ROBUST STRATEGIC TRANSFER LEARNING IN AN UNCERTAIN ENVIRONMENT: FROM SECOND- TO FIRST-PRICE AUCTIONS

NIAN SI, SAN GULTEKIN, JOSE BLANCHET, AND AARON FLORES

Abstract. In this paper we focus on transfer learning in a context in which: a) the target distribution is unknown and the learning task is challenging due to a lack of or limited data; b) there is a related (yet different) environment in which the target distribution can be estimated with an acceptable or even high level of accuracy, and c) the knowledge transfer activity itself induces bias and strategic behavior in the learning tasks and the data generating distributions. While a) and b) are typically studied in transfer learning, c) is a key feature in our study here. We apply our analysis to the task of transfer learning of the optimal (profit-maximizing) bidding strategy from second- to first-price auctions. In this application setting, we achieve distributionally robustness in various game-theoretic frameworks, where optimal equilibrium knowledge transfer bidding policies are explicitly characterized.

1. Introduction

The goal of this paper is to study transfer learning and achieve robustness in a context in which the transfer learning activity itself (through participating agents) introduces bias and strategic behavior in the underlying domains and learning tasks.

We focus on strategic transfer learning with unknown target distributions; see, for example, [30] for a discussion and [21] for a review of transfer learning (also known as knowledge transfer) in general.

Transfer learning with unknown target distributions typically involves two basic elements: a) the learning task under the target distribution is relatively challenging due to a lack of information or limited data, and b) there is a related (yet different) environment and task in which the target distribution can be estimated with an at least acceptably high degree of accuracy. Here we introduce and study a new element, namely, c) the fact that one or several agents will be able to apply knowledge transfer from one activity to another naturally introduces strategic behavior and, therefore, biases in the data generating process and the learning task.

The knowledge transfer setting with unknown distributions presents a challenge for traditional transfer learning methods. This is the case even if the target domain distribution is accessible only through a limited amount of data because standard reweighting procedures may be subject to substantial selection bias; see [7]. Thus, it is important to develop a robust method that performs well in this uncertain environments, which arise in the context of unknown target distributions. To achieve this goal, distributionally robust optimization (DRO) formulations [9, 19, 12, 4] have been applied; see, for example, [30, 10, 11], and [26]. In the distributionally robust optimization, one optimizes the worst case loss, where the underlying distribution is allowed to vary in a distributional uncertainty set centered on a reference measure. In this

Date: This Version: 12$^{\text{th}}$ Feb. 2021

.

work, we also apply distributionally robust optimization, but we are also required to take into account strategic behavior in the transfer learning process.

When combining both transfer learning and strategic behavior, we naturally connect several important areas that are traditionally disconnected, namely, transfer learning in machine learning and robust game theory in economics and operations research. We see this as our first contribution in this paper. The connections between these areas and our work are discussed in Section 4, after we introduce our quantitative results.

In turn, this connection leads to our second contribution, namely, motivated by the need for obtaining less conservative estimates (relative to those derived from robust game theory), we introduce a concept for the study of strategic behavior, which we call "global distributionally robust Nash equilibrium," introduced and discussed in Section 3.2. We believe learning according to this new concept leads to a better performance in the strategic transfer learning context.

Finally, we apply our insights to a particular setting involving optimal bidding policies for transfer learning with unknown distributions from second-price auctions to first-price auctions. In this setting, we are able to show explicit characterizations or closed-form approximations for various forms of distributionally robust optimal bidding policies – which in turn are interpreted as optimal strategic transfer learning strategies.

In addition to being natural and important from an applied standpoint, optimal bidding policies provide an excellent case-in-point for our conceptual framework. To illustrate this, consider the setting of auctions in the online advertising industry. In this setting, publishers sell ad impression opportunities through auctions organized by ad exchanges [20]. Demand-side platforms (DSPs) bid, on behalf of advertisers, for these ad impression opportunities. DSPs need to optimize the bids on behalf of advertisers, and they should try to infer the bid of the highest competitor. This is often done by using data collected from repeated auctions – but it is important to keep in mind that the ad impression opportunities, the auction participants, and the campaigns from which the data is collected are different for various auctions.

For the sake of simplicity, we focus only on two types of auction design mechanisms: first-price auctions and second-price auctions. There are other designs which can be seen as combinations of the two basic types, but we do not consider those to simplify our exposition.

In a second-price auction mechanism, the DSPs are able to elicit much more information than in the first-price auction. This is so because the winning DSP gets to know the second-highest bidding price (i.e., the highest competing bid) after winning. In a first-price auction, on the other hand, the winner of the auction does not get to learn any information about the rest of the bids sent by its competitors – other than that the winning bid was obviously higher than any other bid.

A common scenario in online advertising occurs when many ad impression opportunities across two exchanges (one running a first-price auction and another one operating a second-price auction) may be similar in nature, e.g., from similar websites, the same ad slots, etc. This scenario may seem to fit well within a transfer learning task. But this cannot be done directly and it is where the application of our framework becomes apparent. First of all, as mentioned earlier, the data may not follow the exactly same distribution. And, second, the set of DSPs bidding on these different types of auctions of ad impression opportunities may coincide. However, this creates an incentive for strategic behavior, since DSPs who participate in both auctions will naturally decide to try to inform their bids based on the information gathered from both auctions. This strategic behavior complicates the transfer learning task from the second-price auction domain to the first-price auction domain.

This scenario is modeled in Section 2 by introducing a profit maximization formulation for optimal bidding strategies. We supply several key technical results in Section 3 and discuss our results with the related literature in Section 4. Then, we illustrate our results visually using concrete examples in Section 5. Section 6 discusses results in the heterogeneous setting. We conclude with a discussion on potential future research and applications in Section 7.

2. Preliminaries and Assumptions

We consider that similar ad impression opportunities are run in two exchanges: A and B. Exchange A uses a second-price auction mechanism, where the bidder who bids the highest price wins the auction, but only needs to pay the price offered by the second-highest bidder; while Exchange B uses a first-price auction mechanism, where the highest-price bidder wins the auction and pays the exact amount of his bidding price. Specifically, we denote the bidder's value as u, the bidder's bidding price as b, the reservation price as r, and the highest competing prices as V_A,V_B in Exchanges A and B, respectively (where V_A and V_B are random). Bidders with value $u \geq r$, in Exchange A optimize $\max_{b\in[r,u]} E\left[(u - V_A)\,\mathbb{I}\left\{V_A < b\right\}\right],$ where $\mathbb{I}\{\cdot\}$ denotes the indicator function, and in Exchange B, they optimize $\max_{b\in[r,u]} (u - b)\, P\left(V_B < b\right).$ Bidders with value $u < r$ will not enter the auction, and we simply say they bid $b = 0$ and the payoff is zero.

Recall that bidding truthfully $b = u$ in the second-price auction is the weakly dominant strategy. Therefore, we assume that the highest competing value in the second-price exchange A is precisely the highest competing bid V_A. From the data we collect from Exchange A, we obtain an estimate of the highest competing value. The goal here is to design a robust bidding policy to transfer the knowledge from the second-price exchange A to the first-price exchange B.

Under the first-price mechanism, bidders will discount their bids. Therefore, it is not appropriate to directly use the distribution of the competitive landscape in Exchange A, even if we know that the competition may still be similar. We thus attempt to characterize a transformation, based on an equilibrium argument, which may provide a sensible way to gaining insight into the landscape expected in Exchange B, given the landscape in Exchange A. In order to do this, we first list standard assumptions in the literature (see, for example, [22]).

Assumption 1. (1) *There are $N+1$ bidders in the auction, where bidder i has value $u_i, i = 1, 2, \ldots, N+1$. Bidders can exit with a bid $b_i = 0$, or submit a bid above the reservation price r.*

(2) *The value of bidders U_i are independent and identically distributed, which follows a common distribution $F_U(u)$ with $F_U(\underline{u}) = 0, F_U(\overline{u}) = 1$ and $F_U(u)$ is strictly increasing and has a continuous density $f_U(u)$ over the interval $[\underline{u}, \overline{u}]$, with $\underline{u} < r < \overline{u}$.*

(3) *Each bidder makes a bid b_i, which is a strictly increasing and continuous function of their value u_i, $b_i = s(u_i), i = 1, \ldots, N+1$, with $s(\cdot) \in IC\left([r, \overline{u}]\right),$ where $IC\left([r, \overline{u}]\right)$ denotes the collection of all strictly increasing and continuous functions over $[r, \overline{u}]$.*

Assumption 1.1 captures the key features of the ads bidding auctions in practice; Assumption 1.2 guarantees a tractable Nash equilibrium in the standard non-robust setting; Assumption 1.3 is natural since people tend to bid higher if they believe the opportunity has a larger value.

Suppose we are the bidder $N+1$ and let V denote the highest competing value in Exchange B, i.e.,

$$V = \max\left(U_1, ..., U_N\right).$$

Then, under Assumption 1, V has a strictly increasing CDF $F_V(\cdot)$ and a continuous density bounded away from zero $f_V(\cdot)$. Further, in the symmetric equilibrium, the highest competing bid V_B is $\max\left(s\left(U_1\right), ..., s\left(U_N\right)\right) = s\left(V\right)$.

If the distribution of the highest competing value in Exchange A V_A is known accurately to everyone, and Exchanges A and B have exact the same environments, i.e., $V = V_A$, the knowledge can be transferred perfectly from Exchange A to Exchange B. Therefore, rational bidders in Exchange B will optimize:

$$\max_{b\in\{0\}\cup[r,u]} \left(u-b\right) P\left(s\left(V\right) < b\right).$$

It is well-known that there exists a Bayesian Nash equilibrium of the first-price auction (see, for example, [22]).

Proposition 1. *Under Assumption 1, the symmetric Bayesian Nash equilibrium strategy is*

$$b^*_{\mathrm{NE}}\left(u\right) = u - \frac{\int_r^u F_V\left(x\right)\mathrm{d}x}{F_V\left(u\right)}, \text{ for } u \in [r, \overline{u}]. \tag{1}$$

3. Strategic Transfer Learning from Second to First-price auctions

However, in practice, there are three sources of uncertainties preventing bidders from bidding the Nash equilibrium (1):

(1) The highest competing value cannot be estimated perfectly.
(2) There are small differences between two exchanges. For example, there may be different groups of bidders or slightly different slots in the websites.
(3) Bidders may hold different beliefs and deviate from the equilibrium.

Therefore, we propose our robust bidding policy to perform strategic transfer learning from second to first-price auctions. Suppose $\hat{V}$ is an estimate of the highest competing value V, obtained from the second-price data. To reflect the ambiguity nature, we modify Assumption 1 to Assumption 2.

Assumption 2. (1) *There are many agents participating in the auction. Each agent faces the same distribution of the highest competing value, which is independent of their own value.*

(2) *The estimated highest competing value $\hat{V}$ follows the distribution $\hat{F}_V(u)$ with $\hat{F}_V(\underline{u}) = 0, \hat{F}_V(\overline{u}) = 1$ and $\hat{F}_V(u)$ has a continuous density $\hat{f}_V(u)$ bounded away from zero over the interval $[\underline{u}, \overline{u}]$. We further denote $\hat{P}_V$ as the probability measure induced by random variable $\hat{V}$.*

(3) *The true highest competing value V is inside an uncertainty set around $\hat{P}_V$, i.e.,*

$$P_V \in \mathcal{U}_\delta\left(\hat{P}_V\right) := \{P \in \mathcal{P}\left([0,+\infty)\right) | D(P, \hat{P}_V) \le \delta\},$$

where $D(\cdot,\cdot)$ is some divergence notion in the measure space and $\mathcal{P}([0,+\infty))$ denotes the Borel probability space on $[0,+\infty)$.

(4) *The highest competing bid b is a strictly increasing and continuous function of the highest competing value, i.e. $b = s(u)$, with $s(\cdot) \in IC\left([r, \overline{u}]\right)$, where $r > \underline{u}$ is a reservation price.*

Assumption 2.1 is approximately correct in practice due to reasons that we shall discuss in Section 6; Assumption 2.2 is an analog of Assumption 1.2 to guarantee a tractable solution, but note in Assumption 2.3, we do not require the measures in the uncertainty set $\mathcal{U}_\delta\left(\hat{P}_V\right)$ to be a continuous distribution.

3.1. **Distributionally Robust Nash Equilibrium.** We consider ambiguity averse bidders with maxmin utility [13]. We define a concept of symmetric distributionally robust Nash equilibrium (DRNE) in the first-price bidding setting.

Definition 1 (Distributionally robust Nash equilibrium)**.** *A strategy $b^*(\cdot) = s^*(\cdot)$ is called symmetric distributionally robust Nash equilibrium if*

$$\inf_{P_V \in \mathcal{U}_\delta(\hat{P}_V)} \{u - b^*(u)\} P_V(s^*(V) \leq b^*(u)) \geq \inf_{P_V \in \mathcal{U}_\delta(\hat{P}_V)} \{u - b\} P_V(s^*(V) \leq b),$$

for any $b \geq r$.

Therefore, in the distributionally robust Nash equilibrium, bidders choose the best response with respect to the worst case distribution in the uncertainty set. In other words, bidders robustify against an adversary that is allowed to choose any distributions in the uncertainty set. In this section, we consider the following distance notion:

$$D_p(P, Q) := \left(\int_0^\infty |P(V_1 \leq x) - Q(V_2 \leq x)|^p \, \mathrm{d}x\right)^{1/p}, \tag{2}$$

for $p \geq 1$. When $p = 1$, this distance is known as the optimal transport distance or the Wasserstein distance of order 1 (see, for example, [29]). It turns out that

$$\begin{aligned} D_1(P, Q) = \min\{&E_\pi |V_1 - V_2| : \\ &\pi \text{ is a joint distribution of } (V_1, V_2) \\ &\text{such that the marginal distribution of } V_1 \text{ is } P \text{ and} \\ &\text{the marginal distribution of } V_2 \text{ is } Q\}. \end{aligned}$$

Proposition 2 shows that $D_p(\cdot, \cdot)$ is indeed a metric.

Proposition 2. *$D_p(P, Q)$ is a metric in $\mathcal{P}_1([0, +\infty))$ for $p \geq 1$, where $\mathcal{P}_1([0, +\infty))$ denotes the collection of all probability measure on $[0, +\infty)$ with a finite first moment.*

Proof. Notice that

$$\begin{aligned} D_p(P, Q) &= \left(\int_0^\infty |P(V_1 \leq x) - Q(V_2 \leq x)|^p \, \mathrm{d}x\right)^{1/p} \\ &\leq \left(\int_0^\infty |P(V_1 \leq x) - Q(V_2 \leq x)| \, \mathrm{d}x\right)^{1/p} \leq D_1(P, Q)^{1/p}. \end{aligned}$$

Therefore, for any $P, Q \in \mathcal{P}_1([0, +\infty))$, we have $D_p(P, Q) < \infty$.

Then, it is easy to check $D_p(P, Q) \geq 0, D_p(P, Q) = 0 \Leftrightarrow P = Q$, and $D_p(P, Q) = D_p(Q, P)$. And the triangle inequality is obtained by the triangle inequality for L^P-norm. □

Then, we give a characterization of the DRNE in Theorem 1.

Theorem 1. *Let u_0 be the solution of the equation*

$$\left(\int_{\underline{u}}^{u_0} \hat{F}_V(t)^p \mathrm{d}t\right)^{1/p} = \delta.$$

Under Assumption 2, the symmetric DRNE of first-price auction with $D_p(\cdot,\cdot)$ distance uncertainty sets is any strictly increasing and continuous function in $[r, \overline{u}]$ with

$$b^*_{\mathrm{DRNE}}(u) = u - \int_{\max(r,u_0)}^{u} \hat{F}_V^{\delta}(x)\,\mathrm{d}x / \hat{F}_V^{\delta}(u)$$

for $u \geq \max(r, u_0)$, where $\hat{F}_V^{\delta}(u) = \hat{F}_V(-\lambda_u + u)$ and λ_u is the unique solution of

$$\delta = \left(\int_{u-\lambda_u}^{u} (\hat{F}_V(t) - \hat{F}_V(u-\lambda_u))^p \mathrm{d}t\right)^{1/p}.$$

The proof of Theorem 1 relies on Lemma 1 below. We first denote $\mathcal{U}_\delta^p\left(\hat{P}_V\right) := \{P \in \mathcal{P}([0,+\infty)) \,|\, D_p(P, \hat{P}_V) \leq \delta\}$.

Lemma 1.

$$\inf_{P \in \mathcal{U}_\delta^p(\hat{P}_V)} P(V \leq u) = \begin{cases} 0 & \left(\int_{\underline{u}}^{u} \hat{F}_V(t)^p \mathrm{d}t\right)^{1/p} < \delta \\ \hat{P}_V(V \leq u - \lambda_u) & \left(\int_{\underline{u}}^{u} \hat{F}_V(t)^p \mathrm{d}t\right)^{1/p} \geq \delta \end{cases},$$

where

$$\delta^p = \int_{u-\lambda_u}^{u} \left|\hat{F}_V(t) - \hat{F}_V(u-\lambda_u)\right|^p \mathrm{d}t.$$

Proof. If $\left(\int_{\underline{u}}^{u} \hat{F}_V(t)^p \mathrm{d}t\right)^{1/p} < \delta$, consider the probability measure

$$Q(V < u + \epsilon) = 0 \text{ and } Q(V \leq t) = \hat{P}_V(V \leq t) \text{ for } t \geq u + \epsilon.$$

There exists $\epsilon > 0$ such that

$$\left(\int_0^\infty \left|\hat{P}_V(V_1 \leq t) - Q(V_2 \leq t)\right| \mathrm{d}t\right)^{1/p} = \left(\int_{\underline{u}}^{x+\epsilon} \left|\hat{P}_V(V_1 \leq t) - Q(V_2 \leq t)\right| \mathrm{d}t\right)^{1/p} < \delta.$$

Therefore, $\inf_{P \in \mathcal{U}_\delta(P_n^V)} P(V \leq u) = 0$.

If $\left(\int_{\underline{u}}^{u} \hat{F}(t)^p \mathrm{d}t\right)^{1/p} \geq \delta$, consider the distribution

$$\begin{aligned} Q^{\epsilon_1}(V < t) &= \hat{P}_V(V < t) \text{ for } t \leq u - \lambda_u + \epsilon_1, \\ Q(V < u + \epsilon_1) &= \hat{P}_V(V < u - \lambda_u + \epsilon_1) \\ \text{and } Q(V \leq t) &= \hat{P}_V(V \leq t) \text{ for } t \geq u + \epsilon_2. \end{aligned}$$

For any $\epsilon_1 > 0$, there exists $\epsilon_2 > 0$, such that

$$\begin{aligned} &\left(\int_0^\infty \left|\hat{P}_V(V_1 \leq t) - Q(V_2 \leq t)\right| \mathrm{d}t\right)^{1/p} \\ = &\left(\int_{u-\lambda_u+\epsilon_1}^{u+\epsilon} \left|\hat{P}_V(V_1 \leq t) - Q(V_2 \leq t)\right| \mathrm{d}t\right)^{1/p} \leq \delta. \end{aligned}$$

Therefore, we have

$$\begin{aligned}\inf_{P\in\mathcal{U}_\delta(\hat{P}_V)} P\left(V\le u\right) &\le \inf_{\epsilon_1>0}\hat{P}_V\left(V<u-\lambda_u+\epsilon_1\right)\\ &= \hat{P}_V\left(V\le u-\lambda_u\right).\end{aligned}$$

On the other hand, for any probability measure $Q(V\le u)<\hat{P}_V\left(V\le u-\lambda_u\right)$, we have

$$\begin{aligned}\left(\int_0^\infty\left|\hat{P}_V\left(V_1\le t\right)-Q\left(V_2\le t\right)\right|\mathrm{d}t\right)^{1/p} &> \left(\int_{u-\lambda_u}^u\left|\hat{P}_V\left(V_1\le t\right)-Q\left(V_2\le t\right)\right|\mathrm{d}t\right)^{1/p}\\ &\ge \left(\int_{u-\lambda_u}^u\left|\hat{F}_V(t)-\hat{F}_V(t-\lambda_u)\right|^p\mathrm{d}t\right)^{1/p}=\delta.\end{aligned}$$

□

The remaining proof of Theorem 1 is in Appendix A.1.

Remark 1. *[16, 17] consider this problem with different assumptions. They require all probability distributions in the uncertainty set to be strictly increasing and continuous. Contrary to their results, we relax this assumption and give a closed-form solution corresponding to $D_p\left(\cdot,\cdot\right)$ distance uncertainty sets.*

The distributionally robust Nash equilibrium could be viewed as the Nash equilibrium with the belief $\hat{F}_V^\delta(u)$. However, the distribution $\hat{F}_V^\delta(u)$ is not in $\{P\in\mathcal{P}([0,+\infty)|D(P,\hat{P}_V)\le\delta\}$, which means DRNE is robust against an unrealistic bad distribution. Therefore, DRNE is overly conservative, which may cause a bad performance in the strategic transfer learning task.

3.2. **Global Distributionally Robust Nash Equilibrium.** To address the conservativeness of DRNE, we propose the concept of global distributionally robust Nash equilibrium (GDRNE). Before stating the definition, we make a further assumption about the distribution of bidders' own value.

Assumption 3. *The agent knows the distribution of their own value distribution $F_U(\cdot)$ with no uncertainty. Assume $F_U(\underline{u})=0$, $F_U(\overline{u})=1$, and the associated density $f_U(u)$ is continuous differentiable over the interval $[\underline{u},\overline{u}]$.*

Assumption 3 is natural since in practice the bidders have a full control of their value prediction algorithms and are able to estimate their own value distribution without any uncertainties. Instead of only focusing on a one-time bidding problem in Section 3.1, we consider a long-time average performance of the bidding system in a DSP. Suppose there are M auctions with our value U_i, $i=1,2,\ldots,M$, for each auction and we maximize the average performance under an unknown distribution P of the highest competing value, i.e.,

$$\begin{aligned}&\frac{1}{M}\sum_{i=1}^M\max_{b_i\in\{0\}\cup[r,U_i]}\left(U_i-b_i\right)P_V\left(s(V)<b_i\right)\\ \approx& E\left[\max_{b(U)\in[r,U]}\{U-b(U)\}P_V(s(V)<b(U)|U)\mathbb{I}\left\{U\ge r\right\}\right].\end{aligned}$$

This global nature of the learning task motivates us to propose a new concept of global distributionally robust Nash equilibrium.

Definition 2 (Global distributionally robust Nash equilibrium). *A strategy $b^*(\cdot) = s^*(\cdot)$ is called symmetric distributionally robust Nash equilibrium if*

$$\begin{aligned}&\inf_{P_V \in \mathcal{U}_\delta(\hat{P}_V)} E_U\left[\{U - b^*(U)\} \times P_V\left(s^*(V) \le b^*(U) | U\right) \mathbb{I}\{U \ge r\}\right] \\ &\ge \inf_{P_V \in \mathcal{U}_\delta(\hat{P}_V)} E_U\left[\{U - b(U)\} \times P_V\left(s^*(V) \le b(U) | U\right) \mathbb{I}\{U \ge r\}\right],\end{aligned}$$

for $b(\cdot) \in IC([r, \overline{u}])$ and $b(r) \ge r$.

Note that in Definition 1, the bidders try to robustify against a strong adversary, which is allowed to vary his/her distributions for every single auction; while in Definition 2, the adversary is only able to choose the same distribution among all different ads opportunities.

Below we consider a new type of distance which depends on s and b.

$$D_p^{s,b}(P, Q) := \left(\int_0^\infty \left|P\left(s(V_1) \le b(x)\right) - Q\left(s(V_2) \le b(x)\right)\right|^p \mathrm{d}x\right)^{1/p}. \tag{3}$$

At the first glance, the distance in (3) seems complicated, but in the equilibrium, $b(\cdot) = s(\cdot)$, and $D_p^{s,b}(\cdot,\cdot)$ reduces to the distance $D_p(\cdot,\cdot)$. In this case, it is hard to solve for a closed-form, but we give an asymptotic expansion in Theorem 2, when $\delta \to 0$.

Theorem 2. *We assume the uncertainty set is constructed by the distance $D_p^{s,b}(P, Q)$ with $p > 1$ and let $1/q + 1/p = 1$. Under Assumptions 2 and 3, a symmetric GDRNE of first-price auction with $D_p^{s,b}(\cdot,\cdot)$ distance uncertainty sets is*

$$b^*_{\text{GDRNE}}(u) = b_0(u) + b_1(u)\delta + o(\delta), \text{ for } u \ge r,$$

where

$$\begin{aligned}b_0(u) &= u - \frac{\int_r^u \hat{F}_V(x)\,\mathrm{d}x}{\hat{F}_V(u)}, \\ b_1(u) &= \int_r^u h(x) \exp\left(-\int_x^u \frac{\hat{f}_V(t)}{\hat{F}_V(t)} \mathrm{d}t\right) \mathrm{d}x, \text{ and} \\ h(u) &= \frac{b_0'(u)\left((u - b_0(u)) f_U(u)\right)^{q-1}}{\hat{F}_V(u)\left(\int_r^{\overline{u}} ((x - b_0(x)) f_U(x))^q \,\mathrm{d}x\right)^{1-1/q}},\end{aligned}$$

when $\delta \to 0$.

Sketch of Proof. For any strictly increasing and continuous functions, $b(\cdot), s(\cdot)$, we have

$$\inf_{P: D_p^{s,b}(P, \hat{P}_V) \le \delta} E_U\left[\{U - b(U)\} \times P\left(s(V) \le b(U) | U\right) \mathbb{I}\{U \ge r\}\right]$$

$$= E_U\left[\{U - b(U)\} \hat{P}_V\left(s(V_1) \le b(U) | U\right) \mathbb{I}\{U \ge r\}\right] \tag{4}$$

$$+ \inf_{P: D_p^{s,b}(P, \hat{P}_V) \le \delta} \left\{\int_r^{\overline{u}} (u - b(u)) f_U(u) \left(P\left(s(V) \le b(u)\right) - \hat{P}_V\left(s(V) \le b(u)\right)\right) \mathrm{d}u\right\}. \tag{5}$$

By the Hölder inequality, the term (5) is bounded by

$$\begin{aligned}
&\int_r^{\overline{u}} (u-b(u))f_U(u)\left(P\left(s(V)\le b(u)\right)-\hat{P}_V\left(s(V)\le b(u)\right)\right)\mathrm{d}u \\
&\ge -\left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q\mathrm{d}u\right)^{1/q}\times\left(\int_r^{\overline{u}}\left(P\left(s(V)\le b(u)\right)-\hat{P}_V\left(s(V)\le b(u)\right)\right)^p\mathrm{d}u\right)^{1/p} \\
&\ge -\delta\left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q\mathrm{d}u\right)^{1/q}. \qquad (6)
\end{aligned}$$

where $1/p+1/q=1$. We claim the bound is tight. By Theorem 1 in Chapter 7.4 and Lemma 1 in Chapter 7.5 of [18], we have the optimizer must satisfy the condition:

$$(u-b(u))\frac{\hat{f}_V\left(s^{-1}\left(b(u)\right)\right)}{s'(s^{-1}\left(b(u)\right))}-\hat{P}_V\left(s(V)\le b(u)\right)+\delta\frac{\left((u-b(u))f_U(u)\right)^{q-1}}{\left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q\mathrm{d}u\right)^{1-1/q}}=0. \qquad (7)$$

In the symmetric GDRNE, $b^*(\cdot)$ solves the differential equation:

$$(u-b(u))\frac{\hat{f}_V(u)}{b'(u)}-\hat{F}_V(u) \qquad (8)$$

$$+\delta\frac{\left((u-b(u))f_U(u)\right)^{q-1}}{\left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q\mathrm{d}u\right)^{1-1/q}}=0. \qquad (9)$$

Let $b^*(u)=b_0(u)+\delta b_1(u)+o(\delta)$. When $\delta=0$, it reduces to the non-ambiguity Nash equilibrium, i.e.,

$$b_0(u)=u-\frac{\int_r^u\hat{F}_V(x)\mathrm{d}x}{\hat{F}_V(u)}\ \text{and}$$

$$b_0'(u)=\frac{\hat{f}_V(u)\int_r^u\hat{F}_V(x)\mathrm{d}x}{\hat{F}_V(u)^2}=\frac{\hat{f}_V(u)}{\hat{F}_V(u)}\left(u-b_0(u)\right).$$

Then, the two terms in (8) becomes,

$$(u-b^*(u))\frac{\hat{f}_V(u)}{(b^*(u))'}-\hat{F}_V(u)=-\frac{\delta}{b_0'(u)}\left(\hat{F}_V(u)b_1'(u)+\hat{f}_V(u)b_1(u)\right)+o(\delta)$$

Therefore, after ignoring the higher order, (9) becomes a differential equation

$$b_1'(u)+\frac{\hat{f}_V(u)}{\hat{F}_V(u)}b_1(u)=\frac{b_0'(u)\left((u-b_0(u))f_U(u)\right)^{q-1}}{\hat{F}_V(u)\left(\int_r^{\overline{u}}\left((u-b_0(u))f_U(u)\right)^q\mathrm{d}u\right)^{1-1/q}}.$$

By solving the differential equation, we obtain the expression for $b_1(u)$. Finally, we verify that $b^*(\cdot)$ is indeed the GDRNE, when $\delta\to 0$. □

The full proof of Theorem 2 is in Appendix A.2.

Remark 2. *$b_1(u)>0$ for $u\ge r$. Thus, Robustification pushes the bid higher.*

4. Connections with Existing Literature

This work is closely related to many fields in statistics, machine learning, optimization, and economics. We briefly discuss a few of them with their respective focuses and their key differences with us.

Standard transfer learning concerns the setting that a model receives data from one source distribution and are tested on a known target distribution (labeled or unlabeled): the typical approach is to use importance sampling to reweight the source distribution towards the target distribution; see, for example, [25, 27, 14, 3, 21, 31]. In particular, [8] uses optimal transportation to operate the alignment from the source distribution to the target distribution.

However, as mentioned in the bidding problem, we may not be able to access any knowledge about the target distribution. In this case, the aforementioned methods cannot be applied. Furthermore, even if the target domain distribution is accessible, the reweighting procedure could fail due to large selection bias [7]. Therefore, [30, 11, 10] use a distributionally robust optimization approach to minimize the worst-case loss. These methods are similar to ours in the sense that we also use a minimax formulation in the DRO framework. Nevertheless, there is no literature focusing on the scenario in which other agents could behave strategically in the domain adaptation setting.

We also note that the strategic domain adaptation theory has a close connection with robust game theory under ambiguity in the economics and operations research literature. Mechanism design problems under ambiguity are studied in [2, 5, 6, 28, 15]. Furthermore, [1, 24, 16, 17] propose the same notions as our distributionally robust Nash equilibrium (Definition 1), and [23, 16, 17] study the ambiguity-aversed agents' behaviors in the first-price auction. Specifically, as we mentioned in Remark 1, [16, 17] give results similar to those in Theorem 1 under slightly different assumptions. However, we believe the concept of global distributionally robust Nash equilibrium (Definition 2) is new to the literature, and this paper is the first one that applies the techniques in robust game theory to domain adaption settings.

5. Examples and Visual Illustration

In this section, we compute the explicit equilibrium expressions for some sample distributions.

Example 1 (Uniform Distribution)**.** *Let* $\mathrm{unif}([a,b])$ *denote the uniform distribution on the interval* $[a,b]$. *Let* $U \sim \mathrm{unif}([a,b])$ *and* $\hat{V} \sim \mathrm{unif}([a,b])$. *Then, by direct calculation, the Nash equilibrium is*

$$b^*_{\mathrm{NE}}(u) = \frac{u^2 - 2ar + r^2}{2(u-a)}, \ \textit{for } u \in [r,b];$$

the distributionally robust Nash equilibrium is

$$b^*_{\mathrm{DRNE}}(u) = \frac{u^2 - 2u_0 \max(r,u_0) + \max(r,u_0)^2}{2(u-u_0)},$$

for $u \in (\max(r,u_0), b]$, *where*

$$u_0 = a + (p+1)^{1/(p+1)} \left[(b-a)\delta\right]^{p/(p+1)},$$

and for $u \in [r, \max(r, u_0))$, b^*_{DRNE} cannot be determined; then, for $p = q = 2$, the global distributionally robust Nash equilibrium is

$$b^*_{\text{GDRNE}}(u) = \frac{u^2 - 2ar + r^2}{2(u-a)} + \frac{\sqrt{3(b-a)^3}}{4\sqrt{b+3r-4a}(u-a)^3(b-r)^{3/2}} \times \left((u-a)^4 - (r-a)^4 + 4(r-a)^2(u-a)^2 \log\left(\frac{r-a}{u-a}\right)\right)\delta + o(\delta), \text{ for } u \in [r, b].$$

Example 1 shows that for $r > u_0$,

$$\begin{aligned} b^*_{\text{DRNE}}(u) &= b^*_{\text{NE}}(u) + \frac{(u_0-a)(u-r)^2}{2(u-a)(u-u_0)} = b^*_{\text{NE}}(u) + O\left(\delta^{p/(p+1)}\right), \text{ and} \\ b^*_{\text{GDRNE}}(u) &= b^*_{\text{NE}}(u) + O(\delta), \end{aligned}$$

which indicates that $b^*_{\text{DRNE}}(u)$ is overly conservative and pushes the bid higher than $b^*_{\text{GDRNE}}(u)$.

Example 2 (Maximum of N uniform random variables)**.** *We consider $U \sim \text{unif}([a,b])$ and $\hat{V} \sim \max\{U_1, U_2, \ldots, U_N\}$, where $U_i \overset{iid}{\sim} \text{unif}([a,b])$, for $i = 1, 2, \ldots, N$. In this example, we have*

$$b^*_{\text{NE}}(u) = \frac{a + Nu}{1+N} + \frac{(r-a)^{n+1}}{(1+N)(u-a)^n}, \text{ for } u \in [r, b].$$

and we numerically solve DRNE and GDRNE. In Figure 1, we plot NE, DRNE, and GDRNE for $p = 2, a = 1, n = 10, r = 3$. For DRNE, if $u_0 > r$, we plot $b_{\text{DRNE}}(u)$ only for $[u_0, b]$, since $b^*_{\text{DRNE}}(u)$ for $u < u_0$ cannot be uniquely determined. We test the dependence on δ and N of these three equilibrium concepts.

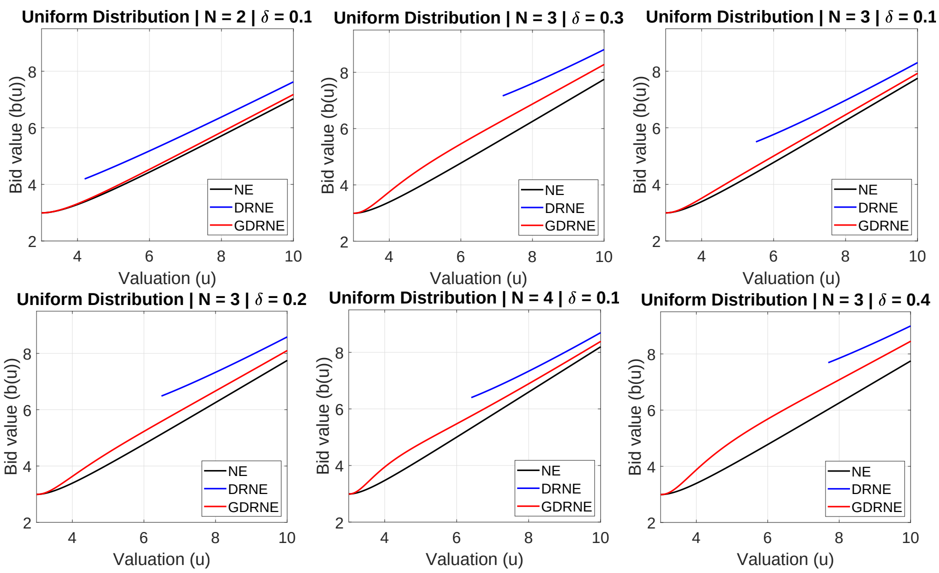

Figure 1: NE, DRNE, and GDRNE for different N and δ, where the nominal distribution is the maximum of N i.i.d. uniform random variables. Left: fixed $\delta = 0.1$ and $N \in \{2, 3, 4\}$. Right: fixed $N = 3$ and $\delta \in \{0.2, 0.3, 0.4\}$.

Example 3 (Maximum of N scaled beta random variables)**.** *We consider* $U \sim \text{unif}([a,b])$ *and* $\hat{V} \sim \max\{B_1, B_2, \dots, B_N\}$*, where* $\{B_i\}_{i=1}^N$ *are i.i.d. samples from the scaled beta distribution* $\text{Beta}^{a,b}(\alpha, \beta)$*, where*

$$\text{Beta}^{a,b}(\alpha,\beta) \overset{d}{=} a + (b-a)\text{Beta}(\alpha,\beta).$$

In this example, we numerically compute NE, DRNE and GDRNE, and plot them for $p = 2, a = 1, n = 10, r = 3, \alpha = 2, \beta = 5$ *in Figure 2. We do for DRNE the same thing that we do in Example 2. We test the dependence on* δ *and* N *of these three equilibrium concepts.*

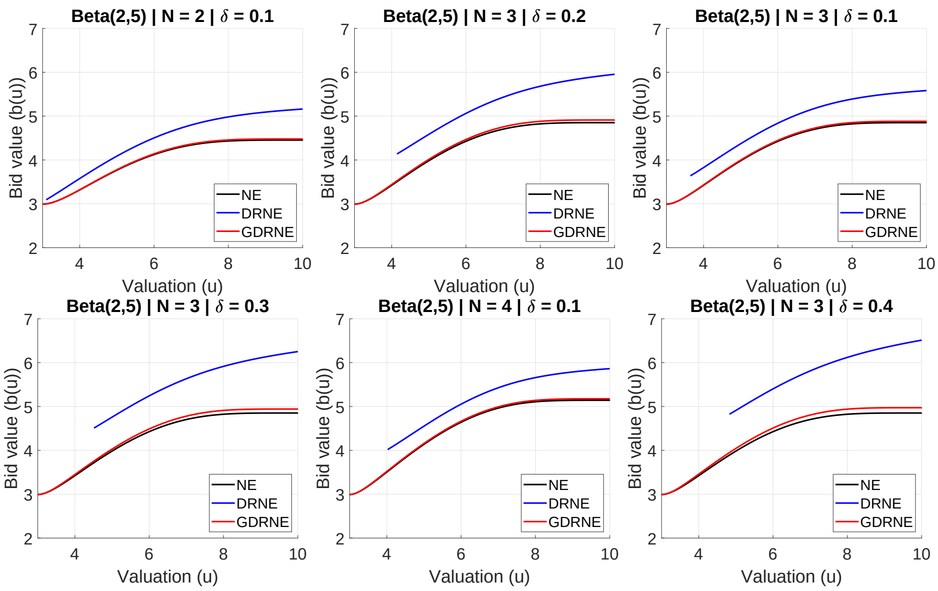

Figure 2: NE, DRNE, and GDRNE for different N and δ, where the nominal distribution is the maximum of N i.i.d. scaled beta random variables. Left: fixed $\delta = 0.1$ and $N \in \{2, 3, 4\}$. Right: fixed $N = 3$ and $\delta \in \{0.2, 0.3, 0.4\}$.

Figures 1 and 2 show that $b^*_{\text{DRNE}}(u) > b^*_{\text{GDRNE}}(u) > b^*_{\text{NE}}(u)$, and $b^*_{\text{DRNE}}(u)$ is far higher than $b^*_{\text{GDRNE}}(u)$, which indicates that bidding according to DRNE may lead to poor performance for the tasks in the target domain. Furthermore, another thing that should be noted is that $b^*_{\text{DRNE}}(u)$ can only be computed for $u > u_0$. Therefore, there is no clear guidance on how to bid for $u < u_0$, and this issue becomes more severe when the competing landscape is more competitive, i.e., other bidder's values are high or there is a large number of bidders; see, for example, $N = 4$ case in top-right plot in Figure 1, where u_0 is larger than 6. In fact, the practice is usually very competitive: there are many bidders participating in ad exchanges, and winning only $1\% - 5\%$ of the auctions is a typical scenario faced by demand-side platforms. Therefore, DRNE may not be a useful concept here. On the other hand, GDRNE provides us a robust bidding policy for all $u > r$, which is easy-to-use for practitioners.

See Section Appendix B for the full set of plots.

6. Discussions on Heterogeneous Agents

The previous sections hinge on the homogeneous-agent assumption, which may not be true in practice. However, in this section, we show that Assumption 2 is approximately true and

Theorem 1 can also be applied in practice when the winning probability of each agent is relatively small.

We denote V_{-i} is the highest competing value, i.e.,

$$V_{-i} = \max\{U_j : j \neq i \text{ and } j = 1, 2, \dots, N+1\}.$$

We suppose that the winner probability of each agent is bounded by α, i.e., $P(U_i > V_{-i}) < \alpha$. In practice, α is usually small. Then, for any $i \neq j$, we have

$$P(V_{-i} \neq V_{-j}) > 1 - P(U_i > V_{-i}) - P(U_j > V_{-j}) > 1 - 2\alpha.$$

Therefore, when α is smaller than 5% (a typical scenerio in practice), the highest competing value each agent faced are usually more than 90% similar. We can thus safely say Assumption 2 is approximately true and apply Theorem 1 directly.

However, for GDRNE, the results of Theorem 2 may not be a good approximation in the heterogeneous setting because the equilibrium depends on each agent's own value distribution, which may vary drastically. Therefore, solving GDRNE (approximately) in the heterogeneous remains an open problem.

7. Conclusions and Future Work

We believe that robust transfer learning with strategic behavior introduced in this work is a new area with many research opportunities. A specific setting which, we believe, connects traditional transfer learning questions and strategic behavior involves information dissemination across different social networks platforms in which different agents wish to transfer knowledge with the goal of achieving competing effects (e.g., a group of agents apply transfer learning from network A to network B with the objective of maximizing competing goals, even involving the network manager). We expect that ideas such as our GDRNE will be particularly useful in these and other types of settings.

References

[1] Michele Aghassi and Dimitris Bertsimas. Robust game theory. *Mathematical Programming*, 107(1-2):231–273, 2006.

[2] Chaithanya Bandi and Dimitris Bertsimas. Optimal design for multi-item auctions: a robust optimization approach. *Mathematics of Operations Research*, 39(4):1012–1038, 2014.

[3] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144, 2007.

[4] Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.

[5] Aaron L Bodoh-Creed et al. Ambiguous beliefs and mechanism design. *Games and Economic Behavior*, 75(2):518–537, 2012.

[6] Subir Bose and Arup Daripa. A dynamic mechanism and surplus extraction under ambiguity. *Journal of Economic theory*, 144(5):2084–2114, 2009.

[7] Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weighting. In *Advances in neural information processing systems*, pages 442–450, 2010.

[8] Nicolas Courty, Remi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.

[9] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3):595–612, 2010.

[10] John Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *arXiv preprint arXiv:1810.08750*, 2018.

[11] John C Duchi, Tatsunori Hashimoto, and Hongseok Namkoong. Distributionally robust losses against mixture covariate shifts. *Under review*, 2019.

[12] Rui Gao and Anton J Kleywegt. Distributionally robust stochastic optimization with Wasserstein distance. *arXiv preprint arXiv:1604.02199*, 2016.

[13] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.

[14] Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex J Smola. Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pages 601–608, 2007.

[15] Çağıl Koçyiğit, Garud Iyengar, Daniel Kuhn, and Wolfram Wiesemann. Distributionally robust mechanism design. *Management Science*, 66(1):159–189, 2020.

[16] Dan Levin and Emre Ozdenoren. Auctions with uncertain numbers of bidders. *Journal of Economic Theory*, 118(2):229–251, 2004.

[17] Kin Chung Lo. Sealed bid auctions with uncertainty averse bidders. *Economic Theory*, 12 (1):1–20, 1998.

[18] David G Luenberger. *Optimization by vector space methods*. John Wiley & Sons, 1997.

[19] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, Sep 2018. ISSN 1436-4646. doi: 10.1007/s10107-017-1172-1. URL https://doi.org/10.1007/s10107-017-1172-1.

[20] S Muthukrishnan. AdX: a model for ad exchanges. *ACM SIGecom Exchanges*, 2009.

[21] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 2009.

[22] John G Riley and William F Samuelson. Optimal auctions. *The American Economic Review*, 71(3):381–392, 1981.

[23] Ahtia Salo and Martin Weber. Ambiguity aversion in first-price sealed-bid auctions. *Journal of Risk and Uncertainty*, 11(2):123–137, 1995.

[24] Alexander Shapiro. Distributionally robust stochastic programming. *SIAM Journal on Optimization*, 27(4):2258–2275, 2017.

[25] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.

[26] Nian Si, Fan Zhang, Zhengyuan Zhou, and Jose Blanchet. Distributionally robust policy evaluation and learning in offline contextual bandits. In *International Conference on Machine Learning*, pages 8884–8894. PMLR, 2020.

[27] Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert MÃžller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8 (May):985–1005, 2007.

[28] Alfredo di Tillio, Nenad Kos, and Matthias Messner. The design of ambiguous mechanisms. *The Review of Economic Studies*, 84(1):237–276, 2016.

[29] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

[30] Junfeng Wen, Chun-Nam Yu, and Russell Greiner. Robust learning under uncertain test distributions: Relating covariate shift to model misspecification. In *ICML*, pages 631–639,

2014.
[31] Han Zhao, Shanghang Zhang, Guanhang Wu, Joao P. Costeira, Jose M. F. Moura, and Geoffrey J. Gordon. Adversarial multiple source domain adaptation. In *Advances in Neural Information Processing Systems*, 2018.

Appendix A. Proofs of Main Results

Appendix A.1. **Proof of Theorem 1.**

Proof of Theorem 1. The optimal bids satisfy $b \leq u$. Since $s(\cdot)$ is strictly increasing and continuous, we have

$$\begin{aligned} & \inf_{P \in \mathcal{U}_\delta(\hat{P}_V)} (u-b)\, P\left(s\left(V\right) \leq b\right) \\ = & \ (u-b) \inf_{P \in \mathcal{U}_\delta(\hat{P}_V)} P\left(V \leq s^{-1}\left(b\right)\right) \\ = & \ (u-b)\, \hat{P}_V\left(V \leq s^{-1}\left(b\right) - \lambda_{s^{-1}(b)}\right). \end{aligned}$$

Consider the measure with the CDF

$$\hat{F}_V^\delta(u) = \left\{ \begin{array}{cc} 0 & u \leq u_0 \\ \hat{F}_V(-\lambda_u + u) & u > u_0 \end{array} \right. .$$

Then, by Lemma 1, we have

$$\inf_{P \in \mathcal{U}_\delta(\hat{P}_V)} (u-b)\, P\left(s\left(V\right) \leq b\right) = (u-b)\, \hat{F}_V^\delta(s^{-1}\left(b\right)).$$

When $u > u_0$, it is just the Nash equilibrium with belief $\hat{F}_V^\delta\left(\cdot\right)$. We claim $\hat{F}_V^\delta\left(u\right)$ is strictly increasing with continuous density over $[u_0, \overline{u}]$. Since λ_u is the unique solution of

$$\delta^p = \int_{u-\lambda_u}^{u} (\hat{F}_V(t) - \hat{F}_V(u-\lambda_u))^p \mathrm{d}t, \tag{A.1}$$

we have $u - \lambda_u$ is strictly increasing with respect to u.Then, by differenting both sides of equation (A.1), we have

$$0 = (\hat{F}_V(t) - \hat{F}_V(u-\lambda_u) + p\left(1 - \frac{\partial \lambda_u}{\partial u}\right) \int_{u-\lambda_u}^{u} (\hat{F}_V(t) - \hat{F}_V(u-\lambda_u))^{p-1} \mathrm{d}t,$$

which means $\frac{\partial \lambda_u}{\partial u}$ is continuous. Therefore, we have

$$\frac{\partial \hat{F}_V^\delta\left(u\right)}{\partial u} = \hat{f}_V\left(-\lambda_u + u\right)\left(1 + \frac{\partial \lambda_u}{\partial u}\right),$$

is continuous. Then, by [22], we have the symmetric DRNE is that

$$b^*\left(u\right) = u - \int_r^u \hat{F}_V^\delta\left(x\right) \mathrm{d}x / \hat{F}_V^\delta\left(u\right) \ \text{ for } u \geq \max(r, u_0).$$

Then, for $u \leq u_0$ and for any $b \leq u$, we have $(s^*)^{-1}\left(b\right) \leq u_0$ and thus $\hat{F}_V^\delta((s^*)^{-1}\left(b\right)) = 0$. Therefore, any strictly increasing and continuous function with

$$b^*\left(u\right) = u - \int_{\max(r,u_0)}^u \hat{F}_V^\delta\left(x\right) \mathrm{d}x / \hat{F}_V^\delta\left(u\right) \ \text{ for } u \geq \max(r, u_0),$$

is a symmetric DRNE. □

Appendix A.2. **Proof of Theorem 2.** For any strictly increasing and continuous functions, $b(\cdot), s(\cdot)$, we have

$$\inf_{P:D_p^{s,b}(P,\hat{P}_V)\leq\delta} E_U\left[\{U-b(U)\}\,P\left(s(V)\leq b(U)|U\right)\mathbb{I}\{U\geq r\}\right]$$

$$= E_U\left[\{U-b(U)\}\,\hat{P}_V\left(s(V_1)\leq b(U)|U\right)\mathbb{I}\{U\geq r\}\right]+ \tag{A.2}$$

$$\inf_{P:D_p^{s,b}(P,\hat{P}_V)\leq\delta}\left\{\int_r^{\overline{u}}(u-b(u))f_U(u)\left(P\left(s(V)\leq b(u)\right)-\hat{P}_V\left(s(V)\leq b(u)\right)\mathrm{d}u\right\}. \tag{A.3}$$

From the optimization problem, we have $b(\bar{u})>s(\bar{u})$ cannot be the optimizer. Thus, we impose the condition $b(\bar{u})\leq s(\bar{u})$. Further, when $u=r$, the optimal bids satisfy $b(r)=s(r)=r$. By the Hölder inequality, the term (A.3) is bounded by

$$\int_r^{\overline{u}}(u-b(u))f_U(u)\left(P\left(s(V)\leq b(u)\right)-\hat{P}_V\left(s(V)\leq b(u)\right)\right)\mathrm{d}u$$

$$\geq -\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/q}\left(\int_r^{\overline{u}}\left(P\left(s(V)\leq b(u)\right)-\hat{P}_V\left(s(V)\leq b(u)\right)\right)^p\mathrm{d}u\right)^{1/p}$$

$$\geq -\delta\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/q}. \tag{A.4}$$

where $1/p+1/q=1$. We claim the bound is tight. We can pick a sequence of measure Q_ϵ, such as

$$\hat{P}_V\left(V\leq s^{-1}\left(b(u)\right)\right)-Q_\epsilon\left(V\leq s^{-1}\left(b(u)\right)\right)=t\left((u-b(u))f_U(u)\right)^{q/p},\ \text{ for } u\in[r,\overline{u}],$$
$$\hat{P}_V\left(V\leq s^{-1}\left(b(u)\right)\right)=Q_\epsilon\left(V\leq s^{-1}\left(b(u)\right)\right)\ \text{ when } u\geq\overline{u}+\epsilon,\ \text{ and}$$
$$\hat{P}_V\left(V\leq u\right)=Q_\epsilon\left(V\leq u\right),\ \text{ when } u\in[\underline{u},r].$$

with

$$t\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/p}=\delta.$$

Since $b(r)=s(r)=r$ and $b(\bar{u})\leq s(\bar{u})$, the inversion $s^{-1}\left(b(u)\right)$ is always valid and gives a unique value. Furthermore, There exists $\delta_0>0$, such that for $0<\delta<\delta_0$, Q_ϵ is a valid probability measure, since $s(\cdot), b(\cdot)$ are strictly increasing and continuous, and $f_U(u)$ has bounded derivative. Finally, notice that

$$\inf_{\epsilon>0} -\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/q}\left(\int_r^{\overline{u}}\left(P\left(s(V)\leq b(u)\right)-\hat{P}_V\left(s(V)\leq b(u)\right)\right)^p\mathrm{d}u\right)^{1/p}$$

$$=-\delta\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/q},$$

which completes the tightness of the bound (A.4).

By Theorem 1 in Chapter 7.4 and Lemma 1 in Chapter 7.5 of [18], we have the optimizer must satisfy the condition:

$$(u-b(u))\frac{\hat{f}_V\left(s^{-1}\left(b(u)\right)\right)}{s'(s^{-1}\left(b(u)\right))}-\hat{P}_V\left(s(V)\leq b(u)\right)+\delta\frac{((u-b(u))f_U(u))^{q-1}}{\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1-1/q}}=0. \tag{A.5}$$

In the symmetric GDRNE, $b^*(\cdot)$ solves the differential equation:

$$(u - b(u))\frac{\hat{f}_V(u)}{b'(u)} - \hat{F}_V(u) + \delta\frac{((u-b(u))f_U(u))^{q-1}}{\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1-1/q}} = 0. \tag{A.6}$$

Let $b^*(u) = b_0(u) + \delta b_1(u) + o(\delta)$. When $\delta = 0$, it reduces to the non-ambiguity Nash equilibrium, i.e.,

$$b_0(u) = u - \frac{\int_r^u \hat{F}_V(x)\mathrm{d}x}{\hat{F}_V(u)} \text{ and } b_0'(u) = \frac{\hat{f}_V(u)\int_r^u \hat{F}_V(x)\mathrm{d}x}{\hat{F}_V(u)^2} = \frac{\hat{f}_V(u)}{\hat{F}_V(u)}(u - b_0(u)).$$

Then, the first two terms of (A.6) becomes,

$$\begin{aligned}
&(u - b^*(u))\frac{\hat{f}_V(u)}{(b^*(u))'} - \hat{F}_V(u)\\
&= (u - b_0(u) - \delta b_1(u))\frac{\hat{f}_V(u)}{b_0'(u) + \delta b_1'(u)} - \hat{F}_V(u) + o(\delta)\\
&= (u - b_0(u) - \delta b_1(u))\left(\frac{\hat{f}_V(u)}{b_0'(u)} - \frac{\delta b_1'(u)\hat{f}_V(u)}{b_0'(u)^2}\right) - \hat{F}_V(u) + o(\delta)\\
&= -\delta b_1(u)\frac{\hat{f}_V(u)}{b_0'(u)} - \frac{\delta b_1'(u)\hat{F}_V(u)}{b_0'(u)} + o(\delta)\\
&= -\frac{\delta}{b_0'(u)}\left(\hat{F}_V(u)b_1'(u) + \hat{f}_V(u)b_1(u)\right) + o(\delta)
\end{aligned}$$

Therefore, after ignoring the higher order, (A.6) becomes a differential equation

$$b_1'(u) + \frac{\hat{f}_V(u)}{\hat{F}_V(u)}b_1(u) = \frac{b_0'(u)((u-b_0(u))f_U(u))^{q-1}}{\hat{F}_V(u)\left(\int_r^{\overline{u}}((u-b_0(u))f_U(u))^q\,\mathrm{d}u\right)^{1-1/q}}.$$

Denote

$$h(u) = \frac{b_0'(u)((u-b_0(u))f_U(u))^{q-1}}{\hat{F}_V(u)\left(\int_r^{\overline{u}}((u-b_0(u))f_U(u))^q\,\mathrm{d}u\right)^{1-1/q}}$$

With $b_1(r) = 0$, the solution is that

$$b_1(u) = \int_r^u h(x)\exp\left(-\int_x^u \frac{\hat{f}_V(t)}{\hat{F}_V(t)}\mathrm{d}t\right)\mathrm{d}x. \tag{A.7}$$

We remain to show $b^*(\cdot)$ is indeed the GDRNE, when $\delta \to 0$. Namely $b^*(\cdot)$ is an optimizer of the problem

$$\begin{aligned}
&\max_{b(\cdot)\in\mathcal{IC}([r,\overline{u}])} E_U\left[\{U - b(U)\}\hat{P}_V(s^*(V_1) \le b(U)|U)\,\mathbb{I}\{U \ge r\}\right]\\
&-\delta\left(\int_r^{\overline{u}}((u-b(u))f_U(u))^q\,\mathrm{d}u\right)^{1/q},
\end{aligned} \tag{A.8}$$

where $s^*(u) = s_0(u) + \delta s_1(u) + o(\delta)$ and $s_0(u) = b_0(u), s_1(u) = b_1(u)$. Let denote functionals

$$\begin{aligned}\Psi\left(b\left(\cdot\right)\right) &= E_U\left[\left\{U - b(U)\right\}\hat{P}_V\left(s^*(V_1) \le b(U)|U\right)\mathbb{I}\left\{U \ge r\right\}\right]\\ &\quad - \delta\left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q \mathrm{d}u\right)^{1/q},\\ \Psi_1\left(b\left(\cdot\right)\right) &= E_U\left[\left\{U - b(U)\right\}\hat{P}_V\left(s^*(V_1) \le b(U)|U\right)\mathbb{I}\left\{U \ge r\right\}\right],\\ \Psi_2\left(b\left(\cdot\right)\right) &= \left(\int_r^{\overline{u}}\left((u-b(u))f_U(u)\right)^q \mathrm{d}u\right)^{1/q}.\end{aligned}$$

First, consider $\Psi_1\left(b\left(\cdot\right)\right)$ and we have

$$\begin{aligned}&\frac{\partial\Psi_1\left(b\left(\cdot\right)\right)}{\partial b\left(u\right)}\\ &= (u-b(u))\frac{\hat{f}_V((s^*)^{-1}\left(b(u)\right)}{(s^*)'\left((s^*)^{-1}\left(b(u)\right)\right)} - \hat{P}_V\left(s^*(V_1) \le b(u)\right)\\ &= (u-b(u))\left(\frac{\hat{f}_V((s^*)^{-1}\left(b(u)\right)}{s_0'\left((s^*)^{-1}\left(b(u)\right)\right)} - \delta\frac{\hat{f}_V((s^*)^{-1}\left(b(u)\right)\left(s_1\right)'\left((s^*)^{-1}\left(b(u)\right)\right)}{s_0'\left((s^*)^{-1}\left(b(u)\right)\right)^2}\right)\\ &- \hat{P}_V\left(s^*(V_1) \le b(u)\right) + o(\delta)\\ &= (u-b(u))\left(\frac{\hat{F}_V((s^*)^{-1}\left(b(u)\right)}{(s^*)^{-1}\left(b(u)\right) - s_0((s^*)^{-1}\left(b(u)\right))} - \delta\frac{\hat{f}_V((s^*)^{-1}\left(b(u)\right)\left(s_1\right)'\left((s^*)^{-1}\left(b(u)\right)\right)}{s_0'\left((s^*)^{-1}\left(b(u)\right)\right)^2}\right)\\ &- \hat{P}_V\left(s^*(V_1) \le b(u)\right)) + o(\delta).\end{aligned}$$

Notice that

$$\begin{aligned}&(s^*)^{-1}\left(b(u)\right) - s_0((s^*)^{-1}\left(b(u)\right))\\ &= (s^*)^{-1}\left(b(u)\right) - s^*((s^*)^{-1}\left(b(u)\right)) + \delta s_1((s^*)^{-1}\left(b(u)\right)) + o(\delta)\\ &= (s^*)^{-1}\left(b(u)\right) - b(u) + \delta s_1(u)((s^*)^{-1}\left(b(u)\right)) + o(\delta).\end{aligned}$$

After simplification, we have

$$\begin{aligned}&\frac{\partial\Psi_1\left(b\left(\cdot\right)\right)}{\partial b\left(u\right)}\\ &= \hat{F}_V\left(s^*\right)^{-1}\left(b(u)\right)\left(\frac{u-(s^*)^{-1}\left(b(u)\right)}{(s^*)^{-1}\left(b(u)\right) - b(u)}\right) +\\ &- \delta(u-b(u))\left(\frac{\hat{F}_V((s^*)^{-1}\left(b(u)\right)s_1((s^*)^{-1}\left(b(u)\right))}{\left((s^*)^{-1}\left(b(u)\right) - b(u)\right)^2} + \frac{\hat{f}_V((s^*)^{-1}\left(b(u)\right)\left(s_1\right)'\left((s^*)^{-1}\left(b(u)\right)\right)}{s_0'\left((s^*)^{-1}\left(b(u)\right)\right)^2}\right)\\ &+ o(\delta).\end{aligned}$$

When $\delta = 0$, $s^*(u) = s_0(u)$, due to $(s^*)^{-1}(b(u)) - b(u) > 0$ and the monotonicity of $s^*(u)$ we have

$$\hat{F}_V\,(s^*)^{-1}(b(u))\left(\frac{u-(s^*)^{-1}(b(u))}{(s^*)^{-1}(b(u))-b(u)}\right) > 0, \text{ when } b(u) < s^*(u), \text{ and}$$
$$\hat{F}_V\,(s^*)^{-1}(b(u))\left(\frac{u-(s^*)^{-1}(b(u))}{(s^*)^{-1}(b(u))-b(u)}\right) < 0, \text{ when } b(u) > s^*(u).$$

Therefore, $b(u) = b_0(u) = s_0(u)$ is the unique maximizer.

For $\delta > 0$, To emphasis the dependence on δ, we explicitly write $b_\delta(u) = b(u)$ and let $\hat{b}_\delta(\cdot)$ be an optimizer of problem (A.8). By directly examining the optimization problem (A.8), we have $\hat{b}_\delta(u) \geq s_0(u)$. We claim $\hat{b}(u) \to b_0(u)$, when $\delta \to 0$. If not, since $\hat{b}_\delta(u) \in [r, u]$, we can construct a sequence $\left\{\hat{b}_{\delta_i}(u)\right\}_{i=1}^{\infty}$ such as

$$\hat{b}_{\delta_i}(u) \to \check{b}(u) \neq b_0(u).$$

Then,

$$\frac{\partial \Psi_1(b(\cdot))}{\partial \hat{b}_{\delta_i}(u)} \to \hat{F}_V\,(b_0)^{-1}(\check{b}(u))\left(\frac{u-(b_0)^{-1}\left(\check{b}(u)\right)}{(b_0)^{-1}\left(\check{b}(u)\right)-\check{b}(u)}\right) \neq 0,$$

which means there exists $\delta_i > 0$, such as $\frac{\partial \Psi_1(b(\cdot))}{\partial \hat{b}_{\delta_i}(u)} \neq 0$. It is a contradiction.

Now, for any $b_\delta(u) = b_0(u) + o(1)$, we have

$$\begin{aligned}\frac{\partial \Psi_1(b(\cdot))}{\partial b_\delta(u)} &= \hat{F}_V\,(s^*)^{-1}(b_\delta(u))\left(\frac{u-(s^*)^{-1}(b_\delta(u))}{(s^*)^{-1}(b_\delta(u))-b_\delta(u)}\right)\\ &\quad - \frac{\delta}{s_0'(u)}\left(\hat{f}_V(u)s_1(u) + \hat{F}_V(t)s_1'(u)\right) + o(\delta).\end{aligned}$$

Therefore, we have

$$\begin{aligned}&\frac{\partial \Psi_1(b(\cdot))}{\partial b_\delta(u)}\\ &= \hat{F}_V\,(s^*)^{-1}(b_\delta(u))\left(\frac{u-(s^*)^{-1}(b_\delta(u))}{(s^*)^{-1}(b_\delta(u))-b_\delta(u)}\right) + o(\delta)\\ &= \hat{F}_V\,(s^*)^{-1}(b_\delta(u))\left(\frac{u-(s^*)^{-1}(b^*(u)) - \frac{b_\delta(u)-b^*(u)}{(s^*)'(b^*(u))}}{(s^*)^{-1}(b_\delta(u))-b_\delta(u)}\right) + o(b_\delta(u)-b^*(u)) + o(\delta)\\ &= -\frac{\hat{F}_V\,(s^*)^{-1}(b_\delta(u))}{\left((s^*)^{-1}(b_\delta(u))-b_\delta(u)\right)(s^*)'(b^*(u))}(b_\delta(u)-b^*(u)) + o(b_\delta(u)-b^*(u)) + o(\delta).\end{aligned}$$

For $\epsilon > 0$, there exists $\delta_0 > 0$, such that for $0 < \delta < \delta_0$, we have

$$\frac{\partial \Psi_1(b(\cdot))}{\partial b_\delta(u)} > 0, \text{ when } b_\delta(u) < b^*(u) - \epsilon\delta, \text{ and}$$
$$\frac{\partial \Psi_1(b(\cdot))}{\partial b_\delta(u)} < 0, \text{ when } b_\delta(u) > b^*(u) + \epsilon\delta.$$

Therefore, we have the optimizer $\lim_{\delta\to 0}\frac{1}{\delta}\left(\hat{b}_\delta(u) - b^*(u)\right) = 0$, which completes the proof.

Appendix B. Numerical Examples

Appendix B.1. **Maximum of N Uniform Random Variables.** We consider the same parameters in Example 2. Here, we provide a full grid of plots for $(N \times \delta) \in \{1, 2, 3, 4\} \times \{0.1, 0.2, 0.3, 0.4\}$.

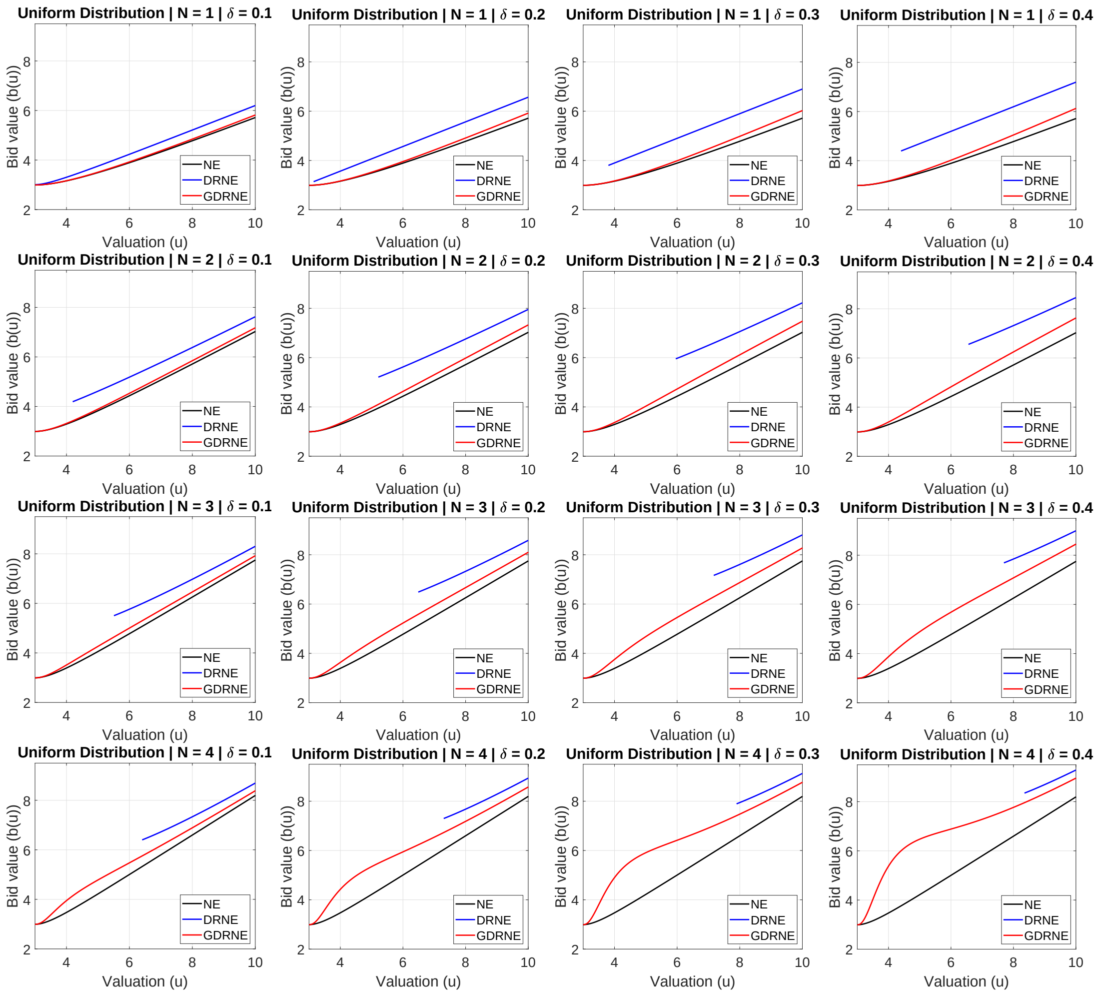

Figure 3: NE, DRNE, and GDRNE for different $N \in \{1, 2, 3, 4\}$ and $\delta \in \{0.1, 0.2, 0.3, 0.4\}$, where the nominal distribution is the maximum of N i.i.d. uniform random variables.

Appendix B.2. **Maximum of N Scaled Beta Random Variables.** We consider the same parameters in Example 3. Here, we provide a full grid of plots for $(N \times \delta) \in \{1, 2, 3, 4\} \times \{0.1, 0.2, 0.3, 0.4\}$.

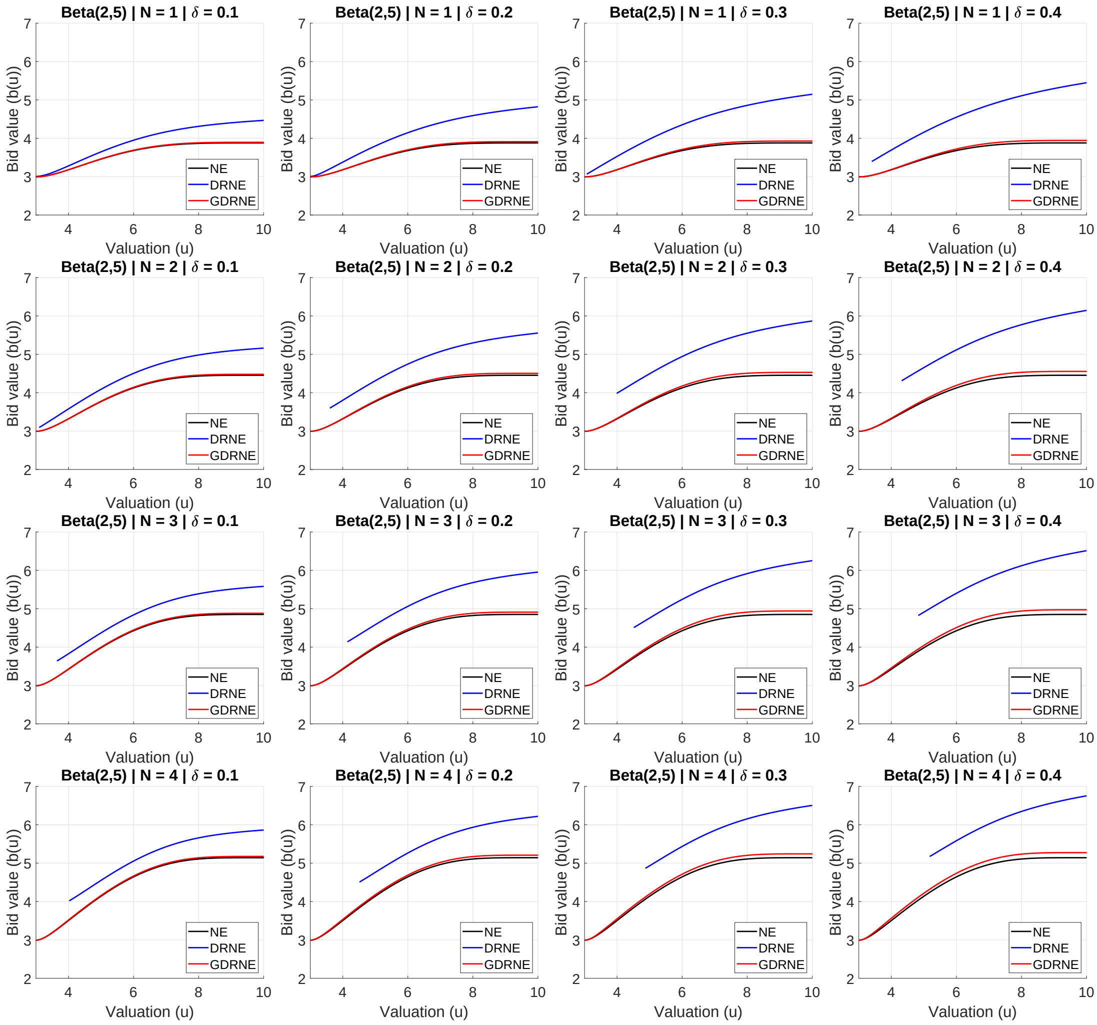

Figure 4: NE, DRNE, and GDRNE for different $N \in \{1, 2, 3, 4\}$ and $\delta \in \{0.1, 0.2, 0.3, 0.4\}$, where the nominal distribution is the maximum of N i.i.d. scaled beta random variables.