Approximation in Mechanism Design

Jason Hartline*

First Draft: October, 2010

Abstract

This paper considers three challenge areas for mechanism design and describes the role approximation plays in resolving them. Challenge 1: optimal mechanisms are finely tuned to precise details of the distribution on agent preferences. Challenge 2: in environments with multi-dimensional agent preferences economic analysis has failed to provide generic characterizations optimal mechanisms. Challenge 3: optimal mechanisms are parameterized by unrealistic knowledge of the distribution of agents' private preferences. The theory of approximation is well suited to address these challenges. While the optimal mechanism may require precise distributional assumptions, there may be approximately optimal mechanism that depends only on natural characteristics of the distribution. While the multi-dimensional optimal mechanism may resist precise economic characterization, there may be simple description of approximately optimal mechanisms. While the optimal mechanism may be parameterized by the distribution of the agents' private preferences, there may be a single mechanism that approximates the optimal mechanism for any distribution. Finally, these approximately optimal mechanisms, because of their simplicity and tractability, may be more likely to arise in practice, thus making the theory of approximately optimal mechanism more descriptive than that of (precisely) optimal mechanisms. This paper surveys positive resolutions to these challenges with emphasis on basic techniques, relevance to practice, and future research directions for approximation in mechanism design.

*Northwestern University, Evanston, IL 60208. Email: `hartline@eecs.northwestern.edu`.

1 Introduction

A mechanism gives the a mapping between the actions of strategic agents and outcomes of the system. Equilibrium theory describes what outcomes will arise in the equilibrium of selfish agent play. Mechanism design then considers the optimization question of what mechanisms have good outcomes in equilibrium.

Optimal mechanism design searches for the best of these mechanisms. The space of all mechanisms is rich and positive results for optimal mechanism design (a) identify a subclass of mechanisms from which an optimal mechanism can be drawn, (b) interpret the salient characteristics of this subclass, and (c) predict the mechanisms that arise in practice. This agenda has a rich and elegant history in the economic literature with many success stories.

But what can a theory of mechanism design say (a) when the only subclass of mechanisms that contains all optimal mechanisms is the full class, (b) when analytical approaches fail to identify salient characteristics of optimal mechanisms, or (c) when the mechanisms in practice are not the ones predicted by optimal mechanism design? To address these and other issues we propose a theory of approximation for mechanism design.

A mechanism is a β-approximation in some setting if its objective performance is within a multiplicative factor of β of that of the optimal mechanism for the same setting. For example, a 2-approximation always obtains 50% of the optimal performance. A subclass of mechanisms is a β-approximation if for every setting there is a mechanism in the subclass that is a β-approximation.

As discussed, the class of all mechanisms is incredibly rich, nonetheless, there are environments, see, e.g., Vincent and Manelli (2007), where any undominated mechanism is optimal for some setting. We face a tradeoff: if we consider only optimal mechanisms we are stuck with the full class from which we can make no observations about what makes a mechanism good; on the other hand, if we relax optimality, we may be able to identify a small subclass of mechanisms that are approximately optimal, i.e., for any setting there is a mechanism in the subclass that approximates the optimal mechanism. This subclass is important in theory as we can potentially observe salient characteristics of it. It is important in practice because, while it is unlikely for a real mechanism designer to be able to optimize over all mechanisms, optimizing over a small class of, hopefully, natural mechanisms may be possible. For instance, a conclusion of this paper is that reserve-price-based mechanisms and posted pricings are approximately optimal in a wide range of settings including those with multi-dimensional agent preferences.

Approximation also provides a lens with which to explore the salient features of an environment or mechanism. Suppose we wish to determine whether a particular feature of a mechanism is important. If there exists an subclass of mechanisms without that feature that gives a good approximation to the optimal mechanism, then the feature is perhaps not that important. If, on the other hand, there is no such subclass then the feature is quite important. For instance, an analysis of this sort easily concludes that transfers are very important for surplus maximization; where as one of the conclusions of this paper is that competition between agents is not that important. Essentially, approximation provides a means to determine which aspect of a setting are details and which are not details. The approximation factor quantifies the relative importance on the spectra between unimportant details to salient characteristics. Approximation, then allows for design of mechanisms that are not so dependent on details of the setting and therefore more robust (e.g., as in the Wilson (1987) doctrine). In particular, often mechanisms based on the monopoly price are sufficient to approximate ones parameterized by the full distribution, moreover, some environments permit a single (prior-independent) mechanism to approximate the optimal mechanism in every setting.

While it is no doubt a compelling success of mechanism design that its mechanisms are so prevalent in practice, optimal mechanism design cannot claim the entirety of the credit. These mechanisms are employed by practitioners well beyond the settings for which they are optimal. Approximation can explain why: the mechanisms that are optimal in ideal settings may continue to be approximately optimal much more broadly. It is important for the theory to describe how broadly these mechanisms are approximately optimal and how close to optimal they are. Thus, the theory of approximation can complement the theory of optimality and justify the wide prevalence of certain mechanisms. For instance, the wide prevalence reserve-price-based mechanisms and posted pricings is corroborated by their approximate optimality.

The primary focus of this paper is to compare and contrast simple approximately optimal mechanisms with optimal ones. We focus on environments where the designer's objective is to maximize profit, though similar results hold for the objective of social surplus. As examples for such an analysis we discuss in detail three distinct environments: *single-dimensional agent environments* such as a single-item auctions, *multi-dimensional agent environments* such as multi-item auctions, and *prior-independent environments* such as multi-unit auctions with "unknown demand". For each environment we describe the optimal mechanism, contrast it to a simple approximately optimal mechanism, and discuss generalizations.

Results and Organization. We start with an example of approximation for a stochastic optimization problem that comes from optimal stopping theory (Section 2). This example allows us to discuss many of the benefits of approximation and also provides a basis for many of the subsequent results. In Section 3 we further define and motivate approximation. Section 4 describes the abstract model for mechanism design that we consider. In Section 5 we describe the role of approximation in single-dimensional environments; the focus is on single-item auctions. In particular, we show that the Vickrey auction (and the VCG generalization) with reserve prices is approximately optimal in many settings. In Section 6 we consider approximation in multi-dimensional environments; the focus is on multi-item item-pricing for a single unit-demand agent. We give a simple approach for identifying approximately optimal item-pricings, and we discuss generalizations of this result to multiple agents where posted pricings continues to give a good approximations. In Section 7 we discuss the design of mechanisms that are less reliant on prior-information. The focus is on multi-unit auctions for a single item. Under standard distributional assumptions, a single random sample is enough to approximate the optimal mechanism; furthermore, this implies that completely prior-independent mechanisms can be good approximations. Concluding remarks will be provided at the end of each section.

Brief literature review. Optimal mechanisms for single-dimensional agent environments can be derived from Myerson (1981), Riley and Samuelson (1981), and Bulow and Roberts (1989). The economics work in multi-dimensional environments that is most relevant to ours is for the setting of additive valuations, i.e., when the value of a bundle of resources is the sum of the value for each individual resource. For environments with independent additive valuations, Armstrong (1996) shows that pricing the grand bundle is optimal in the limit; while Vincent and Manelli (2007) show that any undominated mechanism is optimal in some setting. This latter result can be viewed as an impossibility result for optimal mechanism design as good kinds of mechanisms cannot be distinguished from bad kinds of mechanisms. Unlike the case of single-dimensional environments, optimal mechanisms in multi-dimensional environments may require randomization, see, e.g., Thanassoulis (2004) and Vincent and Manelli (2007), and the resulting mechanisms can be quite complex.

Implementation theory considers design of mechanisms that have desirable equilibria, e.g., see the Maskin and Sjöström (2002) survey. In this context, our question of prior-independent mechanism design would be referred to as non-parametric virtual Bayes-Nash implementation. Non-parametric Bayes-Nash implementations are not parameterized by knowledge of the prior distribution; virtual implementations are allowed to not implement the desired outcome with probability that can be made arbitrarily small. In this context, Choi and Kim (1999) gives a non-parametric Bayes-Nash implementation of ex post efficient and budget balanced mechanisms for public projects; the mechanism requires agents to cross-report their beliefs. This sort of mechanism can also give non-parametric virtual implementation of the (revenue) optimal mechanism. This approach is regarded as being "consequentialist" as the mechanisms suggested are not descriptive of mechanisms in practice (e.g., the Baliga and Sjöström (2008) survey).

In this paper we provide economic context and discussion for a number of technical results from the recent literature on approximation in mechanism design. The single-dimensional approximations that we present come from Hartline and Roughgarden (2009) and Chawla et al. (2010a). Our approximations for multi-dimensional agent environments come from Chawla et al. (2010a). In both environments posted pricings and reserve-price-based mechanisms are approximately optimal. The relationship between mechanism design and prophet inequalities, a technical tool discussed in Section 2 is from Chawla et al. (2010a). Our discussion of prior-independent mechanism design is motivated by the Bulow and Klemperer (1996) theorem. We describe a "single sample mechanism" which was suggested and studied by Dhangwatnotai et al. (2010). Early work on prior-independent mechanism design includes Goldberg et al. (2001); Segal (2003); Baliga and Vohra (2003); Goldberg et al. (2006); Hartline and Roughgarden (2008).

2 Example: A Stopping Game

Consider the following scenario. A gambler faces a series of n games on each of n days. Game i has prize distributed according to F_i. The order of the games and distribution of the game prizes is fully known in advance to the gambler. On day i the gambler *realizes* the value $v_i \sim F_i$ of game i and must decide whether to keep this prize and *stop* or to return the prize and *continue* playing. In other words, the gambler is only allowed to keep one prize and must decide which prize to keep immediately on realizing the prize and before any other prizes are realized.

The gambler's optimal strategy can be calculated by *backwards induction*. On day n the gambler should stop with whatever prize is realized. This results in some expected value. On day $n-1$ the gambler should set a threshold t_{n-1} equal to the expected prize for the last day and stop with any prize bigger than the threshold. On day $n-2$ the gambler should stop with any prize with greater value than the expected value of the strategy thus-far calculated. Proceeding in this manner the gambler can calculate a threshold t_i for each day where the optimal strategy is to stop with prize i if and only if $v_i \geq t_i$.

Of course, this optimal strategy suffers from many drawbacks. It is complicated: it takes n numbers to describe it. It is not robust to small changes in the game, e.g., changing of the order of the games or making small changes to distribution F_i strictly above t_i or below t_i. It does not allow for intuitive understanding of the properties of good strategies. Finally, it does not generalize well to give solutions to other similar kinds of games.

We turn to approximation to give a crisper picture. A *threshold strategy* is given by a single threshold t and requires the gambler to accept the first prize i with $v_i \geq t$. Threshold strategies are clearly suboptimal as even on day n if prize $v_n < t$ the gambler will not stop and will, therefore, receive no prize.

Theorem 2.1 (Prophet Inequality; Samuel-Cahn, 1984) *There exists a threshold strategy such that the expected prize of the gambler is at least half the expected value of the maximum prize. Moreover, one such threshold strategy is the one where the probability that the gambler receives no prize is exactly 1/2; moreover, the bound is invariant on the tie-breaking rule, i.e., which prize the gambler obtains when multiple prizes are above the threshold.*

Unlike our the optimal (backwards induction) strategy the profit inequality theorem provides substantive conclusions. The result is driven by trading off the probability of not stopping and receiving no prize with the probability of stopping early with a suboptimal prize; a good tradeoff point is at a half probability of each event. The suggested threshold strategy is also robust to the order of the games and the precise distribution functions. Notice that the order of the games makes no difference in the determination of the threshold, and if the distribution above or below the threshold changes, nothing on the bound or suggested strategy is affected.

We will refer to the setting where there are more than one prize above the threshold as a *tie*. There is an implicit tie-breaking strategy given by the definition of the game: ties are broken in favor of the earliest game. The aforementioned invariance of the performance bound to the order of the games suggests that the bound might invariant on the tie-breaking rule entirely; the theorem explicitly states this conclusion. The invariance with respect to the tie-breaking rule means that the prophet inequality theorem has broad implications to other similar settings and in particular to auction design and posted pricing, as we will see in subsequent sections.

The prophet inequality is suggesting something quite strong: it is saying that even though the gambler does not know the realizations of the prizes in advance, he can still do as well as a "prophet" who does. While it is nice to know that a simple strategy does as well as the prophet, the comparison is intuitively the wrong one to make. What we should instead care about is our performance relative to the performance of the optimal (i.e., backwards induction) strategy. Such a direct analysis, however, is difficult to make. Just as the optimal strategy did not provide much conceptual understanding, it also does not permit analytically tractable direct comparison. As the prophet's performance, i.e., the maximum of the realized values, is always an upper bound on the optimal strategy's performance, comparison to it implies as a corollary the result that we should have been after in the first place: a bound on the approximation factor of threshold strategies.

Corollary 2.2 *There exists a threshold strategy that is a 2-approximation to the optimal (backwards induction) strategy.*

One should immediately be concerned that the indirect route by which we came to this corollary may have been lossy. In fact, it is often the case in the theory of approximation that with the "right" upper bound, the indirect approach is not lossy. In this respect the prophet's performance is the right upper bound. For threshold strategies it is easy to see by example that the bound of the corollary is tight. Imagine a setting with $n = 2$ days: on day 1, the prize is a high value $h \gg 1$ with probability $1/h$ and otherwise zero, on day 2 the prize is deterministicly $v_2 = 1$. There are only two reasonable threshold strategies $t = 1$ or $t = h$; both give the gambler an expected payoff of 1. Of course, the optimal strategy is to take the first prize if it is high and otherwise the second prize. This has expected payoff $1 + (1 - 1/h)$ which is 2 in the limit as $h \to \infty$.

We conclude with a simple proof of the prophet inequality theorem.

Proof: (of theorem) Define $q_i = 1 - F_i(t)$ as the probability that $v_i \geq t$. Let $\chi = \prod_i (1 - q_i)$ be the probability that the gambler receives no prize. The proof follows in three steps. In terms of t and χ, we get an upper bound on the expected maximum prize. Likewise, we get a lower bound on the algorithm performance. Finally, we plug in $\chi = 1/2$ to obtain the bound. If there is no t with $\chi = 1/2$, which is possible if the F_i are not continuous, one of the t that corresponds to the smallest $\chi > 1/2$ or largest $\chi < 1/2$ suffices.

In the analysis below the notation $(v_i - t)^+$ is short-hand for $\max(v_i - t, 0)$.

1. An upper bound on $\mathrm{OPT} = \mathbf{E}[\max_i v_i]$.

   Notice that regardless of whether there exists a $v_i \geq t$ or not, OPT is at most $t +$

$\max_i (v_i - t)^+$. Therefore,

$$\begin{aligned} \mathrm{OPT} &\le t + \mathbf{E}\big[\max_i (v_i - t)^+\big] \\ &\le t + \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+\big]\,. \end{aligned}$$

2. A lower bound on $\mathcal{A} = \mathbf{E}[\text{prize of gambler with threshold } t]$.

   Clearly, we get t with probability $1 - \chi$. Depending on which prize i is the earliest one that is greater than t we also get an additional $v_i - t$. It is easy to reason about the expectation of this quantity when there is exactly one such prize and much more difficult to do so when there are more than one. We will ignore the additional prize we get from the latter case and get a lower bound.

$$\begin{aligned} \mathcal{A} &\ge (1-\chi)t + \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+ \mid \text{other } v_j < t\big]\,\mathbf{Pr}[\text{other } v_j < t] \\ &\ge (1-\chi)t + \chi \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+ \mid \text{other } v_j < t\big] \\ &= (1-\chi)t + \chi \sum\nolimits_i \mathbf{E}\big[(v_i - t)^+\big] \end{aligned}$$

   The second inequality follows because $\mathbf{Pr}[\text{other } v_j < t] = \prod_{j \ne i}(1 - q_i) \ge \chi$. The final equality follows because the random variable v_i is independent of random variables v_j for $j \ne i$.

3. Plug in $\chi = 1/2$.

   From the upper and lower bounds calculated, if we can find a t such that $\chi = 1/2$ then $\mathcal{A} \ge \mathrm{OPT}/2$.

   For discontinuous distributions, e.g., ones with point-masses, χ may be a discontinuous function of t. There may be no t with $\chi = 1/2$. Let $\chi_1 = \sup\{\chi < 1/2\}$ and $\chi_2 = \inf\{\chi > 1/2\}$. Notice that an arbitrarily small increase in threshold causes the jump from χ_1 to χ_2; let t be the limiting threshold for both these χs. Therefore, as a function of χ, t is constant and we get the same lower bound formula $\mathrm{LB}(\chi) = (1-\chi)t + \chi \sum_i \mathbf{E}[(v_i - t)^+]$ which is linear in χ.

   We know that this function evaluated at $\chi = 1/2$ (which is not possible to implement) satisfies $\mathrm{LB}(1/2) \ge \mathrm{OPT}/2$. Of course this is a linear function so it is maximized on the end-points on which it is valid, namely χ_1 or χ_2. Therefore, one of $\chi \in \{\chi_1, \chi_2\}$ satisfies $\mathrm{LB}(\chi) \ge \mathrm{OPT}/2$. We conclude that either the threshold that corresponds to χ_1 or χ_2 is a 2-approximation.

□

3 The Philosophy of Approximation

The preceding discussion of the gambler's stopping game and prophet inequalities exemplifies a setting where there is little to learn from an optimal strategy and much to learn from a simple, approximately-optimal strategy. In this section we expand on the philosophy of approximation as is relevant for mechanism design.

Our goal is a good theory of mechanism design. Such a theory would ideally satisfy the following three criteria:

Prescriptive: It gives concrete suggestions for how a good mechanism should be designed.

Predictive: It describes how mechanisms are designed.

Conclusive: It identifies the essence of the problem and allows for broad conclusions.

Notice that optimality is not one of the criteria, nor is exactly suggesting a mechanism to a practitioner. Instead, intuition from the theory of mechanism design should help guide the design of good mechanisms in practice. Such guidance is possible through conclusive observations about what good mechanisms do. Observations that are robust to modeling details are especially important.

The question of designing an *optimal mechanism* for a given setting and objective can be viewed as a standard optimization problem. Given incentive constraints, imposed by game theoretic strategizing, and feasibility constraints, imposed by the environment, optimize the designer's given objective. I.e., put all the constraints into a solver and out will come a solution. It is immediately clear that this approach, absent further analysis, is inconclusive. What have learned about the optimal mechanism? Furthermore, the optimization suggested is often computationally intractable. The resulting mechanism is likely to be complex and impractical making the approach neither predictive or prescriptive. Much of the work in economics has focused on environments where the constraints simplify and, for instance, the mechanism design problem can be reduced to a natural optimization question without incentive constraints. E.g., the VCG mechanism for maximizing social surplus, and Myerson's mechanism for maximizing profit in single-dimensional environments. Unfortunately, there are many environments where analysis has failed to simplify the problem and these environments are considered "unsolved".

In environments where optimality is impossible (by any of the above critiques) one should instead try to approximate. The formal definition of an approximation is given below. Notice that approximation factors are always at least one as it is impossible to perform better than optimal; small approximation ratios are good and a large approximation ratios are bad.

Definition 3.1 *A mechanism is a β-*approximation *in a given setting if the ratio of the optimal mechanism's expected performance to its expected performance is at most β. A class of mechanisms is a β-approximation in a given environment, if for any setting in the environment there is a mechanism in the class that is a β-approximation.*

Depending on the problem and the approximation mechanism, approximation factors can range from $(1 + \epsilon)$, or arbitrarily close approximations, to linear factor approximations (or sometimes even worse). Notice a linear factor approximation is one where as some parameter in the environment grows, i.e., more agents or more resources, the approximation factor gets worse.

Compare and contrast an approximation schema which allows for a $(1+\epsilon)$-approximation for any given $\epsilon > 0$ with a 2-approximation. Obviously a practitioner would be happier obtaining 99% of the optimal revenue than 50%. On the other hand, $(1+\epsilon)$-approximations usually arise from brute-force searches over restricted search spaces therefor they provide as little conceptual insight as the optimal mechanisms. On the other hand a simple approach that gives a 2-approximation is usually made possible by some fundamental properties of good mechanisms.

In this paper we will not view super-constant approximation factors as positive results. It has been our experience in mechanism design that super-constant approximation factors are indicative of (a) a bad mechanism, (b) failure to appropriately characterize optimal mechanisms, or (c) an imposition of incompatible modeling assumptions or constraints. However, super-constant lower-bounds for approximation are, as mentioned before, useful for identifying modeling parameters that are not details.

Our conclusion form the above discussion is that constant approximations, e.g., two, are the "sweet spot" where mechanisms are performing fairly well and giving good insight. Nonetheless, if you were approached by a seller (henceforth: principal) to design a mechanism and you returned to triumphantly reveal an elegant mechanism that gives them a 2-approximation to their profit, you would probably find them a bit discouraged. After all, your mechanism leaves have of their profit on the table.

In light of this critique we review the main points of constant, e.g., two, approximations. First, a 2-approximation provides informative conclusions that can guide the design of better mechanisms for specific settings. Second, the approximation factor of two is a theoretical result that holds in a large set of settings, in specific settings the mechanism may perform better. It is easy, via simulation to evaluate the mechanism performance on specific settings to see how close to optimal it actually is. Third, in many environments the optimal mechanism is not understood at all, meaning the principals alternative to your 2-approximation is an ad hoc mechanism with no performance guarantee. This principal is of course free to simulate your mechanism and their mechanism in their given setting and decide to use the better of the two. In this fashion the principal's ad hoc mechanism, if used, is provably a 2-approximation as well. Fourth, mechanisms that are 2-approximations in theory arise in practice. In fact, that it is a 2-approximation explains why the mechanism arises. Even though it is not optimal, it is close enough. If was far from being optimal the principal (hopefully) would have figured this out and adopted a different approach.

Finally, suppose the principal was worried about collusion, risk attitudes, after-market effects, or other economic phenomena that are usually not included in standard ideal models for mechanism design. One option would be to explicitly model these effects and study optimal mechanisms in the augmented model. These complicated models are difficult to analyze and optimal mechanisms may be overly influenced by insignificant-seeming modeling choices. Optimal mechanisms are precisely tuned to details in the model and these details may drive the form of the optimal mechanism. On the other hand, we can consider constant approximations that are robust to various out-of-model phenomena. In such an environment the comparison between the approximation and the optimal mechanism is unfair because the optimal mechanism may suffer from out-of-model phenomena that the approximation is robust to. In fact, this "optimal mechanism" may perform much worse than our approximation when the phenomena are explicitly modeled. For example, the posted pricings that we show are approximately optimal are robust to timing effects and for this reason an online auction house such eBay may want to sellers to switch from auctions to "buy it now" pricings, see, e.g., Wang et al. (2008) and Reynolds and Wooders (2009).

Finally, there is an issue of non-robustness that is inherent in any optimization over a complex set of objects, such as mechanisms. Suppose the designer does not know the setting exactly but can learn about it through, e.g., market analysis. Such a market analysis is certainly going to be noisy and then exactly optimizing a mechanism to it may "over fit" to this noise. Both statistics and machine learning theory have techniques for addressing this sort of overfitting. Approximation mechanisms also provide such a robustness. Since the class of approximation mechanisms is restricted from the full set, for these mechanisms to be good, they must pay less attention to details such as sampling noise.

In our study of mechanism design, we will search for positive results, a.k.a., upper-bounds, which will be in the form of a constant approximation. Where we fail to find positive results, we will search for negative results, a.k.a., lower bounds, in the form of an impossibility of constant approximation.

4 General Environments for Mechanism Design

We will be considering mechanism design questions in independent private value environments with quasi-linear agent preferences. In these environments an agent's preference is given by a valuation function over outcomes and their objective is to maximize their utility which is defined to be the difference between the expected value they derive for the outcome of the mechanism and the expected price they are charged.

It is assumed that the distribution over agent preferences is common knowledge; however, all of the mechanisms we discuss are *ex post incentive compatible*, i.e., truthful reporting is a (weakly) dominant strategy. Nonetheless, the mechanisms we give achieve their guarantees relative to optimal mechanisms in Bayes-Nash equilibrium.

One of the most interesting and important distinctions between environments for mechanism design is between ones with single-dimensional preferences and multi-dimensional preferences. We will consider single-dimensional environments where each agent's private preference is given by a single numeric value for an abstract service. The agent obtains this value from all outcomes in which they receive the service and they have value zero for all other outcomes. Multi-dimensional environments are more general. In the multi-dimensional environments we consider, an agent has distinct private value for any of several abstract services. Much of our focus will be on unit-demand environments, i.e., where each agent desires at most one of these services.

The agents' preferences will drawn from a distribution. We will consider both i.i.d. environments and non-identical environments, but not general correlated environments. We also allow the environment to impose constraints on the distribution as, e.g., the monotone hazard rate condition, or Myerson's (1981) regularity condition. Within any environment we will refer to the specific distribution as the *setting*.

The environment also specifies any feasibility constraints that the designer may face. For instance, in *single-item* environments at most one agent can be served, in k-*unit* environments at most k agents can be served, in *digital good* environments any subset of agents can be served.

In general we can talk about environments where the designer incurs a cost for serving the agents that is an arbitrary function of the set of agents served. We refer to these environments as *general cost* environments. A special case are those where the costs are zero or infinity denoting feasible and infeasible outcomes, respectively. We refer to these as *general feasibility* environments.

An important subclass of general feasibility environments are those with a natural downward-closure property, i.e., where subset of a feasible set is feasible. An example of a *downward-closed* environment is that of combinatorial auctions: given any feasible allocation of items any subset of this allocation is feasible. An example of a non-downward-closed environment is that of (non-excludable) public good provisioning, where either no agents are served or all agents are served.

A important special case of downward-closed environments has the feasible sets given by the *independent sets* of a *matroid set system*. These matroid environments generalize multi-unit auction environments and include some relevant matching problems. Many mechanism design results for multi-unit environments extend to matroid environments, and furthermore matroid environments are often the extent of their generality. For the purpose of our discussion: matroid environments are ones where ordinal optimization suffices, i.e., exact magnitudes of values are unnecessary for optimization.

In all of the environments above, for the objective of social surplus maximization, i.e., maximizing the sum of the values of agents served (less any service cost), the Vickrey-Clarke-Groves (VCG) mechanism is optimal and ex post incentive compatible (Vickrey, 1961; Clarke, 1971; Groves, 1973). In single-item environments VCG is precisely the Vickrey (a.k.a. second-price) auction and we will refer to it as such. In k-unit environments VCG is the k-Vickrey auction, i.e., it sells to the highest k bidders at the $k + 1$st highest bid value.

Throughout the rest of the paper we will consider the objective of maximizing the designers profit, i.e., the sum of the agent payments less any service costs.

5 Single-dimensional Environments

We start our exploration of approximation in mechanism design with environments where each agent's private preference is single dimensional, i.e., the agent has a single private value for an abstract service. The most fundamental example of such environment is that of a single-item auction.

5.1 Optimal Mechanisms

We briefly review optimal mechanism design in single-dimensional environments as developed by Myerson (1981) and Riley and Samuelson (1981), and further refined by Bulow and Roberts (1989).

Each agent i is a risk-neutral quasi-linear utility maximizers with value v_i for service drawn independently from distribution F_i. Considering just this single agent, an important quantity is the *revenue curve*, $R_i(q)$, which is the revenue one can obtain from the agent as a function of the ex ante probability with which the agent is served. Notice, if we offered the agent a price p they should accept with probability $1 - F_i(p)$ and our expected revenue is $p \times (1 - F_i(p))$. Expressing this revenue as a function of probability instead of price, we get $R_i(q) = F_i^{-1}(1-q) \times q$.

Economic intuition suggests that maximizing profit is equivalent to maximizing marginal revenue. Indeed, for mechanism design this result can be derived. Given a revenue curve $R(q)$ denote its derivative with respect to q as $R'(q)$. We summarize with a definition and a theorem.

Definition 5.1 *The* virtual valuation *of agent i is $\phi_i(v_i) = R_i'(1 - F_i(v_i))$.*[1] *The* virtual surplus *of a mechanism is the sum of the virtual values of the agents served less any service costs.*

Theorem 5.1 (e.g., Myerson, 1981) *In Bayes-Nash equilibrium, the expected profit of any mechanism in is equal to its expected virtual surplus.*

Simple intuition suggests that in single-dimensional environments the interim probability to which any agent is served must be monotone non-decreasing in the value of the agent. Otherwise, an agent with a high value could improve their payoff by simulating the strategy of a lower-valued agent. Such monotonicity is in fact a well known characterization.

Theorem 5.2 (e.g., Myerson, 1981) *An interim allocation rule can be (weakly) implemented in Bayes-Nash equilibrium (i.e., with an appropriate payment rule) if and only if it is monotone non-decreasing.*

The two theorems above narrow the question of optimal mechanism design to finding one that maximizes virtual surplus subject to monotonicity. A standard approach to this sort of questions is to ignore monotonicity and optimize the objective virtual surplus and hope that the resultant allocation rule satisfies monotonicity anyway.

Consider a single-item auction. Maximizing virtual surplus is identical to serving the agent with the highest positive virtual value. If the virtual valuation functions are monotone non-decreasing then such an allocation rule monotone: as an agent increases their value, their virtual value does not decrease and therefore, if they were being served with the lower value, they are certainly also served with the higher value. Virtual valuation functions are

[1] This is identical to the more typical formulation, e.g., by Myerson (1981), of $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$.

monotone if and only if the revenue curve is concave, and, thus, the optimal mechanism is identified.

Definition 5.2 *A distribution F is* regular *if its associated revenue curve $R(\cdot)$ is concave. An environment is regular if all distributions are regular.*

Theorem 5.3 (e.g., Myerson, 1981) *For regular, single-item environments, the optimal auction allocates the item to the agent with the highest positive virtual value.*

In the special case where the values of the agents are identically distributed, i.e., according to F; the virtual valuation functions are identical, i.e., $\phi(\cdot)$; and the agent with the highest virtual value is simply the agent with the highest value. This agent has a positive virtual value if their value is at least the inverse virtual value of zero, i.e., $\phi^{-1}(0)$. The Vickrey auction with reserve price $\phi^{-1}(0)$ obtains this outcome in equilibrium and therefore it is optimal. Also notice that $\phi^{-1}(0)$ is exactly the value for which the revenue curve is maximized, i.e., the price a monopolist would charge a single agent with value drawn from the distribution.

Definition 5.3 *The* monopoly price *for a distribution F is $\phi^{-1}(0)$.*

Corollary 5.4 (e.g., Myerson, 1981) *For i.i.d., regular, and single-item environments, the Vickrey auction with monopoly reserve price $\phi^{-1}(0)$ is optimal.*

For irregular environments optimizing virtual surplus does not lead to a monotone allocation rule. We skip a detailed derivation of the optimal mechanism in this case and just state it.

Definition 5.4 *The* ironed revenue curve*, $\bar{R}(\cdot)$, is the concave hull of the revenue curve, $R(\cdot)$. The* ironed virtual valuation function *$\bar{\phi}(v) = \bar{R}'(1 - F(v))$ is the derivative of the ironed revenue curve. The* ironed virtual surplus *is the virtual surplus of the agents served less any service cost.*

Ironed virtual values are monotone non-decreasing. Therefore optimizing ironed virtual surplus results in a monotone allocation rule.

Theorem 5.5 (e.g., Myerson, 1981) *In (possibly non-identical and irregular) single-item environments, the optimal mechanism serves the agent with the highest positive virtual value.*

5.2 Conclusions and Extensions

The characterization of optimal mechanisms in single-dimensional environments as ironed virtual surplus optimizers is quite general.

Theorem 5.6 *In any single-dimensional, general costs environment, the optimal mechanism serves the agents with the maximum ironed virtual surplus.*

Maximizing ironed virtual surplus leads to a reserve pricing mechanism for independent, regular, matroid environments. Recall that matroid environments are ones were optimization is ordinal, i.e., based on the relative order of values and not their exact magnitudes. For regular distributions virtual values are in the same order as values, so the optimal mechanism is VCG with the monopoly reserve price.

Corollary 5.7 *For i.i.d., regular, and matroid environments, the VCG mechanism with monopoly reserve price $\phi^{-1}(0)$ is optimal.*

When the distributions of the agents are not identically distributed or irregular, or the set system is not a matroid (e.g., combinatorial auctions), then reserve-price based mechanisms are not optimal. Instead the optimal auction depends on the exact form of the distributions. While many standard distributions are regular, e.g., uniform, normal, and exponential; many realistic environments are not, e.g., bimodal. Many realistic distributions are non-identical, for instance in eBay buyers have distinct user feedback ratings, and a seller could potentially use these ratings to discriminate. Finally, many important environments for mechanism design are not matroids. Therefore, we should expect the reserve price based mechanisms to be relatively rare.

On the contrary, reserve-price-based mechanism are widely prevalent; eBay, for instance, is a reserve-price based mechanism. The above theory of optimal auctions does not in itself justify such a wide prevalence; importantly, it does not justify their usage in environments where distributions are irregular or non-identical, or in non-matroid environments. Certainly there are numerous reasons these optimal mechanisms are not used. One reason may be that they are too complicated to be practical.

5.3 Approximation

One of the most intriguing conclusions from the preceding subsection is that for i.i.d. regular distributions the optimal single-item auction is the second-price auction with a reservation price. This result is compelling as the solution it proposes is quite simple; therefore, making it easy to prescribe. Furthermore, reserve-price-based auctions are often employed in practice so this theory of optimal auctions is also descriptive. Unfortunately, i.i.d., regular, single-item (or more generally: matroid) environments are hardly representative of the scenarios in which we would like to design good mechanisms. Furthermore, if any of the assumptions are relaxed, reserve-price-based mechanisms are not optimal.

In this section we address this deficiency by showing that while reserve-price-based mechanisms are not optimal, they are approximately optimal in a wide range of environments. These approximately optimal mechanisms are more robust, less dependent on the details of the setting, and provide conceptual understanding that their optimal counterparts do not. They are 2-approximations under worst-case distributional settings. Of course, in any particular setting they may perform better than their worst-case guarantee.

5.3.1 Single-item Auctions

We start with single-item auctions and show that the second-price auction with suitably chosen reserve prices is always a good approximation to the optimal mechanism. The mechanism is the following and is parameterized by agent-specific reserve prices $\mathbf{r} = (r_1, \ldots, r_n)$.

Definition 5.5 *The* Vickrey auction with reserves $\mathbf{r}$, $\mathrm{Vic}_{\mathbf{r}}$ *is:*

1. *reject each agent i with $v_i < r_i$,*
2. *allocate the item to the highest valued agent remaining (or none if none exists), and*
3. *charge the winner their "critical price".*

Our goal in single-item auctions is to select the winner with the highest (positive) ironed virtual value. To do this we draw a connection between the auction problem and the gambler's problem in Section 2. Note that the gambler's problem in prize space is similar to the auctioneer's problem in ironed-virtual-value space. The gambler aims to maximize expected prize while the auctioneer aims to maximize expected virtual value. A constant threshold in the gambler in prize space corresponds to a constant *ironed virtual price* in ironed-virtual-value space. This suggests strongly that constant ironed virtual prices would make good reserve prices in the second-price auction.

Definition 5.6 *A* constant ironed virtual price *is* $\mathbf{p} = (p_1, \ldots, p_n)$ *such that* $\bar{\phi}_i(p_i) = \bar{\phi}_{i'}(p_{i'})$ *for all i and i'.*

Now compare the second-price auction with a constant ironed virtual reserve price with the gambler's threshold strategy. The difference is the tie-breaking rule. The second-price auction breaks ties by value whereas the gambler's threshold strategy breaks ties by the ordering assumption on the games (i.e., lexicographically). Recall though that the tie-breaking rule was irrelevant for our analysis of the prophet inequality. We conclude the following theorem and corollary where, as in the prophet inequality, the constant virtual price is selected so that the probability that the item remains unsold is about 1/2.

Theorem 5.8 (Chawla et al., 2010a) *In (possibly non-identical and irregular) single-item environments, the Vickrey auction with a constant ironed virtual reserve price is a 2-approximation to the optimal revenue.*

Corollary 5.9 *In (possibly irregular) i.i.d. single-item environments, the Vickrey auction with an anonymous reserve is a 2-approximation to the optimal revenue.*

It should be clear that what is driving this result is the specific choice of reserve prices and not explicit competition in the auction. So instead of running an auction imagine the agents arrived in any, perhaps worst-case, order and we made each in turn a take-it-or-leave-it offer of their reserve price? Such a *posted pricing* mechanism is certainly also a 2-approximation.

Theorem 5.10 (Chawla et al., 2010a) *In (possibly non-identical and irregular) single-item environments, a posted pricing with constant ironed virtual prices is a 2-approximation to the optimal revenue.*

5.3.2 Anonymous Reserves

The reserve prices discussed thus far, except for the i.i.d. regular case where the monopoly reserve price is optimal and the i.i.d. irregular case where an anonymous reserve price is a 2-approximation, have been non-anonymous. Different agents face different reserve prices. The question of whether anonymous reserve prices with the Vickrey auction are good approximations in non-identical environments is natural and important. Specifically, anonymous reserve prices are often used in asymmetric environments; can we justify this observation with approximation.

For instance, the eBay auction is essentially a second-price auction with an anonymous reserve. Of course, the buyers are not identical. Some buyers have higher *ratings* and these ratings are public knowledge. The distributions values of agents with different ratings may generally be distinct. Therefore, the eBay auction may be suboptimal. Surely though, if the eBay auction was very far from optimal, eBay may have switched to a better auction. The theorem below justifies eBay sticking with the second-price auction with anonymous reserve.

Theorem 5.11 (Hartline and Roughgarden, 2009) *In (possibly non-identical) regular single-item environments, the second price auction with an anonymous reserve is a 4-approximation to the optimal revenue.*

The proof of this theorem is a non-trivial extension of basic prophet inequality proof that we do not give here. It should be noted that the proof does not make use of competition between agents, i.e., the tie-breaking rule. Therefore, an anonymous price that satisfies the conditions of the theorem is the monopoly price for the distribution of the maximum value. Finally, the bound given in this theorem is not known to be tight. The two agent example with F_1, a point mass at 1, and F_2, the equal-revenue distribution (i.e., $F_2(z) = 1 - 1/z$), shows that the worst-case distributional setting has approximation factor at least 2.

5.3.3 Regular Distributions

Recall that for i.i.d. regular distributions the optimal reserve price is the monopoly price. This reserve price is fairly easy to determine. For the environments considered thus far the suggested reserve prices were not the monopoly price. Furthermore, the calculation of reserve prices were interdependent. They depended on the number of agents and the specific distributions of all agents. It turns out that for (non-identical) regular environments monopoly reserve prices also give a 2-approximation.

Recall that in non-identical environments the optimal auction is not a reserve-price-based mechanism. The optimal auction carefully optimizes agents' virtual values at all points of their respective distributions, whereas the Vickrey auction with reserves monopoly only considers the inverse virtual values of zero. While the reserve-pricing revenue is suboptimal, the following theorem shows that its expected revenue is never too far from the optimal.

Theorem 5.12 (Hartline and Roughgarden, 2009) *In (non-identical) regular single-item environments, the second-price auction with monopoly reserves gives a 2-approximation to the optimal revenue.*

Before giving the proof of this theorem consider the following intuition. The optimal auction, OPT, and Vickrey with monopoly reserves, $\mathrm{Vic}_{\mathbf{m}}$, either have the same winner or different winners. If they have the same winner then they have the same virtual surplus. If they have different winners then winner in OPT is not the agent with the highest value. Of course OPT's winner can pay at most their value, but $\mathrm{Vic}_{\mathbf{m}}$'s winner pays at least the second highest value which is at least the value of OPT's winner. Therefore, in this case the payment in $\mathrm{Vic}_{\mathbf{m}}$ is higher than the payment of OPT. An important observation that is glossed over in this informal argument but necessary for a formal proof is that for regular distributions $\mathrm{Vic}_{\mathbf{m}}$ never sells to an agent with negative virtual value.

Proof: (of theorem) Let I be the winner of OPT and J be the winner of $\mathrm{Vic}_{\mathbf{m}}$. I and J are random variables. Notice OPT and $\mathrm{Vic}_{\mathbf{m}}$ both do not sell the item if and only if all virtual values are negative. In this situation define $I = J = 0$.

We start by simply writing out the expected revenue of OPT as its expected virtual surplus conditioned on $I = J$ and $I \neq J$.

$$\mathbf{E}[\mathrm{OPT}] = \mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J] + \mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J]\,.$$

We will prove the theorem by showing that both the terms on the right-hand side are bounded from above by $\mathbf{E}[\mathrm{Vic}_{\mathbf{m}}]$.

$$\begin{aligned}
\mathbf{E}[\phi_I(v_I) \mid I = J]\,\mathbf{Pr}[I = J] &= \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] \\
&\leq \mathbf{E}[\phi_J(v_J) \mid I = J]\,\mathbf{Pr}[I = J] + \mathbf{E}[\phi_J(v_J) \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&= \mathbf{E}[\mathrm{Vic}_{\mathbf{m}}]
\end{aligned}$$

The inequality in the above calculation follows because regularity of the distribution implies that the virtual value of the winner of $\mathrm{Vic}_{\mathbf{m}}$ is always non-negative. Therefore, this added term is non-negative.

$$\begin{aligned}
\mathbf{E}[\phi_I(v_I) \mid I \neq J]\,\mathbf{Pr}[I \neq J] &\leq \mathbf{E}[v_I \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J \mid I \neq J]\,\mathbf{Pr}[I \neq J] \\
&\leq \mathbf{E}[p_J \mid I \neq J]\,\mathbf{Pr}[I \neq J] + \mathbf{E}[p_J \mid I = J]\,\mathbf{Pr}[I = J] \\
&= \mathbf{E}[\mathrm{Vic}_{\mathbf{m}}]
\end{aligned}$$

The first inequality in the above calculation follows because $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)} \leq v_i$ (since $\frac{1-F_i(v_i)}{f_i(v_i)}$ is always non-negative). The second inequality follows because J is the highest valued agent and I is a lower valued agent and therefore, J's price is at least I's value. The third inequality follows because payments are non-negative so the term added is non-negative. □

This 2-approximation theorem is tight in two senses. First, there is a non-identical regular distribution such that the Vickrey auction with monopoly reserves is a 2-approximation. Second, Vickrey with monopoly reserves is not a 2-approximation for irregular distributions, even i.i.d. irregular distributions. See Hartline and Roughgarden (2009) for these examples.

5.4 Extensions and Conclusions

Monopoly reserve prices provide good approximations more generally than single-item environments. In particular Hartline and Roughgarden (2009) show that the VCG mechanism with monopoly reserves gives a 2-approximation for (a) regular, non-identical, matroid environments and (b) monotone-hazard-rate, non-identical, downward-closed environments. A downward-closed environment is one where any subset of a feasible set is also feasible. Examples include single-minded combinatorial auctions, e.g., from Lehmann et al. (1999); and do not include (non-excludable) public good environments.

In irregular, matroid environments (and some generalizations), posted pricings provide good approximations to the optimal mechanism (Chawla et al., 2010a). Furthermore, posted pricing mechanisms are not susceptible to collusion. This makes posted pricing mechanisms incredibly robust. This robustness may compensate for their non-optimality and justify their wide prevalence in practice. Most resource allocation is done via posted prices. Even eBay, the popular online auction site, is shifting from an auction-based platform a posted-price platform. Part of the reason for this shift is posted pricings do not require all agents to be present at once; instead, agents can arrive and depart as they wish. The revenue guarantees of posted pricings are robust to any exogenous or endogenous ordering on the agents.

These results on approximating optimal mechanisms with reserve prices all rely on the precise understanding we have from optimal mechanism design that shows that optimal mechanisms are virtual surplus maximizers. Essentially all that needs to be done is show that a reserve-price-based mechanism is a virtual surplus approximator. For all of the positive results we have described, this task is greatly simplified by the downward-closure property. In particular, in a downward-closed environment there is never need to serve an agent with negative virtual value. As any subset of a feasible set is feasible, these agents can be ignored; i.e., reserve prices can always be at least the monopoly price. One of the biggest challenges in extending this work to non-downward-closed environments like public good problems is that in these environments optimal mechanisms do serve agents with negative virtual values. Approximation in environments were the objective contains sums of positive and negative numbers is analytically challenging when not impossible.

6 Multi-dimensional Environments

We now turn to environments where the agents' preferences are multi-dimensional. E.g., in a home buyer may have a distinct value for different houses on the market; an Internet user may have distinct values for various qualities of service; an advertiser on an Internet search engine may value traffic for search phrase "mortgage" higher than that for "loan", etc.

For the objective of profit, there are no general descriptions of optimal mechanisms for environments with multi-dimensional agents. Essentially, mechanisms for multi-dimensional environments are complex and optimizing over them does not yield concise or intuitive descriptions, nor does it yield practical mechanisms. In this section we will explore approximation for the objective of profit maximization. In particular, we will show that both VCG with reserve prices and posted-pricing mechanisms can approximate the optimal mechanism. Furthermore, the prices in these mechanisms can be easily calculated.

We will use as a running example in this chapter the multi-item environments. In a multi-item environment there are n agents and m items (e.g., houses). Each agent i has a value v_{ij} for house j. The agents are unit-demand, i.e., each wants at most one house, and the houses are unit-supply, i.e., each can be sold to at most one agents. Agent values are drawn independently at random, e.g., with $v_{ij} \sim F_{ij}$.

6.1 Item Pricing

We start with the special case of the matching markets where there is only one agent, i.e., $n = 1$. In this environment an important optimization question is in identifying revenue optimal pricings. I.e., a pricing $\mathbf{p} = (p_1, \ldots, p_m)$ such that when the agent buys the item that generates the highest positive utility, i.e., the j that maximizes $v_j - p_j$, the revenue of the seller is maximized.

Unfortunately, there is no concise understanding of optimal pricings and their revenue. Therefore, in pursuit of goal approximately optimal pricings, the first hurdle is in finding concise understanding of an upper bound on the revenue of an optimal pricing. Then, if a pricing approximates this upper bound, it also approximates the optimal pricing.

The main idea in obtaining an upper bound is from the thought experiment where we imagine that instead of one agent with unit-demand preferences over the m items that we have m (single-dimensional) agents who each want their specific item, but with the constraint that at most one can be served. In this latter environment the optimal selling mechanism would be the optimal single-item auction discussed in Section 5. Notice that where in the pricing problem the seller can only post a price on each item, in the auction problem, competition between agents can drive the price up. Therefore, intuition suggests that the revenue in the (single-dimensional) auction environment may be an upper bound on the revenue in the (multi-dimensional) pricing environment. This is indeed the case.

Theorem 6.1 (Chawla et al., 2007) *For any product distribution $\mathbf{F} = F_1 \times \cdots \times F_m$, the expected revenue of the optimal single-agent, m-item pricing when the agent's values for the items are drawn from $\mathbf{F}$ is at most that of the optimal single-item, m-agent auction with each agents' value for the item drawn from $\mathbf{F}$.*

Proof: Any item pricing $\mathbf{p}$ can be converted into a single-item auction $\mathcal{M}_{\mathbf{p}}$ such that the expected revenue from the item pricing is at most that of the auction. For convenience define

$v_0 = p_0$ representing the value the seller receives for keeping the item (which is of course at price zero). The auction $\mathcal{M}_{\mathbf{p}}$ assigns the item to the agent j that maximizes $v_j - p_j$. For any fixed values of the other agents, $\mathbf{v}_{-j}$, this allocation rule is monotone in agent j's value and therefore ex post incentive compatible. It is also deterministic, so there is a critical value t_j for agent j which is the infimum of values for which the agent wins the auction; the agent pays exactly this critical value on winning. Of course $t_j \geq p_j$.

Now notice that the allocation rule of the auction $\mathcal{M}_{\mathbf{p}}$ is identical to the allocation rule of the pricing $\mathbf{p}$. For the pricing the agent chooses the item that maximizes $v_j - p_j$; for the auction the winner is selected to maximize $v_j - p_j$. Furthermore, the revenue for the pricing is exactly the p_j that corresponds to this j whereas in the auction it is t_j which, as discussed, is at least p_j. Therefore, the auction $\mathcal{M}_{\mathbf{p}}$ obtains at least revenue of the pricing $\mathbf{p}$.

Therefore, the optimal auction obtains at least the revenue of the optimal pricing. □

With the upper bound from optimal single-item auctions in hand, our goal of approximating the optimal pricing can be refined to approximating this optimal single-item auction revenue. We already saw how the prophet inequality theorem implied that the Vickrey auction with constant ironed virtual reserve prices (Theorem 5.8) is a 2-approximation. The same argument applies here to show that the same constant ironed virtual pricing, used as an item pricing, gives a 2-approximation to the optimal single-item auction revenue. The only difference in the argument is the tie-breaking rule which in item-pricing is by "$v_i - p_i$" instead of by "v_i" as in Theorem 5.8; of course the profit inequality is invariant on the tie-breaking rule.

Theorem 6.2 (Chawla et al., 2010a) *In (possibly non-identical and irregular) multi-item environments, an item pricing with constant ironed virtual reserve price is a 2-approximation to the optimal revenue.*

For single-agent environments item pricings are equivalent to deterministic mechanisms. Therefore, an approximation to the optimal pricing revenue is equivalently an approximation of the optimal deterministic mechanism.

6.2 Reduction: Unit-demand to Single-dimensional Preferences

It should be noted that the construction in the preceding section can be viewed as a reduction from a multi-dimensional unit-demand environment to a single-dimensional environment. We can conclude that from the perspective of approximation, the multi-dimensional unit-demand preferences are similar enough to single-dimensional preferences that a good approach to unit-demand environments is in simulating the outcome of the corresponding single-dimensional environment. We now make that connection and the reduction precise. (Crucial to this connection is the independence of the agents' values.)

Formally, consider the following general unit-demand environment. There are n agents and m services each agent i has value v_{ij} for service j. An outcome is a matching of agents to service (perhaps with some agents and some services left unmatched). We will denote this matching by the indicator $\mathbf{x}$ with $x_{ij} = 1$ if i receives service j and 0 otherwise. There is an arbitrary feasibility constraint over such matchings which we denote, as before, with a cost function $c(\cdot)$ which is zero or infinity for feasibility problems. We assume, without loss of generality, the implicit feasibility constraint that each agent can only receive one service, i.e., $\mathbf{x}$ such that $x_{ij} = x_{i'j} = 1$ for $i \neq i'$ have $c(\mathbf{x}) = \infty$.

A unit-demand environment is thus specified by the product distribution $\mathbf{F}$ indexed by agent-service pairs and the cost function $c(\cdot)$ over outcomes $\mathbf{x}$, also indexed by agent-item pairs.

6.2.1 Single-dimensional Analogy

As in the pricing environment we can define the single-dimensional analog to any unit-demand environment where each unit-demand agent is replaced with single-dimensional representatives. Notice that in the single-dimensional analog the implicit feasibility constraint that a unit-demand agent can receive at most one service is translated to the constraint that only one of its representatives can be served at once.

Definition 6.1 *The* representative environment *for the n agent, m service unit-demand environment given by $\mathbf{F}$ and $c(\cdot)$ is the nm agent single-dimensional environment given by $\mathbf{F}$, $c(\cdot)$, and single-dimensional agents indexed by coordinates ij.*

6.2.2 Upper bound

The restriction that only one representative of each unit-demand can be served at once induces competition between representatives. Intuitively this competition should result in an increased revenue in the optimal mechanism for the representative environment over the original unit-demand environment. In fact it is almost the case. The optimal mechanism for the representative environment (which is deterministic) is an upper bound on the optimal deterministic mechanism for the original (unit-demand) environment.

Theorem 6.3 (Chawla et al., 2010a) *In any unit-demand environment the optimal deterministic mechanism's revenue is at most that of the optimal mechanism for the single-dimensional representative environment.*

The proof of the above theorem follows from similar principles to the proof of Theorem 6.1.

6.2.3 Reduction

We now show that multi-dimensional unit-demand approximation can be reduced to a single-dimensional approximation problem. The techniques from Section 5 then can be applied to solve the single-dimensional agent approximation problem.

Extend the definition of sequential posted pricings to unit-demand environments with multiple agents (i.e., to generalize item prices). A sequential posted pricing is given by prices $\mathbf{p}$ with p_{ij} the price offered to agent i for service j. After the valuations are realized, the agents arrive in sequence and take their utility maximizing service that is still feasible, given the actions of preceding agents in the sequence. The revenue of such a process clearly depends on the sequence and we pessimistically assume the worst-case.

Definition 6.2 *A* sequential posted pricing *is an pricing of services (specialized) for each agent with the semantics that agents arrive in any order and take their favorite service that remains feasible. The revenue of such a pricing is given by the worst-case ordering.*

Consider the sequential posted pricing problem in both the original unit-demand environment and the representative single-dimensional environment. Suppose you had the choice of being the seller in one of these two environments, given the same distribution and costs, which environment would you choose? I.e., which environment gives a higher expected revenue? Whereas when considering auction problems, you would prefer the representative environment because of the increased competition, for sequential posted pricings there is no benefit from competition. In fact, the seller in the representative environment is at a disadvantage because the agents are in a worst case order and there are more possible orderings of the agents in the nm-agent representative environment than the n-agent original environment.

Theorem 6.4 (Chawla et al., 2010a) *The expected revenue of a sequential posted pricing for unit-demand environments is at least the expected revenue of the same pricing in the representative single-dimensional environment.*

Proof: Compare sequential posted pricings for unit-demand environments (i.e., with n unit-demand agents) with sequential posted pricings for their representative environments (i.e., with nm single-dimensional agents). The difference between these two environments with respect to sequential posted pricings is that in the representative environment the nm agents can arrive in any order whereas in the original environment the an agent arrives and considers the prices on services ordered by utility. Thus, the set of orders in which the nm prices are considered in the representative environment contains the set of orders in the original environment. For worst-case sequences, then, the representative environment is worse. □

Corollary 6.5 (Chawla et al., 2010a) *If a sequential posted pricing is approximately optimal in the representative (single-dimensional) environment it is approximately optimal in the original (unit-demand) environment.*

6.2.4 Instantiation

It remains to instantiate the reduction from sequential posted pricing approximation in unit-demand environments to single-dimensional environments. I.e., we need to show that there are good sequential posted pricing mechanisms for single-dimensional environments. Here we will give such an instantiation for multi-item environments, i.e., where the services are items, and each item has only one unit of supply.

The representative environment for unit-demand multi-item environments is one where there are nm agents each agent ij with value $v_{ij} \sim F_{ij}$ desires item j. For any original agent i and all j at most one representative ij can win. For any item j and all i at most one representative ij can win. The optimal mechanism for this environment is just to choose the matching that maximizes virtual surplus; denote this mechanism by OPT

Let q_{ij}^{OPT} be the probability that OPT serves representative ij. Let $p_{ij}^{\mathrm{OPT}} = F_{ij}^{-1}(1-q_{ij}^{\mathrm{OPT}})$ be the corresponding price at which, if posted to representative ij, would be accepted with probability q_{ij}^{OPT}. Now consider the pricing $p_{ij} = F_{ij}^{-1}(1 - q_{ij})$ for $q_{ij} = q_{ij}^{\mathrm{OPT}}/2$. These probabilities and prices can be calculated, for instance, by simulating the optimal mechanism.

Definition 6.3 *The* simulation prices, $\mathbf{p}$, *are the* $p_{ij} = F_{ij}^{-1}(1 - \frac{1}{2}\mathbf{Pr}[\mathrm{OPT}\ \textit{serves}\ ij])$.

We claim that sequential posted pricing with the simulation prices give an 8-approximation to the optimal mechanism's revenue. The theorem is proven in two steps, the first gets an upper bound on the revenue of the optimal mechanism in terms of the above prices and probability, the second gets a lower bound on the sequential pricing revenue in terms of the same. These steps are given by the lemmas below.

Theorem 6.6 (Chawla et al., 2010a) *For regular distributions in the representative matching market environment, the sequential posted pricing with the simulation prices* $\mathbf{p}$ *is an 8-approximation to the revenue of the optimal mechanism.*

This theorem instantiates of the reduction for unit-demand multi-item environments. Similar instantiations can be proven for generalizations including in environments with feasibility structure induced by matroid set systems. Sequential posted pricings are not good approximations in more general downward-closed environments.

Lemma 6.7 *For regular distributions, the expected revenue of the optimal mechanism,* OPT*, is at most* $\sum_{ij} p_{ij}^{\mathrm{OPT}} q_{ij}^{\mathrm{OPT}}$.

Proof: The proof of this lemma follows from a standard approach. Consider an "unconstrained" mechanism that allocates to each representative ij with probability at most q_{ij}^{OPT} but is not constrained by the original feasibility constraints, i.e., that only one representative ij of each agent i is served and that each item j is only allocated to at most on representative ij. In such an unconstrained environment the agents do not interact at all. Furthermore, by regularity and the fact that the original p_{ij}^{OPT} are at least the monopoly price, the optimal unconstrained mechanism simply posts price p_{ij}^{OPT} to each representative ij. Its expected revenue is $\sum_{ij} p_{ij}^{\mathrm{OPT}} q_{ij}^{\mathrm{OPT}}$. Finally, OPT, the optimal mechanism for the constrained environment, is a valid solution to the unconstrained environment, therefore the optimal unconstrained mechanism revenue gives an upper bound on its revenue. □

Lemma 6.8 *For regular distributions, the expected revenue from the sequential posted pricing of the simulation prices is at least* $\frac{1}{8}\sum_{ij} p_{ij}^{\mathrm{OPT}} q_{ij}^{\mathrm{OPT}}$.

Proof: If the sequential posted pricing is able to make an offer to agent ij then the expected revenue is $q_{ij}p_{ij} \geq q_{ij}^{\text{OPT}} p_{ij}^{\text{OPT}}/2$. This inequality follows because the $q_{ij} = q_{ij}^{\text{OPT}}/2$ and $p_{ij} \geq p_{ij}^{\text{OPT}}$ (since prices only increase with a lower selling probability). We now show that the probability that the sequential posted pricing is able to make the offer to representative ij is at least 1/4. As a consequence the expected revenue from representative ij is $q_{ij}^{\text{OPT}} p_{ij}^{\text{OPT}}/8$; and summing over all representatives ij gives the lemma.

To show that the probability that it is feasible to offer service to representative ij is at least 1/4, consider the worst-case ordering for this probability, i.e., where representative ij is last. Representative ij can be served if for all $j' \neq j$ representatives $j'i$ are not served, and for all $i' \neq i$ representatives $i'j$ are not served. The first event certainly happens if $v_{i'j} < p_{i'j}$ for all $i' \neq i$ and the second if $v_{ij'} < p_{ij'}$ for all $j' \neq j$. Each happens with probability at most 1/2 because $\sum_{j'} q_{ij'} \leq 1/2$ since the optimal mechanism allocates to one of these ij' representatives with probability at most one (by the feasibility constraint). Since the above events are independent the probability that both occur is the product of the probability that each occurs which is, therefore, at least 1/4. $\square$

6.3 Extensions and Conclusions

One of the difficulties of multi-dimensional preferences over single-dimensional preferences is that for multi-dimensional preferences optimal mechanisms are not necessarily deterministic. I.e., it can be optimal for an agent to receive a probability distribution over outcomes. In the single-agent case we refer to such a randomized mechanism as a lottery pricing.

An interesting question is in quantifying the relative difference between optimal randomized mechanisms and optimal deterministic mechanism. Essentially, are deterministic mechanisms good approximations to randomized mechanisms? It turns out the answer to this question is quite different in environments when an agent's preferences over distinct services are independent or correlated. For general correlations, deterministic mechanisms do not approximate randomized mechanisms to any factor (Briest et al., 2010); for product distributions in many unit-demand environments, including multi-item auctions, deterministic mechanisms give are constant approximations (Chawla et al., 2010b). Therefore, for product distribution environment, the results discussed above imply that posted pricing mechanisms approximate the optimal (possibly randomized) mechanism.

The biggest direction for future study in approximation in multi-dimensional environments is in moving beyond unit-demand, product-distribution environments. The most natural next step would be to address additive valuations, e.g., where the value an agent derives for a bundle of services is the sum of their values for each service in the bundle. For these environments (Armstrong, 1996) shows that pricing the grand bundle is approximately optimal in the limit were the sum of the agents independent random values are concentrated. The direction we pose here is to understand bundle pricings before the limiting behavior takes effect.

7 Prior-independent Environments

In this section we consider approximation in environments where designer does not know the prior distribution. For reasons to be enumerated below, we will be proposing approximation mechanisms that are ex post incentive compatible, meaning, the agents do not need to know the prior distribution either. A good mechanism is then one that, for every possible distribution over values, obtains a good approximation to the optimal mechanism for that distribution.

7.1 Motivation

Consider from where the designer may have learned the prior-distribution. There are two most logical candidates. The first is, as alluded to above, is from the designer's history in interacting with these or similar agents. The problem with this point of view is that the earlier agents may strategize so that information about their preferences is not exploited by the designer later. In fact, if a monopolist cannot commit not exploit the agents using information from prior interaction then the socially efficient (i.e., surplus maximizing) outcome is the only equilibrium, e.g., via the Coase Conjecture.

The second candidate is *market analysis*. The designer can hire a marketing firm to survey the market and provide distributional estimates of agent preferences. This mode of operation is quite reasonable in large markets. However, in large markets each agent will have little impact and usually this enables asymptotically optimal mechanisms, see, e.g., Segal (2003). Prior-independent mechanisms are most interesting in small, a.k.a., *thin*, markets. Contrast the large market for automobiles to the thin market for space crafts. There may be five organizations in the world in the market for space crafts. How would a designer optimize a mechanism for selling space crafts? First, even if the agents' values do come from a distribution, the only way to sample the distribution is to interview the agents themselves. Second, even if we did interview the agents, the most data points we could obtain is five. This is hardly enough for statistical approaches to be able to estimate the distribution of agent values. This strongly motivates a question related to prior-independent mechanism design which is how many samples from a distribution are necessary to design a mechanism that can approximate the optimal mechanism for the distribution.

There are other reasons to consider prior-independent mechanism design besides the questionable origin of prior information. The most striking is the frequent inability of a designer to redesign a new mechanism for each scenario they wish to run a mechanism in. This is not just a concern, in many settings it is a principle. Consider the standard Internet routing protocol TCP/IP. This is the protocol responsible for sending emails, browsing web pages, steaming video, etc. Notice that the workloads for each of these tasks is quite different. Emails are small and can be delivered with several minutes delay without issue. Web pages are small, but must be delivered immediately. Comparably, video streaming requires a high responsiveness and a large bandwidth. There is not the flexibility to install new protocols in Internet routers each time a new network usage pattern arises. Instead, a good protocol, such as TCP/IP, should work pretty well in any setting, perhaps ones well beyond the imaginations of the original designers of the Internet.

The final motivation we will discuss for prior-independent mechanism design is that the solution of Bayesian optimal (or approximate) mechanisms is incomplete. It solves the problem of what a designer should do who knows the prior-distribution, but in many real situations a designer may not have such knowledge. Requiring the designer to acquire distribution information from outside "the system", therefore, does not completely solve the designer's problem.

7.2 Optimal Mechanisms

The economics literature on Bayes-Nash implementation has studied the question of what a designer can implement without any knowledge of the setting. For instance, a partial implementation of the optimal mechanism is available by the following uninteresting solution: as the agents to report the distribution, shoot them if they disagree, and otherwise run the optimal mechanism for the reported distribution. More sophisticated approaches enable full virtual implementation of the optimal mechanism, i.e., there is a mechanisms that approximates the revenue of optimal mechanism arbitrarily closely in every Bayes-Nash equilibrium. This mechanism, in a similar fashion to the aforementioned partial implementation, relies on agents reporting the distribution.

Allowing cross-reporting mechanisms such as the above as a solution to prior-independent mechanism design begs the question. Furthermore, there are serious practicality and robustness issues for these cross-reporting mechanisms, e.g., see Bergemann and Morris (2005). In our discussion of prior-independent mechanisms we deliberately rule out such solutions in attempt to explore what is possible without distributional knowledge, either a priori or from cross-reports. Such a restriction is certainly with loss, but this makes the positive results we discuss only stronger. Below we identify simple, ex post incentive compatible, prior-independent approximation mechanisms.

7.3 "Resource" Augmentation

Consider the classic result of Bulow and Klemperer (1996) which states that in i.i.d., regular, single-item environments the Vickrey auction with one more agent (from the distribution) obtains a higher revenue than the optimal auction (without the additional agent). This result is often interpreted as a critique on exogenous entry or a statement about competition being better for revenue than reserve prices.[2] The Bulow-Klemperer result is in fact suggesting a prior-independent strategy for approximating the revenue of the optimal mechanism: recruit one more agent. Dhangwatnotai et al. (2010) provide a nice discussion of this viewpoint.

Theorem 7.1 (Bulow and Klemperer, 1996) *In i.i.d., regular, single-item environments,*

[2] Recall, that in Sections 5 and 6 we came to (approximately) the opposite conclusion, i.e., that reserve- or posted-prices, without competition, are enough to guarantee good revenue. Both viewpoints are correct and interesting. The Bulow-Klemperer result provides useful intuition for competitive environments, where as the Chawla et al. result provides useful intuition for environments where collusion is an issue.

the expected revenue of the Vickrey auction on $n + 1$ agents is at least the expected revenue of the optimal auction on n agents.

Unfortunately the "just add a single agent" result fails to generalize beyond single-item auctions. When k units of an item are auctioned to $n+1$ agents with the k-Vickrey auction, the revenue does not approximate that of the optimal k-unit n-agent auction. In order to beat the optimal auction, k additional agents must be added.

Theorem 7.2 *In i.i.d., regular, k-unit environments, distributions in matroid environments, the expected revenue of k-Vickrey on $n+k$ agents at least the expected revenue of the optimal auction on the original n agents.*

Consider the extreme case where $k = n$, a.k.a., that of *digital goods*. Notice that the Vickrey auction in this environment obtains no revenue. All items are given away for free. Unfortunately, the Bulow-Klemperer result in such an environment seems less actionable as prior-independent approach to mechanism design; to obtain at least the revenue of the optimal mechanism we would need to double the size of the market!

7.4 Single-sample Mechanisms

Dhangwatnotai et al. (2010) show that the Bulow-Klemperer result can be approximately extended to multi-unit environments and simultaneously address the question of how large a marketing sample must be to provide good enough statistical information for the design of an approximately optimal mechanisms. Their answer: one. Suppose instead of recruiting an additional agent to participate in the mechanism, we find a single agent for market analysis, and then run a Vickrey auction with this agent's reported value as a reserve price. In i.i.d., regular, multi-unit environments this auction is a 2-approximation.

Definition 7.1 *The* single-sample auction *is the draws a single-sample from the distribution and runs the Vickrey auction with the sampled value as a reserve price.*

The following lemma can be seen as a corollary of the $n = 1$ agent special case of the Bulow-Klemperer result. We give an alternative geometric proof of it that is due to Dhangwatnotai et al. (2010).

Lemma 7.3 (Dhangwatnotai et al., 2010) *For a single-agent with value drawn from regular distribution F, the revenue from a random take-it-or-leave-it offer $r \sim F$ is at least half the revenue from the (optimal) monopoly offer.*

Proof: Let $R(q)$ be the revenue curve for F. Let q^* be the quantile corresponding to the monopoly price, i.e., $q^* = \operatorname{argmax}_q R(q)$. The expected revenue from such a price is $R(q^*)$. Recall that drawing a random value from the distribution F is equivalent to drawing a uniform quantile $q \sim U[0, 1]$. The revenue from such a random price is $R(q)$. In Figure 1 the area of region A is $R(q^*)$. The area of region B is $\mathbf{E}_q[R(q)]$. Of course, the area of C is

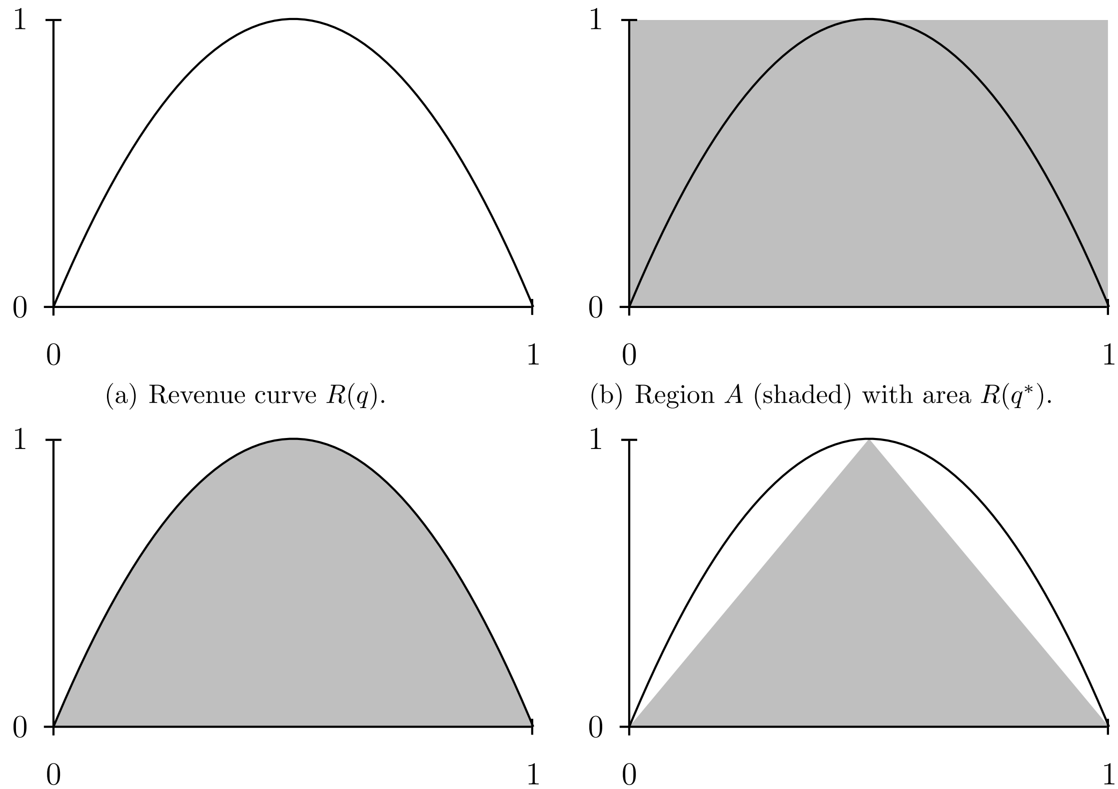

(a) Revenue curve $R(q)$.
(b) Region A (shaded) with area $R(q^*)$.
(c) Region B (shaded) with area $\mathbf{E}[R(q)]$.
(d) Region C (shaded) with area $A/2$.

Figure 1: In the geometric proof of the that a random reserve is a 2-approximation to the optimal reserve, the areas of the shaded regions satisfy $A \geq B \geq C = A/2$.

less than the area of B, by concavity of $R(\cdot)$, but at least half the area of A, by geometry. The lemma follows. $\square$

It is now a simple exercise to generalize Lemma 7.3 to multi-unit auctions to give the following theorem. Essentially, a random reserve is approximately as good as the monopoly reserve. This theorem can also be generalizes to matroid and downward-closed environments with a VCG-based single-sample mechanism.

Theorem 7.4 (Dhangwatnotai et al., 2010) *For any i.i.d., regular, multi-unit environment, the single-sample auction is a 2-approximation to the optimal auction.*

7.5 Prior-independent Mechanisms

To design approximately optimal mechanisms without any prior information, as observed by Goldberg et al. (2001); Segal (2003); Baliga and Vohra (2003), we can use the reports of some agents for market analysis on other agents. For example, for digital goods, the mechanism that pairs the agents and runs a Vickrey auction on each pair is a 2-approximation. This result follows as each agent faces a random reserve from the distribution and Lemma 7.3 implies that such a reserve is a 2-approximation to the monopoly reserve (which is optimal).

Theorem 7.5 *For i.i.d., regular, multi-unit environments, the* pairing auction*, i.e., that randomly pairs the n agents in $n/2$ Vickrey auctions is a 2-approximation to the optimal revenue.*

This approach can be extended to the same environments as the single-sample auction as follows. Simulate both the VCG mechanism and the pairing auction in parallel, but serve only the winners of both mechanisms (at the higher of their prices).

7.6 Conclusions and Extensions

We have exhibited a few simple approaches for designing prior-independent mechanisms. These approaches make strong usage of the regularity of the distribution, the symmetry of identical distributions, and downward-closure of the feasible set system. The main conclusion is that market analysis can be done on-the-fly by the mechanism as it is run; the resulting mechanism is often a good approximation.

The regularity assumption can be relaxed if more samples from the distribution are available. Of course, in i.i.d. environments there for each agent there are $n-1$ other samples from the distribution. The following prior-independent mechanism gives a good approximation fairly generally. Partitioning the agents in to two parts, estimate the distribution for one part, and run the optimal mechanism for the estimated distribution on opposite part. In some environments this can be done symmetrically for both parts. See, e.g., Goldberg et al. (2001); Baliga and Vohra (2003); Devanur and Hartline (2009).

The distributional symmetry (of the i.i.d. assumption) can be relaxed. If the agents are a priori distinguishable by publicly observable attributes, and there are at least two agents with each attribute, then agents can be paired with other agents with the same attribute. See, e.g., Balcan et al. (2008); Dhangwatnotai et al. (2010).

One final note, the viewpoint presented here on prior-independent mechanisms is one where there is a prior but the designer just does not know it. As the mechanisms under discussion are ex post incentive compatible, the prior is not needed for equilibrium. It is possible, then, to dispense with the prior completely; however, in doing so it is not clear to what we should compare to for an approximation. Hartline and Roughgarden (2008) suggest a prior-free benchmark that for each valuation profile is the supremum over i.i.d. distributions of the optimal auction's revenue on the valuation profile. This benchmark has the nice property that a prior-free mechanism that approximates it simultaneously approximates the optimal mechanism for any i.i.d. Bayesian environment. With some minor tweaks, the partitioning mechanisms described above gives a good approximation to this benchmark (Devanur and Hartline, 2009).

Returning to our viewpoint as approximation as a lens by which we can distinguish details from the salient features of the model, the conclusion of this section is that knowledge of the prior is a detail and good, mechanisms can be designed without it. These mechanisms can be more robust than their prior-dependent counter parts. The prior-free mechanisms discussed in the preceding paragraph which obtain their performance guarantee in worst-case over all valuation profiles are especially robust.

References

Armstrong, M. (1996). Multiproduct Nonlinear Pricing. *Econometrica*, 64(1):51–75.

Balcan, M.-F., Blum, A., Hartline, J., and Mansour, Y. (2008). Reducing mechanism design to algorithm design via machine learning. *J. Comput. Syst. Sci.*, 74(8):1245–1270.

Baliga, S. and Sjöström, T. (2008). Mechanism design: new developments. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan.

Baliga, S. and Vohra, R. (2003). Market research and market design. *Advances in Theoretical Economics*, 3.

Bergemann, D. and Morris, S. (2005). "robust mechanism design". *Econometrica*, 73:1771–1813.

Briest, P., Chawla, S., Kleinberg, R., and Weinberg, S. (2010). Pricing randomized allocations. In *SODA*, pages 585–597.

Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86:180–194.

Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *The Journal of Political Economy*, 97:1060–90.

Chawla, S., Hartline, J., and Kleinberg, R. (2007). Algorithmic pricing via virtual valuations. In *Proc. 9th ACM Conf. on Electronic Commerce*.

Chawla, S., Hartline, J., Malec, D., and Sivan, B. (2010a). Sequential posted pricing and multi-parameter mechanism design. In *Proc. 41st ACM Symp. on Theory of Computing*.

Chawla, S., Malec, D., and Sivan, B. (2010b). The power of randomness in bayesian optimal mechanism design. In *ACM Conference on Electronic Commerce*, pages 149–158.

Choi, J. and Kim, T. (1999). A nonparametric, efficient public good decision mechanism: Undominated bayesian implementation,. *Games and Economic Behavior*, 27(1):64–85.

Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:17–33.

Devanur, N. and Hartline, J. (2009). Limited and online supply and the bayesian foundations of prior-free mechanism design. In *Proc. 11th ACM Conf. on Electronic Commerce*, pages 41–50.

Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ECOM10*.

Goldberg, A. V., Hartline, J. D., Karlin, A., Saks, M., and Wright, A. (2006). Competitive auctions. *Games and Economic Behavior*, 55:242–269.

Goldberg, A. V., Hartline, J. D., and Wright, A. (2001). Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM.

Groves, T. (1973). Incentives in teams. *Econometrica*, 41:617–631.

Hartline, J. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Proc. 39th ACM Symp. on Theory of Computing*.

Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*.

Lehmann, D., O'Callaghan, L. I., and Shoham, Y. (1999). Truth revelation in approximately efficient combinatorial auctions. In *Proc. 1st ACM Conf. on Electronic Commerce*, pages 96–102. ACM Press.

Maskin, E. and Sjöström, T. (2002). Implementation theory. In Arrow, K., Sen, A., and Suzumura, K., editors, *Handbook of Social Choice and Welfare*, volume 1, chapter 5, pages 237–288. Elsevier.

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

Reynolds, S. and Wooders, J. (2009). Auctions with a buy price. *Economic Theory*, 38:9–39.

Riley, J. and Samuelson, W. (1981). Optimal auctions. *American Economic Review*, 71:381–392.

Samuel-Cahn, E. (1984). Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216.

Segal, I. (2003). Optimal pricing mechanisms with unknown demand. In *American Economic Review 93*.

Thanassoulis, J. (2004). Haggling over substitutes. *J. of Economic Theory*, 117:217–245.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37.

Vincent, D. and Manelli, A. (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185.

Wang, X., Montgomery, A., and Srinivasan, K. (2008). When auction meets fixed price: a theoretical and empirical examination of buy-it-now auctions. *Quantitative Marketing and Economics*, 6:339–370.

Wilson, R. (1987). Game theoretic analysis of trading processes. *Advances in Economic Theory*.