INCOMPLEX ALTERNATIVES TO MIXED BUNDLING*

Chenghuan Sean Chu
Stanford University

Phillip Leslie
Stanford University & NBER

Alan Sorensen
Stanford University & NBER

May, 2006

*Yes, "incomplex" is a word. We looked it up. We are grateful to Amitay Alter for research assistance, and to TheatreWorks for providing the data. Thanks also to Lanier Benkard, Garth Saloner and Andy Skrzypacz for helpful discussions.

1 Introduction

Under mixed bundling a firm that sells K products can set different prices for all $(2^K - 1)$ bundled combinations of its products. For two products there are three possible prices: two component prices and a price for the bundle of both products. For five products there are 31 possible prices, and for 10 products there are 1,023 possible prices. Perhaps unsurprisingly, observed pricing strategies of multiproduct firms tend to be much simpler than what is theoretically possible. In this study we examine the profitability of incomplex alternatives to mixed bundling, which have the benefit of being much simpler to implement in practice. We show that incomplex alternatives often yield profits within a few percentage points of mixed bundling. Hence, small computational or menu costs may explain the prevalence of these simplified pricing schemes.

We study the relative profitability of five alternative pricing strategies for a multiproduct monopolist with K products.[1] Mixed bundling (MB), as described above. Pure bundling (PB) in which the firm offers only the full bundle of K products at a single price. Uniform pricing (UP) in which the firm sets a single price which is the same for all products and there is no bundling. Component pricing (CP) in which the firm sets a different price for each product and there is no bundling. Lastly, we introduce an alternative pricing strategy that has not previously been discussed in the bundling literature: quantity discounts (QD) in which prices depend only on the number of products in the bundle.[2] In other words, under QD all bundles that contain the same number of products have the same price, but prices may vary for bundles of different sizes.

The bundling literature has focused on the relative profitability of CP, PB and MB. Of course MB always yields weakly higher profits than any the above alternative schemes, because each alternative is a restricted version of MB. More importantly, McAfee, McMillan and Whinston (1989) show that MB strictly dominates CP under rather general circumstances. But for firms with more than a few products, determing optimal prices under MB is a difficult problem—we can't even compute optimal MB prices for a firm that has more than 6 products with any reasonable degree of accuracy. This naturally leads to the question of how well do incomplex alternatives perform? One extreme is to consider pricing strategies with a single price: PB and UP. In fact Armstrong (1999) and Bakos and Brynjolfsson (1999) show that with zero marginal cost PB is equivalent to MB as the number of goods approaches infinity.[3] It is possible that PB

[1] As in the bundling literature, we assume there are no demand-side or supply-side complementarities that may also benefit bundling.

[2] We apply the assumption used throughout the bundling literature: individual consumers purchase at most one unit of each product, but can purchase any number of products.

[3] Fang and Norman (2006) compare the profitability of PB and UP when the number of products is relatively

may also perform well for a finite number of goods. But in this case it is conceivable that CP or QD, with the same number of prices as there are products, may do better than either PB or UP, and may even do well relative to MB.

Our goal is to characterize the circumstances under which each incomplex alternative yields profits close to that of MB. When is it more profitable to implement PB, CP or QD, and when is this a good approximation to MB? We focus on two aspects in our characterization: the number of products (K), and the joint distribution of consumers' tastes for these products. We find that ...

We also include an empirical example in the second part of the paper. The data come from a regional theatre company with an eight-play season. The firm currently offers three prices: a uniform price for each play, a discount for seeing any five plays, and a discount for the bundle of all eight plays. Importantly, we observe which five plays each consumer chooses under the five-play option. This allows us to identify the covariances in the joint distribution of consumers' tastes, which is a key determinant of profitability under alternative pricing schemes. We estimate a flexible demand system and then compute profits under counterfactual pricing strategies.

The example illustrates the kind of data needed to measure the profitability of MB. We also propose a novel demand specification that is well-suited to computing profits under the different pricing strategies. We find that ...

The remainder of the paper is organized as follows. In Section 2 we summarize the relevant prior literature. Section 3 contains our theoretical analysis of the various pricing strategies. The empirical example is presented in Section 4. Section 5 is the conclusion.

2 Prior Theoretical Literature

The bundling literature explores the idea that a multi-product monopolist can increase profits by selling goods in bundles, even when there are no demand-side complementarities or supply-side economies of scope. If a firm sells two products, and consumers vary in their willingness-to-pay for each product, then Stigler (1963) shows by example that selling these two products as a bundle (PB) may yield higher profit than if sold separately (CP). Adams and Yellen (1976) compare CP, PB and MB. They argue that the effectiveness for each these pricing strategies to

small. They show that either may dominate the other, depending on the underlying demand system.

enhance profit depends on the degree to which each satisfies the following conditions.[4] *Complete extraction:* no individual realizes any consumer surplus on his purchases. *Exclusion:* no individual consumes a good if the cost of that good exceeds his reservation price for it. *Inclusion:* any individual whose reservation price for a good exceeds its cost in fact consumes that good.

As Adams and Yellen note, these conditions are met by perfect price discrimination, which defines the maximum possible profit. Roughly speaking, CP is more likely to violate *inclusion* and PB is more likely to violate *exclusion*. MB provides the flexibility to optimally trade-off violations of *inclusion* and *exclusion*.[5] Adams and Yellen also argue that the greater the elasticity of demand for the bundled product, the less likely that PB violates *exclusion* or *inclusion*.

The analysis of Adams and Yellen is based upon judiciously chosen numerical examples. Two subsequent papers show that bundling (PB or MB) dominates CP in a wide variety of circumstances. First, Schmalensee (1984) expands the analysis to demand systems where consumers obtain product valuations from a bivariate normal.[6] Due to the limited computer power at the time, Schmalensee does not compute optimal MB prices, instead focusing on CP and PB. His main finding is that PB can be more profitable than CP even when the correlation of consumers' valuations is non-negative.[7] Second, McAfee, McMillan and Whinston (1989) extend the prior results by showing that MB strictly dominates CP under rather general circumstances.[8]

All the above papers analyze two product monopoly problems. A few prior papers study bundling with more than two goods. Bakos and Brynjolfsson (1999) focus on the profitability of PB as the number of goods (K) goes to infinity. They show that if goods have zero marginal cost, then as K goes to infinity PB approximates perfect price discrimination.[9] This finding is particularly interesting in our context, since it provides an example of an incomplex alternative to MB that closely approximates the profitability of MB in certain circumstances (i.e., large K).

[4]We restate their conditions verbatim. See Adams and Yellen (1976), page 481.

[5]Adams and Yellen (1976) also show that bundling may either increase or decrease social welfare relative to CP.

[6]A concern with this approach is that the bivariate normal implies negative valuations for some consumers which would impact the analysis in non-trivial ways, as noted by Salinger (1995). In all of the analysis in our study we allow for free disposal.

[7]The numerical examples in Stigler (1963) and Adams and Yellen (1976) somehow suggest the importance of negative correlation, as noted by Schmalensee (1984).

[8]McAfee, McMillan and Whinston (1989) also distinguish between firms that can monitor purchases or not. With monitoring, the firm can charge a price for the bundle of two that is higher than the sum of component prices. We limit our analysis to the no-monitoring case.

[9]As already noted, the assumption of zero marginal cost rules out violations of the *exclusion* criteria. Bakos and Brynjolfsson (1999) also show that, under certain conditions, increasing the number of goods under PB, monotonically increases profit. Geng, Stinchcombe and Whinston (2005) extend the analysis of Bakos and Brynjolfsson to incorporate diminishing marginal utility.

Armstrong (1999) provides a more general but similar result to Bakos and Brynjolfsson (1999). Armstrong shows that a two-part tariff, in which consumers are charged a fixed fee and can then purchase any products at marginal cost, achieves approximately the same profit as perfect price discrimination if the number of the products approaches infinity. In the special case of zero marginal cost the two-part tariff is equivalent to PB. The focus on settings with large numbers of products may be quite relevant for some firms, such as booksellers or supermarkets. But clearly these results are of questionable relevance to firms with 5 products, say.

Fang and Norman (2006) also examine the profitability of PB with more than two goods. In contrast to Armstrong (1999) and Bakos and Brynjolfsson (1999) they focus on finite K, and they seek to determine under what circumstances is PB an attractive pricing strategy. They confirm the prior intuition that increasing marginal cost tends to favor CP over PB. Fang and Norman also use numerical simulations to show that increasing the number of goods may favor PB over UP. We discuss their simulations in more detail below.

Spence (1980) and Hanson and Martin (1990) propose methods for computing optimal MB prices for $K \geq 2$ goods. Both papers assume a finite number of consumer types (defined as combinations of valuations for the K products). Naturally, the computational burden increases with both the number of goods and the number of consumer types. However, it is important to allow for a high degree of consumer heterogeneity, since one of the key benefits to bundling is heterogeneity reduction. If there is limited heterogeneity in the first instance, any pricing strategy that exploits heterogeneity reduction (MB, PB and QD) will have limited impact on profitability. Also, neither of these papers considers the relative profitability of CP and MB, and neither mentions QD.

To summarize the relevant prior results about monopoly bundling: MB yields strictly higher profits than CP or PB in a wide range of circumstances, but if the number of products is large and there is zero marginal cost then PB obtains approximately the same profit as MB.

The focus on two good examples in the prior literature makes a lot of sense. It is surely the case that the results of McAfee, McMillan and Whinston (1989) would generalize to $K > 2$ goods. Our starting point is that solving for optimal MB prices is a highly complex problem even for relatively small number of products. Furthermore, while it is possible to analyze the relative profitability of UP, PB, CP, QD, and MB in a two good setting, our claim is that the relative profits systematically depend on K. Hence, in evaluating the appeal of these incomplex alternatives, it is particularly important to analyze examples with $K > 2$.

Lastly, previous authors have offered various explanations for real-world firms' apparent lack of sophistication in pricing strategies. Blinder Canetti, Lebow and Rudd (1998) and Kahneman, Knetsch and Thaler (1986) discuss the possibility that firms avoid complex pricing schemes because consumers perceive such schemes to be unfair. Metrick and Zeckhauser (1999) propose that firms in some instances may be concerned that differential prices will be perceived by consumers as quality signals, discouraging demand for the lower-priced products. McMillan (2005) considers the possibility that menu costs alone can explain why firms implement UP instead of CP in some circumstances.

3 The Multiproduct Pricing Problem

In principle, multiproduct firms can employ a wide variety of pricing schemes. For a firm with K products, the optimal MB strategy requires setting $(2^K - 1)$ prices.[10] PB and UP require only one price to be set: the price for the bundle of all K products (in the PB case), or the per-product price (in the UP case). In between these extremes are CP, by which we mean setting K different prices for the K different products, and QD, by which we mean setting K prices that depend on the number of products purchased. Note that MB nests all the simpler pricing strategies as special cases, so it will always be weakly more profitable than any of these alternatives. Similarly, CP nests UP as a special case, and QD nests both UP and PB as special cases. CP and QD are non-nested alternatives.

In practice, multiproduct firms tend to use a wide variety of different pricing/bundling strategies. Consider baseball teams, for example, which have 81 home games (products).[11] For the 2006 season, the Los Angeles Dodgers offer several bundles of specific games, a discount for choosing any 27 games, and equal prices for all individual games. In contrast, for 2006 the San Francisco Giants do not offer any bundles or quantity discounts, but the Giants do vary prices by day of week and by opponent. Variation in pricing strategies is also evident in settings with fewer products. Consider the Steppenwolf Theatre in Chicago that produces a 5-play season. In 2006-07, they offer a discount for the 5-play bundle at a variety of prices that vary by time-of-week, and equal prices for individual shows (also varying by time-of-week). In 2006-07 the San Francisco Opera has a 10-opera season and offers 37 bundles (combinations of specific operas and time-of-week), and equal prices for individual shows (also varying by time-of-week). These examples highlight the dramatic differences in pricing strategies implemented by different firms in similar settings. We have been unable to find an example of MB being used in practice.

[10] We subtract 1 because the firm does not set the price for the outside good.

[11] Under MB, with 81 products a firm would choose 2.4×10^{24} prices.

CP and QD are of particular interest because the number of prices equals the number of products, which is a reasonable benchmark for practical pricing strategies. However, there are many other potential pricing strategies that also involve K prices, which are arbitrary subsets of MB. The problem in these cases is that it is ex-ante unclear which subset of K prices to choose. Ideally, a firm would compute the profits from all possible combinations that involve K prices, then implement the profit maximizing combination. This burdensome complication could be avoided if we knew in advance the circumstances under which each combination closely approximates the profitability of MB. As a starting point, we examine the circumstances under which QD is more profitable than CP, and to determine whether either of these strategies can be expected to yield profits "close" to the fully optimal MB profits.

In our analysis we focus on settings with zero marginal cost. This is not innocuous. Prior authors have shown that increasing marginal costs tend to reduce or eliminate the benefit of MB relative to CP.[12] The reason for this focus is partly because we include an empirical example which happens to have zero marginal cost, and partly because it simplifies the analysis.[13] Nevertheless, with zero marginal cost it is still possible for CP to dominate QD, or the reverse, and it is interesting to investigate the circumstances under which each strategy is preferred by firms. Examples of products for which bundling is relevant, and which may have zero marginal cost, include information goods, cable television, music, software, sports teams, and performing arts organizations.

Ideally, one would like to answer such questions with analytic solutions derived from a general model of demand. There are at least two obstacles to such an approach. First, as K increases, the region of integration becomes highly complex. Second, an important determinant to the profitability of bundling is the covariance of consumers' tastes for the K products. We are unaware of distributions with analytic integrals that allow for non-zero covariance (even if the regions of integration are simple). We therefore begin by offering a heuristic analysis of QD and CP in a two-good model with general demand, and then follow Schmalensee's (1984) approach by relying on numerical analysis of a specific demand model designed to capture the key features of the problem. Obviously, the results of our numerical exercises cannot be guaranteed to extend beyond the class of examples we consider. However, our setup captures the essential advantages and disadvantages of the alternative pricing strategies, and we suspect the patterns that emerge would hold in a much broader class of models

The analysis proceeds in two stages. First, we identify features of demand that favor QD

[12]With zero marginal cost, the *exclusion* condition is never violated. Consequently, on the benefits to CP is suppressed. See Fang and Norman (2006).

[13]In the theoretical analysis we also explore how increasing marginal cost may impact our findings.

or PB and features that favor CP, for a given number of products. Second, we analyze how increasing the number of products affects the profitability of CP, QD and PB.

3.1 Profitability of Different Pricing Strategies with Two-Goods

A useful starting point for identifying the features of demand that favor different pricing strategies is to examine a two-good setting. Even with only two goods, the five pricing schemes we consider differ meaningfully from each other. In order to build intuition about the factors that determine their relative profitability, we analyze a very simple setup with discrete consumer types. The firm sells two goods, both at zero marginal cost, to consumers whose valuations for each good can be either 1, 2, or 3. The optimal pricing strategy depends on the composition of consumer types, which can be depicted by plotting consumers in a 3×3 grid.

Figure 1 shows six examples, where in each case we illustrate a different composition of 9 consumers across the 9 consumer types, and above each grid we list a rank ordering of pricing strategies. In Example A, consumers' valuations for the two products are perfectly negatively correlated. This example illustrates the canonical rationale for bundling: while the consumers' valuations for the individual goods are heterogeneous (between 1 and 3), their valuations for the bundle of both goods are completely homogeneous (i.e., 4). The implication is that bundle pricing (via MB, QD, or PB) can extract all the surplus, while CP and UP cannot. Example B shows the opposite extreme, where consumers valuations are perfectly positively correlated. In this example, *all* of the pricing strategies we consider are equivalent. Bundling confers no advantage over uniform pricing because it does nothing to reduce heterogeneity in willingness to pay, and component pricing confers no advantage because demand for the two goods is symmetric.[14]

While negatively correlated tastes favor bundling strategies like PB and QD, asymmetry in demand favors CP. Example C shows that even with perfect negative correlation in tastes, CP can be more profitable than QD if the products' demands are asymmetric enough. In this example, PB and QD are equivalent: both involve selling the bundle to the $(v_1 = 1, v_2 = 3)$ types for a price of 4, but selling nothing to the $(v_1 = 2, v_2 = 1)$ types. QD does not outperform PB because the additional single-good price is of no use: any single-good price that induces purchase by the $(v_1 = 2, v_2 = 1)$ types would not be incentive compatible with the $(v_1 = 1, v_2 = 3)$ types purchasing the bundle. In contrast, by setting component prices of 1 and 3, CP extracts all the surplus from the high types (like PB and QD), but without abandoning the low types altogether.

[14]Note that it doesn't matter how we distribute consumers along the "diagonal" of the grid: for example, if there were 5, 3, and 1 consumers along the diagonal instead of 3, 3, and 3, the pricing strategies would still all be equivalent.

(The $(v_1 = 2, v_2 = 1)$ types purchase good 1.) Note that conditional on having only these two types of consumers, CP is more profitable than QD only if a large fraction of consumers are the (1,3) type. As shown in Example D, reversing the weighting of consumer types also reverses the profitability ranking of CP and QD. In Example D, QD sells the bundle to *all* consumers at a price of 3, while CP sets prices of 2 and 1 and leaves more surplus to the high types.

Both Examples C and D exhibit demand asymmetry in some sense. And in both cases, demand is perfectly negatively correlated. But in one case CP beats QD, and in the other case QD beats CP. One way to describe the difference between these two examples is that demand is “top-heavy” in Example C. That is to say, there is a high concentration of consumers with a high valuation for one of the products. In the next subsection we propose a summary measure of “top-heavy asymmetry”, and we show that this is an important determinant of profitability under CP.

The equivalence of PB and QD in the first four examples is not surprising, because bundling is the main purpose of both. In principle, however, QD should have an advantage over PB: the additional flexibility to set a one-good price allows QD to capture surplus from consumers with a high valuation for one good but a relatively low valuation for the bundle. The demand configuration shown in Example E illustrates this advantage. In this configuration, PB sells the bundle to all types at a price of 4. QD does better by selling the bundle for a price of 5 and single goods for a price of 3, which extracts the full surplus of the seven high types while still inducing purchase by the $(v_1 = 1, v_2 = 3)$ types.[15] Nevertheless, in this example QD is less profitable than CP, which extracts almost all of the potential surplus with component prices of 2 and 3. Example F shows a five-type configuration where QD is strictly more profitable than both PB and CP.

Notice that in all but one of the six examples shown in Figure 1, the three-price mixed bundling outcome can be achieved with a simpler two-price alternative. In fact, of the 24,310 possible configurations of 9 consumers into this 3×3 grid, MB is strictly more profitable than both CP and QD in only 2,890 cases. Interestingly, the average difference in profit (over all possible configurations) between MB and QD is 0.56, while the average difference between MB and CP is 1.85.[16] Moreover, MB profits never exceed QD profits by more than 4, but can exceed CP profits by as much as 14. Similarly, when CP is preferred to QD, the profit difference is typically small (never more than 4), but when QD is preferred to CP the difference can be large

[15] For simplicity we assume that if a single-good purchase and a bundle purchase deliver the same utility, the consumer will purchase the bundle. We could instead say the QD prices are 3 and 5-ϵ to get the same outcome.

[16] For reference, total profits are at least 18 and at most 54 in this setup.

(as much as 14).[17]

3.2 The Importance of Correlation and Asymmetry

The examples in Figure 1 suggest that: (i) negative correlation is good for bundling (PB, QD and MB); and (ii) asymmetry of a particular kind is good for CP. In this subsection we provide a more formal analysis of the role of correlation and asymmetry in our 3×3 framework with two goods.

A benefit of examples with two goods is that the correlation in tastes between the two goods provides an obvious summary of correlation.[18] Again using our simple setup with 9 possible consumer types, we compute the correlation in tastes for each of the 24,310 possible demand configurations. Figure 2 shows the average profitability of MB, QD and CP as a function of demand correlation.[19] To construct this figure, we compute the average profit under each pricing regime for each decile of the distribution of the correlation coefficient. Hence, for any given level of correlation, there is variance in the profitability of each pricing strategy, and we depict only the mean.

Figure 2 shows an interesting pattern that emerges even from this simple setup. The average profitability of QD and MB is highly dependent on the degree of negative correlation. The figure also verifies that bundling may still dominate CP even with positive correlation, as previously shown by Schmalensee (1984). Correlation, at least according to this measure, is not the only determinant of the profitability of MB and QD. But the figure suggests that correlation is a primary determinant of the profitability of MB and QD.

In contrast to QD, Figure 2 shows that the profitability of CP is roughly invariant to changes in correlation. The reason is because profitability under CP depends only on the marginal distributions of consumers' valuations for each product. In other words, the firm's profit maximization problem is equivalent to K separate individual profit maximization problems. There is no additional benefit to jointly solving the K optimal prices. Hence, correlation in tastes across products is irrelevant to CP.

[17]Note that we consider all 24,310 possible configurations of 9 consumers, even though some of these are redundant due to the symmetry of our setup. The above statements still hold even if we do not double-count symmetric cases.

[18]For three or more goods the correlations between all possible pairs needs to somehow be aggregated.

[19]Note that profits are normalized by the total surplus, since the total surplus varies across the 24,310 demand configurations.

What factors affect the profitability of CP? As noted above, Examples C and D suggest a particular form of demand asymmetry favors CP: asymmetry that is "top-heavy". We propose a measure of the degree to which there is top-heavy demand asymmetry, allowing us to distinguish demand systems "like" Example C from demand systems "like" Example D. Define $n_j(v)$ as the number of consumers whose for value good j equals v. The total surplus for good j is defined as $s_j = \sum_{v=1}^{3} n_j(v)v$. Our proposed measure of asymmetry is given by

$$\Lambda = w_1 \sum_{v=1}^{3} \left[\frac{n_1(v)v}{s_1} \right]^2 + w_2 \sum_{v=1}^{3} \left[\frac{n_2(v)v}{s_1} \right]^2 ,$$

where $w_j = s_j/(s_1 + s_2)$, the relative importance of each good to total surplus. Note that $n_1(v)v$ equals the surplus for good 1 that is attributable to consumers whose value for good 1 is equal to v. Hence, Λ is a weighted sum of the surplus attributable to each product for each possible consumer valuation.

Λ is higher the more concentrated is demand in one product at a high value. Λ is lowered by: (i) similarity in demand (symmetry in the extreme) across products; (ii) variance in valuations for a given product; and (iii) concentration in demand for a given product at low values. As an illustration, in Example C we obtain a value of $\Lambda = 0.74$, and in Example D we obtain $\Lambda = 0.66$. Hence, Example C exhibits greater top-heavy asymmetry than Example D.

Figure 3 is the analog to Figure 2 based on our measure of asymmetry instead of correlation. In this case the profitability of all three pricing strategies is increasing in Λ. It would be ideal to construct a measure that is correlated with the profits of CP and uncorrelated with the profits of QD and MB. While this measure does not do that, it is nevertheless correlated with the relative profitability of CP and QD or MB. As asymmetry increases, CP does better relative to QD, and for high enough asymmetry CP is more profitable than QD.

The appeal of Λ is that it captures the notion of top-heavy asymmetry. To emphasize the importance of the top-heavy aspect of this characterization, we compare the performance of Λ to a couple of obvious alternatives that measure asymmetry without taking into account "top-heaviness". One alternative is to measure asymmetry in demand by the difference in mean valuations of the two products. A second alternative is the maximum surplus of a single good as a fraction of the total surplus of the two goods. The performance of Λ is compared with these two alternatives in Figure 4. The metric for comparison is the profitability of CP relative to QD. A useful measure of asymmetry would identify when CP does well relative to QD. Figure 4 shows that both alternatives are also positively correlated with the profitability of CP. But it is clear that Λ is more informative. For example, Λ is the measure that identifies demand systems for which CP obtains higher profit than QD, in these experiments.

3.3 Numerical Analyses for $K \geq 2$

Comparing the profitability of alternative pricing strategies is impossible without making specific assumptions about consumer demand (i.e., the probability distribution of consumers' willingness to pay), and the setup described in the previous section is just one with particularly stark assumptions. Although that setup offers some useful insights, we want to understand how the pricing strategies perform in more general settings. Ideally we want to allow for richer heterogeneity in consumer types, and we especially want to know what happens as the number of goods increases. Unfortunately, even after specifying a probability distribution, the optimal prices and profits under the various pricing strategies typically cannot be obtained in closed form, because calculating the demand for a bundle of two or more products involves integrating probability distributions over non-rectangular regions.[20] Our approach will therefore be to use numerical analyses of a variety of demand specifications, with the idea that these analyses will reveal patterns that might be expected to hold more generally.

As a starting point, consider a demand model in which consumer i's utility from purchasing bundle j is equal to $V_i' D_j - p_j$, where V_i is a $K \times 1$ vector of valuations for the firm's K products, D_j is a $K \times 1$ vector of binary indicators for which of the K products are included in bundle j, and p_j is the price of bundle j.[21] Each consumer's problem is to choose the offered bundle that maximizes her utility. Consumers' product valuations are heterogeneous: V_i is drawn from a (multivariate) distribution F. Given this setup, our goal in this section is to evaluate the profitability of alternative pricing strategies for various distributions (F) and for different numbers of products (K).

Table 1 summarizes the simulated profit outcomes for a number of different distributional assumptions. The first panel is the case with independent exponentials. Specifically, we perform the following simulation exercise:

Step 1: Randomly draw the K products' mean valuations from the uniform distribution: $\theta_k \sim U[0, 1]$.

Step 2: Simulate the valuations of n consumers by taking independent draws from exponential

[20] An exception is the two-good case with exponentially distributed tastes; for that case, there are closed-form solutions for the optimal prices and profits under CP, QD, and MB even though the regions of integration are non-rectangular. The integrals are still analytic as you increase the number of goods, but the regions of integration get very complex, so the solutions are quite messy.

[21] Note that by assuming additive preferences we are ruling out consumption complementarities as a motivation for bundling. As with the bundling literature more generally, our focus is on the price discrimination motive. See Crawford (2005) for a nice discussion of the various motives for bundling.

distributions: $v_{ik} \sim \text{Exp}(\theta_k)$.

Step 3: Given these consumer valuations, numerically calculate optimal prices and profits under the various pricing strategies.

We repeat steps 1–3 many times and then summarize the outcomes. To calculate the optimal prices and profits in step 3, we initially consider the case in which marginal costs are zero. We consider simulations with positive marginal costs in the next subsection below.

Before summarizing the outcomes of these simulations, it is important to acknowledge the obvious limitations of our approach. Although we attempt to cover a large space of parameter values, the results clearly depend on how we define that parameter space and how we sample from it. Further, there is no way for us to know whether we are under- or over-sampling the relevant (i.e., empirically plausible) combinations of parameters. So, for example, when we report "average" outcomes, these should certainly *not* be interpreted as outcomes that should be expected in an empirical sense; they should be interpreted narrowly as the averages of the outcomes we simulated.

4 An Empirical Example

The numerical simulations verify the importance of the joint distribution of consumers' tastes for the K products in determining the relative profitability of UP, PB, CP, QD and MB. In this section we address the problem of how to estimate the joint distribution of consumers' valuations from available data. In contrast to standard econometric models of demand, it is important that such an approach allow for non-zero covariances in tastes, and the ability for individuals to purchase multiple products. We estimate such a demand model using data from a theatre company that offers an 8-play season. Based on these estimates, we compute the profitability of each pricing strategy.

We also measure the impact of each pricing strategy on consumer welfare in this market. This is interesting because bundling, like price discrimination more generally, has ambiguous affects on consumer welfare relative to uniform pricing.[22] To the best of our knowledge, there is only one prior empirical analysis of bundling. Crawford (2005) tests the hypothesis that consumers' demand for a bundle of cable channels becomes less heterogeneous as more channels

[22]See Leslie (2004) for a similar empirical analysis of the welfare effects of price discrimination, which also happens to be in the context of theatre ticket pricing.

are added to the bundle, which he finds to be the case. Based on a calibrated demand model, Crawford argues that adding a top-15 cable channel to a bundle and re-optimizing prices leads to 5.5% lower consumer surplus, and 6.0% higher profit.

In the following subsections we summarize the data, present the demand model, and then implement the estimation and counterfactual pricing experiments.

4.1 Data Summary

The data for our empirical analysis come from TheatreWorks, a theatre company based in Palo Alto, California. We observe all ticket sales for TheatreWorks' 2003–2004 season, which consisted of 229 performances of 8 different plays or musicals. Table 2 provides summary information for each of the 8 plays.

Consumers may purchase tickets to individual plays at a uniform price. TheatreWorks also offers three subscription packages: (i) the full 8-play season; (ii) any combination of 5 plays; or (iii) a pre-specified bundle of 3 plays.[23] These subscriptions were offered at discounted prices, in the sense that the per-play price was significantly lower for subscriptions than for ordinary box office sales for individual plays. Table 3 summarizes the purchase options and their average prices.[24]

As shown in the table, a substantial portion of total ticket sales came through subscriptions. A total of 103,078 tickets were sold to the eight plays. Of these, 88,833 were tickets purchased as part of a subscription package. There were 6,101 buyers of the 8-play bundle, 2,794 buyers of a 5-play bundle, and 245 buyers of the 3-play bundle. The data is rich enough that we also observe non-subscribers that buy tickets to multiple plays. In this case, we observe 8,131 buyers of a single play, 1,409 buyers of two plays, 555 buyers of three plays (different to the pre-set bundle of three) and 224 buyers of four plays.

The popularity of the flexible 5-play subscription is a particularly important feature of the data. Observing which five plays these subscribers' selected allows us to identify the covariance of tastes across plays—e.g., if we observe that two plays tend to be included disproportionately often in the five-play combination, we know that tastes for those two plays are positively corre-

[23]The pre-specified bundle consisted of the only 3 plays that were performed at TheatreWorks' secondary venue, a smaller theater in Palo Alto, CA.

[24]In fact prices also vary by time-of-week (but not by play). We therefore report simple (unweighted) averages of these prices. Note also, prices do not vary by seat quality. This is because the venues are small enough that the variation in seat quality is fairly minor.

lated. Conversely, if another pair of plays is rarely included in the same bundle, we can infer that tastes for those two plays are negatively correlated. Table 4 summarizes the correlations implied by the *pick-five* purchases: it reports the difference between the empirical correlations of the choices and the correlations that would be expected if tastes were independent.[25] The patterns make intuitive sense. For example, tastes for *Bat Boy*, described in the brochure as a "wacky new musical," are positively correlated with tastes for *Memphis*, described as a "rafter-rattling musical comedy." Conversely, tastes for *Bat Boy* are negatively correlated with *All My Sons*, a classic Arthur Miller drama billed as an "intense, compelling tale of love, greed, and personal responsibility."

The set of products for each consumer to pick from includes each individual play, the preset bundle of three, the full bundle of eight, and 56 possible combinations of five plays. However, an individual consumer may create other bundles by adding individual plays. For example, a consumer can bundle any four plays, although there is no discount for doing so. Or a consumer can purchase a bundle of five, and add a sixth play at the regular single play price. Defining products to include every conceivable bundle implies 255 choices. In fact we observe zero sales of bundles of six or seven plays. We therefore ignore these combinations in the estimation. Hence, we model the demand for 219 different bundles, plus an outside alternative, giving a total of 220 possible choices. Lastly, we note that capacity constraints are infreqently binding—27 of the 229 performances were sold out—leading us to abstract from their impact in the estimation and counterfactual analyses.

4.2 Empirical Model

The empirical specification is based on an underlying model of individual consumer utility maximization. The firm offers $j = 1, ..., J-1$ bundles containing combinations of the $k = 1, ..., K$ products. There is also a J^{th} option for consumers which is the outside alternative. We assume the net utility to consumer i from option j is given by

$$u_{ij} = \begin{cases} V_i' D_j - \alpha_i p_j - \beta_i q_j & : \quad j = \{1, ..., J\} \\ 0 & : \quad j = J \end{cases}$$

where V is an $K \times 1$ vector of valuations for the individual plays, D_j is an 8×1 vector of indicators for whether each play in included in bundle j, p_j is the price of the bundle, $q_j = \sum_k D_j(k)$ is the quantity of plays in the bundle, and $\alpha_i > 0$ and $\beta_i > 0$ are the individual's sensitivity to price

[25]Since each consumer selected five plays of eight, the pairwise correlations will be nonzero even if tastes are independent. The expected correlation if plays are chosen independently is -1/7.

and quantity, respectively. As mentioned above, we estimate demand for 219 bundle options, plus an outside alternative with zero utility for all consumers. Hence, in our notation $J = 220$.

Following the approach taken in the theoretical bundling literature, we assume there are no demand-side complementarities from consuming particular plays together.[26] However, we depart from the standard utility function in the bundling literature in two ways that are important for our empirical setting. First, we incorporate heterogeneity in consumers' sensitivity to price. Second, we allow for diminishing marginal utility, and we also allow consumer heterogeneity in this parameter. In the next subsection we discuss the identification of these elements of the model, and hence why they are important in our empirical setting.

Individual consumers are defined by their own combination of V_i, α_i and β_i. We estimate the population distribution of each component, and we assume the following parametric distributions:

$$\begin{aligned} V &\sim \mathrm{N}\left(\mu_v, \Sigma_v\right), \\ \alpha &\sim \mathrm{LN}\left(\mu_\alpha, \sigma_\alpha\right), \quad \text{and} \\ \beta &\sim \mathrm{Exp}\left(\mu_\beta\right). \end{aligned}$$

Note that V has support from $-\infty$ to $+\infty$. Hence, to incorporate free disposal, we assume $V(k) = \max\{V(k), 0\}$, $\forall k$. We also impose the restriction that the variance for each marginal distribution of V equals one: $\Sigma_v(k,k) = 1$, $\forall k$. The reason is because the conditional means of V appear not to be empirically separately identified from the covariances. We place no restrictions on the 28 remaining covariance terms in Σ_v, or any of the other parameters in the model. We therefore estimate 39 parameters. Let Θ denote the set of parameters to be estimated.

Let n_j denote the observed number of consumers choosing option j. To determine the number of consumers choosing the outside alternative, we assume a market size of 50,000.[27] It is impossible to know the correct market size, but this seems reasonable based on the population of Palo Alto and the surrounding suburbs, and the fact that consumers of TheatreWorks' shows tend to be elderly.[28]

The model's parameters are estimated via simulated maximum likelihood. We simulate the probability of a consumer choosing each option. For a given set of parameters, Θ, we draw ns

[26] On the supply side, we will also assume zero marginal cost, which also precludes supply-side complementarities.

[27] We check robustness of our estimates to alternative values of market size.

[28] In 2000, the population of Palo Alto was 58,598. Two nearby suburbs are Mountain View (population 70,708) and Los Altos (population 27,693). The combined population is 156,999, and we consider one third of this to be a reasonable measure of the market of potential consumers (mainly elderly people).

simulated consumers from the above distributions of consumer heterogeneity. For each simulated consumer we compute the optimal bundle choice, implying predicted market shares for all 220 choices: s_j, such that $\sum_{j=1}^{J} s_j = 1$. The log-likelihood function is given by:

$$l(\cdot, \Theta) = \sum_{j=1}^{J} n_j \log s_j(P, \Theta),$$

where P is the $(J-1)$-dimensional vector of observed prices. The vector of estimated parameters is the value of θ that maximizes the log-likelihood function.

As is well known, simulated maximum likelihood estimators are biased due to the convexity of the log operator, but are consistent as long as the number of simulation draws (ns) scales up in proportion to $\sqrt{N}$, the squareroot of the sample size. In practice, we try values of ns ranging from 50,000 to 500,000.

Another problem is that when the number choices is large and mean utilities for different choices are very different, a simple frequency simulator will often produce zero probabilities for some choices, leading to an undefined log-likelihood. This problem could be mitigated using advanced techniques for drawing from joint densities, such as Gibbs sampling or the Metropolis-Hastings algorithm. We rely on a simpler method. If there exist choices such that $s_j(\Theta, \cdot) = 0$, we simply remove a very small probability measure δ from the choices with $s_j > 0$ (in proportion to the shares of those elements) and redistribute it equally amongst the zero-share choices. The resulting shares still sum to one and are all strictly positive.[29]

4.3 Identification

What variation in the data serves to identify each parameter of the demand model? The mean terms in the joint distribution of consumers valuations, μ_v, are identified by the sales to individual ticket buyers for each play. The covariance terms in the variance-covariance matrix of V, Σ_v, are identified mainly by the combinations of choices made by the *pick-five* buyers. These subscribers reveal which plays consumers like putting together and which plays they don't like putting together. As noted above, we assume the variance terms in Σ_v all equal one.

Price variation in the data comes in the form of differences in per-play prices across bundles. This variation serves to identify consumers' mean sensitivity to price. A standard concern with demand estimation is the possibility that observed prices are correlated with unobserved demand

[29]In practice, we take δ to be 10^{-10}, and then remove a measure δs_j from each non-zero choice.

shifters, which may bias parameter estimates. However, in the estimation we integrate over all unobserved demand components—there is no remaining error term that may be correlated with observed prices. A consequence of this approach is that we rely on sampling error to explain discrepencies between predicted and actual outcomes.[30] This suggests adding an error term to the demand model, to reduce the likelihood of rejecting the estimated model. But given the rich specification of consumer heterogeneity that is already present in the model, it is unclear how to interpret any additional unobserved demand component.

The empirical model includes two features that are non-standard in the utility functions that appear in the theory literature on bundling: heterogeneity in price sensitivity, and diminishing marginal utility. We include them because they help us to explain some important features of the data.

First, recall from Table 3 that a large fraction of consumers choose the bundle of all eight plays. Consumers that buy this bundle, have relatively high valuations for all eight plays, or at least for most of the plays. To get a large number of consumers with high valuations for almost all plays requires strong positive covariances in consumers' valuations. However, this is at odds with the large number of consumers that attend only one play. If a consumer has a high enough valuation to attend one play, with strong positive correlations for all plays this consumer is likely to have high valuations for many more plays. How can we explain so many single-play buyers?

To explain this pattern of sales, we introduce heterogeneity in the marginal utility of attending more than one play: β_i. The higher is β_i, the less likely a consumer will want to attend multiple plays. This allows us to explain why a consumer would choose to attend one or few plays, even though they may also have high gross valuations for other plays. Hence, given the covariances that are identified from the *pick-five* buyers, the distribution of β_i (i.e. μ_β) is identified by the difference between the actual number of single-play buyers, and the lower number that would be predicted if there was no diminishing marginal utility.

Since we assume that β_i has an exponential distribution, we expect fewer buyers for bundles with successively more plays. But this pattern is not evident in the data. Table 3 shows that the number of 2-play and 3-play buyers is lower than the number of 5-play and 8-play buyers. Some of this may be explained by the fact that a per-play discount is offered for 5- and 8-play bundles, but not for 2-play bundles, and a discount is offered for only one specific 3-play bundle. However, there is also tension between the covariances implied by the choices of the *pick-five* buyers, and the fairly large number of 8-play subscribers that suggest high positive covariances

[30]Ordinarily the difference between predicted and actual outcomes would be explained by an error term in the demand model which has not been integrated out.

for all plays. We explain these aspects of the data by the incorporating heterogeneity in price sensitivity. Consumers with low price sensitivity tend to buy bundles with more plays. In this way, we are able to identify the variance of the distribution of price sensitivity: σ_α.

A final comment on the flexibility of the model. Notice from Table 2 that the rank-ordering of play popularity is different for the single-ticket buyers than it is for subscribers (which is driven by the tastes of *pick-five* buyers because full season buyers attend all plays). Our specification can explain this difference in the following way. The mean quality terms, μ_v, explain the relative popularity of plays among the single-ticket buyers. The covariance terms in Σ_v explain the popularity of certain pairings by the *pick-five* buyers, which also helps to explain differences in the popularity of plays that is contrary to the ranking by the single-ticket buyers.

4.4 Results

4.5 Counterfactual Pricing Experiments

Given our estimates of the demand system, we can calculate the optimal prices and revenues under various counterfactual pricing regimes. Table 5 shows the outcomes.

5 Conclusion

Bibliography

Adams, W.J. and J.L. Yellen (1976): "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics*, 90, 475–98.

Armstrong, M. (1999): "Price Discrimination by a Many-Product Firm," *Review of Economic Studies*, 66(1), Special Issue: Contracts, 151–68.

Bakos, Y. and E. Brynjolfsson (1999): "Bundling Information Goods: Pricing, Profits and Efficiency," *Management Science*, 45(12), 1613–30.

Blinder, A.S., E.R.D. Canetti, D.E. Lebow and J.B. Rudd (1998): "Asking About Prices: A New Approach to Understanding Price Stickiness," *Working Paper, Russell Sage Foundation.*

Crawford, G.S. (2005): "The Discriminatory Incentives to Bundle in the Cable Television Industry," *Mimeo*, University of Arizona.

Fang, H. and P. Norman (2006): "To Bundle or Not to Bundle," forthcoming in *RAND Journal of Economics.*

Geng, X., M.B. Stinchcombe, and A.B. Whinston (2005): "Bundling Information Goods of Decreasing Value," *Management Science*, 51(4), 662–7.

W. Hanson and R.K. Martin (1990): "Optimal Bundle Pricing," *Management Science*, 36(2), 155–74.

Kahneman, D., J.L. Knetsch and R. Thaler (1986): "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review*, 76(4), 728–41.

Leslie, P. (2004): "Price Discrimination in Broadway Theatre," *RAND Journal of Economics*, 35(3), 520–41.

McAfee, R.P., J. McMillan and M. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *Quarterly Journal of Economics*, 114, 371–84.

McMillan, R. (2005): "Different Flavor, Same Price: The Puzzle of Uniform Pricing for Differentiated Products," *Mimeo*, Federal Trade Commission.

Metrick, A. and R. Zeckhauser (1998): "Price Versus Quantity: Market Clearing Mechanisms when Consumers are Uncertain about Quality," *Journal of Risk and Uncertainty*, 17(3), 215–42.

Schmalensee, R. (1984): "Gaussian Demand and Commodity Bundling," *Journal of Business*, 57(1), 211–30.

Spence, A.M. (1980): "Multi-Product Quantity-Dependent Prices and Profitability Constraints," *Review of Economic Studies*, 47(5), 821–41.

Stigler, G.J. (1963): "United States v. Loew's Inc.: A Note on Block Booking," *Supreme Court Review*, 152–7.

Table 1. Simulation results for various distributions of consumer tastes

	K	π_{QD}/π_{CP} min	π_{QD}/π_{CP} median	π_{QD}/π_{CP} max	π_{QD}/π_{MB} min	π_{QD}/π_{MB} median	π_{QD}/π_{MB} max
Exponential (1)	2	1.00	1.11	1.16	0.97	1.00	1.00
	3	1.00	1.19	1.26	0.96	1.00	1.00
	4	1.00	1.26	1.33			
	5	1.09	1.31	1.38			
Extreme Value (2)	2	1.04	1.07	1.09	0.98	0.99	1.00
	3	1.08	1.12	1.15	0.98	1.00	1.00
	4	1.12	1.15	1.18			
	5	1.15	1.18	1.21			
Uniform (3)	2	0.99	1.04	1.12	0.94	0.97	1.00
	3	0.99	1.11	1.17	0.94	0.98	1.00
	4	1.00	1.16	1.23			
	5	1.02	1.20	1.26			
Normal (4)	2	0.99	1.03	1.09	0.96	1.00	1.00
	3	1.01	1.07	1.14	0.98	1.00	1.00
	4	1.04	1.10	1.20			
	5	1.06	1.13	1.22			
Normal (5)	2	0.99	1.03	1.10	0.98	1.00	1.00
	3	1.00	1.05	1.10	0.98	0.99	1.00
	4	1.02	1.08	1.16			
	5	1.03	1.10	1.18			
Normal (6)	2	0.74	1.06	1.11	0.74	0.99	1.00
	3	0.88	1.11	1.17	0.90	0.99	1.00
	4	0.98	1.15	1.22			
	5	1.02	1.19	1.26			

Ratios of optimal QD profits to CP and MB profits, respectively. Each row reports the minimum, median, and maximum value from 1,000 simulations. The details of the preference distributions are as follows: (1) Consumers' valuations for the products are independent exponential random variables, with product means drawn randomly from U[0,1]. (2) Consumers' valuations are independent extreme value random variables, with means drawn randomly from U[1,5] and scale parameter equal to 1 (i.e., standard deviation $= \pi/\sqrt{6}$). (3) Valuations are independent uniform random variables on the interval $[0, a_k]$, with the a_k's drawn randomly from U[0,10]. (4) Valuations are independent normal random variables, with means drawn randomly from U[-1,1] and variances set to 1. (5) Same as (4), but the valuations are positively correlated across products ($\rho = 0.5$ for each pair of products). (6) Valuations are independent normal random variables, with means drawn randomly from U[0,2] *and* variances drawn randomly from U[0,2].

Table 2. Summary of ticket sales for 8 plays

Play	Type	Number of Performances	Average Attendance	Ticket sales (subscription)	Ticket sales (non-subscription)
A Little Night Music	Musical	30	327.63	8,001	1,828
All My Sons	Drama	33	263.88	7,817	891
Bat Boy	Musical	30	297.63	7,793	1,136
Memphis	Musical	30	384.70	7,968	3,573
My Antonia	Drama	26	348.69	7,946	1,120
Nickel and Dimed	Drama	26	380.42	7,757	2,134
Proof	Drama	25	357.68	7,830	1,112
The Fourth Wall	Comedy	29	348.86	8,401	1,716
Total		229	336.34	63,513	13,510

Three plays (*Bat Boy*, *All My Sons*, and *The Fourth Wall*) were performed at the Lucie Stern Theatre in Palo Alto (capacity=428). The remaining five were performed at the Mountain View Center for the Performing Arts (capacity=589).

Table 3. Sales by purchase option

Purchase option	Price per play ($)	Number of consumers
Non-subscription:		
1 play	40.80	8,131
2 plays	40.80	1,409
3 plays	40.80	555
4 plays	40.80	224
Subscription:		
3-play bundle	36.20	245
5-play pick	37.00	2,794
8-play bundle	34.55	6,101

The 3-play subscription bundle was for the specific three plays performed at the (smaller) Lucie Stern Theatre in Palo Alto, which is why the per-play price is lower than the 5-play bundle. Consumers purchasing the 5-play subscription could combine any 5 plays of their choice.

Table 4. Correlation of tastes for pick-5 bundles

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) *A Little Night Music*	.000							
(2) *All My Sons*	-.026	.000						
(3) *Bat Boy*	.067	-.233	.000					
(4) *Memphis*	.072	-.081	.257	.000				
(5) *My Antonia*	.177	.067	-.086	-.037	.000			
(6) *Nickel and Dimed*	-.160	-.009	-.013	-.039	.001	.000		
(7) *Proof*	-.066	.210	-.030	-.057	-.094	.008	.000	
(8) *The Fourth Wall*	-.038	.034	.003	-.117	-.007	.196	.008	.000

This is the difference between the observed correlation matrix and the correlation matrix that would be expected if plays were chosen independently (i.e., no correlation in tastes). It is constructed for the 8,005 purchasers of the flexible 5-play subscription.

```
Estimates
=========

Play means (mu_v):

 0.173   0.018   0.039   0.447   0.061   0.104   -0.083   0.197

Covariance (Sigma_v):

 1.000
 0.022    1.000
-0.020   -0.100    1.000
-0.050   -0.050   -0.014    1.000
 0.125    0.069   -0.007   -0.028    1.000
-0.006   -0.025    0.016   -0.000    0.032    1.000
 0.045    0.072   -0.006   -0.073    0.089    0.084    1.000
 0.030    0.036    0.162   -0.003    0.090    0.190    0.217    1.000

Log Price sensitivity (alpha)

 1.936 (mean)
 0.160 (variance)

Diminishing returns parameter (beta)

 0.606


Fit statistics:
===============
                   Play1  Play2  Play3  Play4  Play5  Play6  Play7  Play8
[Overall sales]
Actual shares       .128   .113   .116   .150   .118   .128   .116   .131
Predicted shares    .125   .117   .120   .147   .120   .125   .112   .135

[Pick-5 sales]
Actual shares       .136   .105   .104   .134   .132   .119   .124   .147
Predicted shares    .125   .107   .112   .137   .125   .129   .114   .150

[Individual sales]
Actual shares       .131   .109   .089   .274   .094   .130   .082   .092
Predicted shares    .133   .102   .094   .298   .090   .116   .060   .108
```

Table 5. Counterfactual pricing

	UP	PB	TW	CP	QD
p_1	15.50		18.81	15.55	19.30
p_2				13.79	16.95
p_3			14.38	14.77	14.88
p_4				15.56	13.07
p_5			11.13	18.53	11.47
p_6				15.60	10.08
p_7				13.99	8.85
p_8		5.71	7.73	15.59	7.77
Revenue	17.80	18.10	18.96	17.97	19.88
CS	13.57	14.00	12.92	13.64	13.87

For UP, p_1 is the optimal uniform price for a single play. For PB, p_8 is the optimal per-play price for the bundle of all eight plays. TW is the pricing scheme currently employed by the theater company: p_1 is the single-play price, p_3 is the per-play price for a specific bundle of three plays, p_5 is the per-play price for any combination of five plays, and p_8 is the per-play price if you buy all eight. For CP, p_1-p_8 are the prices for the eight individual plays, and for QD, p_1-p_8 are the per-play prices for any bundle containing the corresponding number of plays. The revenue and consumer surplus numbers are normalized by the market size—i.e., we report revenue per consumer.

Figure 1.

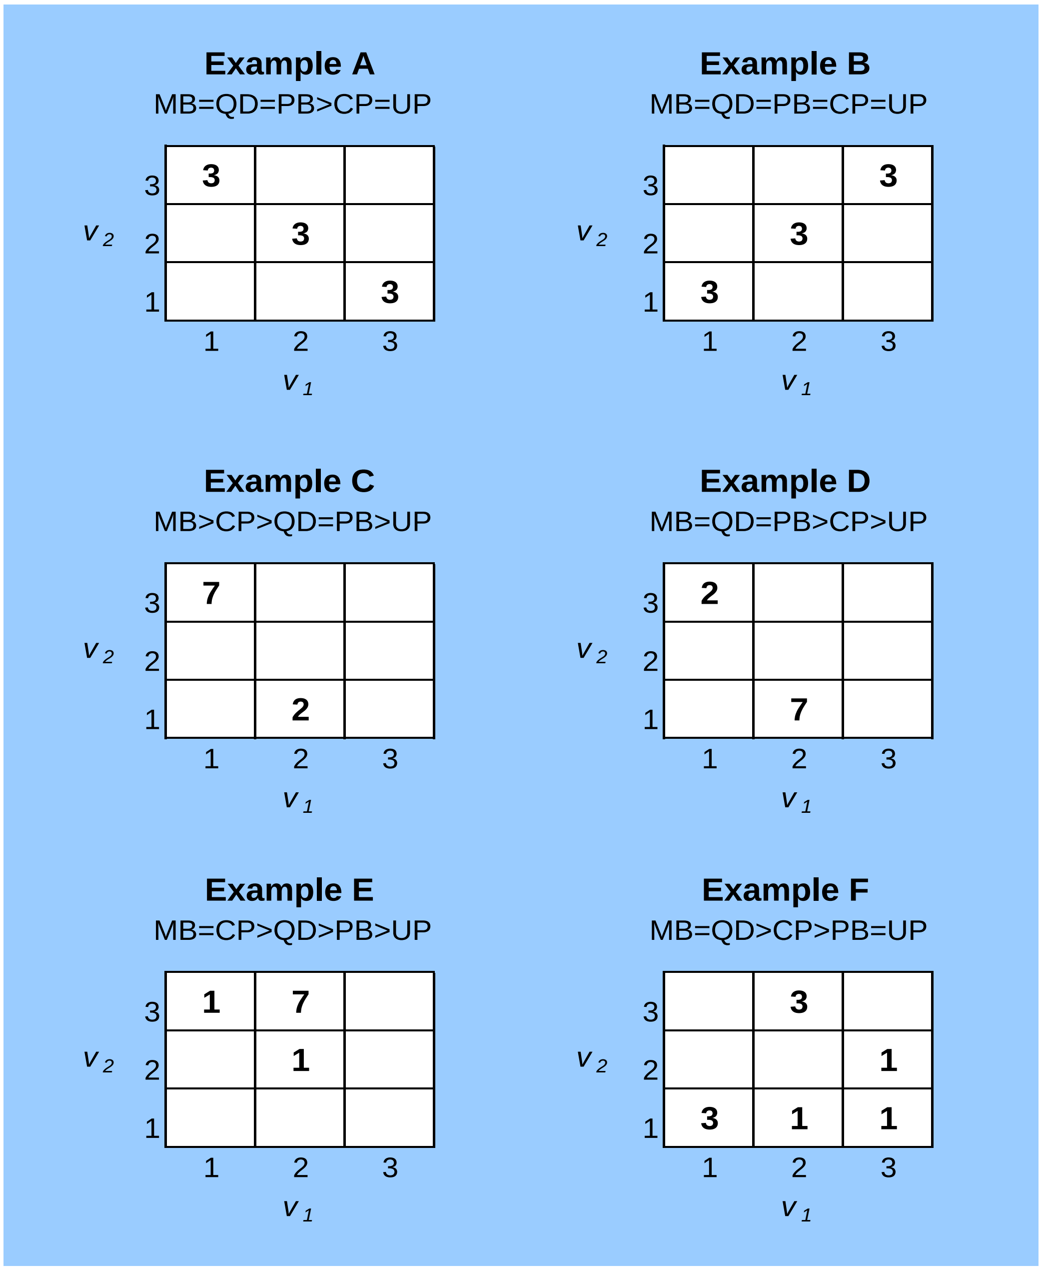
Example A
MB=QD=PB>CP=UP
Example B
MB=QD=PB=CP=UP
Example C
MB>CP>QD=PB>UP
Example D
MB=QD=PB>CP>UP
Example E
MB=CP>QD>PB>UP
Example F
MB=QD>CP>PB=UP
v 1
v 2

Figure 2.

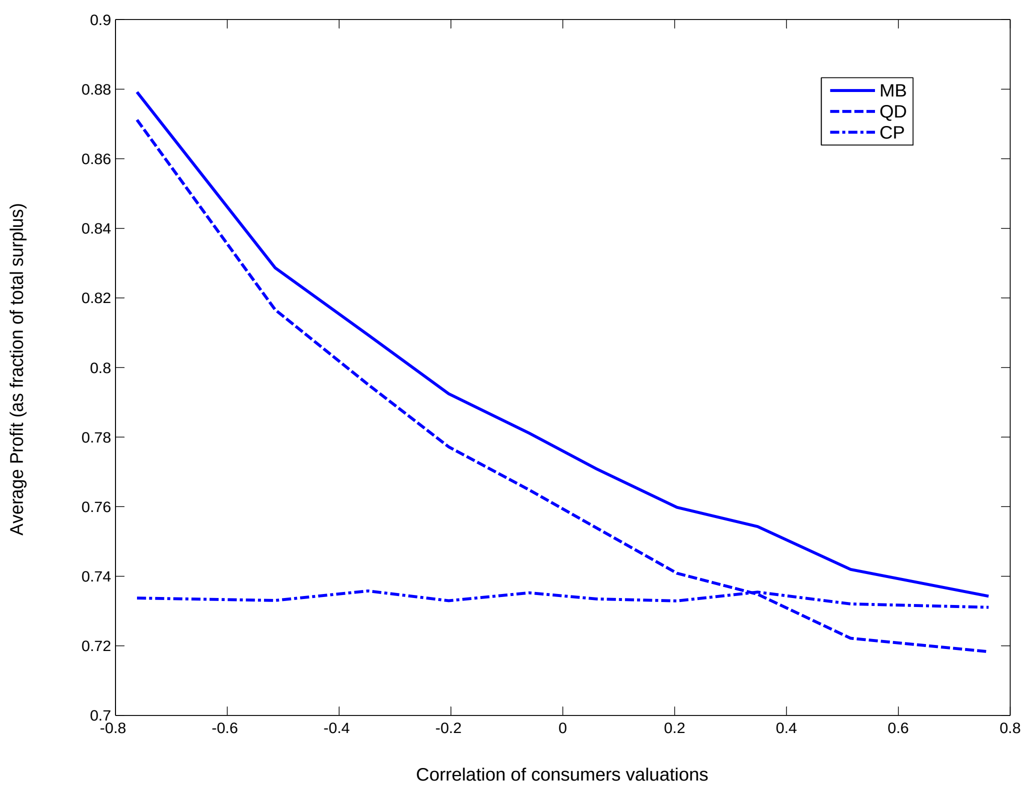
0.9
0.88
0.86
0.84
0.82
0.8
0.78
0.76
0.74
0.72
0.7
-0.8
-0.6
-0.4
-0.2
0
0.2
0.4
0.6
0.8
MB
QD
CP
Average Profit (as fraction of total surplus)
Correlation of consumers valuations

Figure 3.

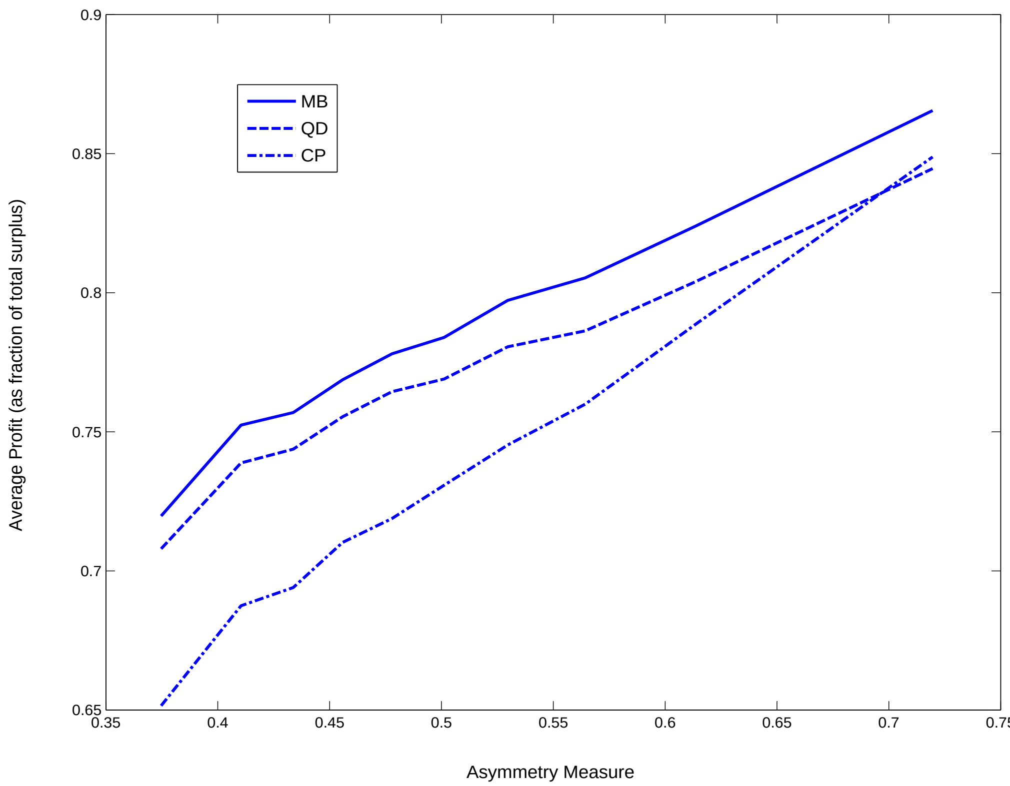

MB
QD
CP
0.9
0.85
0.8
0.75
0.7
0.65
0.35
0.4
0.45
0.5
0.55
0.6
0.65
0.7
0.75
Average Profit (as fraction of total surplus)
Asymmetry Measure

Figure 4.

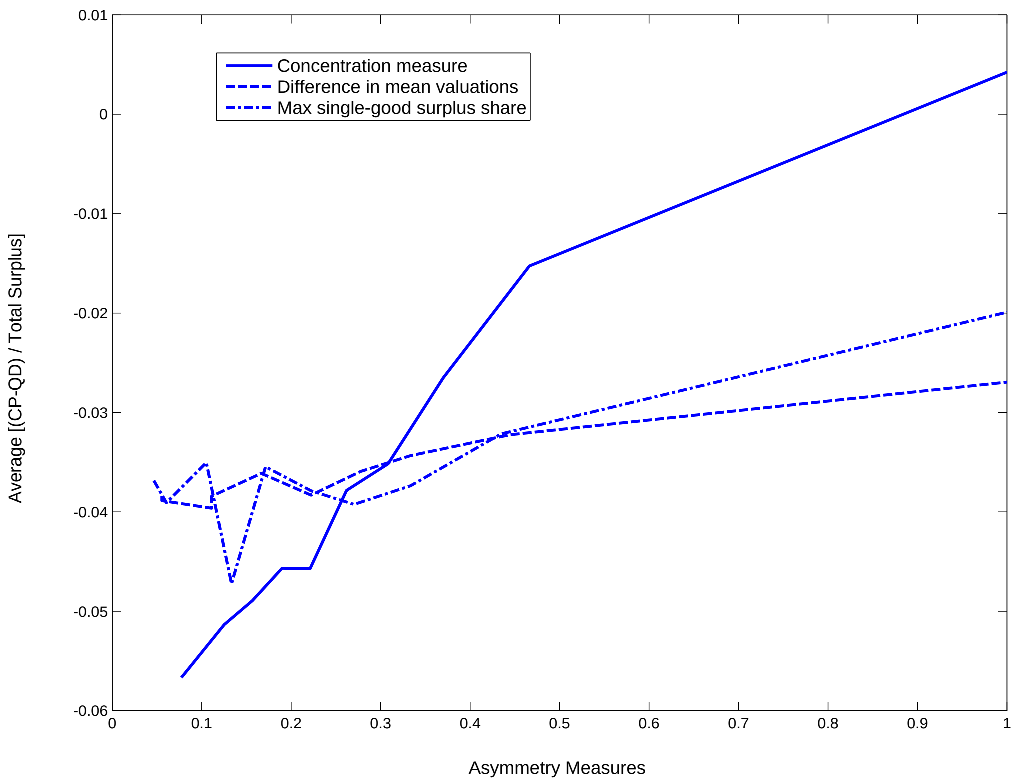

Concentration measure
Difference in mean valuations
Max single-good surplus share
Average [(CP-QD) / Total Surplus]
Asymmetry Measures
0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1