Credibility and Incentives in Gradual Dutch Auctions

Kshitij Kulkarni
UC Berkeley
ksk@eecs.berkeley.edu

Matheus V. X. Ferreira
Harvard University
matheus@seas.harvard.edu

Tarun Chitra
Gauntlet
tarun@gauntlet.network

January 2023

Abstract

Gradual dutch auctions (GDAs) are a class of auctions that have been proposed when an auctioneer would like to sell a batch of illiquid items. They function by making available a fraction of the items for sale at every time, and starting a new auction whose price decays as time passes. This has the effect of allowing a seller to cater to buyers who may not want to purchase the entire batch of items at a single time. We analyze the incentives of participating in GDAs. First, we consider the seller's incentives in running a GDA. We show that the seller can deviate from truthfully running a GDA via an attack in which she initially buys a fraction of the supply available in each dutch auction, forcing buyers to fill their demand with later (more expensive) auctions. This attack is a form of multi-block maximal extractable value (MEV), extending previous work on lending protocols and decentralized exchanges. Next, we consider buyers' incentives in participating in a GDA, and consider an *interdependent* values setting in which the history of the auction is allowed to affect buyers' values in the future. We show conditions in which GDAs are ex post incentive compatible and individually rational for buyers.

1 Introduction

Non-Fungible Tokens (NFTs) are a popular form of digital asset on blockchains that have had a large uptake in usage since late 2020. These assets have been used to represent digital art, collectibles, blockchain representations of real-world assets, and claims on future returns from decentralized finance protocols [2, 47]. One of the key features of NFTs is that a single collection of NFTs can have a fixed, discrete supply that cannot be modified post deployment. This allows for the representation of scarce, indivisible digital objects, which lie in contrast to fungible tokens which can be split and recombined arbitrarily (up to the minimum asset size).

Decentralized trade and exchange mechanisms for these assets, however, can be quite a bit more complex than those of fungible tokens. For fungible tokens, mechanisms such as constant function market makers (CFMMs) provide decentralized means of trade that are synchronized via arbitrage [4, 5, 6]. These mechanisms utilize continuous supply and demand curves, which only make sense for arbitrarily divisible assets, such as fungible tokens.[1] NFTs, much like fine art, are often better suited to direct auction mechanisms versus continuously traded vehicles. This is because the restricted supply of NFTs effectively places a cap on the frequency of transfer of an NFT while also ensuring that there is a well-defined minimum value for a collection of NFTs (the so-called 'floor price').

The lack of continuous trading of NFTs, however, makes price discovery significantly more complex [28, 40, 46]. When a creator of an NFT collection consisting of n items wants to sell these items, they have a number of different auction mechanisms to choose from. For instance, the creator could sell each item individually and in a sequential fashion or try to sell arbitrary sub-bundles as atomic units. As is well-known in combinatorial auctions, this space of possible auctions for multi-item exchange is exponential in the number of items to sell and without a careful choice for the bidding language, could be #P-complete [36, 43]. Moreover, it has been demonstrated that collusion resistant auctions have impossibility results for even single item NFTs auctions [31]. As such, it is also difficult for bidders to accurately participate in the price discovery process (which is perhaps one hypothesis for the excess wash trading found in the NFT market [46]).

Ideally, one would want to prove a stronger property for NFT auctions, that of credibility [3], where it is incentive compatible for the auctioneer to follow the stated auction rules. The results of [31] suggest that it will be difficult to get perfectly credible mechanisms for NFT auctions. However, approximately credible mechanisms and posted price mechanisms have been studied in the blockchain context before when analyzing fee auctions [14, 17, 16]. One goal of this paper is to provide examples of seller incentive compatibility in NFT auctions in order to construct credible (or approximately credible) NFT auctions that have properties similar to those of fee auctions.

Practical Auction Mechanisms. Combined, these results show that designing approximately optimal mechanisms for computationally bounded auctioneers (*e.g.* sellers of an NFT collection) and bidders is tantamount to making price discovery efficient. There have been a number of proposals for such auctions, such as the equilibrium-truthful auction of Milionis, et. al [31] and the Gradual Dutch Auction (GDA) [18]. The former auction relies on traditional auction theory [21], where one makes the assumption that the bidders valuation distributions are independent and that the auctioneer has oracle access to bidders' valuation distributions.[2] In practice, this is a difficult feat to accomplish, even with historical data, as NFT auction data is extremely sparse (especially when compared with online ad auction

[1]We note that while one can fractionalize an NFT (*e.g.* create 100 shares that each own 1% of an NFT), there is often a loss in governance and/or voting rights when this is done. Suppose that an NFT refers to a share of a stock that is allowed to vote on board member selection. Technically, only the address holding NFT has the right to vote and the fractional shareholders often do not get this right passed along to them. As such, price discovery for the fractional shares and the underlying NFT can be quite different, as noted here [34].

[2]Note that there has been some work on inferring the ironed (concave majorant) valuation distributions from sequential data [41]. These results are difficult to apply to NFTs as they have poor sample complexity.

data). As such, auctions with heuristic guarantees, such as the GDA, have been preferably implemented in practice on networks such as Ethereum.[3]

Moreover, another issue with applying traditional auction design to the blockchain world is the lack of true private valuations. The classical Vickrey and/or Myerson optimal auctions assume that users of a system have distinct private valuation distributions that are sampled independently. In the blockchain environment, however, a user's valuation can be far from private as the user has to submit transactions through the network, leaking some information about the their valuation prior to their bid being processed by the smart contract running the auction. This is further complicated by the fact that sequential auctions (like those implemented by GDAs) ask bidders to purchase items over time; a buyer that enters later in the auction therefore may use the public history of the auction to modify their valuation. Milionis, et. al [31] points out this issue for NFT auctions, but provides a model that is analyzed in the private valuation world. Moreover, the fact that NFT valuations often have a common value component [25, 48] suggests that auctions deemed optimal in the private, independent valuation model are insufficient for practice auctions with guarantees.

Finally, blockchain environments are communication and bandwidth constrained. This means that iterative auctions with many rounds of interaction are likely to fail. For instance, MEV bots can snipe early bids (which was mentioned in the original GDA post [18]) and dramatically change the revenue and incentives of an auction. Dutch auctions (and descending price auctions, in general) notably have reduced communication complexity relative to both ascending price and sealed-bid auctions. The communication complexity of allocation mechanisms in which many items need to be allocated to multiple agents (roughly the size of the messages any agent must send to the mechanism designer) can be exponential in the number of items when the allocation mechanism is direct revelation, in which the mechanism designer seeks to elicit truthful reports of agents' valuations) [37]. Even in the case of an ascending price auction, the auctioneer needs to query every bidder in every round to ask whether the bidder would like to increase their bid to stay in the auction, which can be prohibitive in decentralized systems. Therefore, buyers in an ascending price auction may need to speak an unbounded number of times. Alternatively, in descending price auctions, the auctioneer can simply publicly broadcast the current price of the auction to all the bidders. While all auctions can be shown in different communication models to have worst case exponential communication complexity, indirect revelation mechanisms (*e.g.* mechanisms where a buyer's bidding behavior depends on the entire price process and bid sequence) tend to have much lower communication complexity [24, 39]. Moreover, Vickrey-Clarke-Groves optimality results can be extended to descending price auctions [32] and one can view our results in this paper as an extension of this work. Finally, we note that [16] also notes that descending price auctions are significantly safer to design in blockchain environments (and is one reason fee mechanisms resemble descending price auctions).

Gradual Dutch Auctions. The Gradual Dutch Auction (GDA) was introduced in [18] as a mechanism for allowing sellers to not have to price bundled items all at once. As an

[3]There are daily auctions using GDAs with the total auction volume generating $4-5M of revenue for

example, suppose that an NFT collection has 10,000 elements for sale. Simultaneously selling each of the 10,000 items works as a clearing mechanism if there are at least 10,000 items demanded by the market. If there are less than $X < 10,000$ items demanded (as is the case for illiquid assets), then the sequential auction is suboptimal, as the expected clearing price of the auction will decay after X items are sold. In such a case, the seller (auctioneer) will benefit from bundling items, *i.e.* selling sets or bundles of items atomically. However, the question of how to choose the size of the bundle(s) to sell is a difficult question (indeed, at least NP-hard in many scenarios [36]). GDAs can be viewed as a heuristic mechanism that allows *buyers* (as opposed to the seller) to select the size of bundles to purchase.

The main idea behind GDAs is that they auction off each item sequentially, but with a dynamic price that updates as a function of how long a set of items has not been sold. We can view a GDA as constructing a sequence of auctions $A_1, \ldots, A_n$ such that the price tendered by auction A_i is $a_i \in \mathbf{R}_+$ at the time of auction initialization t_i. Moreover, the price of the ith auction at time t, $p_i(t)$, satisfies two conditions:

1. *Price Initialization.* $p_i(0) = a_i$
2. *Strictly Decaying Price.* $p_i(t) < p_i(t')$ for all $t, t' > t_i$ and $t > t'$.
3. *Strictly Increasing Initialization.* $p_i(0) = a_i < a_j = p_j(0)$ if $i < j$

In a sense, the price is monotonically decreasing for an individual auction as $t \to \infty$, similar to a classical dutch auction. However, the initial price of each auction, a_i, can depend on the price of previous auctions $a_1, \ldots, a_{i-1}$ and is usually chosen such that a_i is a strictly increasing sequence of prices.

There are a few different ways to view GDAs that are helpful for constructing some intuition for how they work. First, one can view GDAs as encoding a Bayesian prior of users' valuation and bundle size distributions within the auction mechanism. The particular prices that the auction initializes to and the rate of decay of price over time represents a weak prior distribution on the demand of the bidders in the auction. Once can view the GDA as analogous to a Bayesian posted price mechanism albeit with decaying posted prices (which, for a single auction and single bidder, is not strategyproof [15, 21, 22]). For completeness, we note that descending price clock auctions (which have been proposed for spectrum auctions) have been analyzed [35], but their strategyproofness (or lack thereof) was not established. GDAs are also tightly related to the literature on efficient dynamic auctions, which considers incentive compatibility conditions very similar to ours, albeit with more general price mechanisms [8, 7].

Secondly, we can view a GDA as a form of a descending price candle auction. Candle auctions are auctions with random stopping times [19], where the auction terminates randomly. Since an auction terminates the moment a bid arrives (akin to a posted price mechanism), the stopping time of the auction is equivalent to the bid arrival time. Each descending price auction A_i can be viewed as a candle auction that terminates whenever a bidder places a bid. In the case of random, but strategic bidders, this will be stochastic like the original candle auction. As a simple model for a candle auction, one may assume that buyers arrive at the auction according to a Poisson process, and buy up the cheapest auctions. Then, with high probability, every auction will eventually terminate (since the expected bid arrival time is finite for homogeneous Poisson processes [44]). On the other hand, our attack in §4 shows that the auctioneer can act as an *adversarial* (and not stochastic) player who manipulates the final clearing auction price by buying up early auctions.

Note that the interaction between auctions makes the model of bidders' valuation functions significantly different than a classical dutch auction, and in particular, cause the bidders' valuations to be interdependent, both on each other and on the history of the auction. We will first illustrate this with an example:

Example of interdependent values in GDAs. Consider a GDA with the following pricing:

- $t = 1$: Offering a single item in the first auction with $a_1 = p_1(0) = \$1$
- $t = 2$: Offering a single item in the first auction for $p_1(2) = \$0.10$ and a single item in the second auction for $a_2 = p_2(0) = \$1.50$
- $t = 2.5$: Offering a single item in the first auction for \$0.05 and a single item in the second auction for \$1.05

Now consider two bidders. They have nominal valuations for the items in the first round of \$0.01 and \$1 (that are unknown to each other). The first bidder has the following contingent valuation: if bidder 2 purchases the first auction in round 1 at \$1, then bidder 1 raises his value for the item in round 2 to \$1.05. Else, he keeps his value the same. Bidder 2 has a contingent valuation that if bidder 1 buys in round 1, he will raise his value to \$1.5. In round 1, bidder 2 buys at \$1, and bidder 1 raises to \$1.05. Now, in round 2, neither buyer buys because both values are less than \$1.5. Finally, in round 2.5, bidder 1 buys the second auction at \$1.05. By the virtue of the presence of bidder 2, bidder 1 has incurred a 'regret' of \$0.05. He could have bought the first auction for \$1, but instead chose to buy the second auction for a higher price because his valuation depended on the other party buying.

Auctions with Interdependent Valuations. The theory of auctions with interdependent valuations was first established in the systematic study of Milgrom and Weber in 1982 [30]. Most modern studies of interdependent valuations still follow Milgrom and Weber's model, where each bidder's valuation v_i is viewed as being a function of some signals $s_i \in \mathbf{R}$ for $i \in [k]$, the number of bidders. These signals, in the case of art auctions for a single item, for example, represent private information that each bidder has about the value of the art and the function $v_i(s)$ is the ith bidder's valuation for an item given signals s. In the case where the auctioneer is selling many identical items at different prices over time (as is the case in GDAs), of which buyers can pick bundles (or amounts to buy), this

auctioneers and resales of auctioned items generating \$50-60M in trading volume since November 2022 [38] interdependence can be even more pronounced, as in the above example. In this case, each bidder also sees the history of the auction, and can use this publicly available information to alter her valuation. For instance, in the example above, along with their private signals s_i, each bidder at time 2.5 might see indicators $\hat{s}_j$ of whether the jth auction up to time 2 was purchased or not, *e.g.* $\hat{s}_j \in \{0, 1\}$. In this case, the valuation of bidder 1 for the items, $v_1(s, \hat{s})$ is given by the following:

$$v_1(0.01, 1, 1, 0) = 1.05$$
$$v_1(0.01, 1, 0, 0) = 0.01$$
$$v_2(0.01, 1, 1, 0) = 1.5$$
$$v_2(0.01, 1, 0, 0) = 1$$

The signals encode the conditional dependence between the bidders in an interdependent auction.

Interdependent auctions are known to have unique clearing prices[4] when what is known as the *single-crossing condition* is satisfied [42]. Moreover, with the single-crossing condition and a matroid condition, one can approximate the optimal revenue in an interdependent auction via a modified VCG mechanism [9, 27]. Colloquially, the single-crossing condition states that a user's valuation changes more from their own signals than those of any other bidders. Single-crossing conditions therefore allow the multi-dimensional interdependent valuation setting to be reduced to something akin to the one-dimensional private values world. GDAs turn out to have conditions analogous to single-crossing conditions that depend on the price trajectories $p_i(t)$ (See §2).

This Paper. There are two natural questions that one might wish to ask about GDAs. First, from the perspective of the buyer, is it preferable to purchase an item at the true value that the buyer has for the item? This is the question of *incentive compatibility*. We ask this question in two settings: one, where a buyer is allowed to wait in a GDA for arbitrarily long periods of time, and another, where buyers sequentially arrive to each round of the auction, and are asked to either buy or not buy the currently available items, and then leave the auction. Second, from the perspective of a seller, is it preferable to run a GDA in earnest, without participating in the auction, or otherwise manipulating information to bidders? This is the question of *credibility*. To summarize, we ask two questions about GDAs that we center the subsequent discussion on:

Q1: Are GDAs incentive compatible for buyers?
Q2: Are sellers in GDAs credible?

Our Results. In this paper, we formalize discrete and continuous time GDAs. We first consider the scenario where one is selling many copies of the same item. In such a setting,

[4] Note that [12, 13] are able to analyze interdependent auctions without the single crossing condition but show that single crossing is necessary for an exact optima. They provide means for approximate equilibria

the sequence of GDA auctions should be viewed as creating a supply curve, $p_S(q, T)$ that provides the prices that one would need to pay to get q units of the item. We then first analyze the expected revenue of the auction for the exponential price curves of [18] (which we term *exponential GDAs*).

By analyzing simpler auctions explicitly, like exponential GDAs, we can discern what properties are needed for generic GDAs. This leads to the construction of an attack against exponential GDAs where the auctioneer can artificially inflate the price paid by bidders by purchasing the first r auctions. We construct a profitability condition for this attack (which depends on the demand distribution $p_D(q)$ and demonstrate a geometric condition that relates r to $p_D(q)$. Note that we assume that $p_D(q)$ is time-independent and leave analysis of dynamic demand functions for future work. In particular, we show that the maximum gradient of $p_D(q)$ controls the minimum number of rounds that the auctioneer needs to purchase in order to be profitable. Finally, we conclude by showing that if the auctioneer runs an infinite number of auctions, this attack ceases to be profitable.

We note that to implement our attack in practice, an adversarial auctioneer will need to utilize maximal extractable value (MEV) auctions for multiple blocks to ensure their early transactions are not front run. These multi-block forms of MEV have been studied before in on-chain oracles [29] and on-chain lending [10]. There are additional costs realized by the auctioneer for entering these auctions which are not discussed here and left for future work.

Subsequently, we generalize exponential GDAs (*e.g.* where $p_i(t)$ is a sum of exponentials with increasing initial prices) to the setting of more generic pricing functions $p_i(t)$. These single-dimensional auctions dramatically expand the space of GDAs beyond the exponential pricing setting (which was described in the initial GDA post [18] and in the subsequent post on variable rate GDAs [1][5]. In this setting, we find sufficient conditions for a GDA pricing trajectories $p_i(t)$ to avoid the attack that is described for exponential GDAs. Moreover, we are able to generalize the condition for the number of rounds that a malicious auctioneer would need to cheat to the generic pricing function setting. This expansion of the space of single-dimensional GDAs suggests that the family of single-dimensional GDAs is a tuneable family of interdependent valuation auctions that can provide some credibility guarantees.

With these generalized GDAs, we are able to analyze buyers' incentives in an interdependent values setting, where the history of the auction that a buyer sees may affect the value of the currently available items to that buyer. In particular, we map generalized GDAs to the set of interdependent VCG auctions and utilize results of [9, 27] to show that GDAs achieve ex post incentive compatibility and individual rationality. These conditions involve a single-crossing condition and the construction of sufficient conditions for the set of winners in a GDA to form a matroid.

Our results generalize the known understanding of GDAs, analyze their security (and demonstrate attacks that remove auctioneer credibility), and generalize to the multi-dimensional

without single crossing; this paper will focus on exact and not approximate equilibria.

[5] In the language of the sequel, variable rate GDAs simply correspond to irregular auction start times, *i.e.* $t_i - t_j \neq C|i-j|$. The blog post [1] uses a square-root time scale, e.g. $t_i - t_j = \Theta(\sqrt{i-j})$ and logistic time scale $t_i - t_j = \Theta(\frac{1}{1+e^{-|i-j|}})$, both of which are intended to slow down the frequency of auction while ensuring that the initial price increases.

setting. As a heuristic auction mechanism that is easy to implement in practice, GDAs are promising as a mechanism for blockchain-based auctions. We hope our results expand the universe of mechanisms that can be used safely on-chain.

2 Model and Background

Our setting begins with an auctioneer, and m indivisible identical items for sale. A new buyer enters the auction at every discrete time i. Each buyer has a private signal $s_i \in \mathbf{R}$. We assume s_i is drawn from a distribution F_i. Signals represent the private information a buyer users are able to discern about other participants in order to construct their valuation. Such models of private signals are common in art auctions (auctions which are similar to selling NFTs), where s_i might be the private value of a piece of art to a bidder. Upon realizing that other bidders have a high signal for a piece of art, a bidder infers a higher value for the art.

Buyer Model If bidder i has complete information about the signal $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$ for all other bidders, then bidder i has value $\tilde{v}_i(s_i, s_{-i}) \in \mathbf{R}$ for the item. In practice, bidder i knows s_i, but only observes a proxy for s_{-i}. In GDA, the proxy for s_{-i} will be the price paid by other bidders. To be precise, let $p_{j,k} \in \mathbf{R}$ be the price agent j paid for item k with $p_{j,k} = 0$ if agent j paid nothing. Then the value of bidder i for the item is $v_i(s_i, p)$ where $p \in \mathbf{R}^{n,m}$ is the matrix of payments.

Example 1. *Continuing the above analogy to art auctions, one may be interested in situations where the history of the auction affects the valuation of bidders for items that are identical. In this case, the valuation might take the form* $v_i(s_i, p) = s_i + \frac{1}{n-1} \sum_{j=1}^{n} \sum_{k=1}^{m} p_{j,k}$. *In this case, the value of the bidder for the item equals their private signal plus the average price paid by other bidders.*

Further, the buyer a time n has available to her all the auctions that have started up to time n that have not been purchased by previous buyers. Crucially, as observed in NFT and other art markets [42, 48], we seek to model the fact that a buyer's valuation for the items may depend on both the *public* history of the auction (the purchase prices of items in the past) and the on *private* beliefs that other buyers have about the value of the items. This leads us to consider an interdependent values model for the buyers.

Auctioneer Model The auctioneer will sell the items sequentially, at prices for each item that decay over time. We therefore denote discrete time by $n \in \mathbf{N}$ and continuous time by $t \in \mathbf{R}_+$. At every $n = 1, \ldots, N$, the auctioneer makes available one of the items for sale, with a predefined and public price function for each item $p_n(t)$, which denotes the price of the item t units of time after the discrete time n. The item is available until a buyer purchases it. An example price function is the exponential price that will be studied in the following section, which takes the form:

$$p_n(t) = k\alpha^n e^{-\lambda t}$$

where $k\alpha^n$ is the initial price of the nth auction, and λ is a decay rate. This price function defines a discrete-time GDA. Intuitively, this means that the sequence of auctions get more expensive over time, but each have a price that decays at the same rate as the time since each auction's start passes. One can significantly generalize this price function to generic GDAs, which incorporate a large class of price functions:

Generic GDAs. A discrete time, single-dimensional GDA is defined by a sequence of pricing functions $p_i : \mathbf{R} \to \mathbf{R}_+$, a sequence of auction times $t_i \in \mathbf{N}$ that is distinct (*e.g.* $t_i \neq t_j$ for all $i, j \in \mathbf{N}$). We will define three properties of a sequence of pricing functions that will allow us to prove results similar to the exponential GDA. A sequence of pricing functions $\{p_i\}_{i\in\mathbf{N}}$ is *admissible* to a set of auction times $\{t_i\}_{i\in\mathbf{N}}$ if for all $i \in \mathbf{N}$, $p_i(t) = 0$ for all $t < t_i$ and $p_i(t) < p_i(t')$ for all $t > t' > t_i$. We say that a sequence of pricing functions is *adapted* if for all finite subsets $A, B \subset \mathbf{N}$ such that $A \cap B = \emptyset$ and $\max A \leq \min B$ we have

$$\min_{i\in A} p_i(t) \leq \min_{i\in B} p_i(t)$$

for all $t \geq 0$. Finally, we say that a sequence of pricing functions is *time translation invariant* if there exists a function $\hat{p} : \mathbf{R} \to \mathbf{R}_+$ and values $a_i \in \mathbf{R}_+$ such that for all $i \in \mathbf{N}$ and $t > 0$ we have

$$p_i(t) = a_i \hat{p}(t - t_i)$$

Perpetual Buyers. First, we present a model of buyers that manifestly leads to GDAs that are not incentive compatible, and is hence a negative result for Q1. A single buyer enters a GDA at time T, and has the choice of waiting for arbitrarily long periods of time without purchasing an auction. We call such a buyer *perpetual*. Suppose the buyer has a fixed valuation $v \in \mathbf{R}$ for one item. Then, by the admissibility of prices for generic GDAs, we have $p_T(t') < p_T(t)$ for any $t' > t$. Therefore, exists a time t' large enough such that $p_T(t') < v$. The buyer, having no constraints on how long she is allowed to wait in the auction, can therefore purchase the item for a price strictly smaller than her value. This implies that GDAs are not incentive compatible for perpetual buyers. A potential remedy for this is to have a random stopping time for each item sold in a GDA, which forces buyers to not wait arbitrarily long before purchasing the items. We leave this as future work.

Defining the Mechanism. Having defined the signals and the information environment for the buyers, we now formally define GDAs as mechanisms, comprising of an action space for each buyer, and allocation and pricing mechanisms:

1. *Allocation mechanism.* The functions $x_{ij} : \mathbf{R}^{i\times j} \to [0, 1]$ returns the probability of an item being allocated to the ith bidder in the jth auction. The allocation rule $x_{i,j}(s_i, s_{-i})$ denotes the function with the ith bidder's signal first. An allocation rule is deterministic if $x_{i,j} \in \{0, 1\}$ for all i, j.

2. *Actions.* For any buyer at time n, we define the action space recursively as $A_n =$

$A_{n-1} \cup \{0,1\} \setminus S_{n-1}$ and $A_1 = \{0,1\}^1$, $S_1 = \{0,1\}$ if $x_{i1} = 1$ for some $i \in B_1$ and $S_n = \{0,1\}$ if $x_{in} = 1$ for some $i \in B_n$. That is, bidders at time n have the option to either buy or not buy all the auctions that have started up to time n, which have not been purchased by other bidders in the past.

3. *Pricing Mechanism.* A pricing mechanism $p_{ik}(s_i, s_{-i})$ maps signals to the price tendered to the ith bidder for item k. Note that for GDAs, our pricing mechanism is deterministic and follows a GDA's pricing curve, *e.g.* $p_{ik}(s_i, s_{-i}) = a_{ik}p_{ik}(t(s_i, s_{-i}) - t_i)$, where $t(s_i, s_{-i})$ is the time that the auction is purchased by the ith bidder given the signals $s = (s_i, s_{-i})$.

Desiderata for Sellers. We analyze the seller's incentives in the next section via the *inverse demand function*, denoted by $p_D(q)$ of a buyer. For a single ephemeral buyer that enters a GDA at time T, this function denotes the willingness-to-pay of the buyer for each quantity of the items that the buyer purchases. This function is implicitly a function of the valuation of the bidder for the items available at that time. The question, then, for the seller, is: if the inverse demand of a buyer is fixed and known to the seller, is it always in the seller's best interest to run a GDA in earnest? We provide a negative result by constructing an explicit deviation in which the seller initial buys some of the supply of the items, and leaves the buyer to purchase later items and is able to extract more revenue from the auction than by running a GDA naively.

As a concrete example of $p_D(q)$, suppose there are T items available (one for each auction started since $n = 1$), and the buyer has values $v_i \in \mathbf{R}_+$ for each item, $i = 1, \ldots, T$ (these values are therefore ordered from oldest available auction to the most recent one). Then, we can define the inverse demand function $p_D(q) = \sum_{i=1}^{q} v_i$. In what follows, we do not explicitly specify the dependence of the demand on valuation in a parametric manner. Instead, we analyze the inverse demand function $p_D(q) = p_D(q|v_1, \ldots, v_n)$ only based on global properties such as Lipschitz gradients and/or convexity. The corresponding demand curve is its inverse denoted $q_D(p) = q_D(p|v_1, \ldots, v_n)$. Notationally, we will supress the explicit dependence of the demand function on valuation. We analyze the inverse demand function to simplify the analysis of seller incentives, as the seller's disincentives to participate honestly in the auction are a function of the aggregate behavior rather than individual value functions.

Desiderata for Buyers. An allocation mechanism x_{ij}, valuation functions v_i, and pricing mechanism p_i are said to be *ex post incentive compatible* if

$$x_{in}(s, s_{-i})v_i(s, s_{-i}) - p(s, s_{-i}) \geq x_i(\tilde{s}_i, s_{-i}, \hat{s})v_i(s, s_{-i}) - p_i(\tilde{s}_i, s_{-i})$$

holds for all i, true private and public signals s and $\hat{s}$, and false (adversarially reported) signals $\tilde{s}$. An allocation mechanism is said to be *ex post individually rational* if for all i and signals s the following holds:

$$x_{in}(s, s_{-i})v_i(s, s_{-i}) - p(s, s_{-i}) \geq 0$$

Informally, in our setting, a mechanism is ex post incentive compatible if no agent can increase their welfare (measured by their expected value from allocation less the price they pay) by misreporting a signal and is ex post individually rational if the agent does not incur negative utility as a result of participating in the auction, having seen the prices that items were purchased for in the past. This implies that agents must truthfully report their private values for the items, even if the price history of the past suggests that many bidders have bought at higher prices.

The standard dutch auction is known to be revenue equivalent to the sealed-bid first price and second price auctions (a result going back to Vickrey [45]). However, as mentioned before, the interdependence inherent in buyer valuations in NFT markets on public blockchains and the correlation between the many auctions in a GDA makes this analysis insufficient. One natural question to ask is if it is possible for GDA's to be incentive compatible and individually rational. We will briefly review the necessary auction theory to state this result and then demonstrate that GDAs can be constructed to be *ex post* incentive compatible and individually rational. Further details on the precise definitions and properties of ex post properties can be found in [42]. Note that we will write functions of bidders private data as functions of the vector (s_i, s_{-i}) where s_i is the ith bidder's private data and s_{-i} is the private data of all other bidders.

3 Exponential and Generic GDAs

The GDA with an exponential pricing function was first introduced in [18] and (partially) formally analyzed in [33]. The analyses of these papers is incomplete from an auction theory standpoint as there is no analysis of a malicious auctioneer who manipulates the price by purchasing early auctions. In this section, we formally define the supply and demand curves proffered by a GDA and then analyze such an attack by a malicious auctioneer which deems the exponential GDA auction not credible in the sense of [3].

3.1 Discrete Time

To recap, a discrete time GDA is an auction in which a seller would like to set illiquid items like non-fungible tokens (NFTs) that must be sold in integer quantities. In discrete time GDAs, the auctioneer starts an auction to sell a single item at every discrete time $n = 1, 2, \dots$, and the price of that auction t seconds in the future is given by the price function $p_n(t)$. The exponential price function that will be studied in this section takes the following form:

$$p_n(t) = k\alpha^n e^{-\lambda t}$$

where $k\alpha^n$ is the initial price of the nth auction, and λ is a decay rate. Intuitively, this means that the sequence of auctions get more expensive over time, but each have a price that decays at the same rate as the time since each auction's start passes.

Inverse Demand In this section, we assume that the auctioneer is faced with a single buyer with a fixed and known inverse demand curve, which we denote by $p_D : \mathbf{N} \to \mathbf{R}_+$. This is a function of the quantity of items, q, where $p_D(q)$ denotes the price that the buyer is willing to purchase q items at. Recall from Section 2 that this inverse demand curve can be constructed from the valuation of the buyer for the items. We further make the assumptions that the inverse demand curve is convex, differentiable, non-increasing, and has Lipschitz gradients. These assumptions ensure that if the quantity of items the buyer purchases changes, the price the buyer is willing to purchase the new amount of items at bounded by a linear factor. This ensures that the auctioneer is able to extract sufficient revenue from the auction. We denote the demand curve corresponding to $p_D(q)$ as $q_D(p)$, such that $q_D \circ p_D = p_D \circ q_D = \mathbf{id}$.

Inverse Supply Now, given the the auctioneer's price schedule $p_n(t)$, we can define the inverse supply curve $p_S : \mathbf{N} \times \mathbf{R} \to \mathbf{R}_+$, which denotes the price of each available amount of the items for sale at any time T. Since the buyer has available all the auctions that started since $n = 1$, and will purchase them in order from cheapest (oldest) to most expensive (newest), we can assign the lowest price to $p_S(1, T)$ and ascending prices to more supply. Therefore, the inverse supply curve is an increasing function of the quantity. In particular, from the price function $p_n(t)$, we can write:

$$p_S(q, T) = \sum_{n=1}^{q} p_n(T - n)$$

As an example, if $p_n(t) = k\alpha^n e^{-\lambda t}$, then the inverse supply curve takes the form:

$$\begin{aligned} p_S(q, T) &= \sum_{n=1}^{q} k\alpha^n e^{-\lambda(T-n)} \\ &= \frac{k\alpha e^{\lambda - \lambda T}(\alpha^q e^{\lambda q} - 1)}{\alpha e^{\lambda} - 1} \end{aligned}$$

This shows that as the buyer is interested in purchasing more and more items, the price of the items exponentially grows (as a result of the fact that the oldest auctions have a price that is exponentially smaller than the most recent auction).

Auctioneer's Revenue In order to determine how many auctions the buyer will purchase, and consequently how much revenue the auctioneer will collect, we use the fact that the quantity that the buyer will purchase should set equal the inverse demand curve of the buyer with the inverse supply curve of the available auctions. In words, this means that the buyer will purchase an amount of items such that That is, the market clearing quantity of items the buyer will purchase at time T, $q^\star(T)$, satisfies:

$$p_D(q^\star) = p_S(q^\star, T)$$

Note that $q^\star$ is a always a function of T, but we will drop the time-dependence when it is clear from context. Given this quantity, we define the revenue of the auctioneer at time T as:

$$R(q^\star, T) = \sum_{n=1}^{q^\star} p_n(T-n)$$

This revenue adds up the prices of each of the auctions purchased by the buyer, up to the market clearing quantity $q^\star$. Given our example supply curve $p_n(t) = k\alpha^n e^{-\lambda t}$, and a quantity $q^\star$, the revenue at time T is therefore:

$$R(q^\star, T) = \frac{k\alpha e^{\lambda - \lambda T}(\alpha^{q^\star} e^{\lambda q^\star} - 1)}{\alpha e^\lambda - 1} \tag{1}$$

3.2 Continuous Time

We briefly give an introduction to continuous time gradual dutch auctions (GDA), which occur when a seller wants to sell a batch of fungible tokens, but does not want to make them all available at the same time. We will not analyze the incentive properties of such auctions, although they are very similar to discrete time GDAs. In this case, the seller decides an *emissions rate* ρ, of the number of tokens she will sell per minute starting at time $t = 0$. At every time $t \geq 0$, a new auction starts with a price:

$$p(t) = ke^{-\lambda t}$$

where k is the initial price and λ is the decay constant. Note that this means that at any time T in the future, a buyer will have access to all the auctions started at $t \leq T$, and will be able to buy them at their corresponding price. Concretely, the buyer will receive the lowest price for the tokens by participating in the oldest available auctions. That is, at time T, the buyer can purchase all the auctions from the oldest available auction at $p(T)$, to the auction that corresponds to price $p(T - \frac{q}{\rho})$. Once again, we assume that the seller is faced with a buyer with a fixed and known inverse demand curve $p_D(q)$.

Inverse Supply Once again, given the auctioneer's price function $p(t)$, we can define the inverse supply curve $p_S : \mathbf{R} \times \mathbf{R} \to \mathbf{R}_+$, which is the price of each available amount of items at any time T. Note now that the inverse supply function is a function of a real-valued variable in its first argument, because the items are fungible and no longer sold in integer quantities. Therefore, given the price function $p(t)$, we can write:

$$p_S(q, T) = \int_0^{q/\rho} p(T-t)dt$$

This can be seen by noting that the oldest available (and hence cheapest) auction has price $p(T)$, and at every time, ρ units of the tokens are emitted. Therefore, to buy q total units of tokens, a buyer must purchase all the auctions from $t = 0$ to $t = q/\rho$. This generates the inverse supply curve $p_S(q, T)$.

Auctioneer's Revenue Once again, the revenue the auctioneer collects will depend on the clearing quantity of the auction, which is found by setting equal the inverse demand and supply curves. Therefore, the market clearing quantity once again satisfies:

$$p_D(q^\star) = p_S(q^\star, T)$$

Given this quantity, the revenue of the auctioneer at time T in the continuous-time GDA is once again:

$$R(q^\star, T) = \int_0^{q^\star/\rho} p(T - t)dt$$

Given the example price function $p(t) = ke^{-\lambda t}$, and a quantity $q^\star$, the revenue at time T is therefore:

$$R(q^\star, T) = \frac{ke^{-\lambda T}(e^{\lambda q^\star/\rho} - 1)}{\lambda}$$

Future Work. In the remainder of this paper, we will focus on discrete time results. Note that most of the results carry over to the continuous time setting with some extra technical work (*e.g.* correctly taking limits and adding in extra regularity conditions). Given that this does not provide direct insight into the basic mechanics of GDAs, we leave further continuous time analysis for future work.

3.3 Generic Single-Dimensional GDAs

In the previous section, we focused on exponential pricing mechanisms. These mechanisms have the benefits of allowing for an analytical formulation, allowing for easier practical implementations [33]. However, as we demonstrated, these auctions are not credible in that a malicious auctioneer can deviate from honest behavior and cause a loss to buyers. This section aims to close this gap by defining a large family of single-dimensional GDAs (*e.g.* the auctioneer is selling multiple copies of a single object) that can be formally analyzed. Given the generic GDAs defined in Section 2, for adapted, admissible, and time translation invariant pricing functions, we can define the manipulation cost $C(r)$, the manipulation revenue $R(\hat{q}, t)$ (both defined in the next section), and the equilibrium revenue $R(q^\star, t)$ analogously to the exponential case:

$$C(r) = \sum_{i=1}^{r} a_i$$

$$R(\hat{q}, t) = \sum_{i=r+1}^{\hat{q}+r+1} a_i \hat{p}(t - t_i)$$

$$R(q^\star, t) = \sum_{i=1}^{q^\star} a_i \hat{p}(t - t_i)$$

The quantities $\hat{q}, q^\star, r$ are defined analogously to the exponential scenario. Explicitly, the equilibrium quantity $q^\star$ satisfies:

$$p_D(q^\star) = p_S(q^\star, T)$$

4 Sellers' Incentives in GDAs

Let $p_n(t)$ denote a particular instance of a discrete time GDA and $p_D(q)$ a buyer's demand curve. Then, as above, denote $R(q^\star, T)$ to be the revenue earned the auctioneer by running the auction truthfully, where the buyer purchases $q^\star$ units of items at time T. We now demonstrate that the naive GDAs considered in the previous section are not revenue optimal for the seller, in the sense that the seller can deviate from truthfully running a GDA by participating in the auction and therefore incurring a revenue larger than $R(q^\star, T)$. This attack by the seller proceeds by buying up some amount of the supply of items in the first few auctions, therefore forcing the buyer to buy more of the tokens in later auctions, which come at higher prices. If the inverse demand curve of the buyer is not too sensitive to changes in quantity, we show that this deviation by the seller can be profitable, and thus the seller would choose to not truthfully run the GDA. We note that our deviation is one example of a class of deviations where the seller makes initial items unavailable (for a cost) to force buyers to purchase later items. The deviation can be modified to change the sequence of items that the seller makes unavailable, and also can be made to depend on the valuation of a sequence of buyers, as opposed to a single buyer as we consider subsequently.

The seller's deviation We now demonstrate the seller's deviation. Assume for simplicity that the seller chooses to run a discrete time exponential GDA with a price curve of $p_n(t) = k\alpha^n e^{-\lambda t}$. The seller purchases r of the first auctions at the instant they start. In particular, this means that the buyer is forced to start buying only auctions after $r + 1$, whose initial price is higher. The cost to the seller of purchasing these is $C(r) = \sum_{i=1}^{r} k\alpha^i$. We now show that the seller can set r with respect to the buyer's inverse demand curve $p_D(q)$ such that this attack is profitable provided benign conditions on the rate of decay of the demand.

Now, suppose the buyer buys at time $T > r$ and attempts to fill her demand with respect to the inverse demand curve $p_D(q)$. The auctions that haven't been bought by the auctioneer are available at times $r+1, r+2, ..., T$, and have prices $k\alpha^{r+1}e^{-\lambda(T-r+1)}, k\alpha^{r+2}e^{-\lambda(T-r+2)}, \ldots, k\alpha^T$. In order to construct the supply curve, we see that the following quantities and prices of the items are available, from cheapest to most expensive:

1. Quantity 1 at price $k\alpha^{r+1}e^{-\lambda(T-r+1)}$
2. Quantity 2 at price $k\alpha^{r+1}e^{-\lambda(T-r+1)} + k\alpha^{r+2}e^{-\lambda(T-r+2)}$
3. $\ldots$

4. Quantity $T - r + 1$ at price $\sum_{n=r+1}^{T} k\alpha^n e^{-\lambda(T-n)}$

Therefore, the modified inverse supply curve available to the buyer at time T, $\hat{p}_s(q, T)$ for $q < T$ is the function:

$$\hat{p}_s(q,T) = \sum_{n=r+1}^{\max(q+r+1,T)} k\alpha^n e^{-\lambda(T-n)}$$

Note that when $q = T$, the last term of $p_s(T)$ is $k\alpha^T e^{-\lambda(T-T)} = k\alpha^T$, which is the initial price of the latest, and hence most expensive auction. However, as the first r auctions were purchased by the auctioneer, these are not available to the buyer. The price of the cheapest auction has been increased to $k\alpha^{r+1}e^{-\lambda(T-r+1)}$.

Auctioneer's Revenue To determine the amount of items the buyer is willing to purchase given the new supply curve, and therefore the auctioneer's revenue, we again use the fact that the buyer will attempt to intersect $p_D(q)$ with the modified inverse supply curve $\hat{p}_S(q, T)$. Alternatively, the market clearing quantity $\hat{q}$ satisfies:

$$p_D(\hat{q}(T)) = \hat{p}_S(\hat{q}(T), T)$$

Again, note that $\hat{q}(T)$ is explicitly a function of T, but we will drop the dependence for notational simplification when the context is clear. Correspondingly, the revenue to the auctioneer from devitating to the supply curve $\hat{p}_S(q, T)$ is given by:

$$R(\hat{q},T) = \sum_{i=r+1}^{\hat{q}+r+1} \alpha^i e^{-\lambda(T-i)} = \frac{\alpha e^{\lambda-\lambda T}(\alpha^{\hat{q}+r+1}e^{\lambda(\hat{q}+r+1)} - \alpha^r e^{\lambda r})}{\alpha e^{\lambda} - 1}$$

which allows us to write the PNL to the auctioner of deviating:

$$\begin{aligned}
\mathsf{PNL}(\hat{q},T,r) &= R(\hat{q},T) - C(r) \\
&= \sum_{n=r+1}^{\hat{q}+r+1} k\alpha^n e^{-\lambda(T-n)} - \sum_{n=1}^{r} k\alpha^n \\
&= \frac{k\alpha e^{\lambda-\lambda T}(\alpha^{\hat{q}+r+1}e^{\lambda(\hat{q}+r+1)} - \alpha^r e^{\lambda r})}{\alpha e^{\lambda} - 1} - \frac{k\alpha(\alpha^r - 1)}{\alpha - 1}
\end{aligned}$$

Note that by construction, $\hat{q} < q^\star$ (we formally show this in Appendix B).

When is the deviation profitable? We now show conditions on the buyer's inverse demand curve $p_D(q)$ such that the above deviation is profitable, that is, $\mathsf{PNL}(\hat{q}, T, r) > R(q^\star, T)$. The intuition behind showing profitability is to show conditions on the inverse demand curve such that the demand does not decrease too quickly if the price becomes higher. We first give explicit conditions on how $\hat{q}$ depends on r that guarantees the profitability of the deviation. To lighten notation, define $b = \alpha e^{\lambda}$. Then note that we have:

$$
\begin{aligned}
R(\hat{q}, T) - R(q^\star, T) &= \frac{k\alpha e^{\lambda(1-T)}(\alpha^{\hat{q}+r+1}e^{\lambda(\hat{q}+r+1)} - \alpha^r e^{\lambda r})}{\alpha e^{\lambda} - 1} - \frac{k\alpha e^{\lambda(1-T)}(\alpha^{q^\star}e^{\lambda q^\star} - 1)}{\alpha e^{\lambda} - 1} \\
&= \frac{k\alpha e^{\lambda(1-T)}}{\alpha e^{\lambda} - 1}(\alpha^{\hat{q}+r+1}e^{\lambda(\hat{q}+r+1)} - \alpha^r e^{\lambda r} - \alpha^{q^\star}e^{\lambda q^\star} + 1) \\
&= \frac{k\alpha e^{\lambda(1-T)}}{\alpha e^{\lambda} - 1}(\alpha^r e^{\lambda r}(\alpha^{\hat{q}+1}e^{\lambda(\hat{q}+1)} - 1) - \alpha^{q^\star}e^{\lambda q^\star} + 1) \\
&= \frac{kb^{(1-T)}}{b - 1}(b^r(b^{\hat{q}+1} - 1) - b^{q^\star} + 1)
\end{aligned}
$$

which then yields

$$
\begin{aligned}
R(\hat{q}, T) - R(q^\star, T) - C(r) &= \frac{kb^{(1-T)}}{b - 1}(b^r(b^{\hat{q}+1} - 1) - b^{q^\star} + 1) - \frac{k\alpha(\alpha^r - 1)}{\alpha - 1} \\
&= \frac{kb^{(1-T)}}{b - 1}\left(b^r(b^{\hat{q}+1} - 1) - b^{q^\star} + 1 - \frac{(\alpha^r - 1)(b - 1)}{e^{\lambda(1-T)}(\alpha - 1)}\right)
\end{aligned}
$$

Our attack is profitable if the right-hand side of the above is positive, which holds when

$$
b^{\hat{q}+r+1} \geq b^{q^\star} + b^r - 1 + \frac{(\alpha^r - 1)(b - 1)}{e^{\lambda(1-T)}(\alpha - 1)}
$$

Taking logarithm and simplifying gives the condition

$$
\begin{aligned}
(\hat{q} + r + 1)\log b &\geq \log\left(b^{q^\star} + b^r - 1 + \frac{(\alpha^r - 1)(b - 1)}{e^{\lambda(1-T)}(\alpha - 1)}\right) \\
&= q^\star \log b + \log\left(1 + b^{-q^\star}\Lambda(r, \lambda, T)\right) \qquad (2)
\end{aligned}
$$

where we define

$$
\Lambda(r, \lambda, T) = b^r - 1 + \frac{(\alpha^r - 1)(b - 1)}{e^{\lambda(1-T)}(\alpha - 1)} \qquad (3)
$$

In Appendix A, we show the following claim:

Claim 1. *If $q_D(p)$ is Lipschitz, then $b^{q^\star} \leq C(\lambda, T)$ where C is a constant only dependent on λ and T (not r)*

If we define $C"(\lambda, T) = \frac{\Lambda(r,\Lambda,T)}{b^{q^\star}(\alpha^r - 1)}$ then (2) simplifies to

$$
\hat{q} + r + 1 \geq q^\star + \frac{1}{\log b}\log(1 + (\alpha^r - 1)C'(\lambda, T))
$$

or

$$
r + 1 \geq (q^\star - \hat{q}) + \frac{1}{\log b}\log(1 + (\alpha^r - 1)C'(\lambda, T)) \qquad (4)
$$

Note that $\log(1 + AB) \geq \log A + \log B$ for $A, B \geq 1$ and that $\log(\alpha^r - 1) \geq k(\log \alpha)r$ for a positive constant $k < 1$. These two facts combined give

$$r + 1 \geq (q^\star - \hat{q}) + \frac{k \log \alpha}{\log b} r + \frac{\log C'(\lambda, T)}{\log b}$$

or

$$r \left(1 - \frac{k \log \alpha}{\log b}\right) \geq (q^\star - \hat{q}) + \frac{\log C'(\lambda, T)}{\log b} - 1 \tag{5}$$

Note that $\log b > \log \alpha$, so that $1 - \frac{k \log \alpha}{\log b} > 0$, so our bound on the number of rounds for profitability is always positive.

Interpreting (5). How can we interpret equations (4) and (5)? These equations show that the number of rounds to reach profitability is lower bounded by an affine function of $(q^\star - \hat{q})$. This implies that the larger the deviation from a natural equilibrium $q^\star$, the more we have to compensate by buying earlier rounds (to push the price up enough to be profitable).

Similarly, recall that p_D is convex, differentiable, non-increasing, and has Lipschitz gradients, so that the difference $p_D(q^\star) - p_D(\hat{q})$ has a linear lower bound of the form

$$p_D(q^\star) - p_D(\hat{q}) = \Theta(q^\star - \hat{q})$$

The linear lower bound for r combined with the fact that p_D is non-increasing and $\hat{q} \leq q^\star$ implies that $r = \Omega(p_D(q^\star) - p_D(\hat{q}))$. This can be viewed as saying that the difference quotients and gradients of p_D control the number of rounds that have to be bound. If p_D has small difference quotients and/or gradients (*e.g.* decays gently), then the seller does not have to purchase many initial auctions to be profitable. On the other hand, if demand decays quickly (e.g. the diffence quotients are large), then the auctioneer needs to compensate in a manner proportional the the maximum change in demand over the region of quantities tendered.

Unprofitability as $T \to \infty$. We now consider the limit in which T, the time at which the buyer enters the auction, becomes very large. As this time grows, the buyer has access to more and more auctions whose price is exponentially smaller than the price of the most recent auction, and correspondingly, the auctioneer must purchase more of them to maintain the profitability of the deviation. Formally, we show that in the $T \to \infty$ limit, the above deviation is not profitable for any reasonable inverse demand curve.

The simplest way to prove this is to simply note that in equation (3), $\Lambda(r, \Lambda, T) = \Omega(c^T)$ for some $c > 1$ as $\alpha, b > 1$. This implies that $C'(\lambda, T) = \Omega(c^T)$ so that we equation (4) becomes

$$r + 1 \geq (q^\star - \hat{q}) + K \log(1 + c^T) \geq (q^\star - \hat{q}) + K'T$$

for some positive constants $K, K' > 0$ since $\log(1 + c^T) \geq \log c^T \geq T \log c$. As $T \to \infty$, this implies that the number of auctions that the auctioneer has to buy to be profitable goes to infinity and hence the strategy is not profitable.

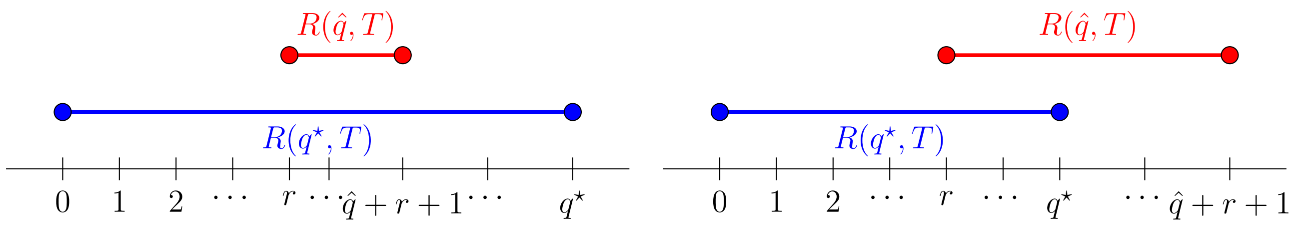

Figure 1: Case 1 (left) and Case 2 (right). The intervals represent the support (*e.g.* set of indices of non-zero terms contributing to R) of the two values. When the red and blue intervals overlap, those terms in the sum cancel. In case 1, the support of $R(\hat{q}, t)$ is contained within the support of $R(q^\star, t)$, which gives us the negative value. On the other hand, Case 2 only has an overlap between r and $q^\star$

Seller Incentives in Generic GDAs In Appendix B, we proved that $\hat{q} \leq q^\star$ for adapted, admissible, time translation invariant pricing function. Moreover, we also know that $r < q^\star$. We will first look at $R(\hat{q}, T) - R(q^\star, T)$. This can be broken down into two cases (see Figure 1):

1. $\hat{q} + r + 1 \leq q^\star$
2. $\hat{q} + r + 1 > q^\star$

For case 1, we have (based on Figure 1):

$$R(q^\star, T) - R(\hat{q}, T) = \sum_{i=1}^{r} a_i \hat{p}(T - t_i) + \sum_{i=\hat{q}+r+2}^{q^\star} a_i \hat{p}(T - t_i) \geq 0 \tag{6}$$

which implies that the attack is never profitable if $q^\star \geq \hat{q} + r + 1$. On the other hand, for case 2, we have for $t > \max_i t_i$

$$\begin{aligned} R(\hat{q}, T) - R(q^\star, T) &= \sum_{i=q^\star+1}^{\hat{q}+r+1} a_i \hat{p}(T - t_i) - \sum_{i=1}^{r} a_i \hat{p}(T - t_i) \\ &\geq ((\hat{q} + r + 1) - q^\star) \left(\min_{q^\star+1 \leq i \leq \hat{q}+r+1} a_i \hat{p}(T - t_i) \right) - r \left(\min_{i \in [r]} a_i \hat{p}(T - t_i) \right) \\ &= r(M_1 - M_2) + ((\hat{q} - q^\star) + 1) M_1 \end{aligned} \tag{7}$$

where

$$\begin{aligned} M_1 &= \min_{q^\star+1 \leq i \leq \hat{q}+r+1} a_i \hat{p}(t - t_i) \geq 0 \\ M_2 &= \min_{i \in [r]} a_i \hat{p}(T - t_i) \geq 0 \end{aligned}$$

The lower bound in eq. (7) is positive if

$$r \geq ((q^\star - \hat{q}) + 1) \frac{M_1}{M_1 - M_2}$$

By the definition of adapted pricing functions, we have $M_1 > M_2$ so that the round complexity is lower bounded by $\hat{q^\star} - \hat{q}$. If we have a positive cost $C(r)$, we have

$$\begin{aligned} \mathsf{PNL}(\hat{q}, t, r) - R(q^\star, t) &= R(\hat{q}, t) - R(q^\star, t) - C(r) \\ &\geq r(M_1 - M_2) + ((\hat{q} - q^\star) + 1)M_1 - C(r) \end{aligned}$$

which is only positive if

$$r \geq \frac{M_1}{M_1 - M_2}(\hat{q} - q^\star) - \frac{M_1 - C(r)}{M_1 - M_2} \tag{8}$$

The two results, equations (6) and (8) demonstrate two key insights into designing pricing functions $p_i(t)$:

- If a mechanism designer knows enough about the demand curve p_D to enforce $\hat{q} + r + 1 \leq q^\star$ by adjusting initial prices a_i, they can ensure that auctioneer deviation is unprofitable
- If a mechanism designer aims to design a mechanism in a prior-free manner (*e.g.* without knowledge of p_D), they can use (8) to increase the cost of deviation by adjusting the pricing function

These two observations demonstrate that optimal design of pricing functions can lead to dramatically better social welfare for participants relative to the exponential GDA.

Unprofitability as $T \to \infty$. In order to recover the unprofitability result of §4, we need to constrain the temporal dependence of the pricing functions, *i.e.* $\hat{p}(t)$. To find a sufficient condition on $\hat{p}(t)$ such that the unprofitability result continues to hold, we need to rewrite (8) as

$$\begin{aligned} r &\geq \frac{M_1}{M_1 - M_2}(\hat{q} - q^\star) + \frac{C(r) - M_1}{M_1 - M_2} \\ &= \frac{M_1}{M_1 - M_2}\left((\hat{q} - q^\star) + \frac{C(r) - M_1}{M_1}\right) \\ &= \frac{M_1}{M_1 - M_2}\left((\hat{q} - q^\star) + \frac{C(r)}{M_1} - 1\right) \end{aligned}$$

Firstly, note that all of the time dependence of this term is in the M_1, M_2 terms. Secondly, this form shows that[6] if $M_1(t) = \omega(1)$ as $t \to \infty$, then $r \to \infty$ as $t \to \infty$. Therefore, our goal is to find sufficient conditions for $M_1(t) = \omega(1)$.

Let $I(q^\star, \hat{q}, r) = \{i : q^\star + 1 \leq i \leq \hat{q} + r + 1\}$, note the following elementary inequality

$$M_1 = \min_{i \in I(q^\star, \hat{q}, r)} a_i \hat{p}(t - t_i) \geq \left(\min_{i \in I(q^\star, \hat{q}, r)} a_i\right)\left(\min_{i \in I(q^\star, \hat{q}, r)} \hat{p}(t - t_i)\right)$$

[6] Recall that a function $f(t) \in \omega(g(t))$ if $\lim_{t \to \infty} \frac{f(t)}{g(t)} = \infty$

Furthermore, since $\hat{p}$ is non-increasing we have

$$\min_{i \in I(q^\star, \hat{q}, r)} \hat{p}(t - t_i) = \hat{p}\left(\max_{i \in I(q^\star, \hat{q}, r)} (t - t_i)\right) = \hat{p}\left(t - \min_{i \in I(q^\star, \hat{q}, r)} t_i\right)$$

Therefore, as long as there exists a function $f(t)$ such that $\lim_{t\to\infty} f(t) = \infty$ such that for all intervals $I(q^\star, \hat{q}, r)$ and time $t \neq t_i$, we have

$$\hat{p}\left(t - \min_{i \in I(q^\star, \hat{q}, r)} t_i\right) \geq \frac{f(t)}{\left(\min_{i \in I(q^\star, \hat{q}, r)} a_i\right)}$$

then we achieve asymptotic unprofitability in a manner analogous to the exponential GDA.

5 Buyers' Incentives in GDAs

We now turn to the incentives of the buyer in a GDA. We seek to understand cases in which a buyer is incentivized to purchase items in alignment with her true value for the items. Having provided a negative result for buyers that allowed to stay perpetually in the auction, we now demonstrate conditions under which GDAs are *incentive compatible* for 'ephemeral' buyers who enter at a single time in the auction, but whose value for the items can depend on private signals of other buyers. That is, these buyers buy items that match their true value, even in the presence of competition at different times. In the case of such ephemeral bidders, we consider an interdependent model of bidder values that has been considered for art auctions (which NFT auctions resemble) [9, 27, 42]. Frequently, in art auctions, the value of a buyer depends on the *resale value* of the art, which depends on how many other buyers are willing to purchase the item. GDAs present an ideal environment to study this kind of interdependence, because many identical copies of the item can be sold over time, which allows potential buyers in the future to adjust their values for an item by observing the history of the auction. For ephemeral buyers, we show conditions under which GDAs are incentive compatible.

Prior Work on Interdependent Auctions. There are a number of papers on analyzing optimal auctions for interdependent auctions. Roughgarden and Talgam-Cohen [42] demonstrated a necessary and sufficient condition for both ex post properties that is analogous to the classical Myerson Lemma. On the other hand, Li [27] found a combinatorial classification for interdependent values and provided an ascending auction that achieved incentive compatibility. Finally, Chawla, et. al [9] provide explicit bounds on revenue maximization. We note that Li's condition for being ex post incentive compatible and individually rational is the most directly connected to GDAs.

There are three concepts we need to state to describe Li's results (for more details, see the original paper [27]). The single crossing condition states that if $t'_i > t_i$ then $v_i(s_i, s_{-i}) \geq v_j(s_i, s_{-i})$ implies $v_i(s'_i, s_{-i}) \geq v_j(s'_i, s_{-i})$. Colloquially, this states that changes to individual agents' valuations are dominated by their signal s_i more than any other signal from another agent s_j. Secondly, recall that a *matroid* on a base set $[n]$ is a set of subsets $\mathcal{M} \subset 2^{[n]}$ such that the following two conditions hold

- *Downward Closed.* If $S \subset T$ and $T \in \mathcal{M}$, then $S \in \mathcal{M}$
- *Basis Expansion.* If $A, B \in \mathcal{M}$ with $|A| < |B|$ then there exists $i \in B - A$ such that $A \cup \{i\} \in \mathcal{M}$

Finally, Li defines a VCG-L mechanism as a standard Vickrey-Clarke-Groves (VCG) mechanism with particular (bidder-specific) monopoly reserve prices $r_i(s_{-i})$ for each user i. The mechanisms used by Roughgarden and Chawla, et. al are subsets of this mechanism for $r_i(s_{-i})$; see [27] for more details. Li proves the following:

Theorem 1. *[9, Lemma 2.2], [27, Thm. 1] Suppose that $\mathcal{M}$ is a matroid on the set of bidders and v_i satisfies the single crossing condition for all i. Then VCG-L are ex post incentive compatible and individually rational.*

We will define VCG-L mechanisms shortly and show that the GDAs can be viewed as a form of VCG-L mechanism. Therefore, if we can construct a sequence of auctions that satisfies the matroid constraint and valuations that satisfy the single crossing condition, we can achieve ex post incentive compatibility for GDAs. In order to utilize this theorem, we will first need to construct signals, valuations, and allocation rules to analyze.

Valuation and Allocation. It is well-known that NFT auctions suffer from the "winner's curse" [48], which is a well-known phenomena within interdependent auctions [20]. These phenomena are often due to the existence of *common values* — my valuation for the item depends on the valuations that others have, and in fact knowing that another player values the item more increases my value for the item. As such, we will model the valuation of a GDA participant as the sum of two components:

$$v_{ij}(s_{i,j}, s_{-i,j}) = v^P_{i,j}(s_{ij}) + v^C_{ij}(s_{ij}, s_{-i,j})$$

where v_i^P is a private valuation drawn from a fixed distribution (independent of v_j^P) and v_i^C is a common value that depends on all agents' signals. We will be focused on constructing v_i^C while leaving v_i^P unspecified as classical auction theory can handle the private components.

At time t, we define the jth auction's price to be $p_j(t) = a_j\hat{p}(t-t_j)\mathbf{1}_{t\geq t_j}$. We can treat the demand function $p_D(q)$ as the average demand of a user, such that if there exists a bidder k with $s_i(t) > p_D(q)$, then bidder i can realize $s_i(t) - p_D(q)$ in value. Since the demand is connected to $p_j(t)$ via the supply-demand clearing equation $\hat{p}_S(q,T) = p_D(q,T)$, we can simply look if the jth auction provides a profit, *e.g.* $s_i(t) \geq p_j(t)$. As such, we can define

$$v^C_{ij}(s_{ij}, s_{-i,j}) = \max(s_i(t) - p_j(t), 0)$$

Therefore, bidder i has positive value if

$$v_i(s_i, s_{-i}) = v_i^P(s_i) + \max(s_i(t) - p_j(t), 0) \geq p_j(t)$$

which implies the condition

$$s_i(t) \geq 2p_j(t) - v_i^P(s_i)$$

This intuitively says that as long as the price is high enough to be worth more than my private valuation but simultaneously low enough that there is another agent who demands the item, then there is positive value. This suggests the following deterministic allocation rule

$$x_{ij}(s_i, s_{-i}) = \mathbf{1}_{s_i(t) \geq 2p_j(t) - v_i^P(s_i)}$$

Finally, we will remark on when these valuations are single-crossing. For differentiable valuations, the single crossing condition becomes $\partial_{s_i} v_{ij}(s) > \partial_{s_k} v_{ij}(s)$ for all $k \neq i$. This holds here as long as $\partial_{s_i} v^P(s_i) > \partial_{s_j} \max(s_j(t) - p_j(t), 0)$ which is a condition that can be enforced by adjustment of maximum posted prices a_i.

Vickrey-Clarke-Groves mechanisms. We will briefly describe the generalized Vickrey-Clarke-Groves (VCG) mechanisms of [9, §2] and [27, §3], which will be used for understanding incentive compatibility. Such a mechanism first relies on constructing a family of subsets of possible winners, $\mathcal{M}$. If we have had T auctions by discrete time $T \in \mathbf{N}$, an element of $\mathcal{M}$ will be an element of $[n] \times [T]$, where n is the number of bidders. More formally, we define a set of valid auction outcomes to be $\mathcal{M} \subset 2^{[n]\times[T]}$. We will slightly abuse notation and define $\mathcal{M} \subset 2^{[T]}$ if there is only one bidder. Finally, note that if $\mathcal{M}$ is a matroid then we can utilize Theorem 1.

Given a set of admissible allocations $\mathcal{M}$, VCG mechanisms do three things:

1. *Construct a set of winners.* $W(s) = \text{argmax}_{A \in \mathcal{M}} \sum_{(i,t)\in A} v_{it}(s)$
2. *Compute threshold signal.* The threshold signal for agent i, i^* is defined as

$$s_i^*(s_{-i}) = \inf_{s_i} \{\exists j \text{ such that } (i,j) \in \operatorname*{argmax}_{A\in\mathcal{M}} \sum_{(l,m)\in A} v_{lm}(s_i, s_{-i})\}$$

3. *Compute the payment.* The VCG payment for agent i is defined as $p_i^{VCG}(s,T) = \sum_{j\in[T]} v_{ij}(s_i^*, s_{-i})$

A VCG mechanism with lazy reserve prices (VCG-L) functions takes as input a vector $r \in \mathbf{R}_+^n$ and performs an auction as follows:

- Agents report signals s_i
- The mechanism chooses a subset $W(s) = \text{argmax}_{A\in\mathcal{M}} \sum_{(i,j)\in A} v_{ij}(s)$
- Each bidder is given a take-it-or-leave-it posted price of $\max(r_i, p_i^{VCG}(s))$

The prices r_i are known as reserve prices and are tuned to provide particular guarantees. In the case of a GDA, we can view the current price represented by the inverse supply curve as a posted price. We take advantage of the fact that the GDA is effectively a posted price mechanism to construct r_i and $\mathcal{M}$ that achieve incentive compatibility.

Incentive compatibility for ephemeral buyers. We now consider the case of an ephemeral bidder that enters the auction at time $T \in \mathbf{N}$ (and must either purchase a subset of the auctions available, or leave the auction). Recall that this buyer has a private signal $s_i \in \mathbf{R}$ and can see the history of the auction, using the public signals p_{ij} which denotes the price that item i was sold for in auction j. We demand ex-post incentive compatibility and individual rationality over the entire sequence of signals $(s_1, \ldots, s_n)$.

Suppose that the set of admissible auctions, indexed by time, is $\mathcal{M} \subset 2^{[T]}$ The subset of auctions purchased by this bidder (who we will denote by i) will be a set $K(s,T) \in \mathcal{M}$ defined as

$$K(s,T) = \operatorname*{argmax}_{A \in \mathcal{M}} \sum_{t \in A} v_{it}(s) - p_t(T)$$

This is analogous to the bidder purchasing a number of auctions that matches supply to demand. Following [9, §2], we define the threshold signal s_i^* for the ith bidder as

$$s_i^* = \inf_{s \geq 0} \left\{ \operatorname*{argmax}_{A \in \mathcal{M}} \sum_{t \in A} (v_{it}(s, s_{-i}) - p_t(T)) \geq 0 \right\}$$

Recall single ephemeral bidder either realizes positive value at time T or drops out and therefore the social welfare of the bidder is $\max(K(s,T) - K(s_i^*, s_{-i}, T), 0)$. Note that this formulation constructs a VCG-L mechanism with reserve prices $r_i(A) = \sum_{(i,j) \in A} p_j(T)$ for a set $A \in \mathcal{M}$. Theorem 1 shows that this mechanism, which a VCG mechanism is incentive compatible if $\mathcal{M}$ is a matroid and $v_{i,j}$ are single crossing. In §5.1, we explicitly construct a matroid that provides conditions that depend on the valuation and the pricing function parameters that achieves this bound. We note that if the auctioneer and/or mechanism designer know the value function $v_{i,j}$, then they can construct GDA initial prices a_j and decay functions $\hat{p}(t)$ such that $A \in \mathcal{M}$ always holds. This holds for an ephemeral buyer at any time, and therefore, the auction is ex post incentive compatible provided that $\mathcal{M}$ is in fact a matroid, which we show next.

5.1 Matroid Construction

In this section, we construct a matroid using the GDA pricing functions that holds only if particular conditions relating the pricing functions to bidder valuations are satisfied. This matroid naturally represents the set of possible winning bundles and the conditions are related to the gap between the buyers' valuation and the current purchase price. Our conditions can likely be relaxed (constructing a larger matroid or even a general downward closed set) and mainly serve to illustrate that it is possible to achieve ex post incentive compatibility and individual rationality with GDAs. Given that the construction of a matroid $\mathcal{M}$ is tantamount to achieving incentive compatibility, the existence of this matroid for a given set of GDAs demonstrates that a mechanism designer can construct price curves $p_j(t)$ to achieve ex post incentive compatibility and individual rationality.

We construct a matroid $\mathcal{M}$ on the base set $[n] \times [T]$ where n is the number of bidders and T is the number of auctions. $\mathcal{M}$ is constructed recursively using pricing functions $\hat{p}(t)$, initial prices a_i, and valuations v_i as follows:

$$\mathcal{M}_1 = \{(i,1) : v_i(s_{i1}) \geq a_1\}$$
$$\mathcal{M}_k = \{S \cup \{(i,\ell)\} : \ell \leq k, S \in \mathcal{M}_{k-1}, \nexists i' \text{ such that } (i',\ell) \in S,$$
$$v_i(s_{ik}) \geq a_\ell \hat{p}(t_k - t_\ell)$$
$$\sum_{(j,r)\in S} v_j(s_{jk}) + v_i(s_{ik}) \geq a_\ell \hat{p}(t_k - t_\ell) + \sum_{(j,r)\in S} a_r \hat{p}(t_k - t_r)$$
$$v_i(s_{ik}) - a_\ell \hat{p}(t_k - t_\ell) \geq \min_{S \subset \mathcal{M}_k} \sum_{(j,r)\in S} a_r \hat{p}(t_k - t_r) - v_j(s_{jk})\}$$

Informally, the conditions in the matroid do the following:

- $v_i(s_{ik}) \geq a_\ell \hat{p}(t_k - t_\ell)$: This ensures that the agent winning the kth auction has non-negative welfare
- $\sum_{(j,r)\in S} v_j(s_{jk}) + v_i(s_{ik}) \geq a_\ell \hat{p}(t_k - t_\ell) + \sum_{(j,r)\in S} a_r \hat{p}(t_k - t_r)$: This ensures that at time t_k (when the kth auction is run), then total social welfare of all elements of the set (including those added previously whose prices have decayed) are still positive
- $v_i(s_{ik}) - a_\ell \hat{p}(t_k - t_\ell) \geq \min_{S \subset \mathcal{M}_k} \sum_{(j,r)\in S} a_r \hat{p}(t_k - t_r) - v_j(s_{jk})$: This condition ensures that the welfare realized at the kth round is at least the minimum total social welfare of the previous round (and is needed for the basis expansion property). Note that the right hand side of this condition can only be positive if the the majority of the values $v_j(s_{jk})$ decayed faster than the fixed prices $a_r \hat{p}(t_k - t_r)$ did (which only occurs if the gradients of v_j are larger than those of $\hat{p}$). Finally, we note that this can be viewed as a 'no-regret' condition (in the sense of the example in the introduction).

In Appendix D, we prove that $\mathcal{M} = \cup_k \mathcal{M}_k$ satisfies the downward closed and basis expansion properties and hence is a matroid. We note that the valuations v_i constructed in the previous section are single crossing under mild conditions (see §2). Therefore, the results of [9, 27] demonstrate that if a GDA satisfies these matroid conditions, then it achieves ex post incentive compatibility and individual rationality. In particular, this matroid construction shows that even if buyers change their valuations based on expectations of future resale value (using the is

6 Conclusion and Future Work

In this paper, we provided the first formal description of gradual dutch auctions and analyzed seller and buyers incentives in these auctions. While showed that the exponential GDA of [18] still has issues with seller and buyer incentives (especially for buyers that remain in the auction over time), there exist other ways to provide auctioneer credibility and incentive compatibility for GDAs. By constructing a more generalized version of the GDA, we were able to provide sufficient conditions for these desiderata. This suggests that the mechanism design space for GDAs is more rich than initially described in [18].

This work only scratched the surface on the close relationship between GDAs, posted prices, and clock auctions. There are a number of open questions about GDAs to answer, especially with regards to their efficiency and/or the trade-off between computational efficiency and revenue achieved. Below we will highlight a few directions for future inquiry.

Connection to Clock Auctions and Deferred Revelation Auctions. Combinatorial clock auctions [26], which are used in spectrum auctions, and deferred revelation auctions [11] both involve price clocks (analogous to our price functions) that buyers respond to. These auctions, which can also be indirect revelation mechanisms, have a number of similarities to GDAs and it is likely that GDAs inherit their strong strategyproofness properties.

Multiple type extensions of Single Crossing. We note that the single-crossing condition for single dimensional signals does not carry over cleanly to signals for multiple types of items. Since our model only auctions one type of items, we were able to construct single dimensional signals for each buyer. We note, however, that there do exist sufficient conditions for auctions with bidders of multiple types and multiple item types. These conditions [23] have not been adapted to the interdependent auction setting and we believe this is a worthwhile future endeavor (partial results for adapting this condition to interdependent auction can be found in [13]).

Multidimensional GDAs. Extending the analysis of this paper to multi item, multidimensional auctions is still an open problem. One could consider a multidimensional extension of GDAs in which a seller wants to sell batches of different items simultaneously. Suppose there are m types of items, each with N_i number of items to be sold, where $i = 1, \dots, m$. The auctioneer sells each of the items at discrete times n, with a price curve $p_n^i(t)$ for the ith type of item. We once again assume that a buyer has an inverse demand curve, $p_D(q)$, where now $q \in \mathbf{R}_+^m$ denotes the vector of types of item purchased by the buyer. Under what conditions do we recover the round complexity lower bound of §4 and the ex post incentive compatibility and/or individual rationality of §2? We suspect that further assumptions on the valuations and the demand are needed to handle this case.

Approximate Revenue Optimality. In [9], the authors demonstrate that VCG-L mechanisms can be approximately revenue optimal. In order to do this, they add an extra constraint upon the value function that is likely satisfies in our scenario. However, we suspect that the precise approximation ratio is worse for the common plus private valuation we use and leave this for future work.

References

[1] Transmission 11, Frankie, and Dave White. 2022. Variable Rate Gradual Dutch Auctions. https://www.paradigm.xyz/2022/08/vrgda

[2] Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. 2021. Uniswap v3 core. *Tech. rep., Uniswap, Tech. Rep.* (2021).

[3] Mohammad Akbarpour and Shengwu Li. 2020. Credible auctions: A trilemma. *Econometrica* 88, 2 (2020), 425–467.

[4] Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. 2022. Constant function market makers: Multi-asset trades via convex optimization. In *Handbook on Blockchain*. Springer, 415–444.

[5] Guillermo Angeris and Tarun Chitra. 2020. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 80–91.

[6] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. 2019. An analysis of Uniswap markets. *arXiv preprint arXiv:1911.03380* (2019).

[7] Dirk Bergemann and Maher Said. 2010. Dynamic auctions: A survey. (2010).

[8] Dirk Bergemann and Juuso Valimaki. 2006. Efficient dynamic auctions. (2006).

[9] Shuchi Chawla, Hu Fu, and Anna Karlin. 2014. Approximate revenue maximization in interdependent value settings. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 277–294.

[10] Tarun Chitra, Peters Erins, and Kshitij Kulkarni. 2022. When do dynamic DeFi rate curves reduce capital efficiency? (2022). https://gauntlet.network/reports/pid

[11] Paul Dütting, Vasilis Gkatzelis, and Tim Roughgarden. 2014. The performance of deferred-acceptance auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 187–204.

[12] Alon Eden, Michal Feldman, Amos Fiat, and Kira Goldner. 2018. Interdependent values without single-crossing. *arXiv preprint arXiv:1806.03865* (2018).

[13] Alon Eden, Michal Feldman, Amos Fiat, Kira Goldner, and Anna R Karlin. 2019. Combinatorial auctions with interdependent valuations: SOS to the rescue. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 19–20.

[14] Meryem Essaidi, Matheus VX Ferreira, and S Matthew Weinberg. 2022. Credible, strategyproof, optimal, and bounded expected-round single-item auctions for all distributions. *arXiv preprint arXiv:2205.14758* (2022).

[15] Michal Feldman, Vasilis Gkatzelis, Nick Gravin, and Daniel Schoepflin. 2022. Bayesian and Randomized Clock Auctions. *arXiv preprint arXiv:2202.09291* (2022).

[16] Matheus VX Ferreira, Daniel J Moroz, David C Parkes, and Mitchell Stern. 2021. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM conference on Advances in Financial Technologies*. 86–99.

[17] Matheus VX Ferreira and S Matthew Weinberg. 2020. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 683–712.

[18] Frankie, Dan Robinson, Dave White, and andy8052. 2022. Gradual Dutch Auctions. `https://www.paradigm.xyz/2022/04/gda`

[19] Sascha Füllbrunn and Abdolkarim Sadrieh. 2012. Sudden termination auctions—An experimental study. *Journal of Economics & Management Strategy* 21, 2 (2012), 519–540.

[20] William N Goetzmann and Matthew Spiegel. 1995. Private value components, and the winner's curse in an art index. *European Economic Review* 39, 3-4 (1995), 549–555.

[21] Jason D Hartline. 2013. Mechanism design and approximation. *Book draft. October* 122, 1 (2013).

[22] Jason D Hartline et al. 2013. Bayesian mechanism design. *Foundations and Trends® in Theoretical Computer Science* 8, 3 (2013), 143–263.

[23] Philippe Jehiel and Benny Moldovanu. 2001. Efficient design with interdependent valuations. *Econometrica* 69, 5 (2001), 1237–1259.

[24] Jayant Kalagnanam and David C Parkes. 2004. Auctions, bidding and exchange design. In *Handbook of quantitative supply chain analysis*. Springer, 143–212.

[25] Peyman Khezr and Vijay Mohan. 2021. Property rights in the Crypto age: NFTs and the auctioning of limited edition artwork. *Available at SSRN 3900203* (2021).

[26] Jonathan Levin and Andrzej Skrzypacz. 2016. Properties of the combinatorial clock auction. *American Economic Review* 106, 9 (2016), 2528–51.

[27] Yunan Li. 2017. Approximation in mechanism design with interdependent values. *Games and Economic Behavior* 103 (2017), 225–253.

[28] Kristof Lommers, Jack Kim, and Mohamed Baioumy. 2022. Market Making in NFTs. *Available at SSRN 4226987* (2022).

[29] Torgin Mackinga, Tejaswi Nadahalli, and Roger Wattenhofer. 2022. TWAP Oracle Attacks: Easier Done than Said? *Cryptology ePrint Archive* (2022).

[30] Paul R Milgrom and Robert J Weber. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* (1982), 1089–1122.

[31] Jason Milionis, Dean Hirsch, Andy Arditi, and Pranav Garimidi. 2022. A Framework for Single-Item NFT Auction Mechanism Design. In *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security*. 31–38.

[32] Debasis Mishra and David C Parkes. 2009. Multi-item vickrey–dutch auctions. *Games and Economic Behavior* 66, 1 (2009), 326–347.

[33] Pablo Misirov, Devansh Batham, and Hrishikesh Bhat. 2022. Spearbit Art Gobblers Security Review. `https://github.com/spearbit/portfolio/blob/master/pdfs/ArtGobblers-Spearbit-Security-Review.pdf`

[34] Mayukh Mukhopadhyay and Kaushik Ghosh. 2021. Market Microstructure of Non Fungible Tokens. *arXiv preprint arXiv:2112.03172* (2021).

[35] Tri-Dung Nguyen and Tuomas Sandholm. 2014. Optimizing prices in descending clock auctions. In *Proceedings of the fifteenth ACM conference on Economics and computation*. 93–110.

[36] Noam Nisan. 2006. Bidding languages. *Combinatorial auctions* (2006), 400–420.

[37] Noam Nisan and Ilya Segal. 2002. The communication complexity of efficient allocation problems. *Draft. Second version March 5th* (2002), 173–182.

[38] User: @pandajackson42. 2023. Gobblers NFT: General Stats. `https://dune.com/pandajackson42/art-gobblers`

[39] David Christopher Parkes. 2001. *Iterative combinatorial auctions: Achieving economic and computational efficiency*. University of Pennsylvania.

[40] Andrei-Dragos Popescu. 2021. Non-Fungible Tokens (NFT)–Innovation beyond the craze. In *5th International Conference on Innovation in Business, Economics and Marketing Research*.

[41] Tim Roughgarden and Okke Schrijvers. 2016. Ironing in the dark. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. 1–18.

[42] Tim Roughgarden and Inbal Talgam-Cohen. 2013. Optimal and near-optimal mechanism design with interdependent values. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 767–784.

[43] Tuomas Sandholm. 2000. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce* 4, 3 (2000), 107–129.

[44] Galit Shmueli, Ralph P Russo, and Wolfgang Jank. 2007. The BARISTA: A model for bid arrivals in online auctions. *The Annals of Applied Statistics* 1, 2 (2007), 412–441.

[45] William Vickrey. 1961. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* 16, 1 (1961), 8–37.

[46] Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, and Omri Ross. 2022. NFT Wash Trading: Quantifying suspicious behaviour in NFT markets. *arXiv preprint arXiv:2202.03866* (2022).

[47] Qin Wang, Rujia Li, Qi Wang, and Shiping Chen. 2021. Non-fungible token (NFT): Overview, evaluation, opportunities and challenges. *arXiv preprint arXiv:2105.07447* (2021).

[48] Z John Zhang. 2022. Cryptopricing: Whence Comes the Value for Cryptocurrencies and NFTs? *International Journal of Research in Marketing* (2022).

A Proof of Claim 1

From the defining equation $p_D(q^\star) = p_S(q^\star, T)$ and the smoothness of p_D, we can use the implicit function theorem to write

$$q^\star = q_D(p_S(q^\star, T)) = p_D^{-1}(K(\lambda)(b^{q^\star} - 1)) \leq LK(\lambda)(b^{q^\star} - 1)$$

The first equality used the fact that $p_S(q^\star, T) = \frac{kb^{1-T}}{b-1}(b^{q^\star} - 1) = K(\lambda)(b^{q^\star} - 1)$ from eq. 1 and the inequality used the fact that Lipschitz $p_D^{-1}(p)$ is Lipschitz for some constant $L > 0$ and $\lim_{p\to\infty} q_D(p) = 0$. The above equation can be rewritten as

$$Ab^{q^\star} - q^\star \geq A$$

where $A = LK(\lambda)$. In terms of the Lambert W function, the solution for $Ae^{q \log b} - q = A$ this can be written as

$$q^\star = -\frac{W\left(-\frac{A \log b}{b^A}\right) + A \log B}{\log b} = -\frac{W\left(-\frac{A \log b}{b^A}\right)}{\log b} + A$$

Thus provided that $-\frac{A \log b}{b^A}$ is in the domain of the Lambert W function, which holds immediately from $b \geq 1, A > 0$, we have the conclusion of the claim.

B Proof of $\hat{q} < q^\star$ for Generic GDAs

We show here that for generic GDAs with price functions $p_n(t) = a_i \hat{p}(t - t_i)$, the modified clearing quantity under the attack $\hat{q}$ and the equilibrium quantity $q^\star$ satisfy $\hat{q} < q^\star$. Note that $\hat{q}$ and $q^\star$ satisfy:

$$p_D(q^\star) = \sum_{n=1}^{q^\star} p_n(T-n)$$
$$p_D(\hat{q}) = \sum_{n=r+1}^{\hat{q}+r+1} p_n(T-n)$$

where for brevity we denote $p_n(T-n) = a_i \hat{p}_n(T-n)$. Assume for the sake of contradiction that $\hat{q} > q^\star$. This implies that $p_D(\hat{q}) - p_D(q^\star) > 0$. We have two cases: $\hat{q}+r+1 \leq q^\star$ and $q^\star < \hat{q}+r+1$. In the first case, when $\hat{q}+r+1 \leq q^\star$. Then:

$$\begin{aligned} p_D(\hat{q}) - p_D(q^\star) &= \sum_{n=r+1}^{\hat{q}+r+1} p_n(T-n) - \sum_{n=1}^{q^\star} p_n(T-n) \\ &= -\sum_{n=\hat{q}+r+1}^{q^\star} p_n(T-n) - \sum_{n=1}^{r} p_n(T-n) < 0 \end{aligned}$$

which is a contradiction. On the other hand, when we have $q^\star < \hat{q}+r+1$, when we have:

$$\begin{aligned} p_D(\hat{q}) - p_D(q^\star) &= \sum_{n=r+1}^{\hat{q}+r+1} p_n(T-n) - \sum_{n=1}^{q^\star} p_n(T-n) \\ &= \sum_{n=q^\star+1}^{\hat{q}+r+1} p_n(T-n) - \sum_{n=1}^{r} p_n(T-n) \end{aligned}$$

We can bound this as

$$\sum_{n=q^\star+1}^{\hat{q}+r+1} p_n(T-n) - \sum_{n=1}^{r} p_n(T-n) \leq p_{\hat{q}+r+1}(T-\hat{q}+r+1)(\hat{q}-q^\star+r+1) - r p_1(T-1)$$

where we use the fact that a_i is non-decreasing and $\hat{p}(t)$ is decreasing in t, so $p_n(T-n)$ is non-decreasing in n. As the right hand side is less than or equal to zero (which has to be true since p_D is non-increasing and by assumption $\hat{q} > q^\star$) this implies that

$$\hat{q}-q^\star < \frac{r p_1(T-1)}{p_{\hat{q}+r+1}(T-\hat{q}+r+1)} - r - 1 = r\left(\frac{p_1(T-1)}{p_{\hat{q}+r+1}(T-\hat{q}+r+1)} - 1\right) - 1 < r\left(\frac{a_1}{a_{\hat{q}+r+1}} - 1\right) - 1$$

Since the a_i are non-decreasing, the right hand side is negative, implying that $\hat{q} - q^\star < 0$ or $\hat{q} < q^\star$.

C Proof that $r < q^\star$ for general GDAs

We now show that $r < q^\star$, having proved in the last section that $\hat{q} < q^\star$. We immediately see that if $r > \hat{q}$, then

$$\begin{aligned}\mathsf{PNL}(\hat{q}, T) &= R(\hat{q}, T) - C(r) \\ &= \sum_{n=r+1}^{\hat{q}} p_n(T-n) - \sum_{n=1}^{r} a_n < 0\end{aligned}$$

because $\sum_{n=r+1}^{\hat{q}} p_n(T-n) = 0$ due to the fact that $r+1 > \hat{q}$. That is, if $r > \hat{q}$, the revenue from the deviation is immediately negative, and thus no rational seller would choose this. Therefore, $r < \hat{q} < q^\star$.

D $\mathcal{M}$ is a matroid

We now show that the pair $([n] \times [T], \mathcal{M})$ is a matroid. First, note that that $\mathcal{M}$ is downward closed by construction. This is because we incrementally add admissible auction winners to existing sets. We will show basis expansion via contradiction. Suppose that the basis expansion property doesn't hold and there are sets $A, B \in \mathcal{M}$ such that $|B| > |A|$ but there is no (i, ℓ) such that $A \cup \{(i, \ell)\} \in \mathcal{M}$. Since $(i, \ell) \in B \in \mathcal{M}$, we know that $v_i(s_{i1}, \dots s_{ik}) \geq a_\ell \hat{p}(t_k - t_\ell)$. Since $A \cup \{(i, \ell)\} \notin \mathcal{M}$, we must have $\sum_{(j,r)\in A} a_r \hat{p}(t_k - t_r) - v_j(s_{j1}, \dots, s_{jk}) > v_i(s_{i1}, \dots, s_{ik}) - a_\ell \hat{p}(t_k - t_\ell)$..But since $A \in \mathcal{M}_k$, we have

$$\sum_{(j,r)\in A} v_j(s_{j1}, \dots, s_{jk}) \geq \sum_{(j,r)\in A} a_r \hat{p}(t_k - t_r)$$

so $v_i(s_{i1}, \dots, s_{ik}) - a_\ell \hat{p}(t_k - t_\ell) \leq 0$, a contradiction for the first condition.

To see that the second condition cannot be violated, note that since $(i, \ell) \in B - A$ we have

$$\sum_{(j,r)\in B} v_j(s_{j1}, \dots, s_{jk}) \geq \sum_{(j,r)\in B} a_r \hat{p}(t_k - t_r)$$

or

$$\begin{aligned}v_i(s_{i1}, \dots, s_{ik}) - a_\ell \hat{p}(t_k - t_\ell) &\geq \sum_{(j,r)\in B - \{(i,\ell)\}} a_r \hat{p}(t_k - t_r) - v_j(s_{j1}, , \dots, s_{jk}) \\ &\geq \min_{S \subset \mathcal{M}_k} \sum_{(j,r)\in S} a_r \hat{p}(t_k - t_r) - v_j(s_{j1}, \dots, s_{jk})\end{aligned}$$

since $\mathcal{M}$ is downward closed. Therefore there must exist $(i, \ell) \in B$ such that $A \cup \{(i, \ell)\} \in \mathcal{M}$ and $\mathcal{M}$ is a matroid.