learning and enables a more effective pricing decision. But on the other hand, a finite sales horizon means that launch delays subtract from the time available to make sales in the mass market. We focus on closed-form characterizations of insights, in order to build intuition on how online auctions might best be used in this novel context.

The firm's decision about when to enter the mass market is modeled as a dynamic program, one trading off better demand information versus faster market entry. The two key elements of the model – the demand information accumulation process (i.e., auction format and associated information updating structure), and the mass-market sales process (i.e., shape of the mass-market sales trajectory), are left general, provided the information updating process is martingale. General insights are derived, and applied to infer implications for specific instances (such as Bayesian information updating, or diffusion-based market saturation processes).

The demand information gathering process informs the firm's predictions about the consumer's probability of purchase at N different price points. There is an appealing structure for the optimal policy: once the expected return of one price p (essentially, price times probability of purchase) emerges as sufficiently attractive relative to all other prices, the firm sets p as the fixed price and enters the mass market. This is captured by Proposition 2. The upshot is that, while it is still possible that the firm chooses "incorrectly" (the optimal decision under limited information may not be the same as under full information), as more demand information is gathered the firm will eventually make a choice once there is a pricing option that looks particularly dominant.

While the firm waits to discover which price becomes particularly attractive, it must also weigh the implications of the sales trajectory pattern for the mass market. When sales are anticipated to "snowball" dramatically in the time from mass-market launch, the firm may wish to forego prolonged demand information gathering and simply jump into the mass market quickly in order to allow sales sufficient time to ramp up prior to the end of the horizon. However, when sales growth is expected to need less ramping-up, or even decrease over time, the firm has less to lose by delaying its mass-market launch and auctioning for demand learning is more attractive. In fact, this latter case prevails for the classic and well-tested Bass model of new product diffusion, suggesting that for practical applications auctioning at the outset of a short selling horizon may indeed be attractive. These results are captured by Proposition 3 and Corollary 3. Additional results, Proposition 4 and Corollaries 4-5, characterize how entry decisions depend on other factors of the market trajectory. For example, when applied to the Bass model these results suggest that innovator effects encourage auctioning, while imitator effects can discourage it. In practice, because innovator effects may be possible to strengthen through activities such as increased advertising, a larger promotional budget could be seen as one way to "buy time" for the firm who wishes to delay mass-market launch to allow for further demand learning via auctions.

Finally, this paper also introduces a novel auction format, the second price sealed-bid auction with rebate, for use in this market research setting. Rather than a traditional auction design approach which takes per-auction revenue maximization as its objective, this auction design seeks to simplify the inverse problem between bids and willingness to pay. This approach facilitates purchase probability estimations at the various prices considered by the firm, and accounts for the presence of the future fixed price. While this auction format is shown to be incentive compatible in this setting, the stopping time analyses do not depend upon the use of this particular auction format nor the assumptions made (see §3) when proving its incentive compatibility. Rather, our main results only require that the purchase probabilities for the N prices can somehow be updated based on bids, and this updating process is martingale. The design and testing of other auction formats for market research in such settings is an attractive future direction for research.

Mass-market fixed pricing helps focus our results on the main tradeoff between learning demand through auctions on one hand, and delaying market entry on the other (this tradeoff can become rather complex to characterize, see, for example, the proof of Corollary 1). Future work could perhaps extend the results to include additional complexities such as enhanced revenue capture through flexible posted pricing during mass-market selling. Flexible posted pricing would also extend the firm's learning opportunities into the mass-market phase. While in this paper we have focused on auctions as the demand-learning forum, examining the use of auctions in tandem with one or more other demand learning strategies, such as strategically adjusting posted prices to observe demand impacts during the mass-market phase, is an interesting, but daunting, opportunity for future work.

References

Aviv, Y., A. Pazgal. 2005. A Partially Observed Markov Decision Process for Dynamic Pricing. *Management Sci.* **51**(9) 1400–1416.

Balvers, R. J., T. F. Cosimano. 1990. Actively Learning about Demand and the Dynamics of Prices Adjustment. *Econ. J.* **100**(402) 882–898.

Bar-Yossef, Z., K. Hildrum, F. Wu. 2002. Incentive-Compatible Online Auctions for Digital Goods. *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 964–970.

Bass, F. 1969. A New Product Growth Model for Consumer Durables. *Management Sci.* **15**(5) 215–227.

Becker, G. M., M. H. DeGroot, J. Marschak. 1964. Measuring Utility by a Single-Response Sequential Method. *Behavioral Sci.* **9**(3) 226–232.

Bernanke, B. S. 1983. Irreversibility, Uncertainty, and Cyclical Iinvestment. *Quart. J. Econ.* **98**(1) 85–106.

Caro, F., J. Gallien. 2007. Dynamic Assortment with Demand Learning for Seasonal Consumer Goods. *Management Sci.* **53**(2) 276–292.

Chan, T. Y., V. Kadiyali, Y.H. Park. 2007. Willingness to Pay and Competition in Online Auctions. *J. Marketing Res.* Forthcoming.

Chen, L., E. L. Plambeck. 2005. Dynamic Inventory Management with Learning about the Demand Distribution and Substitution Probability. *Manufacturing Service Oper. Management* Forthcoming.

Cohen, M.A., J. Eliashberg, T.H. Ho. 1996. New Product Development: The Performance and Time-to-Market Tradeoff. *Management Sci.* **42**(2) 173–186.

DeGroot, M. 1970. *Optimal Statistical Decisions*. John Wiley & Sons, Hoboken, NJ.

Dolan, R.J. 1993. *Managing the New Product Development Process*. Addison-Wesley.

eBay.com. 2007. Annual Meeting of Stockholders Presentation, June 14. Retrieved 2 August 2007 from eBay: `http://files.shareholder.com/downloads/ebay/153778008x0x111466/c0b0ae65%-abea-4c2f-9b0d-87d2e2679f6d/25858ebay_ASM.pdf`.

Fox, J. A., D. J. Hayes, J. B. Kliebenstein, J. F. Shogren. 1994. Consumer Acceptability of Milk from Cows Treated with Bovine Somatotropin. *J. Dairy Sci.* **77**(3) 703–707.

Harrison, G. W., R. M. Harstad, E. E. Ruström. 2004. Experimental Methods and Elicitation of Values. *Exper. Econ.* **7**(2) 123–140.

Ho, T. H., S. Savin, C. Terwiesch. 2002. Managing Demand and Sales Dynamics in New Product Diffusion Under Supply Constraint. *Management Sci.* **48**(2) 187–206.

Hoffman, E., D. Menkhaus, D. Chakravarti, R. Field, G. Whipple. 1993. Using Laboratory Experimental Auctions for Marketing Research: A Case Study of New Packaging for Fresh Beef. *Marketing Sci.* **12**(3) 318–338.

Jain, S., J. Srivastava. 2000. An Experimental and Theoretical Analysis of Price-Matching Refund Policies. *J. Marketing Res.* **37**(3) 351–362.

Jensen, R. 1982. Adoption and Diffusion of an Innovation of Uncertain Profitability. *J. Econ. Theory* **27**(1) 182–193.

Johnson, C., B. Tesch. 2005. US Online Auction Sales: 2005 to 2010. Report from Forrester Research, Inc. Retrieved 2 August 2007 from Forrester: `http://www.forrester.com/Research/Document/Excerpt/0,7211,37848,00.html%`.

Johnson, R. 2006. 10-in Cabinet Saws. *Fine Woodworking* (184) 46–51.

Kagel, J. H., A. E. Roth. 1995. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, N.J.

Kalish, S. 1985. A New Product Adoption Model with Price, Advertising, and Uncertainty. *Management Sci.* **31**(12) 1569–1585.

Kalish, S., G.L. Lilien. 1986. A Market Entry Timing Model for New Technologies. *Management Sci.* **32**(2) 194–205.

Kornish, L. J., R. L. Keeney. 2005. Repeated Commit-or-Defer Decisions with an Inventory Deadline: The Influenza Vaccine Composition. *Oper. Res.* Forthcoming.

Kumar, S., J. M. Swaminathan. 2003. Diffusion of Innovations Under Supply Constraints. *Oper. Res.* **51**(6) 866–879.

Lariviere, M. A., E. L. Porteus. 1999. Stalking Information: Bayesian Inventory Management with Unobserved Lost Sales. *Management Sci.* **45**(3) 346–363.

Lin, K. Y. 2006. Dynamic Pricing with Real-Time Demand Learning. *European J. Oper. Res.* **174**(1) 522–538.

Lippman, S. A., K. F. McCardle. 1991. Uncertain Search: A Model of Search among Technologies of Uncertain Values. *Management Sci.* **37**(11) 1474–1490.

Lovejoy, W. S. 1990. Myopic Policies for Some Inventory Models with Uncertain Demand Distributions. *Management Sci.* **36**(6) 724–738.

Lucking-Reiley, D. 2000. Auctions on the Internet: What's Being Auctioned, and How? *The Journal of Industrial Economics* **48**(3) 227–252.

Mahajan, V., E. Muller, F. M. Bass. 1990. New Product Diffusion Models in Marketing: A Review and Directions for Research. *J. Marketing* **54**(1) 1–26.

Mahajan, V., E. Muller, Y. Winds, eds. 2000. *New Product Diffusion Models*. Springer, New York.

Mahajan, V., R. A. Peterson. 1978. Innovative Diffusion in a Dynamic Potential Adoptor Population. *Management Sci.* **24**(15) 1589–1597.

McCardle, K. F. 1985. Infromation Acquisition and the Adoption of New Technology. *Management Sci.* **31**(11) 1372–1389.

Mehler, K. 2005. A Safer Tablesaw Finally Arrives. *Fine Woodworking* (174) 66–67.

Müller, A., D. Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*. John Wiley & Sons, West Sussex, England.

Noussair, C., S. Robin, B. Ruffieux. 2004. Revealing Consumers' Willingness-to-Pay: A Comparison of the BDM Mechanism and the Vickrey Auction. *J. Econ. Psych.* **25**(6) 725–741.

Paarsch, H., H. Hong. 2006. *An Introduction to the Structural Econometrics of Auction Data*. The MIT Press, Cambridge, MA.

Paarsch, H. J. 1997. Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales. *J. Econometrics* **78**(2) 333–357.

Pinker, E. J., A. Seidmann, Y. Vakrat. 2007. Using Bid Data for the Management of Sequential, Multi-Unit, Online Auctions with Uniformly Distributed Bidders Valuations. Simon School, University of Rochester, Research Paper Series, Forthcoming. Available at SSRN: `http://ssrn.com/abstract=251958`.

Roosen, J., J. A. Fox, D. A. Hennessy, A. Schreiber. 1998. Consumers' Valuation of Insecticide Use Restrictions: An Application to Apples. *J. Agr. Resource Econ.* **23**(2) 367–384.

Ruström, E. E. 1998. Home-Grown Values and Incentive Compatible Auction Design. *Int. J. Game Theory* **27**(3) 427–441.

Scarf, H. 1959. Bayes Solution of the Statistical Inventory Problem. *Ann. Math. Stat.* **30**(2) 490–508.

Segal, I. 2003. Optimal Pricing Mechanisms with Unknown Demand. *Amer. Econ. Rev.* **93**(3) 509–529.

Silver, E. A. 1965. Bayesian Determination of the Reorder Point of a Slow Moving Item. *Oper. Res.* **13**(6) 989–997.

Skrzycki, C. 2006. Inventor's Attempt at Saw Safety Cuts Against Industry's Grain. Retrieved 2 August 2007 from Washington Post: `http://www.washingtonpost.com/wp-dyn/content/article/2006/10/09/AR20061%00901172_pf.html`.

Srivastava, J., N. Lurie. 2001. A Consumer Perspective on Price-Matching Refund Policies: Effect on Price Perceptions and Search Behavior. *J. Cons. Res.* **28**(2) 296–307.

Online Appendix for

Auctions to Learn Consumer Demand for a Product with a Short Selling Horizon

Proof of Proposition 1. Let x_i be the valuation (type) of bidder i, let b_j be the bid submitted by bidder j, let b_0 be the secret reserve price, let $y = \max_{j \neq i} b_j$, and let p denote the fixed price, which is unknown at the time of the auction. We show i bidding its true valuation is a (weakly) dominant strategy for any possible value of p. That is, i's expected payoff with bid $b_i = x_i$ is at least as large as that with any other bid.

Suppose $b_i > x_i$. If $x_i \geq y$, i's ex post payoff is $x_i - \min\{y, p\}$, which would be the same payoff if i had instead bid x_i. If $x_i < y$ and $b_i < y$, the payoff is $\max\{x_i - p, 0\}$, which would again be the same had i instead bid x_i. On the other hand, suppose $x_i < y \leq b_i$. If $x_i \geq p$, i's payoff is $x_i - p$, which is again what would have resulted from instead bidding x_i. However, if $x_i < p$, i's payoff of $x_i - \min\{y, p\}$ would be negative, which is worse than the zero payoff which would have resulted had i instead bid x_i.

Next, suppose $b_i < x_i$. If $b_i \geq y$, i's payoff is $x_i - \min\{y, p\}$, which would be the same payoff if i had instead bid x_i. If $b_i < y$, the payoff is $\max\{x_i - p, 0\}$, which is no greater than $\max\{x_i - \min\{y, p\}, 0\}$, which would have been i's payoff if he had bid x_i.

Note that the above incentive compatibility proof also applies to the case in which y, instead of being chosen as the second-highest bid (or as the reserve price, whichever is higher), is simply chosen randomly and every bid of at least y transacts at price y. This type of "auction" is the Becker-DeGroot-Marschak mechanism, and thus the Becker-DeGroot-Marschak mechanism with rebate is incentive compatible. ■

Lemma 1. *The bidding strategy $\beta(x) = x - \int_0^x H(p)dp$ is a symmetric equilibrium in the second price sealed-bid auction without rebate, where H is the distribution of (unknown) future price p.*

Proof. Let D be the number of bidders participating in the auction, and let Y be the highest of $D - 1$ independently drawn values from the bidder valuation distribution $G_0(\cdot)$, which is assumed differentiable and strictly increasing (as is standard in the literature). The cumulative distribution function of Y is given by $G(y) = G_0(y)^{D-1}$ for all y. The bidders know that the item will be available at a later date at some price P which is uncertain but follows distribution H; that is, the expected payoff of a bidder with valuation x who does not purchase at the auction is $\int_0^x (x - p)h(p)dp = \int_0^x H(p)dp$.

Consider a bidder, call him bidder 1, with valuation x who faces $D - 1$ competitors playing strategy β; we show that bidder 1 can do no better than follow β. Let $\Pi(b, x)$ denote bidder 1's expected payoff when bidding b. Since β is strictly increasing in x, a bid of b wins the auction with probability $G(\beta^{-1}(b))$. Thus, defining $G' = g$, we get

$$\begin{aligned}\Pi(b,x) &= G(\beta^{-1}(b))\left(\frac{\int_0^{\beta^{-1}(b)}(x-\beta(y))g(y)dy}{G(\beta^{-1}(b))}\right) + \left(1 - G(\beta^{-1}(b))\right)\int_0^x H(p)dp,\\ &= \int_0^{\beta^{-1}(b)}(x-\beta(y))g(y)dy + \left(1 - G(\beta^{-1}(b))\right)\int_0^x H(p)dp.\end{aligned}$$

Differentiating the above with respect to the bid, b, gives the following first order condition (FOC).

$$\frac{\partial \Pi(b,x)}{\partial b} = (x-b)\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} - \frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}\int_0^x H(p)dp = 0, \tag{A.6}$$

which is satisfied when $b = x - \int_0^x H(p)dp$. It remains to show that this solution is unique. The profit function is now shown to be unimodal since it is strictly concave whenever the first order condition is satisfied.

$$\begin{aligned}\frac{\partial^2 \Pi(b,x)}{\partial b^2} &= \left(x - b - \int_0^x H(p)dp\right)\frac{\partial}{\partial b}\left(\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}\right) - \frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))}, \quad \text{hence}\\ \left.\frac{\partial^2 \Pi(b,x)}{\partial b^2}\right|_{FOC} &= -\frac{g(\beta^{-1}(b))}{\beta'(\beta^{-1}(b))} < 0.\end{aligned} \tag{A.7}$$

The last expression follows since the first term is zero where the first order condition is satisfied. The final expression is negative since β is increasing and $g(\cdot) > 0$. Concavity where the first order condition holds implies the profit function is unimodal and that the first order condition solution is unique. ■

Proof of Proposition 2. To prove existence of the thresholds, it is sufficient to show that, for $\mathcal{B}_t \in \Omega_t$, $p_i w_i(\mathcal{B}_t)M(t) - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$ is nondecreasing in $w_i(\mathcal{B}_t)$ for all t and i (prices). (Increasing $w_i(\mathcal{B}_t)$ can be thought of as shifting from $\mathcal{B}_t$ to $\widehat{\mathcal{B}}_t \in \Omega_t$, where $w_i(\widehat{\mathcal{B}}_t) > w_i(\mathcal{B}_t)$ and $\mathbf{w}_{-i}(\widehat{\mathcal{B}}_t) = \mathbf{w}_{-i}(\mathcal{B}_t)$.) First, it is shown that $p_i w_i(\mathcal{B}_t)M(t) - J_t(\mathcal{B}_t)$ is nondecreasing in $w_i(\mathcal{B}_t)$ by induction on t. Clearly this holds for $t = T$, since $J_T(\mathcal{B}_T) = 0$ and the first term is increasing in $w_i(\mathcal{B}_T)$. Assume that it holds for $t+1$ for some $t \leq T-1$. Now,

$$\begin{aligned}&p_i w_i(\mathcal{B}_t)M(t) - J_t(\mathcal{B}_t)\\ &\qquad = p_i w_i(\mathcal{B}_t)M(t) - \max\Big[p_1 w_1(\mathcal{B}_t)M(t), \ldots, p_N w_N(\mathcal{B}_t)M(t), \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]\Big].\end{aligned}$$

On the righthand side, the result is trivially nondecreasing in $w_i(\mathcal{B}_t)$ if the maximum of the second term is one of the first N expressions. If the maximum of the second term is the last expression, then

$$\begin{aligned}
& p_i w_i(\mathcal{B}_t) M(t) - J_t(\mathcal{B}_t) \\
&= p_i w_i(\mathcal{B}_t) M(t) - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t], \\
&= p_i w_i(\mathcal{B}_t) M(t) - \delta E[p_i w_i(\mathcal{B}_{t+1}) M(t+1)|\mathcal{B}_t] + \delta E[p_i w_i(\mathcal{B}_{t+1}) M(t+1)|\mathcal{B}_t] - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t], \\
&= p_i w_i(\mathcal{B}_t)\big(M(t) - \delta M(t+1)\big) + \delta E[p_i w_i(\mathcal{B}_{t+1}) M(t+1) - J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]. \qquad \text{(A.8)}
\end{aligned}$$

The first term of (A.8) is derived by noting that $E[w_i(\mathcal{B}_{t+1})|\mathcal{B}_t] = w_i(\mathcal{B}_t)$ by the martingale assumption, equation (1). This term is nondecreasing in $w_i(\mathcal{B}_t)$ since $M(t)$ is nonincreasing in t. The expression $p_i w_i(\mathcal{B}_{t+1}) M(t+1) - J_{t+1}(\mathcal{B}_{t+1})$ is nondecreasing in $w_i(\mathcal{B}_{t+1})$ by the induction hypothesis, and also clearly is a bounded function. By equation (3), the expectation of this expression, conditional on $\mathcal{B}_t$, is nondecreasing in $w_i(\mathcal{B}_t)$, and thus it has been shown that $p_i w_i(\mathcal{B}_t) M(t) - J_t(\mathcal{B}_t)$ is nondecreasing in $w_i(\mathcal{B}_t)$. The existence of the threshold function follows immediately from the fact that $p_i w_i(\mathcal{B}_t) M(t) - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$ is nondecreasing in $w_i(\mathcal{B}_t)$, which was just shown. Next, it is shown that the threshold function, $h_{it}(w_{-i}(\mathcal{B}_t), \Omega_t, \mathbf{p})$, is nondecreasing in $w_{-i}(\mathcal{B}_t)$.

The threshold function for price i must be nondecreasing in $w_j(\mathcal{B}_t)$, $i \neq j$. To see this, the threshold occurs when $p_i w_i(\mathcal{B}_t) M(t)$ is greater than or equal to all $p_j w_j(\mathcal{B}_t) M(t)$, $i \neq j$ and when $p_i w_i(\mathcal{B}_t) M(t)$ is greater than $\delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$. In the first case, the threshold must be nondecreasing in $w_j(\mathcal{B}_t)$. In the second case, it is necessary that the expectation, $E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$, is nondecreasing in $w_j(\mathcal{B}_t)$ for all j, which follows from a simple induction argument and equation (3). The behavior of the threshold in p_j, $j \neq i$, is similar. ■

Proof of Corollary 1. Multidimensional Beta-multinomial updating satisfies the martingale assumption (1) by the definition of Bayesian updating and conditional expectation. On page 12 we discussed a sufficient statistic for multidimensional Beta-multinomial updating. Below we will use a slightly different sufficient statistic, using terminology from page 12. To this end, let $S_{j,t} \triangleq \sum_{i \geq j} \gamma_{i,t}$. Let $\mathbf{w} \triangleq (w_1, w_2, \ldots, w_N)$, and let $\mathbf{w}_{-i}$ be the vector excluding w_i.

Note that $\mathbf{S}_t \triangleq (S_{0,t}, S_{1,t}, \ldots, S_{N,t})$ is a sufficient statistic, and $w_i(\mathcal{B}_t = \mathbf{S}_t) = S_{i,t}/S_{0,t}$. Thus for two histories, $\mathcal{B}_t = \mathbf{S}_t$ and $\widehat{\mathcal{B}}_t = \hat{\mathbf{S}}_t$, having the same number of observations $S_{0,t}$, with $w_i(\widehat{\mathcal{B}}_t) \geq w_i(\mathcal{B}_t)$ and $\mathbf{w}_{-i}(\widehat{\mathcal{B}}_t) = \mathbf{w}_{-i}(\mathcal{B}_t)$, it must be that $\hat{S}_{i,t} \geq S_{i,t}$ and $\hat{S}_{j,t} = S_{j,t}$ $j \neq i$. In words, compared to $\mathbf{S}_t$, $\hat{\mathbf{S}}_t$ has some observations (bids) shifted from interval $[p_{i-1}, p_i)$ to $[p_i, p_{i+1})$.

Note that the joint prior distribution of $w_0 \geq w_1 \geq \cdots \geq w_{N+1}$ is multidimensional beta,

$$f(\mathbf{w}|\mathcal{B}_t) = \Gamma(S_{0,t}) \prod_{j=1}^{N} \frac{(w_j - w_{j+1})^{S_{j,t} - S_{j+1,t} - 1}}{\Gamma(S_{j,t} - S_{j+1,t})}, \qquad \text{where } S_{N+1,t} \triangleq 0. \qquad \text{(A.9)}$$

Integrating (A.9) over $w_i \in [w_{i+1}, w_{i-1}]$ shows that the marginal distribution, $f(\mathbf{w}_{-i}|\mathcal{B}_t)$, depends only on $\mathbf{S}_{-i,t}$, which from the arguments above equals $\hat{\mathbf{S}}_{-i,t}$ (the subscript $-i$ refers to all but the i^{th}

element). Thus, we conclude that $f(\mathbf{w}_{-i}|\mathcal{B}_t) = f(\mathbf{w}_{-i}|\widehat{\mathcal{B}}_t)$. However, since

$$f(w_i|\mathcal{B}_t) = w_i^{S_{i,t}-1}(1-w_i)^{S_{0,t}-S_{i,t}-1}\frac{\Gamma(S_{0,t})}{\Gamma(S_{i,t})\Gamma(S_{0,t}-S_{i,t})}, \tag{A.10}$$

it is easy to see that $\hat{S}_{i,t} > S_{i,t}$ implies $f(w_i|\mathcal{B}_t) \neq f(w_i|\widehat{\mathcal{B}}_t)$. To show that equation (3) holds (our desired result), we will use the notion of likelihood ratio ordering, denoted "$\leq_{LR}$" (see p12 of Müller and Stoyan 2002). By definition, for two random variables X and Y, $X \leq_{LR} Y$ if and only if $f_X(v)f_Y(u) \leq f_X(u)f_Y(v)$ for all $u \leq v$ (see p12 Müller and Stoyan 2002). It is easy to check using (A.10) that $[w_i|\mathcal{B}_t] \leq_{LR} [w_i|\widehat{\mathcal{B}}_t]$.

Towards showing (3) holds, we first characterize properties of $E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t]$. Define $p_0 \triangleq 0$, $p_{N+1} \triangleq \infty$, $w_0 = 1$ and $w_{N+1} = 0$. Let K_j, $j = 0,\ldots,N$ be the number of bids received in the $t+1^{\text{st}}$ auction that are greater than or equal to price p_j. Note that K_0 is the total number of bids received in the auction. (Thus K_0 is the same as b_t in the statement of Corollary 1.) Let $\mathcal{B}_{t+1} \triangleq (S_{0,t}+K_0, S_{1,t}+K_1,\ldots,S_{N,t}+K_N) \triangleq \mathbf{S}_t + \mathbf{K}$ denote the updated sufficient statistic after the $t+1^{\text{st}}$ auction. We now show that $E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t, K_0, \mathbf{w}]$ is nondecreasing in w_i for all K_0, $\mathbf{w}$, and $\mathcal{B}_t$. Let $0 \leq k_j \leq K_0$ be the number of bids in interval $[p_j, p_{j+1})$. Then

$$\begin{aligned}
E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t, K_0, \mathbf{w}] &= \sum_{\substack{k_0,\ldots,k_N \\ s.t. \sum_{j=0}^N k_j = K_0}} K_0! \prod_{j=0}^{N} \frac{(w_j - w_{j+1})^{k_j}}{k_j!}\phi(w_i(\mathbf{S}_t+\mathbf{K})), &\text{(A.11)}\\
&= \sum_{R=0}^{K_0} \sum_{\substack{k_j,\ j\neq i-1,i \\ s.t. \sum_{j\neq i-1,i} k_j = K_0 - R}} K_0! \prod_{\substack{j=1,\ldots,N \\ j\neq i-1,i}} \frac{(w_j - w_{j+1})^{k_j}}{k_j!} &\text{(A.12)}\\
&\quad \times \sum_{k_{i-1}=0}^{R} \frac{(w_{i-1}-w_i)^{k_{i-1}}}{k_{i-1}!}\frac{(w_i - w_{i+1})^{R-k_{i-1}}}{(R-k_{i-1})!}\phi(w_i(\mathbf{S}_t+\mathbf{K})),
\end{aligned}$$

where for the second equality we have rewritten the terms using $R = k_{i-1} + k_i$. We differentiate the above with respect to w_i using the product rule.

$$\begin{aligned}
\frac{dE[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t, K_0, \mathbf{w}]}{dw_i} &= \sum_{R=0}^{K_0} \sum_{\substack{k_j,\ j\neq i-1,i \\ s.t. \sum_{j\neq i-1,i} k_j = K_0 - R}} K_0! \prod_{\substack{j=1,\ldots,N \\ j\neq i-1,i}} \frac{(w_j - w_{j+1})^{k_j}}{k_j!} \\
&\sum_{k_{i-1}=0}^{R-1} \frac{(w_{i-1}-w_i)^{k_{i-1}}}{k_{i-1}!}\frac{(w_i - w_{i+1})^{R-k_{i-1}-1}}{(R-k_{i-1}-1)!} \\
&\times\Bigg(\phi(w_i(\mathbf{S}_t + (K_0,\ldots,K_{i-1},\ K_{i-1}-k_{i-1},\ K_{i+1},\ldots,K_N))) \\
&-\phi(w_i(\mathbf{S}_t + (K_0,\ldots,K_{i-1},\ K_{i-1}-k_{i-1}-1,\ K_{i+1},\ldots,K_N)))\Bigg),
\end{aligned}$$

which is nonnegative since ϕ is nondecreasing and $w_i(\mathbf{S}_{t+1}) = S_{i,t+1}/S_{0,t+1}$. Finally, using the Corollary's assumption $\Pr(K_0 = j|\mathcal{B}_t) = \Pr(K_0 = j|\widehat{\mathcal{B}}_t)$, and for convenience writing this probability as simply $\Pr(K_0 = j)$, we have

$$\begin{aligned}
& E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t] \\
& = \int_{\mathbf{w}} f(\mathbf{w}|\mathcal{B}_t) \sum_{j=0}^{\infty} E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) d\mathbf{w} \\
& \leq \int_{\mathbf{w}} f(\mathbf{w}|\mathcal{B}_t) \sum_{j=0}^{\infty} E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) d\mathbf{w} && \text{(A.13)} \\
& = \int_{\mathbf{w}_{-i}} f(\mathbf{w}_{-i}|\mathcal{B}_t) \int_{w_i|\mathbf{w}_{-i}} f(w_i|\mathcal{B}_t, \mathbf{w}_{-i}) \sum_{j=0}^{\infty} E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) d\mathbf{w}, \\
& = \int_{\mathbf{w}_{-i}} f(\mathbf{w}_{-i}|\mathcal{B}_t) \sum_{j=0}^{\infty} \int_{w_i|\mathbf{w}_{-i}} f(w_i|\mathcal{B}_t, \mathbf{w}_{-i}) E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) d\mathbf{w}, && \text{(A.14)} \\
& = \int_{\mathbf{w}_{-i}} f(\mathbf{w}_{-i}|\mathcal{B}_t) \sum_{j=0}^{\infty} E_{w_i|\mathcal{B}_t, \mathbf{w}_{-i}} \left[E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) \right] d\mathbf{w}, \\
& = \int_{\mathbf{w}_{-i}} f(\mathbf{w}_{-i}|\widehat{\mathcal{B}}_t) \sum_{j=0}^{\infty} E_{w_i|\mathcal{B}_t, \mathbf{w}_{-i}} \left[E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) \right] d\mathbf{w} && \text{(A.15)} \\
& \leq \int_{\mathbf{w}_{-i}} f(\mathbf{w}_{-i}|\widehat{\mathcal{B}}_t) \sum_{j=0}^{\infty} E_{w_i|\widehat{\mathcal{B}}_t, \mathbf{w}_{-i}} \left[E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) \right] d\mathbf{w} && \text{(A.16)} \\
& = \int_{\mathbf{w}} f(\mathbf{w}|\widehat{\mathcal{B}}_t) \sum_{j=0}^{\infty} E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) d\mathbf{w} \; = \; E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t].
\end{aligned}$$

The inequality in (A.13) follows from substituting $\hat{\mathbf{S}}_t$ for $\mathbf{S}_t$ in equation (A.11) and noting that $\phi(w_i(\mathbf{S}_{t+1}))$ is nondecreasing in $S_{i,t+1}$. The interchange of the integral and the infinite summation in (A.14) is allowed by the dominated convergence theorem since

$$\left| \lim_{n\to\infty} \sum_{j=0}^{n} f(w_i|\mathcal{B}_t, \mathbf{w}_{-i}) E[\phi(w_i(\widehat{\mathcal{B}}_{t+1}))|\widehat{\mathcal{B}}_t, K_0 = j, \mathbf{w}] \Pr(K_0 = j) \right| < f(w_i|\mathcal{B}_t, \mathbf{w}_{-i}) Q$$

and $f(w_i|\mathcal{B}_t, \mathbf{w}_{-i})Q$ is measurable. The equality in (A.15) follows by $f(\mathbf{w}_{-i}|\mathcal{B}_t) = f(\mathbf{w}_{-i}|\widehat{\mathcal{B}}_t)$, which was established above. To understand the inequality in (A.16), first note that $[w_i|\mathcal{B}_t] \leq_{LR} [w_i|\widehat{\mathcal{B}}_t]$ (also established above) implies $[w_i|\mathcal{B}_t, \mathbf{w}_{-i}] \leq_{st} [w_i|\widehat{\mathcal{B}}_t, \mathbf{w}_{-i}]$ (where $\leq_{st}$ denotes stochastic dominance, see p13 Müller and Stoyan 2002). Since we showed $E[\phi(w_i(\mathcal{B}_{t+1}))|\mathcal{B}_t, K_0, \mathbf{w}]$ is nondecreasing in w_i for all K_0, $\mathbf{w}$, and $\mathcal{B}_t$, the inequality in (A.16) follows by the definition of stochastic dominance. ■

Proposition 8. (Fewer past observations encourages auctioning for one extra period with multidimensional Beta-multinomial updating.) *Suppose the firm updates its priors using the* *multidimensional Beta-multinomial information structure, and $\mathcal{B}_t$ contains x_0 observations of bidder valuations, while $\widehat{\mathcal{B}}_t$ only contains $y < x_0$ observations. If $w_i(\mathcal{B}_t) = w_i(\widehat{\mathcal{B}}_t) = \bar{w}_i$ for all $i = 1, \ldots, N$ and at time t and sufficient statistic $\mathcal{B}_t$ the firm prefers continuing the auction phase for exactly one additional period (entering at time $t+1$) to entering immediately (entering at time t), the same would also be preferred at time t with sufficient statistic $\widehat{\mathcal{B}}_t$.*

Proof. We first show that for any $s > 0$,

$$E[\max_j\{p_j w_i(\mathcal{B}_{t+s})\}|\mathcal{B}_t] \;\leq\; E[\max_j\{p_j w_i(\widehat{\mathcal{B}}_{t+s})\}|\widehat{\mathcal{B}}_t]. \tag{A.17}$$

Let d be the number of bids received during $(t, t+s]$, and let k_i be the number of such bids in the interval $[p_i, p_{i+1})$, where $p_0 \triangleq 0$ and $p_{N+1} \triangleq \infty$. Hence, $d = \sum_{j=0}^{N} k_j$. For shorthand let $\mathbf{k} \triangleq (k_0, \ldots, k_N)$. Set $\bar{w}_0 = 1$ and $\bar{w}_{N+1} = 0$. We establish (A.17) by showing that, for all d, $E_{\mathbf{k}}[\max_j\{p_j w_i(\mathcal{B}_{t+s})\}|\mathcal{B}_t, d] \;\leq\; E_{\mathbf{k}}[\max_j\{p_j w_i(\widehat{\mathcal{B}}_{t+s})\}|\widehat{\mathcal{B}}_t, d]$.

$$\begin{aligned}
\text{Let} \qquad j^*(\mathbf{k}, x_0) &\triangleq \arg\max_j \left\{ \frac{p_j \bar{w}_j x_0 + p_j \sum_{m=j}^{N} k_m}{x_0 + d} \right\}, \\
\text{and let} \qquad g(x) &\triangleq E_{\mathbf{k}} \left[\frac{p_{j^*(\mathbf{k},x_0)} \bar{w}_{j^*(\mathbf{k},x_0)} x + p_{j^*(\mathbf{k},x_0)} \sum_{m=j^*(\mathbf{k},x_0)}^{N} k_m}{x + d} \right].
\end{aligned}$$

Expanding the expectation we can write

$$g(x) = \frac{d!}{x+d} \sum_{\substack{k_0,\ldots,k_N \\ s.t.\ \sum_{j=0}^{N} k_j = d}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \left[p_{j^*(\mathbf{k},x_0)} \bar{w}_{j^*(\mathbf{k},x_0)} x + p_{j^*(\mathbf{k},x_0)} \sum_{m=j^*(\mathbf{k},x_0)}^{N} k_m \right].$$

For any constants c and d, $\frac{\partial}{\partial x} \frac{p_j \bar{w}_j x + p_j c}{x+d} = \frac{d p_j \bar{w}_j - p_j c}{(x+d)^2}$. Letting

$$h_m(\mathbf{k}) = \begin{cases} p_{j^*(\mathbf{k},x_0)} \bar{w}_{j^*(\mathbf{k},x_0)} - p_{j^*(\mathbf{k},x_0)} & \text{if } m \geq j^*(\mathbf{k}, x_0), \\ p_{j^*(\mathbf{k},x_0)} \bar{w}_{j^*(\mathbf{k},x_0)} & \text{otherwise,} \end{cases}$$

and letting $\mathbf{e}^i$ be the vector of all zeros except for a 1 in i^{th} position, we can write

$$\begin{aligned}
\frac{\partial g(x)}{\partial x} &= \frac{d!}{(x+d)^2} \sum_{\substack{k_0,\ldots,k_N \\ s.t.\ \sum_{j=0}^{N} k_j = d}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \sum_{m=0}^{N} k_m h_m(\mathbf{k}), \\
&= \frac{(d-1)!}{(x+d)^2} \sum_{\substack{k_0,\ldots,k_N \\ s.t.\ \sum_{j=0}^{N} k_j = d-1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \sum_{i=0}^{N} (\bar{w}_i - \bar{w}_{i+1}) \left[h_i(\mathbf{k} + \mathbf{e}^i) + \sum_{m=0}^{N} k_m h_m(\mathbf{k} + \mathbf{e}^i) \right].
\end{aligned}$$

We now re-write the terms involving the summation over m:

$$\begin{aligned}
&\frac{(d-1)!}{(x+d)^2} \sum_{\substack{k_0,\ldots,k_N \\ s.t. \sum_{j=0}^N k_j = d-1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \sum_{i=0}^{N} (\bar{w}_i - \bar{w}_{i+1}) \sum_{m=0}^{N} k_m h_m(\mathbf{k} + \mathbf{e}^i) \\
&= \frac{(d-1)!}{(x+d)^2} \sum_{i=0}^{N} \sum_{m=0}^{N} \sum_{\substack{k_0,\ldots,k_N \\ s.t. \sum_{j=0}^N k_j = d-1,\ k_m \geq 1}} \prod_{\substack{l=0,\ldots,N \\ l \neq i,m}} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \\
&\quad \cdot \frac{(\bar{w}_i - \bar{w}_{i+1})^{k_i+1}}{(k_i+1)!} \frac{(\bar{w}_m - \bar{w}_{m+1})^{k_m-1}}{(k_m-1)!} (k_i+1) h_m(\mathbf{k} + \mathbf{e}^i)(\bar{w}_m - \bar{w}_{m+1}),
\end{aligned}$$

where we have used the fact that the terms involving $k_m = 0$ are zeroed out. Substituting $\tilde{k}_i = k_i + 1$, $\tilde{k}_m = k_m - 1$, $\tilde{k}_j = k_j$ for $j \neq i, m$, we can re-write the terms as

$$\begin{aligned}
&\frac{(d-1)!}{(x+d)^2} \sum_{i=0}^{N} \sum_{m=0}^{N} \sum_{\substack{\tilde{k}_0,\ldots,\tilde{k}_N \\ s.t. \sum_{j=0}^N \tilde{k}_j = d-1,\ \tilde{k}_i \geq 1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{\tilde{k}_l}}{\tilde{k}_l!} \tilde{k}_i h_m(\tilde{\mathbf{k}} + \mathbf{e}^m)(\bar{w}_m - \bar{w}_{m+1}) \\
&= \frac{(d-1)!}{(x+d)^2} \sum_{m=0}^{N} \sum_{\substack{\tilde{k}_0,\ldots,\tilde{k}_N \\ s.t. \sum_{j=0}^N \tilde{k}_j = d-1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{\tilde{k}_l}}{\tilde{k}_l!} (d-1) h_m(\tilde{\mathbf{k}} + \mathbf{e}^m)(\bar{w}_m - \bar{w}_{m+1}),
\end{aligned}$$

where we have used the fact that $\sum_{i=0}^{N} \tilde{k}_k = d-1$ and terms for which $\tilde{k}_i = 0$ zero out. Hence,

$$\begin{aligned}
\frac{\partial g(x)}{\partial x} &= \frac{(d-1)!}{(x+d)^2} \sum_{\substack{k_0,\ldots,k_N \\ s.t. \sum_{j=0}^N k_j = d-1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \sum_{i=0}^{N} (\bar{w}_i - \bar{w}_{i+1}) \left[h_i(\mathbf{k} + \mathbf{e}^i) + (d-1) h_i(\mathbf{k} + \mathbf{e}^i) \right], \\
&= \frac{d!}{(x+d)^2} \sum_{\substack{k_0,\ldots,k_N \\ s.t. \sum_{j=0}^N k_j = d-1}} \prod_{l=0}^{N} \frac{(\bar{w}_l - \bar{w}_{l+1})^{k_l}}{k_l!} \sum_{i=0}^{N} (\bar{w}_i - \bar{w}_{i+1}) h_i(\mathbf{k} + \mathbf{e}^i).
\end{aligned}$$

We next show that $\sum_{i=0}^{N} (\bar{w}_i - \bar{w}_{i+1}) h_i(\mathbf{k} + \mathbf{e}^i) \leq 0$ for all fixed $\mathbf{k}$ such that $\sum_{j=0}^{N} k_j = d-1$ (hence $\partial g / \partial x \leq 0$). For such a fixed $\mathbf{k}$, define

$$j_q \triangleq \arg\max_j \left\{ \frac{p_j \bar{w}_j x_0 + p_j \sum_{m=j}^{N} k_m}{x_0 + d - 1} \right\}.$$

Price p_{j_q} may or may not remain optimal once the d^{th} arrival is considered. If the d^{th} arrival has a valuation below p_1, the optimal choice of price will not change from p_{j_q}. Thus, $h_0(\mathbf{k} + \mathbf{e}^0) = p_{j_q} \bar{w}_{j_q}$. If the d^{th} arrival has a valuation above p_1 but below p_2, the optimal price will either change to p_1 or remain at p_{j_q}. If the d^{th} arrival has a valuation above p_2 but below p_3, the optimal price will either change to p_2 or remain at p_1 or p_{j_q} (whichever price was optimal for an arrival between p_1 and p_2). Continuing in this manner, define the indices, $j_1, \ldots, j_{q-1}$, as the $q-1$, $q \geq 1$, indices for which the optimal price changes for arrivals with valuations less than p_{j_q}. Define the indices, $j_{q+1}, \ldots, j_{q+r}$, as the r, $r \geq 0$, times the optimal price changes for arrivals with valuations greater than p_{j_q}. Let $l^* = \arg\max_{l=q,\ldots,q+r}\{p_{j_l}\bar{w}_{j_l} - p_{j_l}\}$, and set $\bar{w}_{j_{r+q+1}} = 0$. We can then write

$$\begin{aligned}
\sum_{i=0}^{N}(\bar{w}_i - \bar{w}_{i+1})h_i(\mathbf{k}+\mathbf{e}^i) &= (1-\bar{w}_{j_1})p_{j_q}\bar{w}_{j_q} + \sum_{l=1}^{q+r}(\bar{w}_{j_l} - \bar{w}_{j_l+1})(p_{j_l}\bar{w}_{j_l} - p_{j_l}) \\
&\leq (1-\bar{w}_{j_1})p_{j_q}\bar{w}_{j_q} + \sum_{l=1}^{q-1}(\bar{w}_{j_l} - \bar{w}_{j_l+1})(p_{j_l}\bar{w}_{j_l} - p_{j_l}) + \bar{w}_{j_q}(p_{j_{l^*}}\bar{w}_{j_{l^*}} - p_{j_{l^*}}) \\
&\leq (1-\bar{w}_{j_1})p_{j_{l^*}}\bar{w}_{j_{l^*}} + \sum_{l=1}^{q-1}(\bar{w}_{j_l} - \bar{w}_{j_l+1})(p_{j_l}\bar{w}_{j_l} - p_{j_l}) + \bar{w}_{j_q}(p_{j_{l^*}}\bar{w}_{j_{l^*}} - p_{j_{l^*}}), \\
&= (\bar{w}_{j_q} - \bar{w}_{j_1})p_{j_{l^*}}\bar{w}_{j_{l^*}} + \sum_{l=1}^{q-1}(\bar{w}_{j_l} - \bar{w}_{j_l+1})(p_{j_l}\bar{w}_{j_l} - p_{j_l}) + p_{j_{l^*}}(\bar{w}_{j_{l^*}} - \bar{w}_{j_q}) \leq 0.
\end{aligned}$$

The second inequality follows since $l^* \geq q$ implies $p_{j_{l^*}} \geq p_{j_q}$ and hence $p_{j_{l^*}}\bar{w}_{j_{l^*}} \geq p_{j_q}\bar{w}_{j_q}$. The final inequality follows since $\bar{w}_i \geq \bar{w}_j$ for $i \leq j$, and $\bar{w}_i \leq 1$ for all i.

Now, having established that $\partial g/\partial x \leq 0$, we can write

$$\begin{aligned}
E_{\mathbf{k}}[\max_j\{p_j w_j(\mathcal{B}_{t+s})\}|\mathcal{B}_t, d] = g(x_0) \leq g(y) &= E_{\mathbf{k}}\left[\frac{p_{j^*(\mathbf{k},x_0)}\bar{w}_{j^*(\mathbf{k},x_0)}y + p_{j^*(\mathbf{k},x_0)}\sum_{m=j^*(\mathbf{k},x_0)}^{N}k_m}{y+d}\right] \\
&\leq E_{\mathbf{k}}\left[\max_j\left\{\frac{p_j\bar{w}_j y + p_j\sum_{m=j}^{N}k_m}{y+d}\right\}\right], \\
&= E_{\mathbf{k}}[\max_j\{p_j w_j(\widehat{\mathcal{B}}_{t+s})\}|\widehat{\mathcal{B}}_t, d].
\end{aligned}$$

Because the above holds for all d, we have established (A.17), from which the proposition follows easily. ■

Proof of Proposition 3. We begin with part 1. Suppose that at time $t-1$ and sufficient statistic $\mathcal{B}_{t-1} = \mathcal{B}$, stopping the auction phase and entering the mass market is optimal. We show that stopping would also be optimal were the time instead t. Let $t+K$ be the optimal stopping time over horizon $[t, T]$ for a firm with sufficient statistic $\mathcal{B}$ at time t. If during horizon $[t-1, T]$ the firm pretends to have started from time t instead of $t-1$, it can do no better than if it optimizes its stopping time relative to its true starting point of $t-1$. Thus if stopping is optimal at time $t-1$, then

$$\sum_{k=1}^{T-t}\Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1} = \mathcal{B}]M(t-1+k) \leq \max_i\{p_i w_i(\mathcal{B})\}M(t-1).$$

Dividing both sides by $M(t-1)$ yields

$$\sum_{k=1}^{T-t}\Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1} = \mathcal{B}]\frac{M(t-1+k)}{M(t-1)} \leq \max_i\{p_i w_i(\mathcal{B})\}.$$

For now, suppose that $M(t-1+k)/M(t-1) \geq M(t+k)/M(t)$ for all $k \in \{1, \ldots, T-t\}$. Thus,

$$\sum_{k=1}^{T-t} \Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1}=\mathcal{B}]\frac{M(t+k)}{M(t)} \;\leq\; \max_i\{p_i w_i(\mathcal{B})\}.$$

Since the data updating process is assumed to be stationary, we can write

$$\sum_{k=1}^{T-t} \Pr(K=k)\delta^k E_{\mathcal{B}_{t+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t+k})\}|\mathcal{B}_t=\mathcal{B}]\frac{M(t+k)}{M(t)} \;\leq\; \max_i\{p_i w_i(\mathcal{B})\}. \tag{A.18}$$

The LHS of (A.18) is the payoff (divided by $M(t)$) of continuing auctioning at time t when following optimal stopping policy $t+K$. Thus, (A.18) implies that stopping the auction phase is optimal at time t and sufficient statistic $\mathcal{B}_t = \mathcal{B}$, as long as we can show $M(t-1+k)/M(t-1) \geq M(t+k)/M(t)$. Clearly the inequality holds at $k = T-t$, since $M(T) = 0$. For $k \in \{1, \ldots, T-t-1\}$ it is sufficient to show that

$$\frac{\partial}{\partial s}\left(\frac{M(s+k)}{M(s)}\right) \leq 0$$

for $s \in [t-1, t]$. Since $M(s+k) > 0$ for $k = 1, \ldots, T-t-1$, this is equivalent to

$$\frac{\frac{\partial}{\partial s}(M(s+k))}{M(s+k)} \leq \frac{\frac{\partial}{\partial s}(M(s))}{M(s)},$$

which holds if $M(\cdot)$ is log-concave over $[t-1, T-1]$.

Proof of part 2 is analogous to that of part 1. Suppose that at time $t-1$ and sufficient statistic $\mathcal{B}_{t-1} = \mathcal{B}$, continuing the auction phase is optimal. We show that continuing the auction phase would also be optimal were the time instead t. Let $t-1+K$ be the optimal stopping time if following an optimal policy from time $t-1$. Since by assumption $t-1+K \leq T-2$ with probability one, if continuing is optimal at time $t-1$, then

$$\sum_{k=1}^{T-t-1} \Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1}=\mathcal{B}]\frac{M(t-1+k)}{M(t-1)} \;\geq\; \max_i\{p_i w_i(\mathcal{B})\}.$$

For now, suppose that $M(t-1+k)/M(t-1) \leq M(t+k)/M(t)$ for all $k \in \{1, \ldots, T-t-1\}$. Thus,

$$\sum_{k=1}^{T-t-1} \Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1}=\mathcal{B}]\frac{M(t+k)}{M(t)} \;\geq\; \max_i\{p_i w_i(\mathcal{B})\}.$$

Since the data updating process is assumed to be stationary, we can write

$$\sum_{k=1}^{T-t-1} \Pr(K=k)\delta^k E_{\mathcal{B}_{t+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t+k})\}|\mathcal{B}_t=\mathcal{B}]\frac{M(t+k)}{M(t)} \;\geq\; \max_i\{p_i w_i(\mathcal{B})\}. \tag{A.19}$$

The LHS of (A.19) is the payoff (divided by $M(t)$) of continuing auctioning at time t and stopping at $t+K$ per the optimal policy for the firm starting from time $t-1$. When continuing auctioning from time t, the true optimal stopping policy in $[t+1,T]$ performs at least as well as when following the stopping time policy $t+K$. Thus, (A.19) implies that continuing the auction phase is optimal at time t and sufficient statistic $\mathcal{B}_t = \mathcal{B}$, as long as we can show $M(t-1+k)/M(t-1) \leq M(t+k)/M(t)$. It is sufficient to show that

$$\frac{\partial}{\partial s}\left(\frac{M(s+k)}{M(s)}\right) \geq 0$$

for all $k \in \{1,\ldots,T-t-1\}$, $s \in [t-1,t]$. Since $M(s+k) > 0$ for $k \in \{1,\ldots,T-t-1\}$, this is equivalent to

$$\frac{\frac{\partial}{\partial s}(M(s+k))}{M(s+k)} \geq \frac{\frac{\partial}{\partial s}(M(s))}{M(s)},$$

which holds if $M(\cdot)$ is log-convex over $[t-1,T-1]$. ■

Proof of Proposition 4. We first show that $\partial/\partial x(M(s)/M(t))\big|_x \geq 0$ for all $t \leq s \leq T$, $x \in I$, implies the stopping region shrinks as the market size parameter, x, increases within I. Let $x_1, x_2 \in I$. Suppose that at time t, with sufficient statistic $\mathcal{B}_t$ and market size parameter x_1, continuing the auction phase is optimal. We show that continuing the auction phase would also be optimal were the market size parameter $x_2 \geq x_1$. Let $M_k(t)$ be the market size at time t with parameter x_k, and S_k be the optimal stopping time if following an optimal policy as if the market size function is M_k, $k \in \{1,2\}$. If continuing is optimal with parameter x_1, then

$$\sum_{s=t+1}^{T} \Pr(S_1 = s)\delta^{s-t} E_{\mathcal{B}_s|S_1=s}[\max_i\{p_i w_i(\mathcal{B}_s)\}|\mathcal{B}_t]\frac{M_1(s)}{M_1(t)} \;\geq\; \max_i\{p_i w_i(\mathcal{B}_t)\}.$$

When the market parameter is x_2, the optimal stopping policy, $S_2 \in [t+1,T]$, performs at least as well as when following the stopping time policy, S_1, which is optimal under market parameter x_1. By assumption we have $M_1(s)/M_1(t) \leq M_2(s)/M_2(t)$. Thus, when continuing the auction phase is optimal at time t, sufficient statistic $\mathcal{B}_t$, and market size parameter x_1 we have

$$\begin{aligned}
&\sum_{s=t+1}^{T} \Pr(S_2 = s)\delta^{s-t} E_{\mathcal{B}_s|S_2=s}[\max_i\{p_i w_i(\mathcal{B}_s)\}|\mathcal{B}_t]\frac{M_2(s)}{M_2(t)} \\
&\geq \sum_{s=t+1}^{T} \Pr(S_1 = s)\delta^{s-t} E_{\mathcal{B}_s|S_1=s}[\max_i\{p_i w_i(\mathcal{B}_s)\}|\mathcal{B}_t]\frac{M_2(s)}{M_2(t)} \quad \geq \quad \max_i\{p_i w_i(\mathcal{B}_t)\},
\end{aligned} \tag{A.20}$$

and continuing the auction phase is optimal at time t, sufficient statistic $\mathcal{B}_t$, and market size parameter x_2.

The proof of the expanding stopping regions case is analogous. With parameter x_1, if stopping at t with $\mathcal{B}_t$ is optimal then continuing and following stopping time S_2 is worse than stopping. This and the assumption that $M_1(s)/M_1(t) \geq M_2(s)/M_2(t)$ can be used to show that stopping is also optimal with parameter $x_2 \geq x_1$. We omit the details for brevity. Combining the arguments for both of the above cases implies that the stopping regions do not change if $\partial/\partial x(M(s)/M(t)) = 0$ for $s \geq t$. ■

Proof of Corollary 3. Differentiating the natural log of (5) twice with respect to t is easily shown to always be non-negative, immediately yielding the result. ■

Proof of Corollary 4. The result follows by showing that $M(s)/M(t)$ is monotonically increasing in α and that Proposition 4 applies. The result holds trivially for $s = T$. For $t \leq s < T$,

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial \alpha}M(t)}{M(t)}\right) \geq 0 \Rightarrow \frac{\frac{\partial}{\partial \alpha}M(s)}{M(s)} \geq \frac{\frac{\partial}{\partial \alpha}M(t)}{M(t)} \Rightarrow \frac{\partial}{\partial \alpha}\left(\frac{M(s)}{M(t)}\right) \geq 0.$$

Accordingly, we will show that

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial \alpha}M(t)}{M(t)}\right) \geq 0$$

for all $\alpha, \beta > 0$ and all $t \in [0, T]$. Let $\gamma \triangleq (\alpha + \beta)(T - t)$. The derivative,

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial \alpha}M(t)}{M(t)}\right) = \frac{(\alpha+\beta)e^{-\gamma}\Big(e^{-\gamma}(\alpha - 3\beta) + \beta e^{-2\gamma}(\gamma+2) + \alpha\gamma - \alpha + \beta\Big)}{(\alpha + \beta e^{-\gamma})^2(1 - e^{-\gamma})^2},$$

will be non-negative if and only if

$$D(t) \triangleq e^{-\gamma}(\alpha - 3\beta) + \beta e^{-2\gamma}(\gamma + 2) + \alpha\gamma - \alpha + \beta \geq 0.$$

Since $D(T) = 0$ for all α, β, if we can show that $D(t)$ is decreasing on $[0, T]$, it will follow that $D(t)$ must be non-negative. Differentiation yeilds

$$\frac{\partial D}{\partial t} = (\alpha + \beta)\Big(e^{-\gamma}(\alpha - 3\beta) + \beta e^{-2\gamma}(2\gamma + 3) - \alpha\Big).$$

$D(t)$ will be decreasing if and only if

$$F(t) \triangleq e^{-\gamma}(\alpha - 3\beta) + \beta e^{-2\gamma}(2\gamma + 3) - \alpha \leq 0.$$

Since $F(T) = 0$, $F(t)$ will be non-positive if either (i) $F(t)$ is increasing on $[0, T]$, or (ii) $F(t)$ is unimodal on $[0, T]$ (decreasing then increasing) with $F(0) \leq 0$. Taking the derivative of $F(t)$,

$$\frac{\partial F}{\partial t} = (\alpha + \beta)e^{-\gamma}\Big(4\beta e^{-\gamma}(1 + \gamma) + \alpha - 3\beta\Big).$$

Define

$$G(t) \triangleq 4\beta e^{-\gamma}(1 + \gamma) + \alpha - 3\beta, \quad \text{with} \quad \frac{\partial G}{\partial t} = 4\beta(\alpha + \beta)\gamma e^{-\gamma} \geq 0.$$

The derivative of $G(t)$ is non-negative, which implies that $G(t)$ is increasing on $[0, T]$. This also implies that the derivative of $F(t)$ can change signs at most one time. Also, since $G(T) = \alpha + \beta > 0$,

$F(t)$ is increasing at $t = T$. Thus, $F(t)$ is either (i) increasing on $[0, T]$ or (ii) decreasing then increasing on $[0, T]$.

Case 1. $G(0) \geq 0$. Since $G(t)$ is increasing, $G(t) \geq 0$ for all $t \in [0, T]$, which implies that $F(t)$ is increasing for all $t \in [0, T]$. Thus, $F(t) \leq 0$ and $D(t) \geq 0$ for all $t \in [0, T]$.

Case 2. $G(0) < 0$. Since $G(t)$ is increasing and $G(T) > 0$, we know that $F(t)$ is unimodal (decreasing then increasing) on $[0, T]$. To finish the proof, it is necessary to show that $F(0) < 0$ when $G(0) < 0$. To this end, note that $e^{-\gamma}G(t) < 0 \iff G(t) < 0$. Furthermore,

$$\begin{aligned} F(t) - e^{-\gamma}G(t) &= e^{-\gamma}(\alpha - 3\beta) + \beta e^{-2\gamma}(2\gamma + 3) - \alpha - 4\beta e^{-2\gamma}(1 + \gamma) - e^{-\gamma}(\alpha - 3\beta), \\ &= -\beta e^{-2\gamma}(2\gamma + 1) - \alpha < 0. \end{aligned}$$

Thus, if $G(0) < 0$, then $F(0) < 0$ and $F(t)$ is decreasing then increasing on $[0, T]$, implying that $F(t) \leq 0$ and $D(t) \geq 0$ for all $t \in [0, T]$. ■

Proof of Corollary 5. The proof is similar in spirit to that of Corollary 4. First, note that $\partial/\partial\beta(M(s)/M(t)) \leq 0$ holds trivially if $s = T$. For fixed β, we show there exists a $\hat{t}(\beta) < T$ such that

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial\beta}M(t)}{M(t)}\right) < 0$$

for all $t \in [\hat{t}(\beta), T)$. For readability, where convenient we will suppress the argument β when writing $\hat{t}$. Let $\gamma \triangleq (\alpha + \beta)(T - t)$. We have

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial\beta}M(t)}{M(t)}\right) < 0 \iff D(t) \triangleq (-3\alpha + \beta)e^{-\gamma} - (\beta\gamma - \alpha + \beta)e^{-2\gamma} - \alpha(\gamma - 2) > 0.$$

It is easy to verify that $D(T) = 0$. We show there exists $\hat{t} < T$ such that $\partial D(t)/\partial t < 0$ for all $t \in [\hat{t}, T)$. First,

$$\frac{\partial D(t)}{\partial t} > 0 \iff F(t) \triangleq (-3\alpha + \beta)e^{-\gamma} - (2\beta\gamma - 2\alpha + \beta)e^{-2\gamma} + \alpha > 0 \qquad \text{and}$$

$$\frac{\partial F(t)}{\partial t} > 0 \iff G(t) \triangleq -3\alpha + \beta - 4(\beta\gamma - \alpha)e^{-\gamma} > 0.$$

$G(T) = \alpha + \beta > 0$ implies $F(t)$ is strictly increasing at $t = T$. Since $F(T) = 0$, we have that there exists $\hat{t} < T$ such that $F(t) < 0$ for all $t \in [\hat{t}, T)$. Hence, $\partial D(t)/\partial t < 0$ for all $t \in [\hat{t}, T)$ and $D(t) > 0$ for all $t \in [\hat{t}, T)$.

Next, we show that $\hat{t} = 0$ for α sufficiently small. In particular, for α sufficiently small,

$$\frac{\partial}{\partial t}\left(\frac{\frac{\partial}{\partial\beta}M(t)}{M(t)}\right) < 0 \qquad \text{for } 0 \leq t \leq T - 1.$$

It is easy to check that if $\alpha = 0$, $D(t) > 0 \iff \beta(T-t) > \ln(\beta(T-t)+1)$, which holds for all $t \le T-1$ and $\beta > 0$. For the moment, we find it convenient to think of D as a function of both t and α. Since D is continuous in t and α, D is uniformly continuous over compact domain $(t, \alpha) \in [0, T-1] \times [0, 1]$, where, without loss of generality, for compactness we have chosen 1 as an upper bound on α. Thus, there exists $\alpha_0 > 0$ such that $0 < \alpha < \alpha_0$ implies $D(t) > 0$ for all $t \in [0, T-1]$.

Fixing α, if $\hat{t}(\beta_0) \le T-1$, setting $t_0 = \hat{t}(\beta_0)$ implies $\partial/\partial\beta(M(s)/M(t))\big|_{\beta=\beta_0} < 0$ for $t_0 \le t \le s \le T-1$. Since D is continuous in t and β, $D(t)$ is uniformly continuous over compact domain $(t, \beta) \in [t_0, T-1] \times [0, \beta_0 + 1]$, where, without loss of generality, we have chosen $\beta_0 + 1$ as an upper bound on the β. Thus, there exists $0 < \delta < 1$ such that $\beta \in (\beta_0 - \delta, \beta_0 + \delta)$ implies $D(t) > 0$ for all $t \in [0, T-1]$. Finally, if $\hat{t}(\beta_0) > T-1$, set $t_0 = T$ (the trivial case). Applying Proposition 4 completes the proof. ■

Proof of Proposition 5. Without loss of generality, we will assume that the fixed cost c is paid at the end of each auctioning period, that is, (2) becomes

$$J_t(\mathcal{B}_t) = \max\Big[\max_j\{p_j w_j(\mathcal{B}_t)M(t)\}, -\delta c + \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]\Big]. \tag{A.21}$$

The proof of Proposition 2 showed that for all $\mathcal{B}_t \in \Omega_t$, $p_i w_i(\mathcal{B}_t)M(t) - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$ is nondecreasing in $w_i(\mathcal{B}_t)$ for all t and i. Since δc is just a constant, the same proof can be used to show that $p_i w_i(\mathcal{B}_t)M(t) + \delta c - \delta E[J_{t+1}(\mathcal{B}_{t+1})|\mathcal{B}_t]$ is nondecreasing in $w_i(\mathcal{B}_t)$ for all t and i, and the result follows.

To show that part 1 of Proposition 3 holds, we use an argument very similar to the original proof of part 1 of Proposition 3, whose notation we reuse here. Suppose that at time $t-1$ and sufficient statistic $\mathcal{B}_{t-1} = \mathcal{B}$, stopping the auction phase and entering the mass market is optimal. We show that stopping would also be optimal were the time instead t. Let $t+K$ be the optimal stopping time over horizon $[t, T]$ for a firm with sufficient statistic $\mathcal{B}$ at time t. If during horizon $[t-1, T]$ the firm pretends to have started from time t instead of $t-1$, it can do no better than if it optimizes its stopping time relative to its true starting point of $t-1$. Thus if stopping is optimal at time $t-1$, then

$$\sum_{k=1}^{T-t} \Pr(K=k)\delta^k E_{\mathcal{B}_{t-1+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t-1+k})\}|\mathcal{B}_{t-1} = \mathcal{B}]\frac{M(t-1+k)}{M(t-1)}$$
$$-\sum_{k=1}^{T-t} \Pr(K \ge k)\delta^k \frac{c}{M(t-1)} \;\le\; \max_i\{p_i w_i(\mathcal{B})\}.$$

Log-concavity of M over $[t-1, T-1]$ implies that $M(t-1+k)/M(t-1) \ge M(t+k)/M(t)$ for all $k = 1 \ldots T-t$. Together with $M(t-1) \ge M(t)$ and stationarity of the data updating process, this implies

$$\sum_{k=1}^{T-t} \Pr(K=k)\delta^k E_{\mathcal{B}_{t+k}|K=k}[\max_i\{p_i w_i(\mathcal{B}_{t+k})\}|\mathcal{B}_t=\mathcal{B}]\frac{M(t+k)}{M(t)} - \sum_{k=1}^{T-t} \Pr(K\geq k)\delta^k \frac{c}{M(t)} \leq \max_i\{p_i w_i(\mathcal{B})\}. \quad \text{(A.22)}$$

The LHS of the inequality above is the payoff of continuing auctioning at time t when following optimal stopping policy $t+K$. Thus, stopping the auction phase is optimal at time t and sufficient statistic $\mathcal{B}_t = \mathcal{B}$. ■

Proof of Proposition 6. The argument is similar in spirit to that used in the proof of Proposition 4, whose notation we reuse here. We first show that $\partial/\partial x(M(s)/M(t))\big|_x \geq 0$ and $\partial/\partial x M(t)\big|_x \geq 0$ for all $s \geq t$, $x \in I$, implies the stopping region shrinks with market size parameter, x, for $x \in I$. Let $x_1, x_2 \in I$, and $x_1 \leq x_2$. Suppose that at time t, with sufficient statistic $\mathcal{B}_t$ and market size parameter x_1, continuing the auction phase is optimal. We show that continuing the auction phase would also be optimal were the market size parameter $x_2 \geq x_1$. Let $M_k(t)$ be the market size at time t with parameter x_k, and S_k be the optimal stopping time if following an optimal policy as if the market size function is M_k, $k \in \{1, 2\}$. If continuing is optimal with parameter x_1, then

$$\sum_{s=t+1}^{T} \Pr(S_1=s)\delta^{s-t} E_{\mathcal{B}_s|S_1=s}[\max_i\{p_i w_i(\mathcal{B}_s)\}|\mathcal{B}_t]\frac{M_1(s)}{M_1(t)} - \sum_{s=t+1}^{T} \Pr(S_1\geq s)\delta^{s-t}\frac{c}{M_1(t)} \geq \max_i\{p_i w_i(\mathcal{B}_t)\}.$$

Note that this is the same equation as derived in the proof of Proposition 4, except for the term involving c, the cost of auctioning for another period. By assumption we have $M_1(s)/M_1(t) \leq M_2(s)/M_2(t)$, and $M_1(t) \leq M_2(t)$ implies $c/M_1(t) \geq c/M_2(t)$. Thus,

$$\sum_{s=t+1}^{T} \Pr(S_1=s)\delta^{s-t} E_{\mathcal{B}_s|S_1=s}[\max_i\{p_i w_i(\mathcal{B}_s)\}|\mathcal{B}_t]\frac{M_2(s)}{M_2(t)} - \sum_{s=t+1}^{T} \Pr(S_1\geq s)\delta^{s-t}\frac{c}{M_2(t)} \geq \max_i\{p_i w_i(\mathcal{B}_t)\}. \quad \text{(A.23)}$$

When the market parameter is x_2, the true optimal stopping policy S_2 in $[t+1, T]$ performs at least as well as when following the stopping time policy S_1 that would be optimal were the market parameter instead equal to x_1. Thus, (A.23) implies that continuing the auction phase is optimal at time t, sufficient statistic $\mathcal{B}_t$, and market size parameter x_2. Similar reasoning applies for the other cases of the proposition. ■

Proof of Proposition 7. The argument follows immediately from noting that the LHS of (A.22) decreases in c. ■