

Aleck C. Johnsen – Research Statement (March 2020)

The theories of economics and computer science are inherently and essentially related. On one side solutions for models and games within economics require computational tractability to be practically relevant. On the other side there is a reduction from the analysis of robustness for arbitrary algorithms questions to a fundamental two-player game.¹ The result is the broad field of algorithmic game theory, in which my primary research interest is robustness of algorithms.

Robustness is necessary for the study of settings of *incomplete information*, for which there are two major branches. (1) In mechanism design, one or more items are allocated to n strategic agents with *private types*; design of mechanisms is subject to incentivizing the agents to participate, without knowing exactly how their types inform their behavior. (2) In online algorithms, an algorithm designer receives a sequence of n inputs and must commit to a decision before seeing each new individual input, *without knowing the future*. In both settings, an algorithm’s information about the input is structurally limited, and there is no single algorithm that is optimal given all realizations of the unknown inputs. Typically, mechanism design and online algorithms are treated as separate realms of research. My interest is towards building a unified framework, through recognizing themes of a common abstraction of incomplete information.

I expound using the framework of auctions which has motivated most of my research, but models of information extend to online algorithms. The canonical Bayesian auction setting has n agents who each have a private value for one item drawn independently and identically (i.i.d.) from a common probability distribution, with the distribution known by the auction designer. Bayesian auctions were thoroughly solved by Myerson (1981).

Consider alternatively a spectrum of *information settings* to describe the values of the agents, with “solutions” across the spectrum representing auction design with varying robustness. For different lines of research, I consider prior independent and prior free value settings. In studying robustness, a number of natural perspectives arise during analysis and lead to entire fundamental sub-topics. My research considers the following:

1. **Foundations of Benchmark Design.** *Benchmark functions* describe target performance (to be approximated pointwise). The benchmark can theoretically be set too low, or too high. Benchmarks themselves can and should be subject to scientific inquiry.
2. **Adversarial Play as Information Design.** Benchmark approximation is measured against a worst “test case” as if chosen by an omniscient *adversary*, playing against the algorithm in a zero-sum game. Independent of algorithm setting, there is a prior independent lower-bounding technique from analysis of the adversary’s play, with analogy to *information design* and with solutions as a special case of *non-unique tensor decomposition*.
3. **Scalability of Inputs.** Any distribution can be re-scaled by composing its CDF with $f_k(x) = x/k$. When inputs come from the unbounded set of (positive) reals, algorithms must handle all scales; versus, algorithms can take advantage of inputs having bounded support.
4. **Econometric Inference in Auctions.** Loosely related to robustness, there is a question of *econometric inference*, i.e., identifying auction inputs from its rules and its outcomes. *Counterfactual estimation* considers the question offline, and *dashboard mechanisms* online.

¹Specifically, the resulting game is a two-player *zero-sum* game, a class which has a number of nice properties making these games relatively easy to analyze.

Below is a summary of my past work on these topics and next objectives for future work. I am a co-author on **bolded** paper references.

Benchmark Design

A formal approach to benchmark design will seek out the existence of guiding principles to inform and measure the design of the benchmark, motivated by the heuristic narrative that intuitively we must guard both against designing benchmarks that are too small or too big. A common benchmark across both mechanism design and online algorithms is to use the *offline optimal* performance.

Hartline and Roughgarden (2014) propose a property of benchmarks called *normalization*, which ties prior free benchmark design back to Bayesian settings in a way that benchmarks are justifiably not too small. Hartline (2017) proposes a second property *resolution* for measurement of normalized benchmarks to describe intuitively: how big they are – versus – how big they need to be to simply meet the definition of normalization. Of particular relevance, resolution is a measurement that can be subjected to *optimization*.

As the main result of **Hartline et al. (2020a)** – for a general algorithms setting and not restricted to auctions – we connect the prior independent and prior free settings together. Technically, consider a *benchmark design problem* (BDP) as the following min – max operation: for use in a prior free setting, choose the “optimal” benchmark as the normalized one for which there exists a mechanism achieving smallest worst-case resolution, with resolution measured pointwise on vectors of inputs. For BDP, the solution is “the same” as the solution for the question of the optimal prior independent algorithm – the optimal benchmark is set by scaling up the performance of the optimal prior independent algorithm to be big enough to meet the definition of normalization.

Future work. Philosophically, the BDP embeds an “accommodative” design choice, by choosing a benchmark within the context of existence of a good algorithm for approximation – benchmark and algorithm are optimized together. By contrast, a disinterested benchmark designer might do optimization in a vacuum, and leave algorithm designers to optimize responses against the resulting benchmark in a second, independent step. Define an *information benchmark design problem* (IBDP) as follows: choose the benchmark minimizing worst-case resolution, measured over distributions, in expectation over inputs drawn i.i.d. from a distribution. The term “information” in the problem name comes from a connection to information design, following from a nice symmetry inherent in its constraints. Further, constraints of the IBDP optimization are weaker than for BDP, so the value of the IBDP program is a lower bound on optimal approximation in the prior independent setting. The IBDP would be another nice benchmarks problem to solve, along with a characterization of when the IBDP and BDP programs have the same value. As discussed further below, the tools involved would find extended application in related topics.

A second, orthogonal direction for next research on benchmark fundamentals is the following. The combination of the normalization and resolution properties is elegant, but has known drawbacks. For example, there can be inputs where the optimal benchmark is set to be larger than the offline optimal performance, effectively an unachievable target.² The optimal benchmark is increased (above offline optimal) on inputs where it is not tight to resolution, simply in order to help meet the requirement of normalization. The question is to identify what pertinent new property on benchmarks can address this most effectively, and what new scientific benchmarks result.

²For a single-item revenue auction with 2 agents, consider setting the benchmark greater than 1 when its value-profile input is (1, 1).

Adversarial Play

Lower bounds on approximation come from fixed adversary strategies. With an *optimal* such adversarial strategy, we can measure the tight approximation factor of a prior-independent-optimal auction. I developed the following novel technique to prove lower bounds for an arbitrary algorithm question, in the prior independent setting. Assume the unknown prior F is an element of class \mathcal{F} . An adversary strategy in the prior independent setting is the choice of a distribution $\delta(\mathcal{F})$; draw $F \sim \delta(\mathcal{F})$, then n inputs are drawn i.i.d. from F . The adversary sets a prior-independent benchmark as the expectation of optimal algorithm performances – each knowing its local distribution when realized – over its distribution, denoted $\text{OPT}_{\delta(\mathcal{F})}$.

The technique is to fix a distribution over adversary strategies and show there must exist a gap between $\text{OPT}_{\delta(\mathcal{F})}$ and any one fixed algorithm as can be chosen by the designer. A distribution over symmetric product distributions results in a symmetric correlated distribution, called a *blend*. Assume in fact there exist two distinct adversary mixed strategies (as blends) to describe the same correlated distribution, namely δ_1 and δ_2 . I.e., the same correlated distribution is described by n inputs drawn either i.i.d. as $v_i \sim (F_1 \sim \delta_1(\mathcal{F}))$, or drawn i.i.d. as $v_i \sim (F_2 \sim \delta_2(\mathcal{F}))$.

Intuitively, here is what happens. The adversary picks δ_1 as one side of the description of the correlated distribution, thereby setting the benchmark in the prior independent setting to $\text{OPT}_{\delta_1(\mathcal{F})}$. The designer knows exactly what the adversary has chosen, and sees the resulting correlated distribution. However, the designer can not help but believe that the correlated distribution was generated by its second description δ_2 , with the performance of *any* algorithm in fact upper bounded by $\text{OPT}_{\delta_2(\mathcal{F})}$. The ratio of $\text{OPT}_{\delta_1(\mathcal{F})}/\text{OPT}_{\delta_2(\mathcal{F})}$ is a necessary lower bound on approximation of the original prior independent problem.

The technique is information design of a different flavor than is standard. The adversary fully reveals the choice of δ_1 but the designer is powerless to use it directly. Instead the correlated distribution is like a continuous-dimension tensor, symmetric, with factors as probability distributions and weights on factors also as a probability distribution.³ In this stylized tensor setting, it is the non-uniqueness of decomposition of the tensor that leads to necessary gaps in performance between any algorithm and the benchmark.

Future work. There are many research directions following from blends, such that I cover them only at a high level. The blends technique itself ties into the dual program description of the IBDP problem from the Benchmarks section. A first burning question within prior independent analysis is to determine for what classes of algorithms questions does there exist a lower bound from blends that is tight to the optimal prior independent approximation. The blends technique applies directly for $n = 2$ inputs, but for more inputs the technique only extends indirectly. Individual algorithms settings have the potential to realize first-ever or improved lower bounds on approximation from the blends technique, after which characterizations of classes of algorithms can be studied. The

³I give a quick illustration of such a tensor. For greatest simplicity, I generalize to let the total weight on factors be infinite, but this can be corrected. A continuous tensor is defined on $\mathbf{v} = (v_1, v_2) \in (0, \infty)^2$ by $g(\mathbf{v}) = 1/(\max_i v_i^2)$. I give two decompositions of g . (1) Paretos: for $z \in (0, \infty)$, a factor (in one dimension, and then symmetric) is given by $f_{o,z}(x) = z/x^2$ on $[z, \infty)$ with weight $o_z = 2/z \cdot dz$. (2) Uniforms: for $z \in (0, \infty)$, a factor (same usage) is given by $f_{\omega,z}(x) = 1/z$ on $(0, z]$ with weight $\omega_z = 2/z \cdot dz$. To confirm decompositions, add up all of the “density” at an index $(v_1, v_2 \leq v_1)$ (ordering without loss of generality but needed to determine integral end points). We get:

$$\int_0^{v_2} f_{o,z}(v_1) \cdot f_{o,z}(v_2) \cdot o_z = \int_0^{v_2} \frac{z}{v_1^2} \cdot \frac{z}{v_2^2} \cdot \left(\frac{2}{z} dz\right) = \frac{1}{v_1^2} = g(\mathbf{v}) = \int_{v_1}^{\infty} \frac{1}{z} \cdot \frac{1}{z} \cdot \left(\frac{2}{z} dz\right) = \int_{v_1}^{\infty} f_{\omega,z}(v_1) \cdot f_{\omega,z}(v_2) \cdot \omega_z$$

technique can be framed as a special case of the standard information design model, and connections to existing work on information design can be strengthened. Finally, it will give rise to tactical questions in the study of non-uniqueness of tensor decomposition, for example, there is a nice class of blends solutions whereby its elements are characterized as having their two sides of blends generated each by one of a pair of distributions which are “inverse-distributions” to each other.

Scale-Invariant Design (or lack thereof)

As a mechanism-design-specific result in **Hartline et al. (2020a)**, we solve a simple variant of the longstanding open question of the optimal prior-independent truthful auction for 2 agents with unbounded value support, a revenue objective, and all regular distributions as the adversary’s choice set. Even with the restriction, our result improves on recent work by Fu et al. (2015), and by Allouah and Besbes (2018) who gave first results for the variant setting.

Specifically, the variant adds a restriction to only consider *scale-invariant* mechanisms, i.e., those with performance that is linear in the input vector (of values). Removing the restriction to scale-invariance is harder than it sounds, and remains as a tantalizing, fundamental open question. Remarkably, the description of the optimal prior independent auction is not the result of some complicated equation, as might be feared. It is a simple mix over the second price auction, and an auction posting approximately $2.447\times$ the second price to the largest-valued agent.

Future work. For both prior independent and prior free analysis, an argument is sorely needed that it is sufficient to restrict attention to scale-invariant algorithms; and further, the argument should hopefully exist from first principles, and apply very broadly to algorithms questions with ideally no assumptions on the setting. Our auction solution is specifically for 2 agents. First attempts to extend to $n > 2$ agents show that further technical tools will be needed for a solution.

Scale-invariance of the optimal mechanism is likely to be a sufficient assumption when values are unbounded, and becomes possibly its most prevalent feature. Two follow up questions in the same prior independent space become quite interesting. First, note that within the context of scale-invariant auctions, the optimal auction described above only posts two prices to the largest agent, conditioned on the second price: a price “marked up” by factors of 1 or ~ 2.447 . At this point, the auction is taking quite specific actions with small total support. It makes sense to additionally consider the robustness theme of *minimizing regret* with respect to these markup factors. Related, note that the standard question of prior independence allows the benchmark performance to be set by an optimal auction that knows the distribution. In the case that we solved in **Hartline et al. (2020a)**, the optimal auction knows the scale of the distribution. The first follow up perspective considers changing the benchmark to only be set by the optimal performance of a respectively scale-invariant mechanism. The distribution’s scale is meaningless and only its shape matters.

The second perspective considers an assumption of hard bounds on the support of value space, for example $[1, h]^n$. In this case, the idea of scale-invariant design has no relevance, as inputs observed to be close to the boundaries will have material, scale-dependent effects on optimal design. An immediate observation is that the structure of the optimal auction should be different than in the scale-invariant setting (above) – obviously it will not commit *a priori* to sometimes marking up the second price by $\sim 2.447\times$ when the second price could be in $[h/2, h]$. Of note, lower bounds on approximation from the blends technique of the Adversarial Play section might be directly relevant.

Applications of Inference in Auctions

I summarize another direction of my research in auction theory. The predominant setting is prior free: agents’ values are fixed constants, unknown by the designer. The theme is *econometric inference*, including results for both “offline” estimation and “online” implementation. Generally implicit in the setting is a repeated auction: observations of historical agent behavior inform decisions by the auction designer in the future.

Hartline et al. (2020b) addresses a first question of offline econometric inference – we want to identify the agents’ private values. The setting is truthful single-item auctions with proportional weights allocation rules, and charging winner-pays prices to (stochastic) winners. We take the perspective of an *outsider* to the auction who can not observe the agents’ reports, but rather can observe the payments made by winners. Our two main results are that: (1) the payment function is one-to-one and therefore it can be theoretically inverted; and (2), we give an algorithm to compute the inversion arbitrarily to within ϵ .⁴ As part of the supporting proofs, we give new results that are of independent interest. In particular, we extend uniqueness of pure Nash equilibrium in “concave games” of Rosen (1965), using a related definition of concavity based on Gale and Nikaido (1965).

Hartline et al. (2019) explores the potential use of a *bidding dashboard* by a designer in a repeated *non-truthful* auction setting to facilitate agents’ understanding and computation of optimal play. For simplicity, assume a benevolent designer who wants to maximize the total surplus of the auction. Having assumed a non-truthful setting, the VCG-mechanism is not available.

Two key complications are introduced with non-truthful settings. First, equilibrium is likely to have a complicated description of agents’ jointly-optimal strategies, which implies that real world agents are unlikely to be able to compute and execute the bids of theoretical equilibrium, regardless of other considerations. Second, there might be multiple equilibria, with the *price of anarchy* ratio (Koutsoupias and Papadimitriou, 1999) between the offline-optimal welfare and welfare of the worst equilibrium being possibly quite large. So even if agents are able to self-direct themselves towards an equilibrium, it might have decidedly poor welfare performance.

Bidding dashboards are estimated *price-allocation* curves published by the platform in advance of an auction, with one individually-tailored curve given privately to each agent. Their use is considered in a many-round analysis, for which our main result is to show that a *sequential-non-truthful* auction gets almost exactly the same performance as a comparative *sequential-truthful* auction, in the long-term.

Future work. **Hartline et al. (2020b)** solved price inversion under a number of restrictive assumptions (single-item, proportional weights allocation, fixed agent values), each of which can potentially be generalized. **Hartline et al. (2019)** describes a mechanism that could potentially run in the real world; its theoretical analysis would benefit from empirical study of its implementation – even initially with offline simulation – to help understand its inherent convergence properties and parameter tuning. The work related to Rosen’s concave games connects to the emerging field of “learning in games as a computational tool” (e.g., Daskalakis and Panageas (2018), Abernethy et al. (2020)) and some of its related pursuits like fast convergence (Syrkanis et al., 2015) and the possibility of convergence of learning in games to smaller classes of equilibrium than is guaranteed by the existing literature (Blum et al., 2008).

⁴The algorithm has running time that is polynomial in the number of agents n , the upper bound of the support of value space h , $\ln 1/\epsilon$, and the number of calls to the agents’ weights functions (embedding dependence on the complexity of computing them).

References

- Abernethy, J., Lai, K., and Wibisono, A. (2020). Last-iterate convergence rates for min-max optimization. *Working Paper*. arxiv.org/abs/1906.02027.
- Allouah, A. and Besbes, O. (2018). Prior-independent optimal auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 503–503. ACM.
- Blum, A., Hajiaghayi, M., Ligett, K., and Roth, A. (2008). Regret minimization and the price of total anarchy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC 08, page 373382, New York, NY, USA. Association for Computing Machinery.
- Daskalakis, C. and Panageas, I. (2018). The limit points of (optimistic) gradient descent in min-max optimization. In *Advances in Neural Information Processing Systems 31*, pages 9236–9246. Curran Associates, Inc.
- Fu, H., Immorlica, N., Lucier, B., and Strack, P. (2015). Randomization beats second price as a prior-independent auction. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC 15, page 323, New York, NY, USA. Association for Computing Machinery.
- Gale, D. and Nikaido, H. (1965). The Jacobian matrix and global univalence of mappings. *Mathematische Annalen*, 159(2):81–93.
- Hartline, J. (2011-2017). Mechanism design and approximation. *Working manuscript*. jasonhartline.com/MDnA/.
- Hartline, J., Johnsen, A., and Li, Y. (2020a). Benchmark design and prior-independent optimization. *Working paper*. arxiv.org/abs/2001.10157.
- Hartline, J., Johnsen, A., Nekipelov, D., and Wang, Z. (2020b). Inference from auction prices. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms*, SODA ’20. SIAM.
- Hartline, J., Johnsen, A., Nekipelov, D., and Zoeter, O. (2019). Dashboard mechanisms for online marketplaces. In *Proceedings of the 2019 ACM 20th Conference on Economics and Computation*, ACM EC ’19. Association for Computing Machinery.
- Hartline, J. D. and Roughgarden, T. (2014). Optimal platform design. *CoRR*, abs/1412.8518.
- Koutsoupias, E. and Papadimitriou, C. (1999). Worst-case equilibria. In Meinel, C. and Tison, S., editors, *STACS 99*, pages 404–413, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.
- Rosen, J. B. (1965). Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534.
- Syrkkanis, V., Agarwal, A., Luo, H., and Schapire, R. E. (2015). Fast convergence of regularized learning in games. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS15, page 29892997, Cambridge, MA, USA. MIT Press.