Melange: Reinvent the World with Markets (v1.1)

Melange Technologies, Inc.

"A beginning is the time for taking the most delicate care that the balances are correct."

— from "Manual of Muad'Dib" by the Princess Irulan

I. ABSTRACT

Melange is a prediction markets protocol. Its first iteration intends to solve the problem that market makers (colloquially known as "the House") and professional gamblers make billions of dollars at the expense of retail gamblers each year. If there were some way to coordinate transactions between these retail gamblers in a truly peer-to-peer fashion, then the cost of betting would be virtually eliminated. We propose the Continuous Order Book (COB), a market design of our own invention. This market design will revolutionize live betting in three ways. First, by making bets nearly costless, it enables high frequency recreational swing trading. Second, by decentralizing the provision of liquidity, COB turns any live event anywhere into a potential functioning market. Third, by entirely reimagining the user experience of interacting with markets, it will bring prediction markets to a wholly new audience of people who have never placed a bet before. Built on Solana, the Continuous Order Book will be among the first generation of blockchain-based consumer applications — and lay the foundations for a protocol that will one day service every bet on the planet.

II. INTRODUCTION

A prediction market is a market that enables people to trade on a set of mutually exclusive future outcomes. In this Whitepaper, we define these prediction markets to trade between $1 (when an outcome occurs) and $0 (when the outcome does not occur). Defined as such, the price of any contract implies a probability of the outcome in question.

The beauty of a prediction market is that whenever a market misprices a contract, any informed market participant aware of this mispricing is given a financial incentive to make a trade and correct the mistake. Because this monetary incentive to share information is always at play, we say that markets are efficient — in other words, the prices implied by prediction markets truly reflect humanity's best estimate of the probability of a given event. It is for this reason that many economists dream of mastering the future with prediction markets. These economists dream of markets on every event, real or hypothetical. They imagine a world where every person on the planet interacts on a daily basis with prediction markets — a world where these people check the markets for the odds of future events the same way they check a weather app for the probability of rain.

Yet, in spite of these beautiful stories told by romantic economists, the past decade of prediction markets has largely been a story of failure. As we've watched revolutions in markets for financial and digital assets, prediction markets have lagged behind. Today, the primary application of prediction market technology is still the sportsbook, an outdated product with exorbitant prices, limited options, and a confusing and unappealing user interface.

This Whitepaper explores our theory of why prediction markets have failed and how we intend to improve them. Put succinctly, our insight is that prediction markets are a consumer product, and as such, the only thing that matters is user experience. Steve Jobs once said that building a good consumer product means starting with the customer experience, and then work backwards to a technology. "You can't start with the technology then try to figure out where to sell it." At Melange, we are guided by a single question: What is the ideal way to experience a prediction market?

III. LIQUIDITY, MY PROBLEM CHILD

Any substantive discussion about prediction markets must begin with a discussion about liquidity. Liquidity is the lifeblood of every market, and as such, before our question may be answered, we have some ground to cover.

Liquidity is the idea that when someone wants to buy something, there must simultaneously exist someone willing to sell, or else there is no trade. A market with liquidity is one where it's quick and cheap to find a counterparty and make a trade.

Because it is unfeasible to expect two peers to simultaneously agree on prices and trade against one other, traditional markets function on the basis of liquidity providers. These liquidity providers, commonly referred to as "market makers", make markets by offering to buy and sell contracts at all times. In all likelihood, if you have ever interacted with a market, you have traded against one of these market makers.

The market maker derives profit from buying low and selling high. When the state of the world remains the same, and the flow of orders on each side of the market is balanced, the market maker is guaranteed to make money. However, when there is a change in the value of the underlying asset — when an adverse weather event affects the underlying value of corn futures or a regulatory ruling affects the underlying value of Bitcoin — market participants who receive the new information first may use the information to trade at a profit until the market maker has updated its priors. When trades occur on the basis of this imbalance of information, we call it "adverse selection".

A market maker only contributes liquidity when they believe they can offset their bad trades with good trades. This means that for a market to trade, it must be profitable to provide liquidity. As such, the quality of any market design is best judged by how it protects liquidity providers from adverse selection.

Today, there exist two approaches for providing liquidity to prediction markets: automated market making and centralized market making. In what follows, we describe these two approaches, and examine what deep theoretical limitations they have.

A. Approach #1: Automated Market Makers

An automated market maker (AMM) is an autonomous computer program that defines the price of digital assets and holds different assets in separate homogeneous pools. The ratio of assets between pools determines the exchange rate that users can receive when trading one asset for another.

We say that liquidity provision under the AMM model is decentralized because anyone can provide liquidity to these pools by depositing both assets in proportion to their exchange rate. AMMs incentivize the provision of liquidity by automatically crediting liquidity providers with a small "LP fee" on all transactions executed by market takers. Provision of liquidity under this incentive mechanism is commonly referred to as "yield farming".

Problem: Permanent Loss

AMMs determine the prices of the contracts they sell purely via supply and demand. When users purchase a lot of one asset from the AMM, the AMM automatically increases the price of that asset according to a bonding curve. The problem is that when there is a change in the underlying value of an asset traded by an AMM, the AMM "discovers" the new price by making a series of bad trades. While this isn't a problem for stable assets such as USDC/USDT where the relative underlying value almost never changes, this is devastating for prediction markets. To give an example, suppose Tom Brady throws an impressive touchdown pass, raising the Buccaneers' odds of winning from 50% to 60%. The AMM, initially selling contracts at $0.50, now becomes adversely selected against by market participants. In order to adjust price from $0.50 to $0.60, it must sell contracts at $0.50, then $0.51... all the way up to $0.60. Those bad sales represent losses to the AMM.

Conventionally, we call these losses "impermanent loss" — the idea being that, assuming the value fluctuates randomly, something could occur in the game that reverts the market back to $0.50, recouping all losses. In prediction markets, however, which always resolve to a binary outcome, this is never the case. If Tom Brady wins, for example, the AMM eventually sells contracts at $0.98, then $0.99. Because, after resolution, a market never reverts in price, what is called impermanent loss on an AMM exchange between Bitcoin and Ethereum is actually a permanent loss in prediction markets. To put the extent of these permanent losses in perspective, if a user posted $100 worth of liquidity when the pools were 50:50, then the value of the user's share will be worth less than $20 when the pools are 99:1 (assuming a constant product marketmaker).

Because AMM-based prediction market platforms do not have the trading volume to generate enough transaction fees to offset permanent loss, it is not profitable to provide liquidity to such platforms. Predictably, these platforms suffer enormous liquidity problems. And while venture capital money can subsidize the liquidity provision to these markets temporarily, it cannot forever.

B. Approach #2: Centralized Market Makers

A centralized market maker is a company or an individual that quotes buy and/or sell prices that market participants can take. This centralized liquidity model comes in two forms:

Sportsbooks. The vast majority of bets placed each year are placed on sportsbooks like DraftKings, Bet365, and Pinnacle. Sportsbooks are explicitly centralized as the casino is the only market maker. When a user bets, the user trades instantaneously at a price decided upon by this casino.

Exchanges. A betting exchange allows anyone in the world to make and take markets using a traditional order book. While these platforms brand themselves as "peer-to-peer", the technical competence required to make markets profitably precludes most peers. It is by no coincidence that Susquehanna International Group (SIG), the largest financial options market making firm in the world, also makes markets for sports on BetFair, the largest betting exchange in the world.

Problem: Adverse Selection

Trading is a zero sum game. If I have made money on a trade, then I know that my counterparty has lost. Ultimately, this makes trading a skill-based game about information. If I am more skilled and thus have more information than you, then on average, I can make profit betting against you.

The dilemma of the market maker is twofold. First, when they make a market, their liquidity is public, meaning that anyone anywhere in the world can trade against them. Therefore, if it happens that anyone anywhere in the world has more information than the market maker, then the market maker loses. Second, there is an asymmetry in the information that can be gained during the course of betting. Whereas market takers can see the actions of makers before the trade has been made, the market makers can only see the actions of the takers after the trade has been made.

While the market maker attempts to build models that identify and respond to new information as soon as possible, they can never be sure of certain of their model's veracity and can therefore never completely avoid adverse selection. As a result, market makers price the risk of adverse selection into their spread to maintain profitability. By jacking up the fees on their markets (sometimes up to 15% of a bettor's principal for sporting events like Formula 1), they simultaneously lower the chance of being adversely selected by informed actors and collect more profits from uninformed actors. In other words, they pass the costs of adverse selection onto retail gamblers.

Problem: Model Cost

As bad as all of this sounds, it's worth noting that the centralized market maker is strictly better at dealing with adverse selection than the automated market maker. This is because the centralized market maker builds models to augment price discovery whereas an AMM can only change its price in response to the flow of trading. By observing the event or details of the market dynamics, centralized market makers can change their prices without incurring losses. To reuse the example where Tom Brady's pass raises the Buccaneers' odds of winning from 50% to 60%, we observe that the centralized market maker need not make a series of bad trades at every price point from $0.50 to $0.60 as the AMM does. Rather, if market maker has a model of the event, it can directly lift its prices from $0.50 to $0.60, thereby avoiding these bad trades.

However, these models are not easy to build. Modeling a live League of Legends match being streamed on Twitch, for instance, requires a designated modeler to watch the match in real time, updating prices according to their mental evaluation of the state of play. As a result, live League of Legends markets are very rare (not to mention expensive, given that market makers are going to be quite scared of being adversely selected when the way they model the event is by eyeballing a LoL match), despite the fact that people would almost certainly bet on them.

As it turns out, the two heuristics that make up the most engaging prediction market events, volatility and uniqueness, are also two heuristics that define the difficult and expensive events to model. As such, the offerings on sportsbooks and exchanges are bland and paltry.

IV. CONTINUOUS ORDER BOOK

In our Introduction, we pose the question: What is the ideal way to experience a prediction market? Our solution is a market structure designed specifically for live prediction markets that we call the Continuous Order Book (COB). Every technical detail of the COB's architecture is optimized for three value propositions: 1) price improvement for retail gamblers, 2) diversity of market offerings, and 3) ease of use.

While the protocol is mathematical, we have decided to omit almost all formal description and derivation — not only for the sake of legibility and storytelling, but also because what formulas we currently have are arbitrary. What exact equations we eventually arrive at will largely depend on play-testing and experimentation as they will be ultimately dictated by what feels most natural to our users. Thus, this section provides a high level overview of our protocol, saving the specifics for future iterations of this Whitepaper when interoperability and implementation become larger concerns.

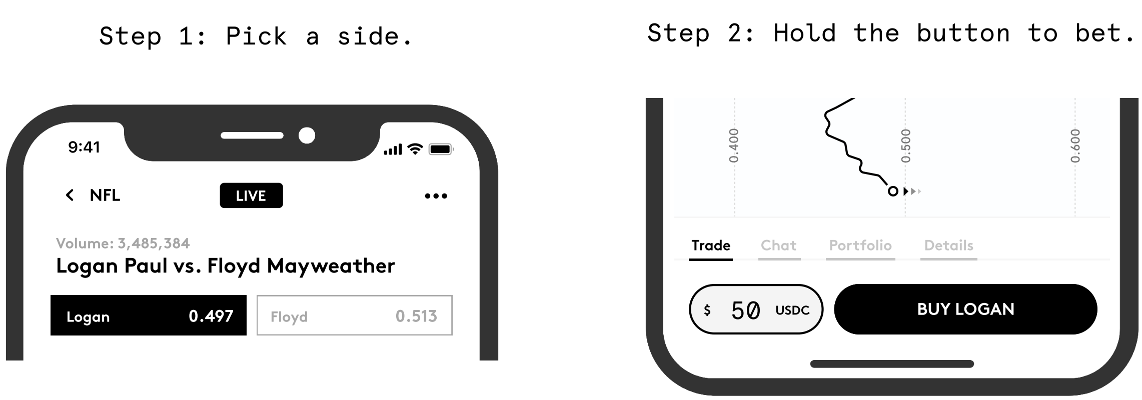

Figure 1: User Interaction

A. The Experience

The Continuous Order Book is a P2P exchange whose prices are automatically adjusted according to supply and demand within the market. The COB achieves this by streaming a single price to all users. If two users want to trade opposite sides of the contract at this price, then the COB instantly matches those users together. Conversely, if users want to trade more of one side than the other, then the COB incrementally changes the price. Because this process occurs in continuous time, we call it the Continuous Order Book.

To see how this works, suppose that one day, Alex Jones attempts to sneak into Bohemian Grove with a GoPro strapped to his head, and Melange implements a prediction market to go alongside the event that resolves to "Yes" if Jones touches the effigy before it is burned.

At home, you tune into the livestream. Checking the Melange app, you see the contract trading at $0.10 — but as you intently watch Jones explain his plan, you are stunned by his logic. You suspect the probability might be higher than 10% so you send an order to buy $100 of "Yes" at $0.10. If, at that very moment, someone had sent an order to sell $100 of "Yes" (in other words, they elected to buy $900 of "No" at $0.90), then your trade would execute peer-to-peer. However, because it's the case that everyone else watching the stream is thinking the same thing as you — that this Alex Jones guy has some really good points — you have no peers. As such, the price climbs to $0.11, then $0.12, then $0.13. As this price climbs, you trade against a time weighted automated market maker, meaning that a small portion of your order executes every second. When the contract climbs to $0.20, supply and demand finally balances. Users enter the market on the "No" side, instantaneously the filling remainder of your order.

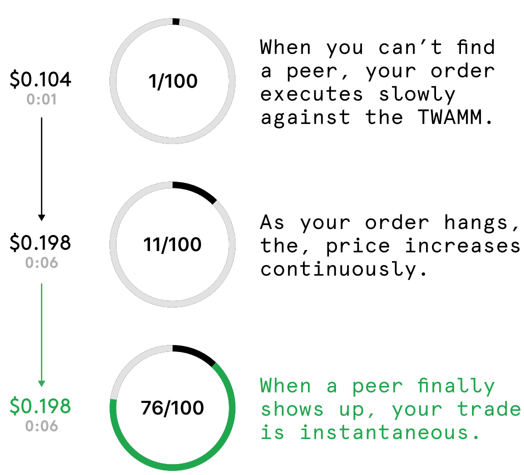

Figure 2: Userflow

Over the course of the next thirty minutes, you watch Jones make good progress. He wades through brush, sneaks past some guards, and we even hear some chanting. As this all happens, the price climbs up, and you bet on both "Yes" and "No" all throughout this. Then, about an hour in, the contract is trading at $0.55 when Jones spots the effigy. The effigy is in the middle of a field with a dozen masked figures standing around it. As soon as you see it, you press the button to buy "Yes", and the price begins shooting up because everyone else has seen it too.

But as Jones makes his advance, a guard appears out of nowhere. The guard tackles Jones. And because you don't get their finger off the "Yes" button fast enough, you trade "Yes" at $0.62 against all the people who have started to press "No" even though Jones has already been caught. So as you watch Jones get hogtied to a big stick and carried off to the fire, you watch in horror as the price of the contract dwindles to $0.00.

B. Price Delta

The Continuous Order Book follows a simple logic. If no one wants to sell you a contract at some price, then the COB assumes this price is too low. Consequently, it bumps the price up to see if anyone will sell it to you at a higher price. This process continues until a clearing price has been found.

In designing the price delta function which dictates the speed at which the COB increases and decreases the price, we follow three high level constraints:

- **Bounds.** The price must be bounded between 0 and 1.
- **Deterministic.** The price movement must be deterministic and easily solvable. Because Solana does not allow for continuous transactions, we simulate continuous transactions using lazy computation.
- **Tail Speed.** The price must move slower in absolute terms at the bounds. This is because a move from 0.02 to 0.03 represents a 50% increase in price while a move from 0.50 to 0.51 represents only a 2% increase in price.

Our speed equation is a variation on the logistic function, giving a price at some future time t based on the current price p and the speed constant s. There are four cases depending on whether the price is increasing or decreasing and whether the price is above or below 0.50. Although this equation might appear complex at first glance, it was chosen because it both fulfills all three criteria while also mimicking the shape of a constant product market maker (something which will become important later in the Whitepaper):

$$p(t) = \begin{cases} \frac{p(1+st)^2}{p((1+st)^2-1)+1} & \text{if } p \geq 0.50, p'(t) > 0 \\ \frac{p(1-st)^2}{p((1-st)^2-1)+1} & \text{if } p \geq 0.50, p'(t) < 0 \\ 1 - \frac{(1-p)(1-st)^2}{(1-p)(1-st)^2+p} & \text{if } p < 0.50, p'(t) > 0 \\ 1 - \frac{(1-p)(1+st)^2}{(1-p)(1+st)^2+p} & \text{if } p < 0.50, p'(t) < 0 \end{cases} \tag{1}$$

This price move equation has the quirk that each particular case ceases to solve for the correct value when the price crosses over 50/50. Consequently, whenever the price is approaching 50/50, we compute the amount of time it will take to reach 50/50, and then run a check to see whether the case should be switched once that amount of time has passed.

Figure 3 is a graphical representation of the price over time. Despite the unruliness of the piece-wise function, observe that the price over time simply resembles a logistic curve, flattening out when the price approaches 0 or 1.

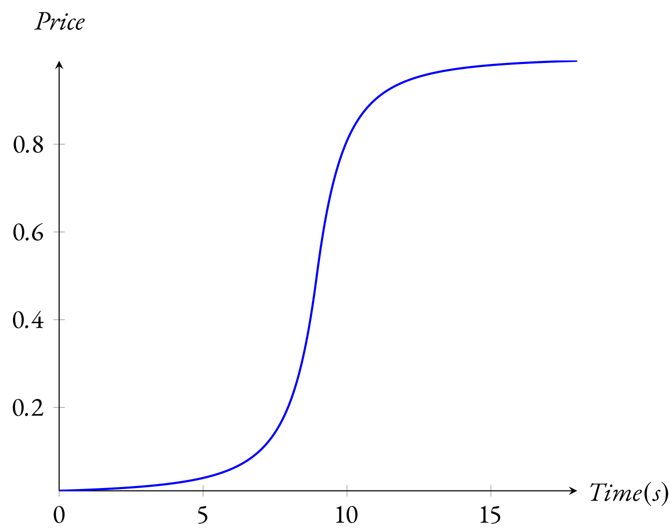

Figure 3: Price vs Time with s = 1, k = 1

C. Time Weighted Automated Market Maker

In order to ensure that our first users always have liquidity even when there are few peers on the market, we introduce a time-weighted automated market maker (TWAMM) that continuously trades against users when there is an imbalance in order flow.

As explained earlier, we do not expect the TWAMM to be profitable for liquidity providers. Rather, the TWAMM should be understood as a market subsidy used to generate initial incentives to bet and build network effects, as well as to reward early adopters who participate in the protocol. Fortunately, as more peers enter the market, by definition, the proportion of all volume traded P2P vs. via the TWAMM increases — and as such, the TWAMM becomes increasingly obsolete and we can remove the TWAMM. At scale, therefore imagine our markets to be a fully P2P.

Formally, our TWAMM is built from a constant product market maker (CPMM) which holds two assets, "yes" outcome shares and "no" outcome shares, denoted x and y respectively. Let R_x be the reserves of asset x and R_y denote the reserves of asset y. A CPMM is defined such that the product of these two reserves always equals a constant k:

$$R_x R_y = k$$

When a user trades against the CPMM, that user deposits some of one asset and withdraws the amount of the other asset that maintains the constant product. Figure 4 is a graphical representation of the CPMM. Where k equals 100, all possible allocations of x and y are charted by the curve:

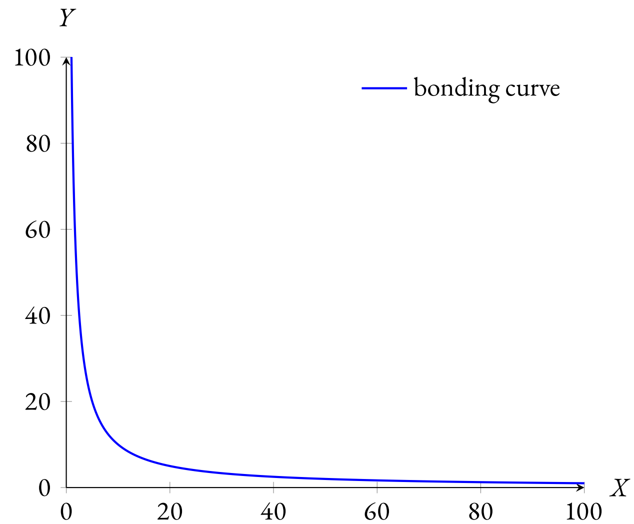

Figure 4: CPMM with k=100

The derivative of this graph at any point represents the relative price between the two assets at that point. Because we want to represent prediction markets as probabilities — i.e., as cents on a dollar — we first convert relative price into dollar price. Let p denote the dollar price of x. Leveraging the observation that the probabilities sum to 1 as well as some algebra, we get the following result:

$$p = 1 - \frac{R_x}{R_x + R_y} = \frac{R_y}{R_x + R_y} \tag{2}$$

Here, we find justification for the complexity of the price delta equation (1). Observe that the price delta equation was derived from the following equation:

$$p(t) = \begin{cases} \frac{(R_y + st)^2}{(R_y + st)^2 + k} & \text{if } R_y \geq R_x, p'(t) > 0 \\ \frac{(R_y - st)^2}{(R_y - st)^2 + k} & \text{if } R_y \geq R_x, p'(t) < 0 \\ 1 - \frac{(R_x - st)^2}{(R_x - st)^2 + k} & \text{if } R_y < R_x, p'(t) > 0 \\ 1 - \frac{(R_x + st)^2}{(R_x + st)^2 + k} & \text{if } R_y < R_x, p'(t) < 0 \end{cases} \tag{3}$$

In other words, equation (1) is derived by setting the larger reserve to 1 and solving for k in equation (3). The implication of this is that when the TWAMM is activated on a market, price equation (3) replaces price equation (1) without changing the experience to the user.

D. Ordering of P2P Fills

Filling occurs on a pro-rata basis, meaning that orders are filled in proportion to their size. In other words, if you place a $100 bet on top of $900 of existing orders, then for every trade made against the TWAMM or filled by another user, you receive 10%. We've chosen this system as opposed to price-time ordering because it always encourages new users to make markets without regard to how late they are.

V. IMPLICATIONS

Robinhood's thesis ten years ago was deeply contrarian. Robinhood saw the potential of a product before anyone else. They claimed there existed an untapped market of millions of young people, and that the key to unlocking this market was to repackage the online brokerage. So, by lowering the price of trading; by offering more things to trade; and by revolutionizing look, feel, and brand identity of the brokerage, they set out to redefine what the stock market meant to people.

By all accounts, they succeeded. Today, Robinhood is a household name and is routinely the #1 app on the App Store. According to documents prepared for its IPO, of Robinhood's 18m monthly active user accounts, more than half were created by people who had never owned a brokerage account before.

At Melange, we intend to reproduce the success of Robinhood on the gambling industry. We believe that the modern sportsbook inhabits the same position that the brokerage inhabited pre-Robinhood: a terrible product with unfathomable untapped potential. By improving upon the user experience via lowering prices, increasing offerings, and precise and intentional design, we will redefine what it is to gamble, uncovering a market much larger than previously imagined.

A. Price

The average spread on a sports bet in the United States is 7.3%. This means that if I place a $100 bet, on average, I should expect to have $92.70 at the end of the bet. This price makes online betting one of the most expensive ways a person can gamble. The innovation of the COB is that it reduces the cost of betting to virtually zero.

Where does price improvement come from?

At a high level, the innovation of the Continuous Order Book is to turn the retail gambler into a market maker, replacing their current role as market taker. We do this by designing our exchange such that making a market is isomorphic to taking a market from the perspective of the user. What results is a new market paradigm, wherein professionals continuously compete against each other to execute retail orders at the cheapest price possible.

At a more granular level, we might consider the following scenario. Suppose an institution like Citadel builds a model of an event and concludes that the fair price of the relevant contract is around $0.53. Let α be the minimal edge required for Citadel to take a market, for some given market and contract size. In this case, the key insight is that even though Citadel would *take* the market at $0.53 - α, it would not *make* the market at the same price due to the risk of adverse selection (as per Section III.B.). As such, under the traditional order book exchange model, where Citadel is the maker, the price of the contract is worse than $0.53 - α.

On the other hand, under the COB model, the contract should trade at exactly $0.53 - α. Suppose, for instance, a retail gambler elects to sell the Patriots at $0.55 and that no peers are willing to take the other side. Here, the price of the contract begins to tick down from $0.55 to $0.54 to $0.53. In this case, as long as Citadel believes that there could be other professionals in the market with a model just as good as theirs, it is game theory optimal for Citadel to act immediately so as to take the order first.

In this way, the price of contracts on the COB are precisely bounded by the minimally profitable fair values of all professional models within the market — the end result of which is an exchange wherein retail gamblers predominately provide liquidity and professional bettors define the spread. Within the bounds, retail gamblers trade so close to the fair value that each bet is essentially free.

Why should I care?

Here, we might observe resemblance that the Continuous Order Book shares with Payment For Order Flow (PFOF), a practice pioneered by Robinhood. Under the PFOF model, when a user places an order on Robinhood, that order is essentially auctioned off to institutional market making firms. Because the market making firms assume that Robinhood's users are uninformed — i.e., there is no reason to believe that there is adverse selection at play with these orders — the market making firms compete for the privilege of *taking* the order by offering better prices.

The introduction of PFOF directly led to what people today call the "retail trading revolution". Whereas brokerages traditionally charged a $5 dollar commission fee on every trade, PFOF meant that on Robinhood, there was no commission at all. By giving users the sense that their trades were free, Robinhood changed user behavior. It made it possible, for instance, to day trade — i.e., buy and sell an asset multiple times throughout the day.

The result is that, today, of Robinhood's 18 million active accounts, 8.5 million use the product daily — and on average, this daily active user on Robinhood checks their account seven times a day. Robinhood realized early on that the key to building the modern brokerage was to optimize for engagement, the same way a social media platform does. In this way, it reimagined the brokerage as a fun product, a high frequency gambling application.

The implications of the Continuous Order Book are therefore obvious. The future of online gambling is not to place a single bet that costs 7.3% of your principal before or during the game. Rather, the future is swing trading during the event. The future is placing multiple trades within a single sports game, each for nearly free.

B. Offerings

One implication of the COB is that it has the potential to turn any livestream into a live prediction market. Whereas it is incredibly rare to find a live market on League of Legends in the status quo (and when you do, the spread is astronomically high), the COB will allow us to run live markets on popular eSports 24 hours a day.

Beyond sports and eSports, the COB gives content creators avenues for monetization beyond donations and advertising. Rather than simply watching streamers attempt challenges, for instance, fans will be able to bet on them. By creating direct incentives for content creators to bring their audiences to trade on our platform, we ensure a diverse and interesting melange of prediction market offerings.

Where do these new markets come from?

In economics, the No-trade Theorem says that if two rational actors share the same utility function for a single good, they should never be willing to trade with each other. This is because, if trading were truly a zero sum game, then there exists no trade that leaves both people better off. Therefore, that prediction markets function at all leaves something to be explained.

The obvious response, of course, is that people have utility functions for things other than profit. Some people derive utility from the act of trading. For the gambler, the experience of gambling is an intrinsic utility, and such, the retail gambler is willing to pay for it.

To recap Section III.B, in the current paradigm, when a retail gambler wants to bet on a sportsbook or exchange, they trade against an institutional market maker. From the perspective of the institutional market maker, the decision to make the market is a very simple and objective calculation about whether it is profitable to do so — i.e., the institutional market maker provides liquidity only when it knows that revenue from collecting the spread will offset the cost of adverse selection and model creation. As we explained, because it is so costly and difficult to model volatile and unique events, it makes no sense for the market maker to make markets on things like eSports livestreams or niche live stunts from content creators.

In Section V.A., we argue that the main innovation of the COB is to enable retail gamblers to make markets, supplanting the role of institutional market makers. What is crucial to understand here is that this flipping of roles does more than reduce the price. It also fundamentally changes the calculus for what markets see liquidity.

This is because whereas the institutional market maker needs to be profitable if they are to provide liquidity, the retail gambler does not have this requirement. Instead, when retail gamblers are the market makers, any event that would be fun to bet on becomes a market. Even when the retail bettor loses to adverse selection due to the difficulty of accurately modelling the event, the market is made nonetheless because profit is no longer the primary consideration.

Why should I care?

Regardless of what Robinhood's PR Department wants you to think, Robinhood is a gambling application. Today, a majority of Robinhood's revenue comes from options and cryptocurrency. In Q2 of 2021, Robinhood made 31% of its revenue from Dogecoin.

Yet, Melange offers a strictly better gambling experience than Robinhood does. Whereas gambling on Robinhood is abstracted — a bundle of variance tied to a financial security that bears no personal relation to the user — on Melange, the gambling is pure. On Melange, we give people what they want: the chance to bet on the things they care about.

C. Design

To design is to communicate meaning. It is to take the logic of a program and to translate it into the human language of symbols. Apple brought personal computing to the masses not through its hardware, but through its design. Marginal advances in computation mattered little to consumers — rather, the move from text-based interfaces to the graphical user interface made computers so intuitive that they became a part of everyday life.

Analogously, the magic of the blockchain and the elegance of the COB is meaningless without proper aesthetic expression. Technology reaches the mind foremost as a feeling. Accordingly, we are sensitive to the observation that most people are repelled by markets not out of illiteracy but from the presentation of markets as asocial and unintelligible black boxes. Through unprecedented attention to visual communication, we aim to create a product that demonstrates the deep social beauty of the market.

In Figure 5 below, we show one very preliminary idea for how we might represent trades on the COB. Here, when a user queues an order, it is represented as a bubble. When a user's order finds a counterparty, we represent the filling of the order as collision of bubbles, enlivening and gamifying the dynamics of the market.

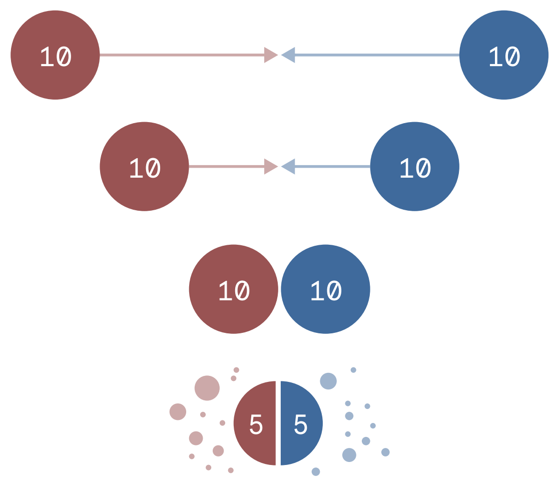

Figure 5: A Visual Representation of a Trade

We include this preliminary idea not because it is in any way final, but rather to signal our general approach to design. Instead of defaulting to a line graph charting price over time as Robinhood does, our approach will be to design engaging and symbolic interfaces that clearly demonstrate the internal logic of our markets.

Why should I care?

What Robinhood understood, that its competitors did not, was the importance of design. They recognized that, at the end of the day, they were selling an experience, and as such, their job was not to create a financial instrument, but to create truly meaningful experiences for their users.

At Melange, we intend to elevate design. For every engineer we hire, we intend to hire two designers because frankly the future belongs to the artist, not the engineer; and the killer app for crypto will be a user experience, not a backend. When the blockchain sees widespread adoption, people will not use a dozen mediocre frontends each doing a single thing, as we do now. Rather, people will demand a precisely designed interface that legibly renders all of the blockchain's complex and variegated applications. Curating and bundling the best consumer applications of the blockchain together, this product will include everything: cryptocurrency, NFTs, stocks, and prediction markets, all traded on one ubiquitous, soulful interface.

VI. ADJUDICATION

The innovation of the Augur protocol was to create a system of escalating disputes that made it nearly impossible to profit from falsely reporting a market outcome. Melange will improve upon Augur's original design by adding additional gamification that encourages more consistent user participation in the dispute process.

Specifically, when Melange cannot resolve a market via a decentralized external oracle such as Chainlink, the Melange protocol will resolve the market via a escalating system of courts similar to the US justice system where, in cases of dispute, evidence will be presented in trials with human justices, plaintiffs, and defendants.

What is written below is by no means final, nor complete. There is a reason why Augur's Whitepaper is completely dedicated to responding to this singular problem. Thus, before we implement this feature, we will think, write, and publish on the subject in a more extensive manner.

Judicial Appointment. Any user can become a justice by putting some amount of governance tokens in escrow in return for a "Judge Badge." Whenever a justice successfully resolves a dispute, a set number of governance tokens will be added to their escrow pool. Conversely, whenever a justice fails, a set number of their tokens will be burned. When a justice's escrow depletes lower than some minimum balance, the badge deactivates until said justice reactivates their badge by posting more tokens to the escrow. Users may also choose to deactivate their "Judge Badge" at any time by withdrawing their governance tokens from escrow.

Resolution. When a market is created, a user (most often the market creator) is identified as responsible for resolving the market. Binary outcome markets can be resolved four ways: in favor of outcome A (e.g. yes), in favor of outcome B (e.g. no), momentarily pausing market (e.g. a tennis game is suspended because of weather conditions), and resolving market to another percentage (e.g. 50/50) in the event of an invalidated market.

Disputes. Whenever a market is created, the resolver is required to post some amount of collateral as specified by the protocol. The amount specified should be enough to discourage users from frivolous disputation and duplication of markets, while not enough to discourage market creators from creating markets at all. (Of course, if the market creator wants to resolve the market instantaneously, the resolver can always post collateral equal to the amount needed to pay out both sides in case of a dispute.) Once a resolver makes a decision as to the outcome of an event, any user on the platform has 24 hours to launch a dispute. If there is no dispute, then the winning side is paid out and the collateral returned to the market creator.

To create a dispute, the user stakes tokens on a challenge. A challenge requires exactly 50% of the market creator's collateral to be launched. Once a challenge is successfully launched, 5 justices are randomly chosen from amongst the available pool.

Trial. Once a dispute is launched, both the challenger and the challenged have 48 hours to submit evidence on their behalf. The evidence will be stored in a secure, decentralized location such as Arweave — at which point the justices of the court will have 48 hours to come to a ruling. Any justice who does not vote will have a portion of their escrowed governance tokens burned and a new justice will be appointed to replace them.

Decision. Once all the justices have voted, any justice who voted for the losing side will have a portion of their escrowed governance tokens burned so long as this ruling isn't successfully challenged in the Supreme Court. Given a 50,000 governance token requirement for the creation of Judge Badge, for instance, deciding for the minority opinion in a trial might lead to 5,000 of the justice's tokens being burned.

If the challenger wins the dispute, then the challenged surrenders their collateral. Half of it is split amongst the judges. The other half is split amongst the challengers. On the other hand, if the challenged wins then all of the challengers' stake will be split amongst the justices. The challenged does not earn anything as this would encourage them to write contracts that might be prone to challenge.

Supreme Court. If people believe that the court made an incorrect ruling they may elevate the dispute to the Supreme Court. To do this, they must first collateralize their challenge by staking the projected court costs. Once the Supreme Court is called, every justice has 7 days to make a ruling. Once the Supreme Court has made its ruling, any justice who voted for the losing side or abstained from voting will also have some of their tokens burned.

VII. CONCLUSION

This Whitepaper outlines a minimal viable version of the Continuous Order Book, but a more maximal account would include features like fees on variance, limit orders, liquidity boundaries to protect LPs, scalar markets, multi-outcome markets. Moreover, it is not a trivial problem that livestream latencies will disadvantage some market participants at the expense of others. Future updates to this Whitepaper may cover these topics, and if they do not, please rest assured that the protocol will and that these features are top of mind.

Nevertheless, the Continuous Order Book is just the beginning. On a longer time scale, we imagine social communities with P2P games like poker and private user generated prediction markets — as well as new long term prediction market structures that solve the same adverse selection problems that the Continuous Order Book solves. We promise these will be interesting — but alas, time is of the essence and the reader must look to the future.

VIII. DISCLAIMER