

Risk Aversion and Auction Design: Theoretical and Empirical Evidence

Shoshana Vasserman* Mitchell Watt[†]

May 28, 2021

Abstract

Auctions are inherently risky: bidders face uncertainty about their prospects of winning and payments, while sellers are unsure about revenue and chances of a successful sale. Auction rules influence the allocation of risk among agents and the behavior of risk-averse bidders, leading to a breakdown of payoff and revenue equivalence and a heightened significance of auction design decisions by sellers. In this paper, we review the literature on risk aversion in auctions, with an emphasis on what can be learned about auction design from theoretical modeling and empirical studies. We survey theoretical results relating to the behavior of risk-averse agents in auctions, the comparison of standard auction formats in the presence of risk aversion and implications for auction design. We discuss standard and more recent approaches to identifying risk preferences in empirical studies and evidence for the significance of risk aversion in auction applications. Finally, we identify areas where existing evidence is relatively scant and ask what questions empirical research might ask given the theory and where further theoretical research may be beneficial given existing empirical results.

*Stanford Graduate School of Business, SIEPR and NBER. Email: svass@stanford.edu

[†]Stanford University. Email: mwatt@stanford.edu.

We are grateful to Zi Yang Kang, Paul Milgrom, an anonymous reviewer and our editor, Leslie Marx, for their helpful comments and support.

1 Introduction

Auctions guide decisions in many economically significant domains, including government procurement, spectrum allocation, electricity generation, online advertising and financial markets. Market designers choose rules for the auction to promote their objectives. In the classic risk-neutral independent private values model, most of these decisions are moot: the revenue and payoff equivalence theorems imply that the only real decision to be made by an auctioneer is the choice of a reserve price. In the optimal auction, characterized by [Myerson \(1981\)](#), there is no role for entry fees, randomized payment or allocation rules, buyout prices, subsidies to losing bidders and a variety of other non-standard auction designs that are observed in practice .

In order to reconcile the use of these non-standard auction design tools, the literature has noted that bidding behavior is often inconsistent with the predictions of the classic model.¹ In this case, the equivalence of different auction designs may no longer hold. Risk aversion—a central example of a force that extends beyond the classic model—has been identified as a relevant factor in economic decision-making in many choice domains ([Holt & Laury, 2014](#)). In auctions, risk-averse agents respond to the uncertainty inherent in the allocation and payment rules, leading to consequential differences in behavior compared to their risk-neutral counterparts. This generates a richer role for auction designers, who may gain by accounting for—and potentially exploiting—this risk-avoiding behavior.

In this survey, we discuss the theoretical literature that analyzes and guides the design of auctions with risk-averse agents, as well as the empirical literature that tests for evidence of risk aversion and estimates the primitives governing risk-averse bidding behavior. A relevant question for empirical study is the extent to which risk aversion substantively affects behavior in real-world auctions. We begin in [Section 2](#) by summarizing different types of empirical evidence that have been offered to argue that risk aversion is a relevant factor in auction settings. Drawing on examples of auctions conducted by governments, on online platforms and in the lab, these studies show that observed bids do not conform to the predictions of a risk-neutral model and are more consistent with some model of risk aversion.

We begin [Section 3](#) by describing a benchmark model of risk-averse bidding in standard auctions. With this, we both fix notation and discuss the classic result, formalized² by [Holt \(1980\)](#), that the revenue of a first-price auction with risk-averse bidders exceeds that of a second-price auction. We survey several papers showing that this ranking is robust to a variety of additional modeling considerations, including endogenous entry and heterogeneity of risk preferences, although not all. We describe revenue-maximizing auctions for risk-averse bidders ([Matthews 1983](#), [Maskin and Riley 1984](#)), which are strikingly different from classic auction formats, including the risk-neutral optimal auction ([Myerson, 1981](#)). Because this mechanism is not observed in practice, we discuss implications of risk aversion for the design of classic auctions, including the choice of reserve prices and entry fees, and a variety of non-standard auction designs that may be rationalized by risk aversion. Finally, we describe a number of further extensions to the model which

¹See [Section 2](#) for evidence.

²The result was foreshadowed in the first game-theoretic analysis of auctions by [Vickrey \(1961\)](#), who noted that his proof of the equivalence of the first- and second-price auctions would not hold in the presence of risk aversion.

may be relevant, including interdependent values, *ex post* uncertainty, related behavioral biases and approximate mechanism design.

In Section 4, we discuss empirical studies of risk aversion in auctions, beginning with the econometric strategies that have been developed to identify the primitives governing bidding behavior observed in risk-averse bidders. Many design decisions, such as the choice of a reserve price or entry fee/subsidy, require knowledge of the distribution of bidders' values and the extent of their risk aversion. A classic non-identification result by [Guerre, Perrigne, and Vuong \(2009\)](#) shows that—in contrast to risk-neutral settings—empirical bid distributions alone are generally insufficient for identifying both value distributions and utility functions of risk-averse bidders. Successful identification strategies thus exploit additional variation from parallel auction formats, exogenous participation, endogenous entry decisions and related strategic actions made by auction participants, such as the portfolio allocation decisions inherent in scaling auctions. We also discuss what has been learned about auction design from the application of these identification strategies to real-world data.

We do not emphasize in this survey the substantial experimental literature on risk aversion in auctions. Much of this literature, including many papers in the excellent surveys of [Kagel \(1997\)](#) and [Kagel and Levin \(2016\)](#), has exploited laboratory studies of auctions to illuminate the risk preferences of human decision makers, rather than for the purpose of informing auction design. Because we are more interested in the implications of risk aversion for the design of auctions, we only mention the key experimental studies that have guided the development of the theoretical and empirical literature on auctions.

We conclude by summarizing the key lessons learned from the theoretical and empirical literature on risk aversion in auctions and discuss what we see as key open questions in the field. Table 1 presents this summary in condensed form. Because many theoretical results in the risk-averse setting depend on the extent of risk aversion and relationships between risk aversion and value distributions, we believe there could be benefits from further empirical studies to understand when and why risk aversion is relevant to auction design. Empirical analysis could also direct theoretical research towards areas where better modeling would be valuable, such as the question of whether simple mechanisms that require less knowledge of the bidder type distribution may be able to approximate optimal results.

2 Empirical Evidence of Risk Aversion in Auctions

Empirical interest in risk-averse bidding has typically stemmed from the observation that bids submitted by real people do not conform to the predictions of a model of risk-neutral agents. Bids submitted in university lab experiments conducted by [Cox, Smith, and Walker \(1985\)](#) were consistently higher than would be expected from risk-neutral bidders with the same value distribution. Building on a literature of similar experiments and using tests of statistical equivalence to model predictions, [Cox, Smith, and Walker \(1985\)](#) conclude that a model of risk-averse bidders with heterogeneous concave utility functions rationalized the bids observed better than alternatives. Examining data from scaling auctions hosted by the US Forest Service, [Athey and Levin \(2001\)](#) note

that the absence of “extreme skewing” in bidding for varieties of timber is suggestive of portfolio diversification by risk-averse bidders. Under the scaling auction format, bidders submit a unit price per unit of each timber variety in a given tract of land. The winner is determined based on the sum of unit bids multiplied by the auctioneer’s estimate of how many of each variety the tract has, but payments are based on how many units are ultimately found and harvested. As such, bidders would maximize expected profits by placing the bulk of their bids on the variety they predict to be most over-estimated by the Forest Service, and this strategy would not harm their chances of beating competitors. [Athey and Levin \(2001\)](#) show that the bidders indeed appeared to predict which varieties were over-estimated, but they spread their bids across all of the varieties. [Bolotnyy and Vasserman \(2021\)](#) make a similar observation in the context of scaling auctions for bridge construction procurement in Massachusetts, adding that bidders’ portfolios place lower stakes on units with higher quantity variance, which further suggests risk-averse portfolio balancing.

Motivated by “reduced form” or observational arguments, the empirical literature has largely focused on model fitting as a key source of evidence for risk-averse bidding. [Bajari and Hortaçsu \(2005\)](#) use data from a university lab experiment to estimate four competing independent private values auction models: a risk-neutral model, a homogeneous constant relative risk aversion (CRRA) model, an adaptive learning model, and a quantal response model. Comparing the goodness of fit, [Bajari and Hortaçsu \(2005\)](#) argue that the CRRA model best rationalizes the bids observed in the data. [Lu and Perrigne \(2008\)](#), [Campo, Guerre, Perrigne, and Vuong \(2011\)](#) and [Campo \(2012\)](#) estimate both constant absolute risk aversion (CARA) and CRRA versions of their models³ and find that risk neutrality is statistically rejected. While [Lu and Perrigne \(2008\)](#) find that CRRA provides a better fit, [Campo et al. \(2011\)](#) weakly prefers the CARA estimate. Applying an endogenous entry model in the same setting, [T. Li, Lu, and Zhao \(2015\)](#) find evidence for decreasing absolute risk aversion preferences⁴. [Campo \(2012\)](#) projects each firm’s estimated risk aversion parameter on its number of years of operation and rejects CARA on the grounds that the fitted model predicts that risk aversion increases with experience, whereas the CRRA model does not. However, not all papers reject risk-neutral bidding. [Grundl and Zhu \(2019\)](#) replicate the exercise from [Bajari and Hortaçsu \(2005\)](#) allowing for unobservable heterogeneity in bidder values and find that the estimated CRRA parameter can no longer be statistically differentiated from zero.

A parallel set of observations argues that many auction design choices made in practice make more sense with risk-averse bidders than with risk-neutral ones. [Ackerberg, Hirano, and Shahriar \(2006\)](#) note that buyout options in auctions such as those hosted by eBay would have no value to the auctioneer if bidders were risk-neutral. [Lu and Perrigne \(2008\)](#) suggest that ascending price auctions in US Timber auctions are used more often than sealed-bid auctions for higher value parcels. Building on this, [Kong \(2020\)](#) finds that the sealed-bid first-price auction format generates 30% more revenue than an equivalent English auction format in auctions for oil and gas basins in New Mexico. Using evidence that bidders in these auctions are uncertain about the number of competitors that they will face, [Kong \(2020\)](#) argues that the superiority of the sealed-bid format

³[Lu and Perrigne \(2008\)](#) and [Campo et al. \(2011\)](#) study U.S. Forest service timber sales, while [Campo \(2012\)](#) studies construction procurement contracts.

⁴Note that CRRA implies decreasing absolute risk aversion, so that this result is not in conflict with the findings of [Lu and Perrigne \(2008\)](#).

is consistent with theoretical predictions of a model with risk-averse bidders. Luo and Takahashi (2019) find that procurement contracts auctioned in a lump-sum format by the Florida Department of Transportation (FDOT) have lower unobserved bidder costs than contracts auctioned using a scaling auction format. Interpreting unobserved costs as risk, Luo and Takahashi (2019) argue that this is consistent with FDOT project managers preferring the scaling format to insure bidders against *ex post* risk in riskier auctions.

Taken together, these studies suggest both that risk aversion is a pertinent factor in many auction settings and that a careful model—that captures the relevant institutional features of each setting—is needed in order to accurately assess the extent of risk aversion and its implications for policy making. In the next section, we survey the different modeling assumptions that have been considered in the literature on risk-averse bidding, highlighting their implications for optimal design. We then survey the literature on empirical identification and estimation, showing how modeling assumptions are used to inform identifying restrictions and counterfactuals of interest.

3 Theoretical literature

3.1 Benchmark model

We introduce a benchmark model of risk aversion in auctions to fix notation for the remainder of this article.

There is a single, indivisible object to be (potentially) sold to one of N interested buyers via an auction.⁵ Each bidder has an identical, twice-differentiable utility function $u : \mathbb{R} \rightarrow \mathbb{R}$ over wealth⁶ that satisfies $u(0) = 0$, $u' > 0$ and $u'' < 0$. Risk aversion is determined by the concavity of the utility function, with the Arrow-Pratt coefficient of absolute risk aversion (ARA) defined as $A(x) = -u''(x)/u'(x)$ and the coefficient of relative risk aversion (RRA) as $R(x) = -xu''(x)/u'(x)$. These measures of risk aversion may be constant with respect to wealth (abbreviated CARA or CRRA, respectively) or exhibit wealth effects: increasing (IARA, IRRA) or, more commonly assumed, decreasing (DARA, DRRA).

Each bidder observes a private signal s_i which is informative of the bidder’s value for the object, and there may also be other value-relevant information contained in a vector s_0 . The bidder’s value is then $v_i = v(s_i, s_{-i}, s_0)$, where v is typically assumed to be increasing in s_i and symmetric and non-decreasing in s_{-i} , the vector of other bidders’ signals. Much of the literature studying risk aversion focuses on the independent private values (IPV) model, where $v_i = s_i$ is drawn independently and identically from some common knowledge distribution F , with support $[\underline{v}, \bar{v}]$. Another well-studied model is the affiliated values model of Milgrom and Weber (1982), where the joint distribution of (s, s_0) is symmetric, affiliated⁷ and has bounded support. Throughout this

⁵In this survey, we mostly describe auctions for the sale of objects, hence we treat “buyers” and “bidders” synonymously and similarly “sellers” and “auctioneers”. Of course, most results hold symmetrically to the case of the procurement auction, in which case roles are reversed.

⁶Implicit in this formulation is the assumption that bidders have an equivalent monetary value for the item being sold, which is not necessarily without loss of generality. Maskin and Riley (1984), Che and Gale (2006), Baisa (2017) and others consider more general preference specifications in a setting with risk aversion, but we will restrict attention to the equivalent monetary value setting in this survey.

⁷A vector-valued random variable X with density $f(\cdot)$ is affiliated if, for all x, y in the support of f , $f(x \wedge y)f(x \vee y) \geq$

survey, results apply to the independent private values setting of this benchmark model, unless otherwise stated.

The auctioneer chooses a mechanism for allocating the object and determining payments to or (more usually) from the bidders. The allocation and payment rules of the mechanism may be random. Bidders are assumed to behave according to the symmetric Bayes-Nash equilibrium of the game induced by this mechanism, which allows us, by the revelation principle, to restrict attention to Bayesian incentive-compatible direct mechanisms. A *standard auction* (Waehrer, Harstad, & Rothkopf, 1998) treats all bidders symmetrically, assigns the asset to the bidder with the highest signal as long as that signal is above some screening level (in equilibrium), and charges the winners a nonnegative payment and losing bidders zero. The first-price auction (FPA), second-price auctions (SPA) and ascending auctions—with and without reserves—fall into the class of standard auctions.

3.2 Ranking classic auctions

In the IPV setting, Holt (1980) first formalized the conjecture of Vickrey (1961) that the FPA has a higher expected revenue than the SPA. The logic for this is simple. In the SPA with IPV, it is a dominant strategy—regardless of risk preferences—for a bidder to bid truthfully, and so the expected revenue of the auction is the same with risk-neutral and with risk-averse agents. In the FPA, risk-neutral agents bid up to the point where the marginal cost of increased payment conditional on winning exactly balances the benefit of increased probability of winning the auction. Risk-averse agents are willing to pay for actuarially-fair insurance and thus bid more aggressively. To see this, note that any increasing and differentiable equilibrium bidding strategy $\beta(\cdot)$ must satisfy $\beta(v) \in \arg \max_b F^{N-1}(\beta^{-1}(b))u(v-b)$, which implies the differential equation,

$$\beta'(v) = \frac{u(v-\beta(v))}{u'(v-\beta(v))} \frac{(N-1)f(v)}{F(v)}, \quad (1)$$

which characterizes the FPA equilibrium under risk aversion. Since $\frac{u(v-\beta(v))}{u'(v-\beta(v))} > v-\beta(v)$ for strictly concave u , and the risk-neutral FPA bidding strategy $\tilde{\beta}(v)$ satisfies $\tilde{\beta}'(v) = (v-\tilde{\beta}(v))\frac{(N-1)f(v)}{F(v)}$, we must have that $\beta'(v) > \tilde{\beta}'(v)$. Because $\beta(v) = v = \tilde{\beta}(v)$, the previous inequality implies that $\beta(v) > \tilde{\beta}(v)$, so that bidding under the FPA is more aggressive in the risk-averse setting. Since the two auction formats are equivalent under risk neutrality, this implies that the FPA earns strictly more revenue than the SPA with risk-averse bidders. This result is easily extended to the case of reserve prices or entry fees, using the ranking lemma of Milgrom and Weber (1982).

Although the expected payments are greater in the FPA, the payment rule is a riskier random variable in the SPA. Consequently, from the point of view of a risk-averse bidder, the ranking of these auctions is not necessarily the reverse of the risk-neutral seller's ranking. Matthews (1987) shows that for a CARA buyer, the increase in the expected payment in the FPA is exactly balanced by the reduced riskiness of the payment rule, so that a CARA buyer is indifferent between partic-

$f(x)f(y)$. Affiliation is a strong form of correlation: bidders who observe a high signal believe that other bidders will also receive high signals, so competition in the auction is likely to be strong.

icipating in a FPA and a SPA with the same reserve.⁸ When bidders have non-constant ARA, the ranking of the two auction formats depends on how their risk preferences change with wealth: DARA bidders prefer the SPA to the FPA, while IARA bidders have the reverse ranking.

3.3 Optimal auctions with risk-averse bidders

Matthews (1983) and Maskin and Riley (1984) both characterize the optimal design of auctions by a risk-neutral seller in the IPV setting with risk-averse bidders. Matthews (1983) focuses on CARA utility functions while Maskin and Riley (1984) consider a more general class of preferences at the expense of reduced specificity of the resulting mechanism's design. The key observation of both papers is that a seller in the presence of risk-averse buyers can, by careful design, influence the allocation of risk among participants in a way that may be profitable. By shielding high-value bidders from risk and exposing low-value bidders to risk, the seller can reduce the information rents earned by high-value bidders and extract more of the expected surplus. Pareto-dominating schemes, where agents are fully insured against risk, result in lower expected revenues for the seller: in fact, Maskin and Riley (1984) show that "perfect insurance" auctions, designed so that the marginal utilities conditional on winning or losing are equal, obtain the same expected revenue as the SPA.

Both papers focus on auctions where bidders may be charged a type-specific entry fee. For tractability of the analysis, the authors consider symmetric, reduced-form allocation rules, where a bidder's *interim* allocation probability depends only on that bidder's type and not those of other bidders. Two technical challenges emerge which do not appear in the risk-neutral analysis of Myerson (1981). First, the reduced-form allocation rule must be chosen so as to ensure that the total probability of assignment of the good is no larger than one. This adds an additional "implementability constraint" to the auction design problem that amounts in the single-object case to the requirement that no subset of the type space be allocated the object with a probability greater than the probability that a bidder from that subset appears in the auction (Border, 1991). Second, the optimal auction may involve random payment rules. Maskin and Riley (1984) determine conditions under which deterministic payment rules are optimal, which includes the case of CARA utility functions.

The revenue-maximizing auction for the seller solves an optimal control problem which generally does not possess a closed-form solution, although Maskin and Riley (1984) derive its properties. In general, the optimal auction fully insures the highest-type bidder and exposes lower types to more risk. As a consequence, buyers strictly prefer to win than lose, even though high-type buyers may receive a subsidy if they do not win the item. The probability of winning and the expected payments of the bidders are typically increasing in the bidder's value, while the lowest types may not enter the auction at all.

In the case of CARA utility, Matthews (1983) obtains a more detailed characterization of the optimal auction. Matthews (1983) first analyzes the case in which the seller has many identical objects available, but buyers have unit demand. The optimal auction then consists of an increasing

⁸This result can be extended more generally: for any auctions with the same allocation rule and where all losers pay zero, CARA bidders obtain the same expected utility (Milgrom, 2004).

payment-probability schedule, where the agent pays for some probability of being able to purchase the item at the seller’s marginal cost.⁹ When there is only a single object available to the seller, the implementability constraint may bind. The optimal mechanism remains similar, but now the seller charges high-type bidders a higher price conditional on winning in exchange for a reduced entry fee or even an entry subsidy. The price charged conditional on winning is the larger of the seller’s marginal cost and the shadow cost of the implementability constraint (when it is binding). If the resulting price rule is non-monotonic, a procedure similar to ironing is applied.¹⁰ These optimal mechanisms approach perfect surplus extraction as the ARA approaches infinity and the Myerson optimal auction as bidders become more risk-neutral.

3.4 Implications of risk aversion for the design of classic auctions

Because the optimal auction in the risk-averse setting is not observed in practice, several papers consider how risk aversion might inform the design of classic auction formats: the FPA and the SPA.

Reserves Hu, Matthews, and Zou (2010) studies the choice of reserves in the FPA and the SPA in an environment with risk aversion. If bidders are risk-averse and the value distribution has the decreasing reverse hazard rate (DRH) property,¹¹ the effect of an increase in the reserve price on the strategy of the bidders in the FPA is less pronounced, the greater the risk aversion coefficient. The intuition for this is that more risk-averse bidders shade their values less and so are not as likely to be affected by the marginal increase in reserve price. Any optimal reserve of the FPA is lower than all optimal reserves of the SPA, so that the FPA is more *ex post* efficient than the SPA. The same conclusion holds when the seller is risk-averse. More risk-averse sellers set lower reserve prices in both auctions to mitigate the risk that the object is not sold, and lower reserves are also preferred for more risk-averse bidders under the DRH assumption.¹² Hu (2011) shows that the optimal reserve price in the FPA is decreasing in the number of bidders. In practice, very low or zero reserve prices, below the value of the object to the seller, are observed more often in practice than would be predicted by optimal auction design in the presence of risk-neutral bidders. Hu, Matthews, and Zou (2019) show that very low or zero reserve prices may be rationalized by a combination of interdependent valuations and sufficiently risk-averse bidders. With interdependent

⁹Baisa (2017), in a setting with non-quasilinearity and risk aversion, proposes a mechanism that directly sells “probability units” of being assigned a good. Risk-averse buyers of a normal good will have a higher willingness-to-pay for the first unit of probability than they will for the item as a whole. In his “probability demand mechanism”, buyers report a schedule of willingness-to-pay for each probability of receiving the good, and a Vickrey-like procedure determines the allocation and payments. Baisa (2017) shows that this mechanism, while non-truthful, has a larger revenue than the standard auction formats and that this revenue approaches the highest attainable revenue from an interim individually rational mechanism as the number of bidders grows large.

¹⁰More on this generalized ironing procedure (which arises in the solution to mechanism design problems with majorization constraints) may be found in Kleiner, Moldovanu, and Strack (2020).

¹¹The valuation distribution F with density f has the decreasing reverse hazard rate property if $f(v)/F(v)$ is a strictly decreasing function on the support of F .

¹²A version of this result was also derived by Riley and Samuelson (1981), who showed that if $A(x) > A'(x)$ for all x , then the optimal reserves satisfy $r \leq r'$. However, this derivation was contemporaneous with an early draft of Hu et al. (2010), despite publication dates differing by several decades.

values, decreasing the reserve (from the seller's true value for the item) may be profitable because it increases the probability that the second-highest bid will be paid instead of the reserve.

Entry fees Chakraborty (2019) extends the analysis of Hu et al. (2010) in the case of CARA utility functions to also consider the optimal choice of entry fees. Entry fees in the FPA increase the risk facing the lower types beneficially for the seller, while the SPA already has sufficient payment risk for low-type bidders that the seller does not benefit from exposing bidders to more risk. As a result, the optimal SPA has a positive reserve price but no entry fee, while the optimal FPA may charge bidders for entry.

3.5 Non-standard auction design tools

The breakdown of revenue equivalence in the case of risk-averse bidders sparked a literature investigating whether risk aversion may help rationalize some non-standard auction designs that are observed in practice. Here we discuss a few: buyout options, hidden reserves and premium auctions.

Buyout options Buyout options (also known as buy prices or buy-it-now prices) are prices advertised by an auctioneer prior to or during an auction, at which any bidder may purchase the item and terminate the auction process. Buyout options may be permanent, in which case the price is available to all bidders throughout the auction procedure, or temporary, where the option is revoked after the first bid is received. Buyout options are used on many online auction platforms.

Budish and Takeyama (2001) analyze an ascending auction with a permanent buyout option, in an IPV model with two bidders and binary types (high or low value). They show that a permanent buyout option can raise the expected revenue of the seller by partially insuring risk-averse high-value bidders. The resulting revenue may be higher than that of a FPA with the same bidders. Mathews and Katzman (2006) study a sealed-bid SPA with a temporary buyout option in an IPV setting. A buyer can choose to purchase the item at an advertised price before the auction, otherwise a standard sealed-bid SPA is conducted. In this setting, a risk-averse seller facing risk-neutral bidders will choose a temporary buyout price low enough to be exercised with some positive probability. This may comprise a Pareto-improvement over the sealed-bid SPA. Reynolds and Wooders (2009) contrast temporary and permanent buyout options in an IPV setting with risk-averse buyers and a risk-neutral seller. Both kinds of buyout options increase expected revenue over the ascending auction, with the permanent buyout option better for sellers in the presence of CARA/DARA bidders. Bidders with DARA utility prefer the temporary buyout option because it reduces risk once bidding commences, but bidders with CARA utility are indifferent between the two formats.

Hidden reserves In the risk-neutral setting, the revenue equivalence theorem implies that it is never beneficial for the seller to keep the reserve price secret, but hidden reserve prices have been observed in practice (Elyakime, Laffont, Loisel, & Vuong, 1994). When buyers are sufficiently risk-averse, H. Li and Tan (2017) show that a secret reserve price may be profitable in a FPA. In their

model, the seller's value for the good is not known by the bidders, but the distribution of possible seller valuations is common knowledge. A hidden reserve price in this setting functions like a competing bid, and risk-averse buyers bid more aggressively as a consequence. [H. Li and Tan \(2017\)](#) show that it is in the best interest for a risk-neutral seller to submit her true valuation as the reserve, but uncertainty in this reserve from the perspective of bidders generates higher profits for the seller.

Premium auctions Premium auctions refer to the observed practice of subsidizing losing bidders in auctions, often by offering the second-highest bidder some share of the auction's revenue. One explanation for this practice is due to [Goeree and Offerman \(2004\)](#) who show that a premium auction may benefit the seller in a setting where several weak bidders compete against a single strong bidder, by increasing the weak bidders' incentives to bid aggressively. [Hu, Offerman, and Zou \(2018\)](#) provide an alternative rationalization for the premium auction related to risk aversion. While risk aversion on the part of bidders alone cannot rationalize the choice of the premium auction,¹³ if *both* bidders and the auctioneer are risk-averse, subsidies to losing bidders may facilitate risk sharing between the seller and buyers. Because risk-averse bidders also prefer the premium auction, this implies that the premium auction may offer an *ex ante* Pareto improvement over the ascending auction.

3.6 Enriching the benchmark model

Risk-averse sellers [Waehrer et al. \(1998\)](#) consider the possibility that auctioneers may themselves be risk-averse and compare the FPA and the SPA with the same reserve under this assumption. In the affiliated-values setting with independent signals (which includes the IPV setting), from the point of view of the seller, the price in the SPA is a riskier random variable than that of the FPA: conditional on the highest realized type, the payment under the SPA is random with mean equal to the *certain* payment received in the symmetric equilibrium of the FPA. This implies, via Jensen's Inequality, that a risk-averse seller strictly prefers the FPA to the SPA. [Waehrer et al. \(1998\)](#) also show that risk aversion on the part of sellers may entail a reduction in the optimal reserve prices and entry fees relative to the choice of risk-neutral sellers, in order to reduce the probability that the object is not sold.

More generally, in the space of all Bayesian incentive-compatible mechanisms, [Eső and Futó \(1999\)](#) show, using the revenue equivalence theorem, that it is without loss of generality for the seller to restrict attention to mechanisms with deterministic revenue. In other words, for any incentive-compatible mechanism, there exist payment rules which implement the same allocation and with certain revenues equal to the expected revenue of the original mechanism. This implies that for auction design, the seller need only worry about risks associated with the allocation rule. A downside of this approach is that the resulting mechanism may violate *ex post* individual ratio-

¹³[Hu, Offerman, and Zou \(2011\)](#) show that the expected revenue of the premium auction is decreasing in the risk aversion parameter of the bidders. Since the expected revenue is the same with or without the premium for risk-neutral bidders, this implies that a risk-neutral seller would be worse off by offering losing bidders subsidies if bidders are risk-averse.

nality and may involve complicated transfers between bidders, so it may be difficult to implement such mechanisms in practice.

Heterogeneity in risk preferences One of the starkest assumptions of the benchmark model (and much of the literature) is that of identical preferences among bidders. A natural dimension along which bidders may differ is their attitudes to risk. [Cox, Roberson, and Smith \(1982\)](#) were the first to analyze a model permitting heterogeneity of risk preferences. One consequence of such a model is that, in equilibrium, the FPA and other standard auctions (although not the SPA) may be inefficient: the bidder with the highest bid may not be the one with the highest value, but may be bidding especially aggressively due to their attitude toward risk. [Cox, Roberson, and Smith \(1982\)](#) restrict attention to bidders with CRRA utility functions, $u_i(x) = x^{r_i}$, where the risk aversion parameter r_i is drawn from a probability distribution on $[0, 1]$. In this case, the equilibrium bidding function for a bidder with risk aversion parameter r_i is

$$\beta_i(v_i) = \underline{v} + \frac{N-1}{N-1+r_i}(v_i - \underline{v}), \quad (2)$$

which is conveniently independent of the distribution of risk parameters among the bidders. This result was extended to characterize the equilibrium of multi-unit pay-as-bid auctions in [Cox, Smith, and Walker \(1982\)](#) and certain other parametrized families of concave utility functions in [Cox, Smith, and Walker \(1988\)](#).

[Che and Gale \(2006\)](#) greatly expanded the scope of revenue comparisons of auctions in settings where bidders have heterogeneous preferences, by introducing a general methodology for characterizing expected revenue when bidder types come from an arbitrary multi-dimensional distribution. Their method involves fixing equilibrium behavior of $N-1$ bidders and constructing a fictitious risk-neutral bidder whose behavior mimics the best response in the actual equilibrium. The resulting distribution of fictitious types may be used to bound the revenue from the auction. [Che and Gale \(2006\)](#) apply their methodology to greatly generalize the revenue ranking of the FPA and the SPA to settings with arbitrary heterogeneity in the risk aversion parameters and allow for the possibility of budget constraints. They also extend these results to bidders with certain non-expected utility preferences.

Interdependent values In the risk-neutral setting with interdependent values, the seminal result of [Milgrom and Weber \(1982\)](#) is that the expected revenue of the SPA is higher than that of the FPA, a reverse of the revenue rankings for risk-averse IPV bidders. [Milgrom and Weber \(1982\)](#) observe that, for models with risk aversion and affiliated values, this implies that the revenue ranking of the FPA and the SPA is ambiguous. [Milgrom and Weber \(1982\)](#) also consider the incentives for a seller to reveal information in such a setting. Unless bidders have CARA utility, it is possible that additional information can *reduce* the willingness-to-pay for risk-averse bidders, another contrast to the risk-neutral setting. With CARA bidders participating in a SPA, they show that full reporting of information results in the highest expected revenue for the auctioneer.

From the vantage point of the buyer, [Matthews \(1987\)](#) shows that buyers with CARA and IARA

preferences still weakly prefer the FPA to the SPA in the presence of affiliated values. With DARA utility functions, the buyers' preferences over auction formats depends upon how the effect of their risk aversion decreases relatively to how strongly their types are affiliated. To the extent that DARA preferences are considered a more realistic assumption in many real-world applications, this result implies that welfare analysis in the presence of risk aversion and interdependent values requires detailed knowledge of the economic primitives.

Esó (2005) studies optimal auction design for risk-averse bidders in the presence of correlated values. Because the classic Crémer and McLean (1985) mechanisms in the risk-neutral correlated values environment typically expose bidders to substantial risk, one might hypothesize that the seller's ability to extract full surplus from the buyers may be hampered by the presence of risk aversion. In a simple binary-type example with CARA bidders, Esó (2005) shows that the optimal auction extracts all buyer surplus whenever the correlation is sufficiently strong, regardless of the buyers' risk aversion coefficients. These mechanisms still expose the seller to substantial risk and thus may not be optimal if the seller is also risk-averse. However, there has not been a general characterization of the optimal mechanism for risk-averse buyers with correlated values beyond these simple examples.

Stochastic and endogenous entry Risk aversion on the part of bidders may have additional effects in models of auctions where entry is stochastic. In models with exogenous random entry or when entry decisions are made strategically on the basis of private information, bidders face an additional source of uncertainty in the number of competing bidders in the auction, which may or may not be resolved prior to bidding. In addition, heterogeneity in risk preferences and costly entry may induce selection into the auction along the dimension of risk preference.

McAfee and McMillan (1987) study an IPV FPA with a known pool of potential bidders, but in which the actual participants of the auction are chosen by an exogenous random process. When bidders in such auctions have CARA preferences, bidders are indifferent between knowing the realized number of participants in the auction and the auctioneer concealing this information. Risk-averse agents bid as if there are more bidders than the objective expected number of bidders, so that a risk-neutral seller earns more expected revenue by concealing the number of bidders in the FPA. Matthews (1987) extends this result to show that DARA bidders in such a setting prefer the number of agents to be revealed, while IARA bidders prefer concealment. The concealment policy thus also results in greater expected revenue in the presence of DARA bidders. Matthews (1987) also consider the affiliated values setting and show that concealment cannot lower expected revenue against CARA or DARA bidders under the additional assumption that bids for all types would be increasing in the number of bidders, if this information was made common knowledge.

Smith and Levin (1996) endogenize entry in the IPV setting and compare the FPA and SPA with reserve prices and entry fees. In their model, bidders' values are drawn after entry decisions have been made, at which time the number of entrants to the auction also becomes common knowledge. Smith and Levin (1996) show that when bidders have CARA or IARA utility functions, the revenue under the FPA with endogenous entry is higher than that of the SPA with the same reserve price and entry fees. This follows from Matthews (1987) since CARA and IARA buyers (weakly) prefer

to participate in the FPA, which also earns higher expected revenue for the seller. The same revenue ranking holds when entry fees are chosen optimally (given fixed reserves). However, when bidders have DARA utility functions, the ranking for sellers *may* be reversed, but this requires the effect of decreased entry into the auction to overpower the increased expected revenue conditional on entry. [Delnoij and De Jaegher \(2020\)](#) augment this model with an earlier stage where competing sellers of identical items choose to run either a FPA or a SPA without reserves. The revenue ranking results of [Smith and Levin \(1996\)](#) imply that *all* sellers are best off selecting the FPA. On the other hand, with DARA, some sellers may prefer to select a SPA for certain value distributions.

[T. Li et al. \(2015\)](#) study a model of endogenous entry where bidders receive a private signal of their value prior to making entry decisions, but do not observe the realized number of bidders in the auction after entry decisions have been made. The distribution of valuations in this model is increasing in the signal in the sense of first-order stochastic dominance, and each bidder's value is realized after the entry decision has been made. In this model, entrance has a cut-off structure: bidders with a sufficiently high signal pay the entrance fee and participate in the auction. [T. Li et al. \(2015\)](#) examine the FPA and the ascending auction in this environment. They show that entry into both auctions is increasing in the number of potential bidders and that the revenue comparison results of [Smith and Levin \(1996\)](#) extends to this setting. Furthermore, participation in the two auction formats depends on the nature of risk preferences: CARA utility functions imply that the two auctions have the same participation rate, while DARA and IARA preferences implying higher and lower participation rates in the ascending auction, respectively.

[Pevnitskaya \(2004\)](#) combines an endogenous entry model with heterogeneity in risk aversion to study the possibility of selection along the dimension of risk preferences. Each potential bidder may either claim an outside option with known value or enter the auction to observe their value for the object. With heterogeneous risk preferences (with a risk parameter drawn from a common knowledge distribution), [Pevnitskaya \(2004\)](#) shows that there is a cut-off rule for entry into the auction where sufficiently risk-tolerant agents enter and more risk-averse agents claim the outside option. The result is that bidding in the auction is less aggressive than would be predicted in a model without an entry decision. The importance of this self-selection effect increases in the opportunity cost of entering the auction and the number of bidders, which implies that the selling price may actually fall as the number of bidders increases.

Ex post risk The benchmark model of risk aversion in auctions assumes that all uncertainty is resolved at the end of the auction, but this is not the case in many applications. [Esó and White \(2004\)](#) study *ex post* risk, which they define to be uncertainty in valuations that is not resolved at the end of the auction. Their key result is that bidders with symmetric, DARA utility functions reduce their bids by more than the increase in the risk premium when noise is added to the valuation, so that, all else equal, DARA bidders prefer auctions for riskier prospects. In the FPA, this is driven by a precautionary savings motive: DARA bidders value their current money holdings more highly, reducing their willingness to bid on the good. Together, this implies that the seller would like to reduce the pure risk faced by buyers, while buyers may prefer to collectively commit not to acquire information about payoff-relevant shocks. [Hu, Matthews, and Zou \(2018\)](#) derive

conditions under which the *ex post* equilibrium of an English auction with *ex post* risk is efficient, so that the winning bidder is the one with the highest willingness-to-pay, which includes the case where bidders have non-increasing ARA. [Skrzypacz \(2013\)](#) surveys a related literature on auctions with contingent payment rules, which may be used when the *ex post* risks are contractible.

Related behavioral biases This survey and most of the auction theory literature has focused on the expected utility approach to modeling risk preferences due to [Bernoulli \(1738\)](#) and [Morgenstern and Von Neumann \(1953\)](#). The behavioral economics literature has proposed alternative ways to capture risk aversion and related behavioral biases in the preferences of a decision maker; in this section, we discuss a few that have been applied to the analysis of auctions.

[Gershkov, Moldovanu, Strack, and Zhang \(2021\)](#) consider the design of auctions when bidders have “dual” utility functionals due to [Yaari \(1987\)](#), in which the probabilities, rather than the payoffs, are distorted by the bidder. They focus on the case of bidders with constant risk aversion, which means that preferences over lotteries are unchanged after adding a sure prospect or multiplying payoffs by a positive constant.¹⁴ In stark contrast to the results of [Matthews \(1983\)](#) and [Maskin and Riley \(1984\)](#), the optimal mechanism in a model of constant risk-averse bidders with dual utility functionals entails *full insurance* for bidders, so that the *ex post* payoff to any bidder who reports truthfully is a constant. This implies that some losing bidders must be compensated, as in the premium auction. With a single buyer, high types receive the object with probability one, while intermediate types receive the object randomly.

[Bose, Ozdenoren, and Pape \(2006\)](#) show that a full insurance auction is also the optimal standard auction when bidders are averse to ambiguity, in the form of maxmin expected utility preferences. Ambiguity-averse bidders also prefer the SPA to the FPA ([Levin & Ozdenoren, 2004](#)) but [di Tillio, Kos, and Messner \(2016\)](#) show that an ambiguous mechanism, in which certain rules may be hidden from bidders, may have higher revenue than any standard auction in such a setting.

[Lange and Ratan \(2010\)](#) study behavior in auctions with loss-averse bidders, using the reference-dependent utility model of [Kőszegi and Rabin \(2007\)](#). Bidders behave differently according to whether the reference point for the decision-maker considers both the object and wealth as separate dimensions (narrow bracketing) or if the reference point is in terms of consumer surplus (wide bracketing). In the narrow bracketing case, loss-averse bidders do not bid truthfully in the SPA, and the resulting expected revenue is lower than the FPA. [Balzer and Rosato \(2021\)](#) extend the analysis of loss-averse bidders in auctions to the case of interdependent values. [Filiz-Ozbay and Ozbay \(2007\)](#) and [Engelbrecht-Wiggans and Katok \(2007\)](#) describe how anticipation of loser regret (felt when bidders miss an opportunity to win at an affordable price) may also explain more aggressive bidding in FPAs than predicted in the risk-neutral model.

Overbidding in FPAs may also be rationalized by non-standard equilibrium concepts. [Goeree, Holt, and Palfrey \(2002\)](#) show that overbidding arises in the quantal response equilibria of FPAs with risk-averse bidders, and [Crawford and Iriberri \(2007\)](#) show that the behavior may also be explained by level-*k* thinking, a form of bounded rationality.

¹⁴In the classic setting, this implies risk neutrality, while in the dual utility formulation, this class includes certain loss-averse, disappointment-averse and mean-dispersion preferences.

Approximate mechanism design In practice, auction designers may not have access to the true distribution of values or the exact form of the bidders’ utility functions. In this setting, the seller may settle for a mechanism which is robust to these details and offers only an approximation of the optimal revenue. Several papers have studied approximately-optimal auctions when either the seller or buyers are risk-averse.

Sundararajan and Yan (2020) show that a certain posted-price mechanism offers a constant-factor approximation¹⁵ to the optimal mechanism for a risk-averse seller, which is robust to the seller’s concave utility function. In the setting where buyers are symmetric and have a certain subclass of risk-averse preferences,¹⁶ Fu et al. (2013) show that the FPA is a constant-factor approximation of the optimal mechanism when the value distribution is unknown. Feng et al. (2019) show that a price-posting mechanism is also a constant-factor approximation in that setting. Fibich and Gavious (2010) show that the revenue differences between the FPA, SPA and all k -price auctions disappear at a rate inversely proportional to the square of the number of bidders, which suggests that risk aversion might not be a primary concern to designers of very large auctions.

4 Empirical literature

4.1 Identification and Estimation

Evidence of risk-averse bidding has significant implications for econometricians and market designers alike. From an econometric perspective, estimates that stem from models that ignore risk aversion in such settings are biased and may be hard to interpret or use. From a design perspective, policy prescriptions that rely on intermediate results that do not hold under risk aversion—such as revenue equivalence—may be misleading. Addressing both concerns requires incorporating risk aversion into the approach for identification and estimation using bidding data. However, as Guerre et al. (2009) explain, empirical bid distributions alone are not generally sufficient to identify both the distribution of bidder values and the confounding risk-averse utility function.

To show this, Guerre et al. (2009) rewrite the equilibrium characterization for IPV FPAs from Equation (1) in terms of the equilibrium bidding function $\beta(v)$, the “fear of ruin” term¹⁷ $\lambda(\cdot) = U(\cdot)/U'(\cdot)$, the value distribution $F(\cdot)$ —which is allowed to depend only on the number of bidders N who participate—and the (observed) equilibrium bid of each bidder i , $b_i = \beta(v_i)$:

$$\beta'(v_i) = (N - 1) \frac{f(v_i | N)}{F(v_i | N)} \lambda(v_i - b_i) \quad \text{for all } v_i \in (\underline{v}(N), \bar{v}(N)]. \quad (3)$$

Following Guerre, Perrigne, and Vuong (2000), they use the monotonicity of the bidding function

¹⁵In the approximate mechanism design setting, a constant-factor approximation is a mechanism which, for some class of problems, guarantees a proportion $0 < \rho < 1$ of the expected payoff of the optimal mechanism on any instance of the problem from that class. For example, Sundararajan and Yan (2020) provide a mechanism which is a constant-factor approximation to the optimal revenue for any concave utility function for the seller, while Fu, Hartline, and Hoy (2013) and Feng, Hartline, and Li (2019) consider mechanisms which are robust to the value distribution of buyers.

¹⁶Namely, Fu et al. (2013) and Feng et al. (2019) assume capacitated utility, where $u(x) = \min\{x, C\}$ for some constant $C > 0$, which is a very specialized form of concave preferences.

¹⁷This terminology is due to Aumann and Kurz (1977).

to change variables, switching from the distributions of values $F(\cdot)$ to the distribution of equilibrium bids $G(\cdot)$.¹⁸ Rewriting Equation (3), they obtain an empirical *inverse-bid* function

$$v_i = b_i + \lambda^{-1} \left[\frac{1}{N-1} \frac{G(b_i | N)}{g(b_i | N)} \right] \equiv \xi(b_i, U, G, N). \quad (4)$$

Whereas under risk neutrality, an observation of the empirical distribution of bids $G(\cdot)$ would be sufficient to identify pseudo-values v from Equation (4), [Guerre et al. \(2009\)](#) note that the additional degree of freedom generated by risk aversion makes this approach insufficient. Indeed, any smooth bid distribution can be rationalized by some utility function, and so the joint structure of the bidders' value distribution and utility function $[U, F]$ cannot be identified by this alone.

Identification can be achieved using additional data or assumptions. [Perrigne and Vuong \(2019\)](#) provide a comprehensive survey of the econometric theory of auction identification strategies. Here, we discuss a few key approaches that are especially relevant to auctions with risk-averse bidders. The main challenge to identification under risk aversion is that the empirical bid distribution $G(\cdot)$ is a function of both the value distribution and the bidders' utility function. If one of those quantities is known, then the other can be backed out from Equation (4). [Lu and Perrigne \(2008\)](#) make use of this by combining data from an auction environment in which both a sealed-bid FPA and an English auction were used. As English auctions are truthful, the distribution of value quantiles can be identified from the empirical distribution of winning bids, as demonstrated by [Athey and Haile \(2002\)](#). This is true with risk-neutral bidders and risk-averse bidders alike. On the other hand, English auction data cannot identify the extent of risk aversion. To make progress, [Lu and Perrigne \(2008\)](#) assume that the distribution of bidders' values is identical across the two auction formats. That is, the same types of bidders participate in both formats. Writing $\alpha = F(v(\alpha)|N)$ for each quantile of the value distribution, they obtain the following invertible form of Equation (4), which may be used to construct an efficient non-parametric estimator for $\lambda(\cdot)$,

$$v(\alpha) = b(\alpha) + \lambda^{-1} \left\{ \frac{1}{N-1} \frac{\alpha}{g[b(\alpha | N) | N]} \right\}. \quad (5)$$

If only data from sealed-bid FPAs are available, then a typical approach employs an additional structural restriction on the the distribution of values. [Guerre et al. \(2009\)](#) assume that the value distribution in a given auction is independent of the number of participants (given observables). A sufficient condition for this exclusion restriction is that participation in an auction is *exogenous* to the realized number of entrants,¹⁹ so that $F(v|N) = F(v)$. Under this assumption, the *quantiles* of the bid distributions for two auctions that have different numbers of bidders must be equal, even though the equilibrium bid functions are different. Applying Equation (5) for two such auctions

¹⁸That is, $G(b | N) = \Pr[B \leq b | N] = \Pr[\beta^{-1}(B) \leq \beta^{-1}(b) | N] = \Pr[V \leq \beta^{-1}(b) | N] = F[\beta^{-1}(b)] = F(v|N)$ and, differentiating, $g(b | N) = f[\beta^{-1}(b) | N] / \beta'[\beta^{-1}(b)|N] = f(v | N) / \beta'(v | N)$.

¹⁹[Guerre et al. \(2009\)](#) provide an extension for settings in which entry is endogenous but good instruments for entry are available.

and equating them, they obtain the following set of equivalences, or *compatibility conditions*, which can be used to non-parametrically trace out $F(v)$ and U .

$$v(\alpha) = b(\alpha; N_2) + \lambda^{-1} \left\{ \frac{1}{N_2 - 1} \frac{\alpha}{g[b(\alpha; N_2) | N_2]} \right\} = b(\alpha; N_1) + \lambda^{-1} \left\{ \frac{1}{N_1 - 1} \frac{\alpha}{g[b(\alpha; N_1) | N_1]} \right\}. \quad (6)$$

An alternative approach to identification applies parametric assumptions in lieu of exclusion restrictions. [Campo et al. \(2011\)](#) restrict bidder preferences to a family of utility functions that can be characterized by a single parameter, such as CARA and CRRA, and show that the non-identification of Equation (4) can be resolved by pinning down a single additional degree of freedom. In this case, the value distribution and risk aversion parameter can be semi-parametrically estimated in three steps. First, the econometrician estimates a (non-parametric) distribution of bid quantiles, as in the risk-neutral case. She then chooses a focal quantile α_0 and fits Equation (7) assuming a parametric utility function governed by parameters θ_u and a flexible parametrization of the quantile function (such as a high-dimensional polynomial or a spline) with parameters θ_v .²⁰

$$g[b(\alpha_0; N) | N] = \frac{1}{N - 1} \frac{\alpha_0}{\lambda [v(\alpha_0; N, \theta_v) - b(\alpha_0; N); \theta_u]}. \quad (7)$$

Finally, the full value distribution is recovered (non-parametrically) from Equation (4).

As the goal of the empirical literature is to adequately capture the data generating process underlying bid data, extensions to the basic model discussed in [Guerre et al. \(2009\)](#) are often considered on the basis of what is needed to rationalize observed bids. In many cases, however, allowing more complex modeling features opens new avenues for identification. [Campo \(2012\)](#) shows that when bidders are allowed to have different risk aversion parameters, one can exploit variation in bids conditional on observables for identification without quantile restrictions at all. Under the assumption that the value distribution in a given auction can be fully parameterized by a vector of observable variables X , the difference between two bids with the same X must be rationalized by the difference in the bidders' risk aversion. As such, [Campo \(2012\)](#) is able to write compatibility conditions akin to Equation (6) by comparing bids that are accompanied by the same observable characteristics and estimate with a simplified procedure.

In most cases, modeling extensions requires additional data or assumptions. [Grundl and Zhu \(2019\)](#) show that a model with unobservable heterogeneity is only identified if the unobservable heterogeneity is separable relative to the bidders' utility function²¹ or if it satisfies a common support assumption for auctions used in compatibility conditions. [Luo, Perrigne, and Vuong \(2018\)](#) consider a model of *ex post* risk in the spirit of [Eső and White \(2004\)](#). Under an assumption of exogenous participation (as in [Guerre et al. \(2009\)](#)), they show that bidders' value distribution and utility function can be non-parametrically identified up to an unidentified threshold risk premium

²⁰All of the estimated quantities in [Campo et al. \(2011\)](#) are conditional on a vector of observables X . We suppress these from the notation in Equation (7) as conditioning on observables can also be done in the other papers in this literature, although this is sometimes less explicit in their exposition.

²¹For example, it can be multiplicatively separable if preferences are CRRA.

π_0 . Under endogenous entry, however, they argue that semi-parametric identification is needed and demonstrate how it can be achieved through a set of “within” compatibility conditions that compare the bid quantiles of two (parametric) bidder types facing a Gaussian ex post shock.

Luo and Takahashi (2019) extend the Luo et al. (2018) setting to consider multivariate ex post uncertainty in the context of scaling procurement auctions. Applying a model of CARA agents who bid on a portfolio of project components that are subject to Gaussian risk, along with a “risk-free asset” that can absorb unlimited slack, they show that bidder values for winning an auction can be projected onto single dimensional “pseudo-cost” types. This allows them to apply identification results similar to Luo et al. (2018) and estimate bidder values under a common utility function through indirect inference, by matching moments of the distribution of total bids.

When more data are available, these approaches can be supported or even supplanted with alternative sources of identification. Chen, Gentry, Li, and Lu (2020) consider an affiliated-signals with risk aversion (AS-RA) model of selective entry and show that data on the pool of potential bidders can be used to point identify the joint distribution of bidder values and signals, along with the risk-averse utility function. Kong (2020) approximates the number of potential bidders in a pool of firms that alternately participate in ascending and sealed-bid auctions to identify selective entry probabilities into each the two formats. Following the approach of Lu and Perrigne (2008), she uses bids from the ascending price format to identify the distribution of bidder values, and bids from the fixed price format to identify bidders’ utility functions. Kong (2020) then accounts for selective entry by modeling the joint probability that a potential bidder participates in an ascending bid auction—subject to a non-negative profit constraint—and the bid value that is incurred upon entry.

Bolotnyy and Vasserman (2021) use historical data on *ex ante* government projections and *ex post* realizations of project needs in scaling procurement auctions to semi-parametrically identify bidder values and risk aversion types without reference to a first order condition like Equation (3). Under the assumption that the realization of ex post uncertainty is independent of who wins the auction, Bolotnyy and Vasserman (2021) show that the *portfolio problem* by which bidders choose to split their total bids across different project components is sufficient for identification. Key to their argument is a notion of separability that is particular to scaling auctions: although the winner of an auction is determined based on their total bid, the value of winning an auction is given by the solution to the portfolio problem conditioned on the winning total. Under Bolotnyy and Vasserman (2021)’s assumptions, the solution to a bidder’s portfolio problem only depends on auction-level factors through their total bid. Therefore, the total serves as a sufficient statistic for all competitive considerations and the moments characterizing the optimal portfolio spread alone can be used to identify heterogeneous value and risk aversion types without restrictions on the distribution of bidders across auctions.

4.2 Counterfactual evidence and policy advice

The primary motivation for empirical work involving risk aversion is to enable accurate policy recommendations on the basis of data. In the majority of the literature to date, this is achieved by virtue of providing empirical evidence that bidders are indeed risk-averse. As we demonstrate in

the theoretical section of this survey, the presence of risk aversion alone is sufficient for a number of policy conclusions. For example, when values are independent, the FPA outperforms the SPA in revenue in the presence of risk aversion, independently of what that level is. However, in other cases, policy recommendations depend on the relative distributions of bidder values and risk aversion. For example, when values are interdependent, the revenue ranking depends on the strength of affiliation versus the degree of risk aversion. Other conclusions, like the preferences of bidders over auction formats and the desirability of reducing *ex post* risk, depend on the magnitude of the risk that bidders are exposed to and the way that bidders' risk attitudes change with wealth.

Although the empirical literature has thus far mostly focused on questions of identification and estimation of bidder types, recent papers have begun to address policy design more directly through counterfactual simulations. These exercises provide valuable context for interpreting the role that risk and risk aversion play in driving welfare outcomes. For instance, [Kong \(2020\)](#) uses a simulated decomposition of auction outcomes to show that uncertainty about the set of competitors explains the majority of a \$25,000 average revenue difference between ascending and sealed-bid auctions in her data. Counterfactual simulations can also inform evaluations of policies that have not yet been enacted. For instance, [Bolotnyy and Vasserman \(2021\)](#) show that a policy to impose price floors on unit bids—which had been rejected by the auctioneer in their setting after protests by bidding contractors—would reduce the excess spending incurred due to *ex post* uncertainty by a fifth. As our discussion in Section 3 demonstrates, the space for beneficial policy design in the settings with risk aversion is substantial and interesting. As this literature continues to mature, we hope that more studies performing counterfactual analysis—guided both by theoretical predictions and by institutional features—will continue to emerge.

5 Conclusion

We summarize the key takeaways for auction design with risk aversion in Table 1. As we discussed in Section 4, design recommendations often depend on details of the auction environment and, in particular, the distributions of values and risk aversion among auction participants. Because empirical studies on risk aversion in real-world auctions are relatively scarce compared to analysis in the risk-neutral setting, we believe there could be significant gains for policy advice through further research.

A fundamental question that remains, in our view, underexplored is the practical salience of risk aversion for the design of auctions. That is, how much value is left on the table by neglecting risk preferences in real-world auction design, and what cost is there in employing standard auctions for such settings, relative to the optimal benchmark? Approximate mechanism design has contributed to our understanding of these questions, but such results are typically focused on worst-case analysis. As a consequence, these questions are largely open for distributions that are actually observed in practice. Furthermore, while the theoretical literature has focused on characterizing optimal mechanisms for a given setting, it is clear that the mechanisms described by [Maskin and Riley \(1984\)](#) and [Matthews \(1987\)](#) are unlikely to be implemented in practice, both due to the complexity of their rules and the precision they require in estimating bidder types.

There is a fertile ground for further research that lies between the optimal full-information benchmark and the worst-case setting of approximate mechanism design. For example, what types of *simple* mechanisms—ones that could be explained persuasively to non-technical domain experts and for which lower-dimensional characterizations of bidder types might suffice—will perform well when in different cases where risk aversion is salient?

The theoretical and empirical literature to date has been mostly responsive, seeking to understand auction data through the lens of risk aversion *post hoc*. But economists are increasingly involved in policy design and experimentation, including in auction design. These active designers typically model auctions under the assumption that bidders are risk-neutral, but this approach may not be appropriate in all settings. To better understand how to structure auctions when risk aversion is salient, theorists and empiricists should don their engineer hats together as in Roth (2002). Auction designers should consider designing adaptively so as to learn about the aspects of risk and risk aversion that bidders are sensitive to. Meanwhile, auction empiricists should motivate their estimation strategies on the drivers of bidding behavior that are most salient for design.

Table 1: Key takeaways from the literature on auction design with risk aversion

Design element	Takeaway from the literature	Related literature
Classic auction formats	FPA has higher expected revenue than the SPA. CARA buyers are indifferent between the formats, while DARA buyers prefer SPA. Affiliated values imply an ambiguous ranking for both sellers and DARA buyers, depending on strength of affiliation vs risk aversion.	Holt (1980), Milgrom and Weber (1982), Matthews (1987), Smith and Levin (1996), Che and Gale (2006), Lu and Perrigne (2008), T. Li et al. (2015), Chen et al. (2020), Kong (2020)
Reserve prices	Optimal reserves are lower than the risk-neutral setting, and lower in the FPA than the SPA, decreasing in the number of bidders. Hidden reserves may be profitable for sufficiently risk-averse bidders. With interdependent values and sufficiently risk-averse bidders, very low reserves may be optimal.	Riley and Samuelson (1981), Hu et al. (2010), Hu (2011), H. Li and Tan (2017), Chakraborty (2019), Hu et al. (2019)
Entry fees	For CARA bidders, optimal SPA has positive reserve but no entry fee, optimal FPA has positive reserves and entry fees. Generally, the optimal auction has type-dependent entry fees.	Matthews (1983), Maskin and Riley (1984), Chakraborty (2019)
Buyout options	Sellers and bidders may both prefer a buyout option. Buyers prefer the option to be revoked after the first bid, while sellers benefit from a permanent option.	Budish and Takeyama (2001), Akerberg et al. (2006), Mathews and Katzman (2006), Reynolds and Wooders (2009)
Subsidies to losing bidders	Optimal auction may involve subsidies to some losing bidders. With risk-averse buyers and sellers, a premium auction format may facilitate risk-sharing.	Matthews (1983), Maskin and Riley (1984), Hu, Offerman, and Zou (2018),
Information about number of bidders	Concealing the number of bidders may result in higher expected revenue against CARA and DARA bidders.	McAfee and McMillan (1987), Matthews (1987), Kong (2020)
<i>Ex post</i> risk	Seller should try to reduce the risk faced by DARA buyers; DARA buyers prefer to bid on riskier prospects.	Athey and Levin (2001), Esó and White (2004), Hu, Matthews, and Zou (2018), Luo et al. (2018), Luo and Takahashi (2019), Bolotnyy and Vasserman (2021)
Risk-averse sellers	Risk-averse seller prefers FPA to SPA and lower reserve prices and entry fees than the risk-neutral seller. Deterministic revenue is theoretically optimal, at expense of <i>ex post</i> individual rationality.	Waehrer et al. (1998), Esó and Futó (1999)

Notes: The independent private values setting is assumed throughout, unless otherwise stated. FPA and SPA are abbreviations for first-price and second-price auctions, respectively.

References

- Akerberg, D., Hirano, K., & Shahriar, Q. (2006). The buy-it-now option, risk aversion, and impatience in an empirical model of eBay bidding. *Unpublished manuscript*.
- Athey, S., & Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6), 2107–2140.
- Athey, S., & Levin, J. (2001). Information and competition in US Forest Service timber auctions. *Journal of Political Economy*, 109(2), 375–417.
- Aumann, R. J., & Kurz, M. (1977). Power and taxes. *Econometrica*, 1137–1161.
- Baisa, B. (2017). Auction design without quasilinear preferences. *Theoretical Economics*, 12(1), 53–78.
- Bajari, P., & Hortacısu, A. (2005). Are structural estimates of auction models reasonable? Evidence from experimental data. *Journal of Political Economy*, 113(4), 703–741.
- Balzer, B., & Rosato, A. (2021). Expectations-based loss aversion in auctions with interdependent values: Extensive vs. intensive risk. *Management Science*, 67(2), 1056–1074.
- Bernoulli, D. (1738). Specimen theoriae novae de mensura sortis. *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5. (Translated by Louise Sommer and reprinted as “Exposition of a New Theory on the Measurement of Risk” in *Econometrica*, 1954, 22:23–36)
- Bolotnyy, V., & Vasserman, S. (2021). Scaling auctions as insurance: A case study in infrastructure procurement. *Working Paper*.
- Border, K. C. (1991). Implementation of reduced form auctions: A geometric approach. *Econometrica*, 1175–1187.
- Bose, S., Ozdenoren, E., & Pape, A. (2006). Optimal auctions with ambiguity. *Theoretical Economics*, 1(4), 411–438.
- Budish, E. B., & Takeyama, L. N. (2001). Buy prices in online auctions: Irrationality on the internet? *Economics Letters*, 72(3), 325–333.
- Campo, S. (2012). Risk aversion and asymmetry in procurement auctions: Identification, estimation and application to construction procurements. *Journal of Econometrics*, 168(1), 96–107.
- Campo, S., Guerre, E., Perrigne, I., & Vuong, Q. (2011). Semiparametric estimation of first-price auctions with risk-averse bidders. *Review of Economic Studies*, 78(1), 112–147.
- Chakraborty, I. (2019). Reserve price versus entry fee in standard auctions. *Economic Inquiry*, 57(1), 648–653.
- Che, Y.-K., & Gale, I. (2006). Revenue comparisons for auctions when bidders have arbitrary types. *Theoretical Economics*, 1(1), 95–118.
- Chen, X., Gentry, M., Li, T., & Lu, J. (2020). Identification and inference in first-price auctions with risk averse bidders and selective entry. *Unpublished manuscript*.
- Cox, J. C., Roberson, B., & Smith, V. L. (1982). Theory and behavior of single object auctions. *Research in Experimental Economics*, 2(1), 1–43.
- Cox, J. C., Smith, V. L., & Walker, J. M. (1982). Auction market theory of heterogeneous bidders. *Economics Letters*, 9(4), 319–325.
- Cox, J. C., Smith, V. L., & Walker, J. M. (1985). Experimental development of sealed-bid auction

- theory; Calibrating controls for risk aversion. *American Economic Review*, 75(2), 160–165.
- Cox, J. C., Smith, V. L., & Walker, J. M. (1988). Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty*, 1(1), 61–99.
- Crawford, V. P., & Iriberry, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica*, 75(6), 1721–1770.
- Crémer, J., & McLean, R. P. (1985). Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2), 345–361.
- Delnoij, J., & De Jaegher, K. (2020). Competing first-price and second-price auctions. *Economic Theory*, 69(1), 183–216.
- di Tillio, A., Kos, N., & Messner, M. (2016). The design of ambiguous mechanisms. *Review of Economic Studies*, 84(1), 237–276.
- Elyakime, B., Laffont, J. J., Loisel, P., & Vuong, Q. (1994). First-price sealed-bid auctions with secret reservation prices. *Annales d’Économie et de Statistique*(34), 115–141.
- Engelbrecht-Wiggans, R., & Katok, E. (2007). Regret in auctions: Theory and evidence. *Economic Theory*, 33(1), 81–101.
- Eső, P. (2005). An optimal auction with correlated values and risk aversion. *Journal of Economic Theory*, 125(1), 78–89.
- Eső, P., & Futó, G. (1999). Auction design with a risk averse seller. *Economics Letters*, 65(1), 71–74.
- Eső, P., & White, L. (2004). Precautionary bidding in auctions. *Econometrica*, 72(1), 77–92.
- Feng, Y., Hartline, J. D., & Li, Y. (2019). Optimal auctions vs. anonymous pricing: Beyond linear utility. In *Proceedings of the 2019 ACM Conference on Economics and Computation* (pp. 885–886).
- Fibich, G., & Gaviious, A. (2010). Large auctions with risk-averse bidders. *International Journal of Game Theory*, 39(3), 359–390.
- Filiz-Ozbay, E., & Ozbay, E. Y. (2007). Auctions with anticipated regret: Theory and experiment. *American Economic Review*, 97(4), 1407–1418.
- Fu, H., Hartline, J., & Hoy, D. (2013). Prior-independent auctions for risk-averse agents. In *Proceedings of the fourteenth ACM Conference on Electronic Commerce* (pp. 471–488).
- Gershkov, A., Moldovanu, B., Strack, P., & Zhang, M. (2021). Optimal auctions: Non-expected utility and constant risk aversion. *Unpublished manuscript*.
- Goeree, J. K., Holt, C. A., & Palfrey, T. R. (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104(1), 247–272.
- Goeree, J. K., & Offerman, T. (2004). Notes and comments the Amsterdam auction. *Econometrica*, 72(1), 281–294.
- Grundl, S., & Zhu, Y. (2019). Identification and estimation of risk aversion in first-price auctions with unobserved auction heterogeneity. *Journal of Econometrics*, 210(2), 363–378.
- Guerre, E., Perrigne, I., & Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3), 525–574.
- Guerre, E., Perrigne, I., & Vuong, Q. (2009). Nonparametric identification of risk aversion in first-price auctions under exclusion restrictions. *Econometrica*, 77(4), 1193–1227.
- Holt, C. A. (1980). Competitive bidding for contracts under alternative auction procedures. *Journal*

- of *Political Economy*, 88(3), 433–445.
- Holt, C. A., & Laury, S. K. (2014). Assessment and estimation of risk preferences. In M. J. Machina & W. K. Viscusi (Eds.), *Handbook of the economics of risk and uncertainty* (Vol. 1, pp. 135–201). Elsevier.
- Hu, A. (2011). How bidder's number affects optimal reserve price in first-price auctions under risk aversion. *Economics Letters*, 113(1), 29–31.
- Hu, A., Matthews, S. A., & Zou, L. (2010). Risk aversion and optimal reserve prices in first-and second-price auctions. *Journal of Economic Theory*, 145(3), 1188–1202.
- Hu, A., Matthews, S. A., & Zou, L. (2018). English auctions with ensuing risks and heterogeneous bidders. *Journal of Mathematical Economics*, 76, 33–44.
- Hu, A., Matthews, S. A., & Zou, L. (2019). Low reserve prices in auctions. *Economic Journal*, 129(622), 2563–2580.
- Hu, A., Offerman, T., & Zou, L. (2011). Premium auctions and risk preferences. *Journal of Economic Theory*, 146(6), 2420–2439.
- Hu, A., Offerman, T., & Zou, L. (2018). How risk sharing may enhance efficiency of English auctions. *Economic Journal*, 128(610), 1235–1256.
- Kagel, J. H. (1997). Auctions: A survey of experimental research. In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (Vol. 1, pp. 501–585). Princeton University Press.
- Kagel, J. H., & Levin, D. (2016). Auctions: A survey of experimental research. In J. H. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (Vol. 2, pp. 563–637). Princeton University Press.
- Kleiner, A., Moldovanu, B., & Strack, P. (2020). Extreme points and majorization: Economic applications. *Unpublished manuscript*.
- Kong, Y. (2020). Not knowing the competition: Evidence and implications for auction design. *RAND Journal of Economics*, 51(3), 840–867.
- Kőszegi, B., & Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review*, 97(4), 1047–1073.
- Lange, A., & Ratan, A. (2010). Multi-dimensional reference-dependent preferences in sealed-bid auctions—how (most) laboratory experiments differ from the field. *Games and Economic Behavior*, 68(2), 634–645.
- Levin, D., & Ozdenoren, E. (2004). Auctions with uncertain numbers of bidders. *Journal of Economic Theory*, 118(2), 229–251.
- Li, H., & Tan, G. (2017). Hidden reserve prices with risk-averse bidders. *Frontiers of Economics in China*, 12(3), 341–370.
- Li, T., Lu, J., & Zhao, L. (2015). Auctions with selective entry and risk averse bidders: theory and evidence. *RAND Journal of Economics*, 46(3), 524–545.
- Lu, J., & Perrigne, I. (2008). Estimating risk aversion from ascending and sealed-bid auctions: The case of timber auction data. *Journal of Applied Econometrics*, 23(7), 871–896.
- Luo, Y., Perrigne, I., & Vuong, Q. (2018). Auctions with *ex post* uncertainty. *RAND Journal of Economics*, 49(3), 574–593.
- Luo, Y., & Takahashi, H. (2019). Bidding for contracts under uncertain demand: Skewed bidding

- and risk sharing. *Available at SSRN 3364708*. (Unpublished manuscript)
- Maskin, E., & Riley, J. (1984). Optimal auctions with risk averse buyers. *Econometrica*, 1473–1518.
- Mathews, T., & Katzman, B. (2006). The role of varying risk attitudes in an auction with a buyout option. *Economic Theory*, 27(3), 597–613.
- Matthews, S. A. (1983). Selling to risk averse buyers with unobservable tastes. *Journal of Economic Theory*, 30(2), 370–400.
- Matthews, S. A. (1987). Comparing auctions for risk averse buyers: A buyer's point of view. *Econometrica*, 633–646.
- McAfee, R. P., & McMillan, J. (1987). Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43(1), 1–19.
- Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press.
- Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. *Econometrica*, 1089–1122.
- Morgenstern, O., & Von Neumann, J. (1953). *Theory of games and economic behavior*. Princeton University Press.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1), 58–73.
- Perrigne, I., & Vuong, Q. (2019). Econometrics of auctions and nonlinear pricing. *Annual Review of Economics*, 11, 27–54.
- Pevnitskaya, S. (2004). Endogenous entry in first-price private value auctions: The self-selection effect. *Unpublished manuscript*.
- Reynolds, S. S., & Wooders, J. (2009). Auctions with a buy price. *Economic Theory*, 38(1), 9–39.
- Riley, J. G., & Samuelson, W. F. (1981). Optimal auctions. *American Economic Review*, 71(3), 381–392.
- Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4), 1341–1378.
- Skrzypacz, A. (2013). Auctions with contingent payments—an overview. *International Journal of Industrial Organization*, 31(5), 666–675.
- Smith, J. L., & Levin, D. (1996). Ranking auctions with risk averse bidders. *Journal of Economic Theory*, 68(2), 549–561.
- Sundararajan, M., & Yan, Q. (2020). Robust mechanisms for risk-averse sellers. *Games and Economic Behavior*, 124, 644–658.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1), 8–37.
- Waehrer, K., Harstad, R. M., & Rothkopf, M. H. (1998). Auction form preferences of risk-averse bid takers. *RAND Journal of Economics*, 179–192.
- Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica*, 95–115.