

Approximation in Mechanism Design

By JASON D. HARTLINE*

A mechanism gives 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 I survey several results from the theory of approximation in 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 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. Below I will motivate the perspective that a, for example, 2-approximation can have important theoretical and practical consequences.

As discussed, the class of all mechanisms is incredibly rich and there are environments, see, e.g., Vincent and Manelli (2007), where any undominated mechanism is optimal for some set-

ting. 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 is approximately optimal. 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, one conclusion I will make is that reserve-price-based auctions and posted pricings are approximately optimal in a wide range of settings.

Approximation also provides a lens with which to explore the salient features of a model or mechanism. Suppose we wish to determine whether a particular feature of a mechanism is important. If there exists a subclass of mechanisms without that feature that gives a good approximation, 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 monetary transfers are very important for surplus maximization; where as a conclusion I will discuss is that competition between agents is not that important. Essentially, approximation provides a means to determine which aspect of the model are details and which are not details (cf., Wilson, 1987). For instance, I will describe how mechanisms based on reserve prices are often sufficient to approximate ones parameterized by the full distribution and how some environments permit a single (prior-independent) mechanism to approximate the optimal mechanism for every prior.

While it is no doubt a compelling success of the theory of mechanism design that the mechanisms it predicts are so prevalent in practice, the optimality of these mechanisms in ideal settings does not provide a complete explanation.

* Hartline: EECS Dept., Northwestern U., Evanston, IL 60208. hartline@eecs.northwestern.edu.

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.

I will illustrate the potential for the theory of approximation in three representative environments: *single-item auctions*, *multi-item pricing*, and *prior-independent digital-good auctions*. I will focus on the revenue objective. For each environment I will describe the optimal mechanism, compare and contrast it to a simple approximately optimal mechanism, and discuss generalizations.

I. 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 independently 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 this threshold. On day $n - 2$ the gambler should stop with any prize greater than the expected payoff 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 . 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 scenarios.

Approximation gives 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. I refer to the prize selection procedure when multiple prizes are above the threshold as the *tie-breaking rule*. The tie-breaking rule implicit in the specification of the gambler's game is by "smallest i ."

Theorem 1 (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 with respect to the tie-breaking rule.*

Theorem 1 is a *prophet inequality*: it suggests that even though the gambler does not know the realizations of the prizes in advance, he can still do half as well as a prophet who does. Unlike our the optimal (backwards induction) strategy this prophet inequality provides substantive conclusions. Most obviously, it is a very simple strategy. The result is clearly driven by trading off the probability of not stopping and receiving no prize with the probability of stopping early with a suboptimal prize. 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. The invariance of the performance bound to tie-breaking rule suggests the bound can be applied to other related scenarios.

PROOF OF THEOREM 1:

Define $q_i = F_i(t) = \Pr[v_i < t]$ and $\chi = \prod_i q_i$ to be the probability that the gambler rejects all prizes. The proof follows in three steps. In terms of t and χ , I get an upper bound on the expected prophet's payoff. Likewise, I get a lower bound on expected gambler's payoff. Finally, assuming that the F_i 's are continuous,¹ I plug in $\chi = 1/2$ to obtain the result.

The prophet's expected payoff is

$$\begin{aligned} P &= \mathbf{E}[\max_i v_i] = t + \mathbf{E}[\max_i (v_i - t)] \\ &\leq t + \mathbf{E}[\max_i (v_i - t)^+] \\ &\leq t + \sum_i \mathbf{E}[(v_i - t)^+] \end{aligned}$$

where $(v_i - t)^+$ denotes $\max(v_i - t, 0)$.

I will split the gambler's payoff into two parts, the contribution from the first t units of the prize and the contribution from the remaining $v_i - t$ units of the prize. The first part is $G_1 = (1 - \chi)t$. To get a lower bound on the second part I consider only the contribution from the no-tie case. For any i , let \mathcal{E}_i be the event that all other prizes j are below the threshold t (but v_i is unconstrained). The bound is:

$$\begin{aligned} G_2 &\geq \sum_i \mathbf{E}[(v_i - t)^+ | \mathcal{E}_i] \Pr[\mathcal{E}_i] \\ &\geq \chi \sum_i \mathbf{E}[(v_i - t)^+]. \end{aligned}$$

The second line follows because $\chi = \prod_i q_i \leq \prod_{j \neq i} q_j = \Pr[\mathcal{E}_i]$ and because the conditioned variable $(v_i - t)^+$ is independent from the conditioning event \mathcal{E}_i . Therefore, the gambler's payoff is at least:

$$G \geq (1 - \chi)t + \chi \sum_i \mathbf{E}[(v_i - t)^+].$$

Plugging in $\chi = 1/2$ gives the result. \square

II. Single-item Auctions

I will start the discussion 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.

I briefly review the theory of optimal single-item auctions as developed by Myerson (1981). Each agent i is a risk-neutral quasi-linear utility maximizers with value v_i for service drawn independently from distribution F_i (density function f_i). Given the distribution, the *virtual value* of agent i is with value v_i is $\phi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)}$. The *virtual surplus* of a mechanism is the virtual value of its winner(s). The most important theorem in optimal auction theory relates virtual surplus to revenue.

Theorem 2 (Myerson, 1981) *In Bayes-Nash equilibrium, the expected virtual surplus and expected revenue of an auction are equal.*

Bayes-Nash equilibrium (BNE) requires that an agent's likelihood of service is monotone in her value. If the virtual valuation functions are monotone (i.e., the distributions are *regular*), the point-wise maximization of virtual surplus satisfies this BNE monotonicity requirement. Therefore, the revenue-optimal auction is the one that serves the agent with the highest positive virtual value. Moreover, when agent values are identically distributed, the virtual valuation functions are identical, and this optimal auction is equivalently the Vickrey auction with a reserve price equal to $\phi^{-1}(0)$ (a.k.a., the *monopoly price*).

The simplicity of the optimal auction when values are i.i.d. and regular should be noted and contrasted to the complexity of the optimal mechanism when values are non-i.i.d. or irregular. The Vickrey auction with reserve price is simple and practical; furthermore, its variants are widely used in practice. When values are non-i.i.d. the optimal auction is not reserve-priced-based, instead it is parameterized by distinct virtual value function for each agent.

This complex mechanism does not arise in practice, even in auction environments with explicit asymmetries. For example, in the eBay auction buyers are distinguished by publicly-observable "reputation", the optimal auction is likely to be a complex, asymmetric mechanism, and nonetheless the eBay auction is essentially the Vickrey auction with an (anonymous) reserve price. Fortunately, this auction is approximately optimal.

¹The theorem is true, modulo a more sophisticated choice of χ , for discontinuous distributions; I omit the discussion necessary to derive this more general result.

Theorem 3 (Hartline and Roughgarden, 2009)

*For regular distributions, the Vickrey auction with an anonymous reserve is a 4-approximation to the optimal single-item auction.*²

With non-identical distributions it may be natural to use non-anonymous reserve prices. With non-anonymous reserves the approximation factor can be improved and the result generalized to allow for irregular distributions. A *uniform virtual pricing* is $\mathbf{p} = (p_1, \dots, p_n)$ that satisfies $\phi_i(p_i) = \phi_j(p_j)$ for all i and j . The following theorem is an immediate consequence of the prophet inequality (Theorem 1) in virtual-value space with tie-breaking by “largest v_i ”.

Theorem 4 (Chawla et al., 2010a) *The Vickrey auction with uniform virtual reserve prices such that the item remains unsold with probability 1/2 is a 2-approximation to the optimal single-item auction.*

In contrast to the optimal auction for non-identically distributed agent values, Vickrey with non-anonymous reserve prices is parameterized by a single number for each agent. These reserve prices are easy to solve for and have an intuitive meaning. The revenue guarantee of these reserve prices is so robust that the “Vickrey” part of the mechanism is unnecessary. Posting these reserve prices as take-it-or-leave-it while-supplies-last offers is also a 2-approximation (by the prophet inequality with tie-breaking by “smallest p_i ”). Therefore, competition plays only a limited role in driving revenue, the performance guarantee is robust to collusion.

The extent to which these results hold beyond single-item auctions is of interest. I will refer to a mechanism for selling multiple identical units to unit-demand buyers as a multi-unit auction. A natural extension of the uniform-virtual-reserve auction generalizes to multi-unit settings where it continues to be a 2-approximation. The anonymous reserve result does not; an anonymous reserve cannot guarantee better than a logarithmic (in the number of units) factor of the optimal multi-unit auction revenue (Chawla et al., 2007).

²This bound is not known to be tight, the true approximation factor lies in [2, 4]. Of course, when distributions are similar then the performance can be much better than the (worst-case over distributions) approximation guarantee.

III. Multi-dimensional Environments

I now turn to environments where the agents’ preferences are multi-dimensional. E.g., 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 revenue, there are no general descriptions of optimal mechanisms for environments with multi-dimensional agents. In fact, Vincent and Manelli (2007) show that any undominated mechanism is optimal for some distribution over agent preferences. 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 contrast I will give simple descriptions of approximately optimal mechanisms for multi-dimensional preferences; furthermore, these mechanisms are based on prices that are easy to calculate.

I will focus on the example of pricing m items to a single unit-demand consumer with value v_i for item i drawn independently from distribution F_i . A revenue-optimal pricing $\mathbf{p} = (p_1, \dots, p_m)$ is 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 expected revenue of the seller is maximized.

Since there is no concise understanding of optimal pricings and their revenue, I will compare the performance of a candidate pricing to an upper bound on the optimal revenue. Compare, as a thought experiment, our single-agent multi-item pricing problem to the multi-agent single-item auction problem of Section II. In particular, keeping the distribution of the i th agent (resp. item) in the auction (resp. pricing) problem the same, which scenario generates more revenue? Intuitively, and this is borne out in the theory, the multi-agent auction problem revenue-dominates the multi-item pricing problem (Chawla et al., 2007). Furthermore, the analysis of Section II gives the optimal auction.

The prophet inequality (Theorem 1) with tie breaking by “smallest p_i ” yields the following.

Theorem 5 (Chawla et al., 2010a) *An item pricing with constant ironed virtual price is a 2-approximation to the optimal item pricing.*

There are a few important conclusions. First, multi-dimensional unit-demand preferences are similar to single-dimensional preferences. Second, the multi-dimensional approximation problem can be reduced to a single-dimensional “optimize virtual price” search space. Finally, this theorem can be generalized to multiple agents (see Chawla et al., 2010a) and therefore one can conclude that widely prevalent posted pricing mechanisms are indeed approximately optimal.

Note that in the result above the agent’s values for distinct items are required to be independent. Without this independence the multi-agent single-item auction does not give a reasonable upper bound on the optimal single-agent multi-item pricing. Additionally, while I have discussed only item pricings, optimal multi-dimensional mechanisms may require randomization, e.g., in the form of “lottery pricings” (Thanassoulis, 2004). Chawla et al. (2010b) show that this benefit of lottery pricing over item pricing is at most a constant factor (for independent distributions).

IV. Prior-independent Environments

In this section I consider approximation in environments where designer does not know the prior distribution. I will give an abbreviated motivation. I will briefly discuss non-parametric implementation. Finally, I will describe results showing that (a) only one sample from the distribution is necessary to design good mechanisms and (b) that a prior-independent mechanism could perform a small distributional analysis on-the-fly as it is being run.

Where does the designer’s knowledge of the prior come from? Perhaps market analysis: then in a thin market, e.g., for space stations, how accurate can this knowledge be? Perhaps from historical transactions: then might agents in the past game the system understanding that their actions will affect future prices (cf., Coase, 1972)?

In many environments optimizing the mechanism for each specific scenario is impossible. For example, eBay is the broker in many buyer-seller exchanges in many rather distinct markets. Instead of optimizing the exchanges for

each market, eBay has essentially one exchange mechanism. This mechanism must perform well in all markets.

The literature of non-parametric Bayes-Nash virtual implementation shows that an ignorant designer can implement essentially the same outcome as an informed designer. Such a mechanism relies on agents making cross reports about the other agents values using their own beliefs (e.g. Maskin and Sjöström, 2002); e.g., agents may be asked to report their entire distributional knowledge. Besides being non-robust (Bergemann and Morris, 2005) and impractical, in a sense, this solution begs the question. Instead, I will focus on mechanisms that approximate the informed designer’s optimal mechanisms in dominant strategies.

Our example problem here will be a digital-good environment. There are n identical units of an item and n agents with identically distributed values. I would like a prior-independent mechanism with expected revenue that approximates the revenue of the optimal auction. Of course the optimal mechanism here is to post the monopoly price $\phi^{-1}(0)$ for the distribution (Section II); this mechanism is not prior independent.

The Bulow-Klemperer (1996) Theorem states that for i.i.d. regular distributions, the Vickrey auction on $n + 1$ agents revenue-dominates the revenue-optimal single-item auction on n agents. Dhangwatnotai et al. (2010) observe that this “recruit another agent” suggestion is essentially a prior-independent strategy. Unfortunately, for this approach to have any bite in digital-good environments we would need to double the size of the market (and refuse to serve half of the agents). Dhangwatnotai et al. further observe that an implication of the $n = 1$ case of the Bulow-Klemperer Theorem is that with a single agent: a random reserve from the distribution obtains at least half the revenue of the optimal reserve. Consequently, a single sample from the distribution is enough to get an approximately optimal auction.

Lemma 1 (Dhangwatnotai et al., 2010)

When agents values are i.i.d. from a regular distribution F , with a single sample $r \sim F$, the digital good auction that posts price r is a 2-approximation to the optimal digital-good auction.

To design approximately optimal mechanisms without any prior information; as observed by Goldberg et al. (2001), Segal (2003), and Baliga and Vohra (2003); it is possible to use the reports of some agents for market analysis on others. For example, by the Bulow-Klemperer Theorem, the mechanism that pairs the agents and runs a Vickrey auction on each pair is a digital-good 2-approximation. The Goldberg et al. (2001) *random sampling auction* is only better, i.e., it is a 2-approximation, for regular distributions when $n = 2$ and its approximation factor approaches one in the limit (e.g. Segal, 2003).

Definition 1 *The random sampling auction randomly partitions the agents into two equal-sized sets and independently offers the agents in each set the monopoly price for the other set.*

These approaches extend to limited supply, position auctions (see, e.g., Varian, 2006), and beyond. I conclude that a single sample from the distribution is enough to get a good approximation in benign (i.e., regular distributions) environments, and that this sample can be obtained on-the-fly as the mechanism is run.

Returning to the viewpoint that approximation is a lens by which details can be distinguished from salient features, the conclusion of this section is that the prior is a detail and good, robust mechanisms can be designed without it.

REFERENCES

- 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.
- Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86:180–194.
- 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.
- Coase, R. (1972). Durability and monopoly. *Journal of Law and Economics*, 15(1):143–149.
- Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *Proc. 12th ACM Conf. on Electronic Commerce*.
- Goldberg, A., Hartline, J., and Wright, A. (2001). Competitive auctions and digital goods. In *Proc. 12th ACM Symp. on Discrete Algorithms*, pages 735–744.
- Hartline, J. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proc. 11th ACM Conf. on Electronic Commerce*.
- 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.
- Myerson, R. (1981). Optimal auction design. *Mathematics of Operations Research*, 6:58–73.
- 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*.
- Thanassoulis, J. (2004). Hagglng over substitutes. *J. of Economic Theory*, 117:217–245.
- Varian, H. (2006). Position auctions. In *Working Paper, UC Berkeley*.
- Vincent, D. and Manelli, A. (2007). Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185.
- Wilson, R. (1987). Game theoretic analysis of trading processes. *Advances in Economic Theory*.