

An Analysis of Intent-Based Markets

Tarun Chitra
Gauntlet
tarun@gauntlet.xyz

Kshitij Kulkarni
UC Berkeley
ksk@eecs.berkeley.edu

Mallesh Pai
Rice University and SMG
mallesh.pai@mechanism.org

Theo Diamandis
MIT
tdiamand@mit.edu

June 7, 2024

Abstract

Mechanisms for decentralized finance on blockchains suffer from various problems, including suboptimal price execution for users, latency, and a worse user experience compared to their centralized counterparts. Recently, off-chain marketplaces, colloquially called ‘intent markets,’ have been proposed as a solution to these problems. In these markets, agents called *solvers* compete to satisfy user orders, which may include complicated user-specified conditions. We provide two formal models of solvers’ strategic behavior: one probabilistic and another deterministic. In our first model, solvers initially pay upfront costs to enter a Dutch auction to fill the user’s order and then exert congestive, costly effort to search for prices for the user. Our results show that the costs incurred by solvers result in restricted entry in the market. Further, in the presence of costly effort and congestion, our results counter-intuitively show that a planner who aims to maximize user welfare may actually prefer to restrict entry, resulting in limited oligopoly. We then introduce an alternative, optimization-based deterministic model which corroborates these results. We conclude with extensions of our model to other auctions within blockchains and non-cryptocurrency applications, such as the US SEC’s Proposal 615.

1 Introduction

Some of the most popular applications used on blockchains are decentralized exchanges, which allow users to trustlessly swap assets without a third-party intermediary. Such exchanges have facilitated over \$2 trillion of transactions since 2018, and their usage continues to increase [Lla23]. However, decentralized exchanges (DEXs) still have worse user experience and price execution than centralized exchanges (CEXs), as they tend to be less capital efficient. That is, it takes more passive capital to facilitate trades on a DEX than on a CEX.

Decreased capital efficiency on DEXs results from their reliance on continuous arbitrage to synchronize prices. The most popular type of DEX, the Constant Function Market Maker (CFMM), was first shown to synchronize prices via arbitrage with an external market in 2019 [Ang+19; AC20]. Recent work has demonstrated that this continuous arbitrage causes liquidity providers to face large losses and adverse selection costs, as their pricing and liquidity are completely public at all times [Mil+22; AEC21; CDM23]. Moreover, the permissionless nature of CFMMs has empirically led to excess liquidity fragmentation [LPZ23] and maximal extractable value, or MEV, which does not have a precise analogue in the centralized world [XFP23; KDC22; Dai+20]. Despite these shortcomings, CFMMs still command substantial trading volume [Blo23], and new mechanisms for improved capital efficiency, such as concentrated liquidity CFMMs [Ada+21], have narrowed the gap between decentralized and centralized exchanges.

Intent-based markets. Recently, there have been numerous proposals for settlement of on-chain transactions collectively known as *intents* that claim to close the gap between centralized and decentralized exchanges. Intent-based mechanisms allow users to specify precise conditions or covenants under which their transactions can be executed. Transactions are then executed on a blockchain by third parties called *solvers*. Cryptographic primitives ensure that solvers cannot violate the user-specified conditions, and the validators of the blockchain verify that these conditions were satisfied. CFMMs, conversely, deterministically quote a price based on public liquidity, and therefore, in contrast to intent-based markets, their execution price is known publicly.

Intents were first described by Anoma [Yin23] as a way of credibly executing asynchronous transactions across multiple blockchains. These mechanisms generalize ‘request-for-quote’ (RFQ) systems [Mor21; OZ21; O’H10]. In particular, intents allow users to express arbitrary covenants, whereas RFQ systems do not. These covenants are guaranteed to hold throughout the execution of the transaction on multiple blockchains; a transaction is only settled after all blockchains have agreed on execution via a communication protocol.

Solvers are assumed to be rational economic actors who seek to satisfy a user’s order only if it benefits them. Importantly, they hold or can acquire private information about the state of external markets that unsophisticated users cannot acquire. The central question is, therefore, whether a solver is incentivized to execute the user’s order at ‘optimal’ prices. In this sense, intent markets resemble principal-agent problems in which the user is a principal who employs an agent—the solver—to execute their order under the intent covenants.

Example intents. To provide concrete examples of user intents, we consider a user who holds 1 unit of token T_1 on blockchain B_1 (for example, the ETH token on the Ethereum blockchain). The user specifies an intent to convert the token into the maximum amount of output token T_2 on the same blockchain (for example, USDC on the Ethereum blockchain). We call this the *maximum output* intent. As another example, consider the same user who now wants to move their token on blockchain B_1 to another blockchain B_2 (for example, from the Ethereum blockchain to the Solana blockchain) only if the realized variance of the

token on a set of predefined venues on chain B_2 is less than a threshold. We call this the *variance limited transfer* intent.

While these two intents are superficially similar, only one can be easily verified after execution. In the first case, the validator is not able to verify that the user did, in fact, receive the maximum available output of token T_2 from all possible external venues. In the second case, however, a validator on chain B_2 is able to verify that the realized variance of the token did meet the conditions specified by the user. In this work, we do not focus on the issue of verifiability, but rather consider the economic incentives of solvers in fulfilling these intents.

Intents in practice: UniswapX. Several intent-based marketplaces are already live in production in 2024. The first example we consider is UniswapX [Ada23]. Since its launch in late 2023, UniswapX has processed over \$3 billion in notional volume [p24]. This protocol incentivizes off-chain agents to provide users with better execution prices than those provided by CFMMs and, therefore, implements the maximum output transfer intent introduced earlier. In UniswapX, solvers participate in a Dutch auction to provide price improvement over the price quoted by a CFMM. Solvers in UniswapX are referred to as ‘fillers’. The solver who wins the auction then commits to filling the order at an execution price equal to their bid. The winning solver has to deliver the specified number of tokens within a fixed-time interval. If the solver delivers the assets within the time-interval, the transaction is settled: the user receives the delivered assets less a payment to the solver for finding a better price. Should the solver be delinquent and not deliver the assets, they are penalized and the user’s trade is sent to a CFMM such as Uniswap. We will use this Dutch auction setup as a starting point for our analysis of solvers’ incentives shortly.

Intents in practice: other protocols. We now review several other intent protocols. CoW Swap [Pro23; Leu21; RGM23], another intent-based protocol that implements the maximum output intent, has processed over \$30 billion since its inception in 2021 [CoW23]. In CoW Swap, solvers participate in an auction to execute user orders at a uniform clearing price, also known as a ‘batch auction’.

Another variant of an intent market is the so-called ‘oracle extractable value’ market [Sen23; Lam24; Ben23]. These markets run auctions that sell the right to execute a trade immediately after a price update is sent on-chain via a price oracle [AC20; LSZ21]. Such markets allow for decentralized finance protocols (who are the principals, akin to users in UniswapX) to outsource trade execution to agents and realize a welfare gain.

While there has been an immense amount of activity in building more general intent-based systems for purposes other than trading (*e.g.* SUAVE, Essential, and Anoma [Cha23]), these systems do not currently have any formal guarantees. Further, the economic incentives of the participants in these systems have not been studied.

Intents without blockchains. Another fulfillment model that is similar to intents but has been proposed for traditional equities markets is the US Security and Exchange Com-

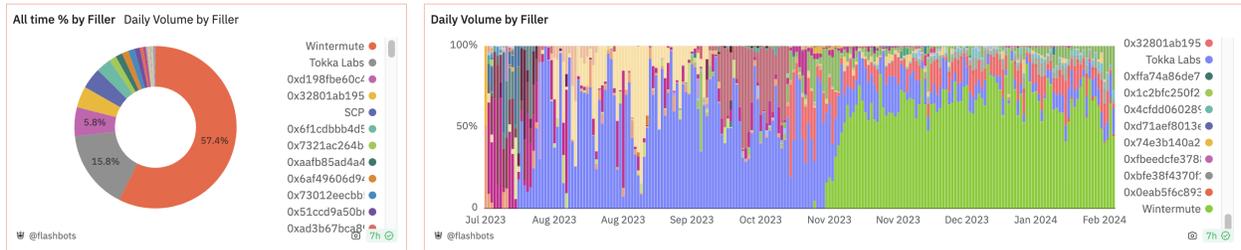


Figure 1: Evolution of competition within UniswapX over time [p24]. Left: the market share of different participants, in terms of the volume of orders filled. Right: the same data evolving over time. The data set has over 2,000 unique addresses that have participated, despite fewer than 20 addresses participating in January and February 2024.

mittee’s Proposal 615 [Sec22]. This rule aims to make Payment for Order Flow (PFOF) relationships more competitive. These relationships, between brokerage firms, who purchase shares on behalf of their users, and wholesalers who execute transactions and settle trades for brokerages, have been controversial [Ang21; BPS23; Wel22].

Currently, private agreements exist between brokerages and wholesalers that diminish competition, as noted in [Sec22]: “Broker-dealers route more than 90% of marketable orders of individual investors in NMS stocks to a small group of six off-exchange dealers, often referred to as ‘wholesalers.’. The wholesaling business is highly concentrated...” The SEC’s rule proposes the construction of a Dutch auction wherein a wholesaler who receives orders from a brokerage must auction right to execute the orders at improved prices before fulfilling it themselves. This extra auction is similar to blockchain intents like UniswapX or CoW Swap, but differs in the precise execution guarantees given to the user. We discuss this auction and the applicability of our model to analyzing it in §4.

Competition. Despite the interest in intent-based systems, these systems are not necessarily guaranteed to perform significantly better than existing atomic mechanisms such as CFMMs. Understanding conditions under which these systems provide improvements to existing mechanisms is therefore crucial to designing welfare-maximizing intent systems. Moreover, reliance on a network of solvers for fulfillment and execution in an intent market raises some questions: first, will the market remain competitive or concentrate to a few specialized actors? And second, how would such concentration affect user welfare?

Empirical signatures of reduced competition in UniswapX. In Figure 1, we see the evolution of competition within UniswapX [p24]. The left panel shows the percentage of volume processed by a particular solver and shows extreme concentration, with the top 3 solvers accounting for $\approx 79\%$ of daily volume transacted. The right panel shows that the number of entrants into the auction has also decreased dramatically. While over 2,000 unique addresses have participated as UniswapX fillers, with most of them participating in the earlier days of the protocol during Summer 2023, only 12 unique addresses participated in January 2024 (which can be viewed on the dashboard [p24]). This concentration of fillers

suggests that the UniswapX intent market suffers from a lack of competition, which may translate to user welfare losses. We note that the filler role is currently permissioned by Uniswap Labs [Ada23], so there is not free entry within this market. However, the heavy concentration to 12 out of 2,000 participating addresses suggests that the market has barriers to competition even among the existing set of entrants.

This paper. In this paper, we formulate two simple models of solvers bidding in an auction for a single user’s intent to receive the best possible price when converting one token into another. We aim to understand the solvers’ economic incentives to fulfill the user’s intent.

Our first model is probabilistic, relying on classical auction theoretic tools, where the standard auction setting is complicated by two features. First, the solvers initially face up-front *entry costs*, which correspond to infrastructure set-up costs that are endemic to any entity participating in intent markets. Second, they face *congestion costs* that are incurred upon entry into the market—these costs potentially depend on the number of other entrants, and are faced while competing with other solvers. For example, the setting may require participants to exert costly effort to find liquidity to satisfy the user’s order in external markets. Further, participants may have to invest in finding this liquidity prior to winning this auction, or indeed may need to expend this effort just to know how much to bid.

We study the effects of these costs in the context of a Dutch auction. We start by assuming there is a universe of n possible solvers who may enter into the auction. For each solver, entry costs are drawn i.i.d. from a distribution F_C . As a function of these costs, an equilibrium number k^* solvers enter the auction, and then search for prices to satisfy the user’s intent, which are drawn i.i.d. from a distribution F . This search is affected by the congestion costs and therefore, so is the solvers’ bid in the Dutch auction. In this model, we have the following results:

First, in the presence of only entry costs but no congestion costs, we show that the equilibrium number of entrants k^* into the solver market is $O(\sqrt{n})$ when the solvers’ prices are distributed exponentially, and $O(n^{1/3})$ when they are distributed uniformly on $[0, 1]$. This means that the percentage of potential solvers who participate in the auction goes to zero as $n \rightarrow \infty$ since $\frac{k^*}{n} \rightarrow 0$. We further show that for sufficiently fat-tailed price distributions, such as the Pareto distribution, one can recover $k^* = n$, although the user may not be able to realize large welfare increases due to significant bid shading.

Second, when congestion costs are present, the distribution of prices is modified to a distribution $F(p, e)$, where e is costly, congestive effort exerted by the solver. Under congestive costs (such as those that might occur while searching for inventory in illiquid external markets), the effort that a solver exerts in the search for better prices for the user degrades, and, therefore, so does the user’s welfare. In practice, this result suggests that intent markets may provide welfare gains over CFMMs in highly liquid markets where congestion costs are small, but they likely do not do so in markets for extremely illiquid assets. Our results therefore connect costs faced by solvers to the welfare gains that intent-based markets can promise to users.

Our second model takes a deterministic approach, which is based on an optimization

framework. This model naturally suggests the Dutch auction as a means to solve the social welfare problem. Although it introduces different modeling assumptions compared to our first model, the deterministic model offers two advantages: it can be easily parameterized, suggesting a fruitful avenue for future empirical work; and it likely has an easier generalization to the multi-asset case, which we also leave to future work.

Our analysis also implies several interesting features of intent markets. Entry costs may cause restricted entry into the markets. Further, when solvers are subject to congestion, limited oligopoly can be better for user welfare than free entry; fewer entrants will have stronger incentives to invest in higher effort due to less congestion. The welfare gains for the user from additional investment from fewer entrants may outweigh the losses to the user from increased bid-shading in an oligopolistic market. Understanding the underlying costs faced by solvers is therefore important for an intent platform to optimally design the market microstructure.

2 An auction-theoretic model

We begin with the simplest setting for an intent market: a single decentralized trade between two tokens. Consider a user who desires to swap 1 unit of token T_1 into as many units of token T_2 as possible (this is an example of the maximum output intent introduced earlier). The user is faced with a choice to trade either with an infinitely liquid public market (an idealization of a CFMM) that quotes a price p^* , or with one of n possible solvers, where each solver i has access to inventory of token T_2 , at a price p_i in token T_1 . The user only trades with a solver if they offer a price improvement above the public market. We assume that p_i are distributed i.i.d. according to a distribution with CDF $F(p)$ and density $f(p)$. This price could be either the price of their inventory or the price of an external market in which they have infinite or arbitrarily large liquidity relative to the user’s demand. The user, in an attempt to elicit the best price for their tokens and improve upon the prices available from the public market, runs an auction (such as the Dutch auction found in UniswapX [Ada23]). We will examine several settings, in increasing order of complexity.

Entry costs, costly effort, and congestion. Solvers face a sequence of choices in attempting to fill a user’s order. First, they must choose whether or not to enter the auction in the first place. This choice is often driven by initial setup and infrastructure costs that are required to bid on and execute a user’s orders. We call these *entry costs*. These entry costs determine an equilibrium number k^* of the universe of n solvers that will enter the auction. Upon making the choice to enter, the solvers then must decide how to bid. This choice, in our model, is driven by the fact that solvers expend costly effort to search for a price to win and satisfy a user’s order, which modifies the distribution of prices to $F(p, e)$, where e is the effort level chosen by the solver. In this case, solvers may try to simultaneously access a set of external markets, and therefore, effort may cause congestion. We seek to understand how the solvers’ bidding behavior changes as a function of these two kinds of costs they face. In summary, we have the following settings and their associated results:

1. *No cost of entry, no effort costs:* No cost of entry implies that the equilibrium number of entrants into the solver market satisfies $k^* = n$. Classical revenue equivalence [Kri09] applies and user welfare is maximized.
2. *Costly entry:* In equilibrium, the intent market has an asymptotically negligible number of entrants k^* (that is, $k^*/n \rightarrow 0$ as $n \rightarrow \infty$) for reasonable distributions like the exponential and uniform distributions. For the Pareto distribution, however, in anticipation of large prices, we may have $k^* = n$ entrants.
3. *Costly effort:* We demonstrate cases where welfare is reduced due to effort causing congestion. This indicates that a planner may in fact want to restrict entry into such a market.

These results demonstrate that entry and effort costs may lead to oligopoly, defined as the combination of reduced competition ($k^* = o(n)$) and reduced welfare or auction revenue. Critically, in the case of effort costs, a social planner may wish to limit entry, as a closed or permissioned solver market may achieve better outcomes from the perspective of a welfare-maximizing planner. Understanding when this occurs in practice, and how such permissioned markets can be implemented is an important design consideration that we leave for future work.

2.1 Dutch auctions for intents

The Dutch auction has been implemented by various protocols, including UniswapX [Ada23], to execute a single user's swap intent in a market with multiple solvers. It is well-known that the Dutch auction is strategically equivalent to a sealed-bid first price auction [Kri09]. Therefore, we analyze the solvers' strategies in the setting of a first price auction, where the highest bidding solver wins at the highest price. Each solver submits a price $\tilde{p}_i \leq p_i$, which is their bid. According to the rules of the first price auction, the highest bid gets filled as long as it is better than p^* . The net profit of a solver from being filled therefore becomes:

$$p_i - \tilde{p}_i.$$

If k of the universe of n possible solvers are present in the auction, revenue equivalence yields that

$$(p_i - \tilde{p}_i)F^{k-1}(p_i) = \int_{p^*}^{p_i} F^{k-1}(x)dx,$$

or, equivalently,

$$\tilde{p}_i = (p_i) - \frac{\int_{p^*}^{p_i} F^{k-1}(x)dx}{F^{k-1}(p_i)}.$$

This formula has a very simple interpretation: in a first price auction, in the absence of entry or congestion costs, each solver shades their bid according to the competition they face in the auction.

Interim profits. In this case, the interim expected profit of a solver with price p_i in a setting with k other solvers can be written as:

$$S(p_i, k) = \int_{p^*}^{p_i} F^k(x) dx.$$

We can therefore write the ex-ante expected profit of a solver as

$$S(k) = \int_0^\infty S(p, k) f(p) dp.$$

Note that:

$$\begin{aligned} S(k) &= \int_0^\infty S(p, k) f(p) dp \\ &= - \int_0^\infty S(p, k) \frac{d(1 - F(p))}{dp} dp \\ &= - \left(S(p, k)(1 - F(p)) \Big|_{p=0}^{p=\infty} - \int_0^\infty S'(p, k)(1 - F(p)) dp \right) \\ &= \int_{p^*}^{\bar{p}} F^k(p)(1 - F(p)) dp. \end{aligned}$$

Therefore, by observation, S is decreasing and satisfies increasing differences in k . Note that $S : \mathbf{N} \rightarrow \mathbf{R}$ can be extended to the positive reals using *e.g.*, linear interpolation between successive integers. The resulting function is convex, and we abuse notation to consider this extended function in the sequel.

Exponential distribution. First suppose that that F is the exponential distribution with rate λ and $p^* = 0$. Then,

$$\begin{aligned} S(k) &= \int_0^\infty (1 - \exp(-\lambda p))^k \exp(-\lambda p) dp \\ &= \frac{1}{\lambda} \int_0^1 x^k dx \\ &= \frac{1}{(k+1)\lambda}, \end{aligned}$$

where the second equality is by substituting $1 - \exp(-\lambda p) = x$.

Uniform distribution. Now suppose that F is the uniform distribution on $[0,1]$ and $p^* = 0$. Then,

$$S(k) = \int_0^1 p^k (1 - p) dp,$$

which in turn implies that

$$S(k) = \frac{1}{(k+1)(k+2)}.$$

Pareto distribution. Finally, suppose that F is the generalized Pareto distribution with CDF $F(x) = 1 - \left[1 + \left(\frac{x-\mu}{\sigma}\right)^{1/\gamma}\right]^{-\alpha}$, where $x > \mu$, and $\sigma, \gamma, \alpha > 0$, and that $p^* = 0$. We consider this distribution in the regime $\gamma/\alpha > 1$, for which its mean is infinity. In this case, we have that

$$\begin{aligned} S(0) &= \int_0^\infty 1 - F(p) dp \\ &= \mathbf{E}_{p \sim f}[p] \\ &= \infty. \end{aligned}$$

We will use this fact shortly to show that in this case, the market admits all n potential solvers as entrants.

2.2 Costly entry

Now, we assume that prior to bidding in the auction, every solver must pay an entry cost of c_i if they choose to bid. These costs correspond to the setup and infrastructure costs considered previously. For now, we consider the case that the solver does not know how many other solvers have entered or their oracle price; they choose to enter only based on their own price and cost. Costs are nonnegative and distributed according to a given distribution F_C . In the subsequent, we will always assume that F_C has a density f_C for whom $f_C(0)$ is finite and nonzero. Under these entry costs, we will show that expected number of entrants in the auction, k^* , is small relative to the universe of possible entrants, n .

Market size as a function of costs. The above setup allows us to study entry into the market by providing a threshold cost \bar{c} such that it is not profitable in expectation for a solver to enter the auction if their cost exceeds the threshold. The maximum cost solver who will choose to enter this market will be one with cost \bar{c} where \bar{c} solves:

$$\sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c}) (1 - F_C(\bar{c}))^{n-k} \right) S(k) = \bar{c}. \quad (1)$$

Recall that solvers enter the auction if their cost is less than the threshold cost \bar{c} . This happens with cumulative probability $F_C(\bar{c})$. Therefore, once the threshold \bar{c} has been found, the expected number of entrants into the solver market can be computed as:

$$k^* = n F_C(\bar{c}).$$

We now show how k^* depends on n for various realistic distributions F .

Exponential distribution. For the case of exponentially distributed prices with rate λ , we have that the left hand side of the above equation (1) can be written as:

$$\sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c}) (1 - F_C(\bar{c}))^{n-k} \right) S(k) = \frac{1}{(n+1)\lambda F_C(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+1}).$$

This equation is derived in Appendix A. So, for exponentially distributed prices we must have that

$$\frac{1}{(n+1)\lambda F_C(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+1}) = \bar{c},$$

which implies

$$\frac{1}{(n+1)\lambda} = \frac{\bar{c} F_C(\bar{c})}{(1 - (1 - F_C(\bar{c}))^{n+1})}.$$

From this equation, we straightforwardly have that \bar{c} must be decreasing to 0 as $n \nearrow \infty$. This is because \bar{c} asymptoting to a constant would result in the left hand side going to 0 while the right hand side stays constant. Now suppose that $\bar{c} = g(n)$ for some decreasing function g and take n large so that we can do Taylor approximations. Then, the above equation becomes

$$\frac{1}{(n+1)\lambda} = \frac{(g(n))^2 f_C(0)}{(1 - (1 - f_C(0)g(n))^{n+1})}.$$

This implies that $g(n) = O(\frac{1}{\sqrt{n}})$ since for this rate of convergence the denominator of the right hand side converges to 1 because $(1 - f_C(0)g(n))^{n+1} \rightarrow 0$ and the numerator matches the left hand side. Recall that the expected number of entrants into the solver market, k^* , equals $nF_C(\bar{c})$. Therefore, for large n , we have $k^* = O(\sqrt{n})$ for exponentially distributed values.

Uniform distribution. Now, for the case of prices uniformly distributed between $[0, 1]$, we have that the left hand side of equation (1) is:

$$\begin{aligned} & \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c}) (1 - F_C(\bar{c}))^{n-k} \right) S(k) \\ &= \frac{1}{(n+1)(n+2)F_C^2(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+2} - (n+2)F_C(\bar{c})(1 - F_C(\bar{c}))^{n+1}) \end{aligned}$$

This equation is derived in Appendix B. Therefore, for uniformly distributed prices, we must have that

$$\frac{1}{(n+1)(n+2)F_C^2(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+2} - (n+2)F_C(\bar{c})(1 - F_C(\bar{c}))^{n+1}) = \bar{c},$$

which implies

$$\frac{1}{(n+1)(n+2)} = \frac{\bar{c} F_C^2(\bar{c})}{(1 - (1 - F_C(\bar{c}))^{n+2} - (n+2)F_C(\bar{c})(1 - F_C(\bar{c}))^{n+1})}.$$

Similar arguments as in the case of exponentially distributed prices now tell us that $\bar{c} = O(n^{-\frac{2}{3}})$ and therefore the expected number of entrants is $k^* = O(n^{1/3})$.

Pareto distribution. Finally, when F is the generalized Pareto distribution (in the regime where $\gamma/\alpha > 1$), recall that this distribution has infinite mean, which means in particular that $S(0)$, the expected interim profits with no other participating solvers, is infinity. This immediately allows us to conclude that \bar{c} in this case is also infinity, as we have:

$$(1 - F_C(\bar{c}))^n S(0) \leq \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c}) (1 - F_C(\bar{c}))^{n-k} \right) S(k) = \bar{c}.$$

Because $S(0) = \infty$, the only \bar{c} that makes the above inequality hold is $\bar{c} = \infty$. Therefore, because $F_C(\bar{c}) \rightarrow 1$ as $\bar{c} \rightarrow \infty$, we have that the total number of entrants into the auction is $k^* = nF_C(\bar{c}) = n$.

Connections to extreme spacings. The quantity $S(k)$, for a given distribution F , is intimately related to the expected difference between the highest and second highest price in a sample of $k + 1$ prices drawn i.i.d. from F . This gap is also known as the *extreme spacing* [MCT09]. For a distribution with CDF F and PDF f , define for any $k \in \mathbf{N}$ the following quantity:

$$ES(k) \equiv \mathbf{E}_{p_i \sim f} [p_{k:k} - p_{(k-1):k}],$$

where $p_{k:k}$ and $p_{(k-1):k}$ denote the largest and second-largest of k i.i.d. draws according to f . We show in Appendix C that if $p^* = 0$, then for any distribution, we have that

$$(k + 1)S(k) = ES(k + 1).$$

For the generalized Pareto distribution, it was shown in [MCT09, Theorem 6.1] that $ES(k) = O(k^{\gamma/\alpha})$, which allows us to conclude $S(k) = O(k^{\gamma/\alpha-1})$. Such bounds may allow one to derive rates for entry into the auction for a wider class of price distributions, which we leave for future work.

Discussion and numerical results. Our first set of results on costly entry demonstrate several interesting features of intent markets. For both the exponential and uniform distributions, entry costs can cause a much smaller fraction of potential solvers to enter the auction than would otherwise enter in the absence of costs. For reference, when $n = 100$, one may have as few as $k^* = 10$ entrants in the case of exponentially distributed prices, and as few as $k^* = 5$ entrants in the case of uniformly distributed prices. These results have a natural economic interpretation: solvers decide to enter the auction depending on whether they believe there is an opportunity to receive net expected profit that exceeds their threshold costs. In the case of the exponential distribution, the probability that a solver receives a price that is much better than another solver decays exponentially. The situation is even worse in the case of the uniform distribution, where there is an upper bound on the best price available to the solvers.

This can be remedied in the case of the Pareto distribution, where even being a single solver in the market yields infinite expected profits. Therefore, in anticipation of these large

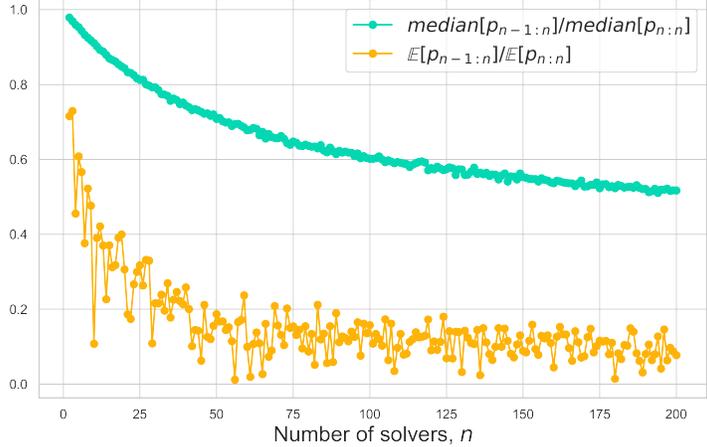


Figure 2: Ratio of the expected revenue to the expected highest price, $\mathbf{E}[p_{n-1:n}]/\mathbf{E}[p_{n:n}]$, for the Pareto distribution, versus the number of solvers, n . The ratio of the medians of $p_{n-1:n}$ and $p_{n:n}$ is plotted for reference.

profits, we can expect n entrants into the market. We note, however, that even in this case, this does not mean that the user is able to realize a large price, as there may be significant bid shading in the auction. We demonstrate this effect in Figure 2, where we plot the ratio of the expected second highest price (equivalent to the expected revenue of the auction, which in turn is the expected price the user receives) to the expected highest price, $\mathbf{E}[p_{n-1:n}]/\mathbf{E}[p_{n:n}]$, for n i.i.d. draws of prices p_i from the Pareto distribution. This ratio tells us how much the price the user can expect to receive from a solver auction with n participants falls short of the theoretical highest price that could be achieved by any solver. The ratio of the medians of $p_{n-1:n}$ and $p_{n:n}$ are also plotted for reference.

Our simulation is done by setting $\sigma = 100$, $\alpha = 0.95$, $\gamma = 1$, $\mu = 0$, which corresponds to a standard Pareto distribution with CDF $F(x) = 1 - \left(\frac{x_m}{x}\right)^\alpha$, with scale parameter $x_m = 100$ and shape parameter $\alpha = 0.95$. For each n , we compute the expectations of $p_{n-1:n}$ and $p_{n:n}$ by averaging over 10000 trials. Code for this simulation can be found at

<https://github.com/kkulk/intent>.

We see that as n gets large, the ratio $\mathbf{E}[p_{n-1:n}]/\mathbf{E}[p_{n:n}]$ approaches zero. This means that even with a large number of solvers present in the auction, the user is able to realize a price that is only a small fraction of the theoretical highest price available to any of the solvers. This suggests that intent markets in assets with heavy tails may fail to realize user welfare that is much greater than that afforded by public mechanisms such as CFMMs.

2.3 Costly effort

So far, we have considered markets where costs of entry were fixed for each potential entrant, independent of the number of entrants, and similarly the distribution of prices for each

entrant is exogenously given. Now, instead, suppose that entrants do not simply receive a price upon entry but have to invest costly effort. This can be viewed as investment into infrastructure or technology in order to be able to participate in an auction—an example: a high-frequency trader who has to spend millions of dollars to achieve a sufficiently low message latency to successfully participate in an auction. Further, this can also correspond to the costs of sourcing liquidity for assets that the solver does not initially hold. As we will see, if these effort costs cause congestion (as they might in low-liquidity assets), user welfare may be harmed.

Effort-dependent prices and costs. Formally, suppose now that k^* entrants have entered the auction (perhaps by the process of costly entry introduced in the previous section). Each entrant i simultaneously chooses an investment level $e_i \in \mathbf{R}_+$. This investment results in the entrant drawing a price $p_i \sim F(p, e_i)$ for a CDF F dependent on the amount of investment. We make the assumption that investments improve values in the sense of first-order stochastic dominance: $F(p, e) \leq F(p, e')$ for any p and $e > e'$.

Effort is both costly and congestive. We model the cost of taking effort e when the total number of entrants is k^* by $c(k^*, e)$ where c is increasing and convex in both arguments. A special case of interest is when $c(\cdot)$ is constant in the first argument, so the cost of effort is not affected by other other entrants, *i.e.*, there is no congestion cost. Note that this setup is a symmetric game between the k^* solvers that enter the auction and thus admits a symmetric equilibrium, which we search for next.

Equilibrium effort. We now solve for the equilibrium investment e^* by each entrant in this model. To do so, we pick a solver i and assume that all $k^* - 1$ competing solvers play the equilibrium strategy of drawing $p \sim F(x, e^*)$. Given this strategy, the expected revenue of solver i with price p is:

$$S(p, k^*, e^*) := \int_0^p F(x, e^*)^{k^*-1} dx$$

This expression corresponds to the probability that solver i 's value is higher than the remaining $k^* - 1$ solvers. Therefore, given that all the other entrants invest e^* , entrant i must pick their level of investment e_i to maximize their expected profit:

$$e_i = \operatorname{argmax}_{e \in \mathbf{R}_+} \int S(p, k^*, e^*) f(p, e) dp - c(k^*, e).$$

Assuming that $e^* > 0$, we must also have e_i satisfy a first order condition at $e = e^*$. This condition, assuming that the densities and costs are differentiable, is:

$$\left(\int S(p, k^*, e^*) \frac{\partial f(p, e)}{\partial e} dp - \frac{c(k^*, e)}{\partial k} \right) \Big|_{e=e^*} = 0. \quad (2)$$

This equation implicitly describes e^* as a function of k^* , which allows us to measure congestion and effort's impacts on net revenue.

Properties of the equilibrium. We now consider the following specific example: $F(p, e) = p^e$ for $p \in [0, 1]$. Note that therefore $f(p, e) = ep^{e-1}$. Informally, an investment of e results in e draws from a uniform distribution, with the entrant's value being the maximum of these. We assume from our results in §2.2 that $k^* = o(n)$ solvers have entered the auction.

Further, suppose $c(k^*, e) = \alpha(k^*)\frac{e^2}{2}$, where $\alpha(\cdot)$ is an increasing function which parametrizes the 'congestiveness' of the underlying market as a function of the number of entrants in the market. This functional form matches one's intuition for inventory cost: if all the entrants have to carry inventory in an intent market, then the existence of other entrants drives up the costs of sourcing and carrying the inventory. In this case we have

$$S(p, k^*, e^*) = \frac{p^{(k^*-1)e^*+1}}{(k^*-1)e^*+1}$$

Substituting this into our first order condition (2) we therefore have

$$\int_0^1 \frac{p^{(k^*-1)e^*+1}}{(k^*-1)e^*+1} (p^{e^*-1} + e^*p^{e^*-1} \ln p) dp - \alpha(k^*)e^* = 0$$

Solving this integral yields the following equation:

$$\frac{1}{(1 + e^*k^*)^2} - \alpha(k^*)e^* = 0 \implies \alpha(k^*)e^*(1 + e^*k^*)^2 = 1.$$

Since $\alpha(\cdot)$ is nondecreasing, we must have that the solution e^* is decreasing in the number of entrants k^* . Note that we can write the left hand side as

$$\frac{\alpha(k^*)}{k^*} (e^*k^*(1 + e^*k^*)^2) = 1$$

Observe that the function $e^*k^*(1 + e^*k^*)^2$ is increasing in e^*k^* . This implies that, viewing e^* as an implicit function of k^* , if $\alpha(k^*)$ is linear (*i.e.* $\alpha(k^*) = \Theta(k^*)$), we have that e^*k^* is constant in k^* or $e^* = c/k^*$ for some constant c . Conversely, if $\alpha(k^*)$ grows faster (slower) than linear, (*i.e.*, $\alpha(k^*)/k^*$ is increasing (decreasing) in k^*), then we have that e^*k^* is decreasing (increasing) in k^* . Since $k^* = \omega(1)$, this implies that if $\alpha(k^*) = \Omega(k^*)$ then $e^* = o(1)$ in k^* (and hence, n). In other words, a larger universe of searchers leads to vanishing expected equilibrium effort as $\frac{e^*}{k^*} \rightarrow 0$.

Welfare as a function of effort. We now analyze the welfare of the user as a function of effort e^* and equilibrium participation k^* . Note that user welfare is maximized when the expected revenue of the auction is maximized, since that is the (expected) price that the user receives. By revenue equivalence, the expected revenue equals the revenue of a second-price

auction, so we can write the revenue as:

$$\begin{aligned}
\text{Rev}(e^*, k^*) &= \int_0^1 pk^*(k^* - 1)(1 - F(p, e^*))f(p, e^*)(F(p, e^*))^{k^* - 2} dp \\
&= \int_0^1 pk^*(k^* - 1)(1 - p^{e^*})(e^* p^{e^* - 1})(p^{e^*})^{k^* - 2} dp \\
&= \frac{e^{*2}(k^* - 1)(k^*)}{(1 + e^*(k^* - 1))(1 + e^*k^*)} \\
&= \frac{e^*(k^* - 1)}{(1 + e^*(k^* - 1))} \times \frac{e^*k^*}{(1 + e^*k^*)} \tag{3}
\end{aligned}$$

We make two claims about this formula:

1. If $\alpha(k^*) = o(k^*)$, then $\text{Rev}(e^*, k^*)$ increases in k^*
2. If $\alpha(k^*) = \omega(k^*)$, then $\text{Rev}(e^*, k^*)$ decreases in k^*

The first claim states that if the congestion cost is sublinear in k^* , then the planner has no reason to restrict entry into the auction, since every new entrant increases their revenue. On the other hand, the second claim states that if the congestion cost is superlinear in k^* , then the user would prefer to *restrict entry*, *i.e.* a limited oligopoly makes the user better off.

We show that the two claims are correct using elementary arguments. The first claim is equivalent to the statement: if $\alpha(k^*) = o(k^*)$, then $\frac{\partial \text{Rev}(e^*, k^*)}{\partial k^*} > 0$. To show this, note that if $\alpha(k^*) = o(k^*)$, then e^*k^* is increasing in k^* and so is $e^*(k^* - 1)$. Therefore both terms in (3) are increasing in k^* . Similarly, the second claim states that if $\alpha(k^*) = \omega(k^*)$ then $\frac{\partial \text{Rev}(e^*, k^*)}{\partial k^*} < 0$. Since $\alpha(k^*) = \omega(k^*)$ implies that e^*k^* is decreasing in k^* and $e^*(k^* - 1)$ is decreasing in k^* . This implies $\frac{\partial \text{Rev}(e^*, k^*)}{\partial k^*} < 0$.

3 An optimization approach

Here, we present an alternative, deterministic model of the intent market, which follows from an optimization framework. Instead of considering prices for each solver, drawn from some distribution, we equip each solver with a utility and cost function. This approach has a number of advantages. First, these functions can be easily parameterized, and, as a result, the approach lends itself to empirical studies and model fitting with real-world data. Second, modeling with utility functions may be more appropriate in some scenarios, particularly when there is a lot of heterogeneity among the solvers. Finally, this approach immediately suggests a Dutch auction-like mechanism, analogous to that considered in the probabilistic model. We also show how, under certain assumptions, the deterministic framework implies similar results as those in the previous section, such as those involving the impact of congestion.

In particular, we turn to convex optimization for our deterministic approach, which is similar in spirit to the linear programs presented in [KP04] in the context of auctions.

We outline how social welfare maximizing allocations may be reached via a primal-dual mechanism, which can be interpreted as a Dutch auction. As in the previous section, we limit our focus to two-asset trades. In the general multiple asset case, one must use a more complicated mechanism, as asset prices may be negatively correlated. We leave this case to future work.

3.1 Social welfare

Solvers. Each solver $k = 1, \dots, n$ is willing to provide, for given amount x of token T_1 , an amount $\alpha_k x$ of T_2 . To do this, the solver incurs a cost $c_k(x)$ and gains utility $u_k(x)$, where the function u_k may, for example, reflect the solver’s relative preference for holding T_1 versus T_2 . We restrict the cost function $c_k : \mathbf{R}_+ \rightarrow \mathbf{R}_+ \cup \{\infty\}$ to be convex, the utility function $u_k : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ to be concave, and both to be defined over the nonnegatives—negative trades are not allowed. Infinite values are used to encode constraints; a trade x such that $c_k(x) = \infty$ is unacceptable (*e.g.*, it is outside solver k ’s budget). We assume that these functions are closed and proper. Thus, the net utility for solver k to tender $\alpha_k x_k$ units of T_2 to the user in exchange for x units of T_1 is given by

$$U_k(x_k) = -\alpha_k x_k - c_k(x_k) + u_k(x_k),$$

where this utility is measured in terms of token T_2 , which we take as the numeraire. This utility is concave by construction.

User. We consider a user who wishes to trade δ units of token T_1 for a maximum amount of token T_2 , using a combination of these solvers and an on-chain CFMMs. The user can find the optimal trade with the network of CFMMs by solving the *optimal routing problem* [Ang+22; Dia+23]. Because this problem is a convex optimization problem with a concave objective, we can partially maximize [BV04, §3.2.5] to define an aggregate forward exchange function $G : \mathbf{R}_+ \rightarrow \mathbf{R}_+$, which maps an input of token T_1 to the maximum output of token T_2 using all of the on-chain CFMMs. Note that this function is concave and nondecreasing. The user’s utility is then given by

$$U_p(x) = G(\delta - \sum_{k=1}^K x_k) + \sum_{k=1}^K \alpha_k x_k,$$

which is concave and nondecreasing.

Social welfare maximization. Putting these utility functions together, we can define the problem of maximizing social welfare:

$$\begin{aligned} \text{maximize} \quad & G(\delta - y) + \sum_{k=1}^n u(x_k) - c_k(x_k), \\ \text{subject to} \quad & y = \sum_{k=1}^n x_k. \end{aligned} \tag{4}$$

with variables $y \in \mathbf{R}$ and $x_k \in \mathbf{R}$ for $k = 1, \dots, n$. Note that the transfer from each solver k to the user, given by $\alpha_k x_k$, has canceled out in the objective.

3.2 Dutch auction

The Lagrangian of the social welfare maximization problem (4) is

$$L(x, \tilde{x}, \nu) = G(\delta - y) + \nu y + \sum_{k=1}^n u_k(x_k) - c_k(x_k) - \nu x_k.$$

The dual function is then

$$h(\nu) = \sup_{y \geq 0} \{G(\delta - y) + \nu y\} + \sum_{k=1}^n \sup_{x_k \geq 0} \{u_k(x_k) - c_k(x_k) - \nu x_k\}.$$

Referring back to the utilities, we can interpret these terms simply: they are exactly the utility for the user and agent respectively when the universal price of the asset is ν . Since the function G is nondecreasing, we can conclude that the price $\nu \geq 0$. Otherwise if $\nu < 0$, letting $y = -t$ we have

$$G(\delta + t) - \nu t \geq G(\delta) + |\nu|t,$$

and the right hand side goes to infinity as $t \rightarrow \infty$ (equivalently, as $y \rightarrow -\infty$).

Optimality conditions. From the Lagrangian, at optimality we have primal feasibility ($y^* = \sum_{k=1}^n x_k^*$) and the dual conditions

$$\begin{aligned} g(\delta - y^*) &= \nu^* \\ u'_k(x_k^*) - c'_k(x_k^*) &= \nu^* \quad \text{for } k = 1, \dots, n \end{aligned}$$

(When u_k or c_k is not differentiable for some k , a similar argument holds using subgradient calculus.) Putting these together, we have that

$$u'_k(x_k^*) - c'_k(x_k^*) = g(\delta - \sum_{k=1}^n x_k^*).$$

In other words, the marginal cost of getting asset 2 from the public CFMMs should exactly equal each participating solver's net marginal utility for providing it. These optimality conditions allow us to also characterize when a solver will not participate for a given price ν . Any solution to the solver subproblem, denoted by \tilde{x}_k , which lies on the interior of the domain ($\tilde{x}_k > 0$) must satisfy

$$u'(\tilde{x}_k) - c'(\tilde{x}_k) = \nu.$$

Thus, if ν is set too high, the solver will simply return the zero trade: $\tilde{x}_k = 0$. Furthermore, this condition implies that when the costs to the solver are zero, it will report its true marginal utility (*i.e.*, its true price), corroborating the results in §2.2.

A dual approach. We can solve this social welfare problem with a mechanism akin to a dutch auction. As a consequence of strong duality, which holds under relatively weak conditions [BV04, §5.2.3], we can solve the primal problem (4) by solving the dual problem. The dual problem is

$$\text{minimize } h(\nu).$$

Since this problem is convex and single dimensional, we can solve it with an iterative method that can be interpreted as a dutch-auction-like mechanism. The first derivative of h is

$$h'(\nu) = \tilde{y} - \sum_{k=1}^n \tilde{x}_k,$$

where \tilde{y} is the optimizer of $G(\delta - y) + \nu y$ and $\{\tilde{x}_k\}$ are the optimizers of the solvers' sub-problems, *i.e.*, the amount they will quote given a price ν . Since this function is unimodal, we can optimize it with a simple strategy: we start with a large ν and query each solver to get $\{\tilde{x}_k\}$, which is exactly the solution to the solver utility-maximization subproblem above. Then, we compute the gradient $h'(\nu)$. If this gradient is greater than zero, we decrease ν and repeat. When the gradient is sufficiently close to zero, we are done. In practice this looks like a Dutch auction. Note that, at termination, the \tilde{y} and $\{\tilde{x}_k\}$ are exactly the solutions to the original social welfare problem (4). The dual setup above can also be extended to the case of n assets with independent solver costs. However, the proposed mechanism does not extend directly, as there is no longer a total ordering on prices. A coordinate descent type mechanism may work in this multiple asset case.

Interpretation. Recall that α_k is the price a solver quotes to the user. It is clear that for any rational solver

$$u_k(\delta) \geq \alpha_k \delta \implies u_k(\delta)/\delta \geq \alpha_k.$$

In other words, $\alpha_k x_k$ is a linear underestimator of the solver's concave utility function u_k . The mechanism above, then, elicits the solver's true price for a given order size, and it terminates when the solver's marginal price is equal to that of the CFMM. However, instead of proposing an order size for each of the n solvers, the mechanism instead can propose one price, which it broadcasts to all solvers. The solvers respond to this price with a proposed order.

3.3 Extension: congestion

From this optimization framework, we also see that positively correlated costs among solvers will reduce the user's welfare. Now assume that the solver cost functions are not independent. We model congestion by introducing a new cost function $\tilde{c}_k : \mathbf{R}_+^n \rightarrow \mathbf{R}_+ \cup \{\infty\}$ which is increasing in all arguments and matches the old cost function c_k when that solver is the only one quoting an order. In other words, we must have that

$$\nabla \tilde{c}(x_1, \dots, x_n) > 0$$

and

$$\tilde{c}_k(0, \dots, x_k, \dots, 0) = c_k(x_k)$$

for any order x_k . These conditions imply that marginal costs are higher in this regime:

$$\tilde{c}'_k(x_1, \dots, x_n) > \tilde{c}_k(0, \dots, x_k, \dots, 0) = c'_k(x_k),$$

where there is at least one solver $k' \neq k$ with $x_{k'} > 0$. As a result, from the dual conditions, we have that

$$\nu^* = u'_k(x_k^*) - \tilde{c}'_k(x_1^*, \dots, x_K^*) \leq u'_k(x_k^*) - c'_k(x_k).$$

Thus, the price ν^* when congestion costs are present is less than the price with independent costs. In other words, the user will receive reduced output from the solvers if their costs are positively correlated. This result mirrors that of the probabilistic model: under congestion costs, user welfare decreases.

4 Model Extensions

As discussed in the introduction, there are both non-blockchain focused intent markets as well as more complex intent-like systems within the blockchain world. We will briefly discuss the applicability of our model to those scenarios and extensions that can be used for analyzing these systems.

Stock order-by-order auctions. The US Securities and Exchange Commission (SEC) recently released a proposal for an order flow execution auction within the US stock market [Sec22]. One rationale for implementing such a system is to improve the fairness of ‘price improvement’ that a retail stock trader receives [Sec22]. The majority of retail, individual stock traders within the US use brokerages such as Robinhood or Vanguard to execute their trades. These brokerages outsource the task of finding an optimal execution price for a user to wholesalers, who aggregate order flow from brokerages and institutions (such as endowments or pension funds) to execute trades. Currently, there is no competitive market for brokerages to engage wholesalers—they instead construct private agreements which are known as payment for order flow [Wel22]. These agreements have come under scrutiny by the US government, as the lack of a competitive market is thought to have led to market failures during the Gamestop panic of 2021 [Ang21].

The SEC’s proposal is to have wholesalers run an auction whose initial price is set at the price tendered by the wholesaler. The wholesaler would run the auction for 100 milliseconds and any regulated entity, such as an accredited investor, could participate in the auction. To win the auction, the wholesaler would have to participate explicitly, as their initial price does not serve as a reserve price. In contrast, the UniswapX Dutch auction guarantees the user at least the amount of output quoted publicly by a CFMM, even if the auction does not clear.

Due to latency constraints, it is assumed that the wholesaler would be running many of these auctions simultaneously. It is an interesting extension of our model to understand the

degree to which auctions in multiple markets, where price movements may be correlated, result in improved prices for users.

Proposer-builder separation. One of the most common auctions settled on blockchains is the so-called ‘proposer-builder separation’ (PBS) auction [PR23; GPR23; BGR24]. The seller in this auction is the proposer: an agent who has a temporary monopoly to propose a valid block within a blockchain. The proposer runs an auction to solicit blocks from agents known as builders who source transactions and construct valid blocks. A number of works have studied the centralizing effect within PBS auctions [BGR24; GPR23]. However, none of the prior works analyzed the impact of congestion on auction performance.

PBS auctions fit into our framework as they are single item first price auctions. Block builders, who source transactions to earn fees or rewards from block construction (*i.e.* the aforementioned MEV [KDC22; BGR23; KDC23]), compete for order flow from exchanges and wallets that generate transactions. Similar to PFOF, these builders often offer incentives to users to send their transactions to them, and may exert costly effort to acquire user order flow. One notable difference is that our model assumes a uniform congestion cost across agents. In the PBS auction, different agents will have varied costs and congestion is more heterogeneous across agents. Modifying our model to include heterogeneity between the agents is future work and can be used to analyze historical auction data to measure reductions in competition.

5 Conclusion

In this paper, we model the incentives that solvers face in intent-based markets. Our first model is a probabilistic model of an auction under two complicating features: (*i*) entry costs incurred by solvers while setting up infrastructure and (*ii*) congestion costs faced by solvers while attempting to search for liquidity at desirable prices. In our single intent model (single item auction model), we demonstrate that entry costs reduce the number of solvers that participate in the market in equilibrium. If, furthermore, participants have to invest costly effort into improving their ability to find prices, and face congestion costs while doing so, we also find the existence of limited oligopoly. These results demonstrate the counterintuitive fact that planners of intent markets may want to limit entry in order to maximize user welfare.

Our second model is an optimization-based deterministic model. This model enables the explicit analysis of the social welfare maximization problem faced by a planner who considers both the user’s and the solvers’ utilities. This approach has several benefits: (*i*) it allows for fitting utility functions to data and (*ii*) suggests a natural Dutch auction-like primal-dual mechanism to solve the welfare maximization problem.

We leave as future work a generalization of the single intent model to more complicated intents which may include transacting in multiple markets. Further, it is an interesting question whether numerical methods may be used to estimate whether solvers face congestion costs in these markets, and the degree to which users face welfare losses. Finally, in our

modeling so far, motivated by practice, we have restricted attention to solvers competing in an auction to fulfill the user’s intent, where the profits from the auction provide the incentives for solvers to enter and invest. However, certain kinds of congestive investment costs may be difficult to observe and even more challenging for users to directly contract [GH92]. Recent developments in the literature provide insight into how a principal can robustly incentivize effort by an agent (*e.g.*, [Car15]). Another direction of future work, therefore, is whether such contracts can be implemented in intent markets.

Our work implies that designers and planners of intent markets need to thoroughly vet the various costs that solvers face. If the goal of an intent network is to maximize welfare while also allowing free entry into the system, designers should be careful to construct auctions such that bidders minimize contention in sourcing inventory.

References

- [AC20] Guillermo Angeris and Tarun Chitra. “Improved price oracles: Constant function market makers”. In: *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 2020, pp. 80–91.
- [Ada+21] Hayden Adams et al. “Uniswap v3 core”. In: *Tech. rep., Uniswap, Tech. Rep.* (2021).
- [Ada23] Hayden Adams. *Introducing the UniswapX protocol*. July 2023. URL: <https://blog.uniswap.org/uniswapx-protocol>.
- [AEC21] Guillermo Angeris, Alex Evans, and Tarun Chitra. “Replicating monotonic payoffs without oracles”. In: *arXiv preprint arXiv:2111.13740* (2021).
- [Ang+19] Guillermo Angeris et al. “An analysis of Uniswap markets”. In: *arXiv preprint arXiv:1911.03380* (2019).
- [Ang21] James Angel. “Gamestonk: What happened and what to do about it”. In: *Georgetown McDonough School of Business Research Paper 3782195* (2021).
- [Ang+22] Guillermo Angeris et al. “Optimal routing for constant function market makers”. In: *Proceedings of the 23rd ACM Conference on Economics and Computation*. 2022, pp. 115–128.
- [Ben23] Burak Benligiray. *Oracle Extractable Value (OEV)*. 2023. URL: <https://medium.com/api3/oracle-extractable-value-oev-13c1b6d53c5b>.
- [BGR23] Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. “Transaction fee mechanism design with active block producers”. In: *arXiv preprint arXiv:2307.01686* (2023).
- [BGR24] Maryam Bahrani, Pranav Garimidi, and Tim Roughgarden. “Centralization in Block Building and Proposer-Builder Separation”. In: *arXiv preprint arXiv:2401.12120* (2024).

- [Blo23] The Block. *DeFi Exchange Data and Charts for DEXs, AMMs, and Swaps*. Dec. 2023. URL: <https://www.theblock.co/data/decentralized-finance/dex-non-custodial>.
- [BPS23] Svetlana Bryzgalova, Anna Pavlova, and Taisiya Sikorskaya. “Retail trading in options and the rise of the big three wholesalers”. In: *The Journal of Finance* 78.6 (2023), pp. 3465–3514.
- [BV04] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [Car15] Gabriel Carroll. “Robustness and linear contracts”. In: *American Economic Review* 105.2 (2015), pp. 536–563.
- [CDM23] Álvaro Cartea, Fayçal Drissi, and Marcello Monga. “Predictable losses of liquidity provision in constant function markets and concentrated liquidity markets”. In: *Applied Mathematical Finance* (2023), pp. 1–25.
- [Cha23] Jon Charbonneau. *Suave, anoma, shared sequencers, & Super Builders*. Apr. 2023. URL: https://dba.mirror.xyz/NTg5FSq1o_YiL_KJrKB0s0kyeiNUPobvZUrLBGceagg.
- [CoW23] Protocol CoW. *CoW Protocol — Dune*. <https://dune.com/cowprotocol1>. Dec. 2023.
- [Dai+20] Philip Daian et al. “Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability”. In: *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2020, pp. 910–927.
- [Dia+23] Theo Diamandis et al. “An Efficient Algorithm for Optimal Routing Through Constant Function Market Makers”. In: *Financial Cryptography* (2023).
- [GH92] Sanford J Grossman and Oliver D Hart. “An analysis of the principal-agent problem”. In: *Foundations of Insurance Economics: Readings in Economics and Finance*. Springer, 1992, pp. 302–340.
- [GPR23] Tivas Gupta, Mallesh M Pai, and Max Resnick. “The centralizing effects of private order flow on proposer-builder separation”. In: *arXiv preprint arXiv:2305.19150* (2023).
- [KDC22] Kshitij Kulkarni, Theo Diamandis, and Tarun Chitra. “Towards a theory of maximal extractable value i: Constant function market makers”. In: *arXiv preprint arXiv:2207.11835* (2022).
- [KDC23] Kshitij Kulkarni, Theo Diamandis, and Tarun Chitra. “Routing MEV in Constant Function Market Makers”. In: *International Conference on Web and Internet Economics*. Springer. 2023, pp. 456–473.
- [KP04] Jayant Kalagnanam and David C Parkes. “Auctions, bidding and exchange design”. In: *Handbook of quantitative supply chain analysis: modeling in the e-business era*. Springer, 2004, pp. 143–212.
- [Kri09] Vijay Krishna. *Auction theory*. Academic press, 2009.

- [Lam24] Hart Lambur. *Announcing oval: Earn protocol revenue by capturing Oracle mev*. 2024. URL: <https://medium.com/uma-project/announcing-oval-earn-protocol-revenue-by-capturing-oracle-mev-877192c51fe2>.
- [Leu21] Felix Leupold. *The batch auction optimization problem*. <https://docs.cow.fi/solvers/in-depth-solver-specification/the-batch-auction-optimization-problem>. June 2021.
- [Lla23] DeFi Llama. *DEXs — DeFi Llama*. <https://defillama.com/dexs>. Dec. 2023. URL: `\url{https://defillama.com/dexs}`.
- [LPZ23] Alfred Lehar, Christine Parlour, and Marius Zoican. “Liquidity fragmentation on decentralized exchanges”. In: *arXiv preprint arXiv:2307.13772* (2023).
- [LSZ21] Bowen Liu, Pawel Szalachowski, and Jianying Zhou. “A first look into defi oracles”. In: *2021 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS)*. IEEE. 2021, pp. 39–48.
- [MCT09] Govind S Mudholkar, Yogendra P Chaubey, and Lili Tian. “Extremes, extreme spacings and tail lengths: an investigation for some important distributions”. In: *Calcutta Statistical Association Bulletin* 61.1-4 (2009), pp. 243–266.
- [Mil+22] Jason Milionis et al. “Quantifying loss in automated market makers”. In: *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security*. 2022, pp. 71–74.
- [Mor21] Fulvia Morales. *0x API swap now supports 0x Protocol v4*. <https://medium.com/@fulviamorales/0x-api-swap-now-supports-0x-protocol-v4-4c53c732d04f>. Mar. 2021.
- [O’H10] Maureen O’Hara. “What is a quote?”. In: *The Journal of Trading (Retired)* 5.2 (2010), pp. 10–16.
- [OZ21] Maureen O’Hara and Xing Alex Zhou. “The electronic evolution of corporate bond dealers”. In: *Journal of Financial Economics* 140.2 (2021), pp. 368–390.
- [p24] phu. *UniswapX — Dune*. <https://dune.com/phu/uniswapx>. Jan. 2024.
- [PR23] Mallesh Pai and Max Resnick. “Structural Advantages for Integrated Builders in MEV-Boost”. In: *arXiv preprint arXiv:2311.09083* (2023).
- [Pro23] CoW Protocol. *Introducing the Programmatic Order Framework, from CoW Protocol*. <https://blog.cow.fi/introducing-the-programmatic-order-framework-from-cow-protocol-088a14cb0375>. Nov. 2023.
- [RGM23] Geoffrey Ramseyer, Ashish Goel, and David Mazières. “SPEEDEX: A Scalable, Parallelizable, and Economically Efficient Decentralized {EXchange}”. In: *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. 2023, pp. 849–875.
- [Sec22] Securities and Exchange Committee. *Proposed Rule: Order Competiton Rule*. Tech. rep. Version Release No. 34-96495. Securities and Exchange Committee, 2022.

- [Sen23] Shayon Sengupta. *Oracles and the New Frontier for Application-owned Orderflow auctions*. 2023. URL: <https://multicoin.capital/2023/12/14/oracles-and-the-new-frontier-for-application-owned-orderflow-auctions/>.
- [Wel22] Ivo Welch. “The wisdom of the Robinhood crowd”. In: *The Journal of Finance* 77.3 (2022), pp. 1489–1527.
- [XFP23] Matheus Venturyne Xavier Ferreira and David C Parkes. “Credible Decentralized Exchange Design via Verifiable Sequencing Rules”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 2023, pp. 723–736.
- [Yin23] Awa Sun Yin. *An overview of Anoma’s architecture*. Sept. 2023. URL: <https://medium.com/anomanetwork/an-overview-of-anoma-s-architecture-26b72e8c9be5>.

A Exponential value distributions

We now derive the left hand side of equation (1) in the case where F is the exponential distribution with rate λ (that is, $F(p) = 1 - \exp(-\lambda p)$ for $p \geq 0$). We have:

$$\begin{aligned}
& \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) S(k) \\
&= \frac{1}{\lambda} \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \frac{1}{k+1} \\
&= \frac{1}{\lambda} \sum_{k=0}^n \left(\frac{n!}{k!(n-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \frac{1}{k+1} \\
&= \frac{1}{\lambda} \sum_{k=0}^n \left(\frac{n!}{(k+1)!(n-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \\
&= \frac{1}{(n+1)\lambda F_C(\bar{c})} \sum_{k=0}^n \left(\frac{(n+1)!}{(k+1)!(n-k)!} F_C^{k+1}(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \\
&= \frac{1}{(n+1)\lambda F_C(\bar{c})} \sum_{k=1}^{n+1} \left(\frac{(n+1)!}{(k)!(n+1-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n+1-k} \right) \\
&= \frac{1}{(n+1)\lambda F_C(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+1})
\end{aligned}$$

B Uniform value distributions

We now derive the left hand side of equation (1) in the case where F is the uniform distribution on $[0, 1]$ (that is, $F(p) = p$ for $p \in [0, 1]$). We have:

$$\begin{aligned}
& \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) S(k) \\
&= \sum_{k=0}^n \left(\binom{n}{k} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \frac{1}{(k+1)(k+2)} \\
&= \sum_{k=0}^n \left(\frac{n!}{k!(n-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \frac{1}{(k+1)(k+2)} \\
&= \sum_{k=0}^n \left(\frac{n!}{(k+2)!(n-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \\
&= \frac{1}{(n+1)(n+2)F_C^2(\bar{c})} \sum_{k=0}^n \left(\frac{(n+2)!}{(k+2)!(n-k)!} F_C^{k+2}(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \\
&= \frac{1}{(n+1)(n+2)F_C^2(\bar{c})} \sum_{k=2}^{n+2} \left(\frac{(n+2)!}{(k)!(n-k)!} F_C^k(\bar{c})(1 - F_C(\bar{c}))^{n-k} \right) \\
&= \frac{1}{(n+1)(n+2)F_C^2(\bar{c})} (1 - (1 - F_C(\bar{c}))^{n+2} - (n+2)F_C(\bar{c})(1 - F_C(\bar{c}))^{n+1})
\end{aligned}$$

C Extreme spacing

In this section, we derive the equality $(k+1)S(k) = ES(k+1)$ for any CDF F and PDF f , where $p^* = 0$. To see this, observe the following chain of equalities:

$$\begin{aligned}
ES(k+1) &= \mathbf{E}_{p_i \sim f} [p_{k+1:k+1} - p_{k:k+1}] \\
&= \int_0^\infty \left((k+1)F(p)^k f(p) - (k+1)kF(p)^{k-1}(1-F(p))f(p) \right) pdp \\
&= (k+1) \int_0^\infty F(p)^{k-1} f(p) (F(p) - k(1-F(p))) pdp \\
&= (k+1) \int_0^\infty F(p)^{k-1} f(p) ((k+1)F(p) - k) pdp \\
&= (k+1) \int_0^\infty f(p) \left((k+1)F(p)^k - kF(p)^{k-1} \right) pdp \\
&= (k+1) \int_0^\infty \frac{d \left(F(p)^{k+1} - F(p)^k \right)}{dp} pdp \\
&= (k+1) \int_0^\infty (F(p)^k - F(p)^{k+1}) dp \\
&= (k+1) \int_0^\infty F(p)^k (1-F(p)) dp,
\end{aligned}$$

where the second to last equality follows from integration by parts.