

An improved marketplace

The goal of this paper is to describe a novel method of price discovery, discuss its tradeoffs, and show that this system, Coordinated Discovery Markets (CDM), is superior to existing systems. This document is organized in nine sections. Section one describes the useful aspects of existing markets which must be preserved. Section two shows the difficulties which must be overcome. In section three the assumptions of economics and human behavior are laid out. In section four the mathematical foundations referenced by this paper are described. Section five details the proposed solution. Section six begins the validation process with qualitative support. Section seven expands on this with scenarios to stress test the system behavior. Section eight moves into quantitative support. Finally section nine provides logical-mathematical proofs.

Section 1: Price discovery the function of markets

Price is vital economic information. It is also abstract information, when a storm system devastates crops or a breakthrough raises yields the consequences positive or negative are communicated through price. It is the aggregation and simplification of the multivariate factors of an economy.

We need this because being simultaneously aware of everything going on isn't possible. In order for individuals to coordinate action prices provide the information and incentives that make it feasible. In addition to the actions of the moment people can think and plan for the future and so prices over time are the distillation of all economically relevant information.

For commodities several methods are used to find and communicate price information, but the most effective and therefore most important are markets. At a high level markets that we currently use work by constantly maintaining a balance between supply and demand and moving the price point so that both sides are in accord. By fixing certain dates contracts to deliver or receive are standardized and the creation and trade of these contracts produce accurate future price sentiment and a historical record of price evolution.

At a similarly high level this paper proposes an inversion of this model where a predetermined price projection seeks to attract equal supply and demand. This sounds strange and it hasn't been done before but this paper will show that it is both practical and economically desirable in many instances.

So marketplaces provide a service and that service is information, but how valuable is that service? To answer that we look at why the service is valuable. The economy is a process of billions of people planning actions globally. The effects of events major and minor can impinge on these plans in myriad ways and so price information has evolved as a mechanism to abstract the news of the day into the most relevant and useful message. Individuals can then assess their plans in light of this data and guide themselves as best they can. So it is by improvement of human decision making that market value is delivered.

It is this coordination of decisions which creates the "invisible hand" Adam Smith remarked upon. The discovery and publishing of price information enables this advance in economic activity which can pay for the service of finding and advertising prices. The value of various market systems can then be assessed by referencing their effects on the economy. The

provision of information as valuable for coordinating human action for less cost than existing methods is consequently to be preferred. This paper will lay out one such system.

Section 2: Problems with current systems

Advances in communication technology stretching back at least to Morse telegraphs have driven markets to speed up. Modern computers have pushed that to levels that fundamentally alter the behavior of the marketplace as information is now communicated too quickly for humans to process.

These microsecond transactions have not resulted in appreciable improvement in either price stability nor lower transaction costs for the entire system. Instead we see deal flow increasing faster than fees fall.

With finer time distinctions more chances for liquidity events occur creating risks in the form of flash crashes and High Frequency Trade (HFT) schemes. These risks raise the transactional burden of the marketplace as a whole.

Existing marketplaces also require liquidity in order to discover prices. This must be provided by persons other than the producers and consumers as if these groups had the spare capital to provide their own liquidity it would be possible to out compete them by operating in the market without those resources. External liquidity must be attracted by return which contributes to transaction costs.

Without trades taking place the marketplace has no price to offer. With market speed driven by computational technology which far outstrips the rate of productivity improvement in the rest of the economy inevitably trade volume outstrips trades for delivery. The complexity of the resulting marketplace degrades the value of the information it provides as the difficulty of evaluating that information increases.

These factors combine to make liquidity return proportional to the marketplace's transaction costs. Within existing designs lowering transaction costs can reduce liquidity return which in turn damages or destroys the market's function. By uncoupling these features a new category of better marketplaces can be implemented.

While this service is vital to the functioning of the economy the costs are deadweight losses so a system that can reduce those costs immediately and directly benefits the potential of the economy. Producers will see immediate gains to their bottom lines, consumers will find products being less expensive. Furthermore as those in finance who are no longer needed find their way to other professions the economy will gain there as well.

So we see two sources of transaction cost in existing markets. First there is the cost of attracting liquidity. This cost is proportional to the average opportunity costs by uninformed traders, a category which includes the producers and consumers on average. Reduction of this cost using existing market design will result in a broken marketplace rather than a competitive advantage, as failing to offer attractive return on capital will cause liquidity to dry up in the market.

Secondly there is the cost of using the information provided by the marketplace. The more complex and ephemeral this data is the less useful it is. This creates a tradeoff for existing market design leading to diverging interests between the price taking users of the market and the informed speculators. Resolving this divergence allows for the creating of markets which function materially better than current designs.

The complexity of the marketplace at scale also brings up regulatory issues. The amounts and importance of deal flow with the marketplace makes them ripe targets for fraudulent, manipulative, or malicious behaviors. Countering these requires independent monitoring which becomes yet another expense of operating the marketplace which ultimately must be borne by the producers and consumers. Simplifications which damage liquidity returns suffer the same fate as markets which attempt to compete by favoring price takers. And so regulation is a complex and technical field which drive costs up.

One difficulty is that many manipulations are known to be plausibly deniable. This is structurally linked to the unification of speculation and trade in existing market designs. Effectively every action is a trade and a speculation and a manipulation, so disambiguating intent is intrinsically difficult. Since both under and over regulating are harmful and fraud is inevitable even when you get it just right this is a hard to measure but important cost imposed by existing markets.

A second issue with the derivative driven model of existing marketplaces is the counter-party issue. While marking to market and reserve requirements make counter-party failure rare a cascading market disaster it is never the less a constant risk which must be assessed at the broader economic level when evaluating market designs.

Section 3: Basic economic assumptions

The supply and demand of a good at a given level of price is a function of economic conditions. In a competitive environment opportunities will be exploited in increasing cost order. Or as generally stated declining marginal returns.

This means that for a given set of economic conditions there is a price at which the consumption value of the marginal item closes in on the production cost of that same item. This is known as the clearing price and it has several important properties. First the rate of production and consumption are equal at the clearing price. Second it is the price level which creates the highest returns for producers or consumers which are subject to market competition. Third it constitutes part of a maximally functioning economy providing the broadest benefit to society. Fourth it is in the individual interest of each producer and consumer to lie about it the benefits are at the aggregate and not individual level.

Each individual whether in production or consumption has their own ends to satisfy, the aggregate behavior of the marketplace is simply the sum of these motives. Markets thrive by maximizing the satisfaction of these desires and so reasoning effectively requires analysis at the individual and societal levels.

Broadly the a market can have two states equilibrium and dis-equilibrium. Equilibrium is characterized by the clearing price. This concept of equilibrium still contains significant

dynamism.

Under conditions of equilibrium individual returns among producers and consumers are determined by the relationship between their personal marginal positions and the aggregate curve. Those whose margins are far from the market price will probably find growth quite profitable while those near or at the margin will find themselves pressed to remain operational. These forces work on both the supply and demand side and drive greater productivity benefiting society. So even in the perfect functioning of a marketplace we still expect and indeed hope for the competitive liquidation of sub-marginal market participants.

Dis-equilibrium must also be considered. To avoid this state as much as possible incentives must exist to correct price errors. If these incentives are stronger than countervailing motives to maintain dis-equilibrium then market forces will minimize or eliminate dis-equilibrium events.

Some events, thankfully rare, are truly unanticipated and can result in market price errors, examples include global scale natural disasters, mass delusions, and war. In these cases the marketplace must have mechanisms for distributing the gains and losses and recovery.

The effects of long term price errors are evident from the historical record on price fixing. Prices fixed above or below the clearing price result in persistent shortages or oversupply respectively. Economies maintaining these conditions suffer vast potential loss.

A related phenomenon can be seen in productivity change. No law of physics has changed in the last five centuries but prices and production levels have altered radically as our awareness and application of them has advanced considerably. These new and advanced forms of production require discovery and investment to apply. So we expect the clearing price to co-evolve with the technology available to the society using the marketplace. So an equilibrium marketplace is quite a dynamic object with both an evolving clearing price and constant ferment of competition of both production and consumption.

The success or failure of competitive markets will be determined by the independent choices of their critical users: producers, consumers and informed speculators. Those choices will be governed by the self-interest of the individual. Marketplace's enjoy network effects which cause them to grow to monopoly conditions. So competition between marketplaces will be a winner take all affair in general. There are two well known economic principles governing such a conflict. First Gresham's Law which states that bad money drives out good money shows that economies will shift quite quickly between an alternative of higher cost to lower cost. The second is the historical experience of black markets which show that considerable and continuous state power must be quite publicly deployed to maintain a less effect marketplace in competition with a more effective one.

Section 4: Mathematical foundations

This will not be a thorough treatise on the mathematical subjects which are exploited here and full mastery of those subjects isn't a requirement of understanding this paper. I will mention the critical features without proof. Some familiarity with these areas will allow this section to be skipped. So firstly game theory and specifically coordination games. Define a coordination

game to be one in which no player can secure a better outcome for themselves by disagreeing with the other players. Coordination games have a strong tendency toward stagnation and a concept called the Schelling point is a naturally occurring point of agreement around which a coordination game can stagnate. Also important are iterated games and evolutionary games. The behavior of a game can alter radically when the same players play it repeatedly. If we combine iteration with allowing many participants to independently play in each of the roles of the game the resulting behavior is governed by evolutionary game theory which essentially states that the system will gravitate towards the best conditions it can manage and then stay there.

Metric spaces comprise an incredibly broad category with only the requirements of a set and a metric function. That metric function has only three strictures: defined for any pair of elements of the set, always non-negative, and the triangle inequality essentially no short cuts.

Information theory is a rich subject and is already quite influential in finance with discoveries like the Kelley Bet. We are going to use the most elemental form with Shannon's work classifying information in a quantifiable way using binary systems.

Section 5: the Solution

The basics of CDM is modifying a coordination game with a metric space of potential Schelling points allowing the players to evolve to a globally maximizing outcome. Structurally that is an iterated two stage game. There are three player roles: (M)aker, (U)ser, (S)peculator. First S proposes a price function $P(t)$ then M and U independently choose a level of participation for $P(\text{now})$ and S is rewarded. The less P changes between iterations and the closer the sequence of $P(\text{now})$'s correspond to the clearing price function the higher the return and reward for S respectively. There are multiple people in each of these roles and for a given iteration it is possible and likely that M and U players will also participate as S players. The governing system and metric spaces can be described as follows:

Let $P(t)$ be the function of price over time in a given market.

Let ΔP represent a speculative price function.

Let the distance between potential price functions $D(P, \Delta P) = \int |\log(P) - \log(\Delta P)| dt$

Let the expected value to the market at any time be $V(t)$

$V(t)$ can be updated empirically as the market functions by calculation as a moving average of volume times the markets total transactional overhead charge.

Let $R(t)$ be the rate of return, I'm proposing using 100% annualized so for a t in years $R(1) = 2$, $R(2) = 4$ and $R(.5) = 1.414..$ but any function is possible.

Let $P\Omega$ = the trading price.

Then the cost of speculation = $V \int |\log(P) - \log(\Delta P)|/R dt$

and the return on speculation is ideally $V \int |\log(P\Omega) - \log(\Delta P)| dt = VD(P\Omega, \Delta P)$

Less elegantly return is a fractional payout of the parimutuel payments for speculation augmented by the portion of V designated for rewarding informed speculation, I'm proposing 90%. The fraction is based on the proportional information provided as measured using Shannon information. This is made straightforward because the metric space corresponds to Shannon information measures.

I will now describe each of the bi-lateral relationships between the four roles: producer,

consumer, speculator, and market. For this description I'm going to assume that the system is functioning as expected for simplicities sake. This provides a general overview of CDM function.

Producer - Consumer:

These two exchange their goods securely at ideal prices. They pay a fixed commission that is lower than whatever average cost they currently encounter.

Producer - Speculator:

Speculators attempt to guess the price function that will fetch sufficient demand to meet supply. For producers specifically honest account of their own marginal costs of production is a nearly ideal speculating strategy.

Producer - Market:

The producer puts his goods up for sale and the market informs him as to the distribution. As delivery is made the market releases the funds in escrow into his account

Consumer - Speculator:

Again speculators attempt to guess the price function that will fetch sufficient demand to meet supply. For consumers specifically honest account of their own marginal value of consumption is a nearly ideal speculating strategy.

Consumer - Market:

The consumer escrows payment for his demands. The market informs him of how those are to be met and as delivery is made the payments are freed to the supplier and the commission payments to the marketplace.

Market - Speculator:

Speculators provide price information to the market by paying for the amount of data they provide (how far they move the price). The market holds their money, adds a portion of the commission from trades influenced by their data and redistributes the resulting money on a pro rata basis of the amount of accurate information contributed by each participant. Speculators exist in a commission free positive sum marketplace with average good investments paying off at a rate higher than anything else known. However if their cash exceeds their knowledge there is no practical way to leverage a capital advantage to find return. Speculator and market income are positively correlated with market volume.

Ultimately all four sides want the same thing, for maximum trades at clearing prices to occur. Competition is confined within groups, so speculators compete to provide better forecasts, producers compete to produce more plentiful (and cheaper) goods, consumers compete to use more plentiful (and more expensive) goods. The market maintains its position through transparency and disinterest and really at the heart of it maintaining trustworthiness by having the computers just run the algorithms.

Taking each perspective into account in a pure sense, neglecting the capacity of individuals to play more than one role, we can also describe each experience with the system.

Producer:

At the logistically dictated pace of your industry you see a calendar of price evolution which is as accurate as the human species is capable of producing. The current price which is fixed for human and logistical time frames and specific to the needs of your industry reveals the current price which you have been watching for quite some time. You make decisions calmly make delivery as expected and get paid securely.

Consumer:

The Consumer's position is symmetric to the Producer's.

Speculator:

At a pace as sedate or active as demanded by the logistical needs of the industry you look at the current price information and the historical speculation which lead to it. Synthesizing this with any other market relevant data you possess you attempt to accurately update the price calendar and determine your monetary commitment to your beliefs.

Marketplace:

During each trading period I gather and secret all supply, demand, and speculation data. Any associated monetary commitments are placed in escrow. As the settlement phase replaces the bidding phase the distribution pattern of buyers to sellers is established. Any miss in supply or demand is distributed, currently I propose a minimization of marginal cost. The escrow is maintained to ensure delivery. Also the speculations are settled and recorded. Any unused speculative commitments are rebated to the speculator. The information distribution for this trading period is calculated and as trades clear the market commission is added to the parimutuel pool division for speculators active in this trading period. My portion of the commissions are collected into the corporate account and all data is then published.

This is a good spot for a computational diversion to characterize the risks and costs being described. Risks in existing markets are generally modeled as normal though the tails are known to be fatter than that. CDM prices are aggregate prices which the entire market are potentially participating in so the central limit theorem will produce normal curves so long as interests aren't correlated. The inclusion of both buyers and sellers with their differing perspectives makes a strong case for normal behavior. The comparative risk for prices posted by CDM and ordinary markets naively will be a function of the number of participants. A CDM with 49 players can expect a price risk reduction of 6/7, with 1024 players 96.875% risk reduction, and with 40,000 (a reasonable number for a global benchmark) a 99.5% reduction in price risk. Moreover clearing risk, the chance that some goods may remain unsold or unavailable, is just the risk that the price is wrong so clearing misses are exactly as unlikely as pricing errors. However these dramatic reductions rely on humans being able to behave no better than random and consequently that all civilization, art, invention, and investments had paid off through simple blind luck. A more realistic model of price comes from a chained Bayesian function with each participant updating priors based on all inputs. The rate at which this will converge will depend on the distribution of priors but the convergent behavior is that the error will halve at each iteration. The next section will include some conservative examples of what it takes to induce price and delivery risk to a CDM.

Another interesting number to compute is what is the cost of propping up an existing market

in the face of CDM competition. First the cheap scenario in which all of the competing speculator, broker, market maker, and liquidity provider entities choose to voluntarily and in lockstep accept less revenue in exchange for continuing to provide the identical service which is the present marketplace. Since a CDM can easily be set up with total costs that are lower than the operational expenses of existing markets. So the ability to maintain will be the amount of time that the least funded critical actors in the existing marketplace can operate at a loss. However CDM will offer some of these participants better lifestyles than the existing market will specifically people that will rapidly bankrupt themselves operating at a loss will tend to see a dramatic improvement from switching their activity to the CDM. So the more realistic and expensive scenario involves a small group with deep pockets funding a subsidy for an existing market to continue. How much does it cost to maintain such a subsidy? Well to be effective the subsidy must be an obvious money loser for those offering it, it has to be constantly available so the end users can avail themselves of the advantage, and it can only be limited by the capital posted for the trades. In other words the cost of providing such a thing is infinite since a simple series of trades will transfer some of the capital in the subsidy to the person executing the trades and leave them otherwise unencumbered and able to repeatedly take money from the subsidy pool until the subsidizers are bankrupt. For those wondering how a marketplace can have no calculable risk and infinite competitive costs the next section will explore these causes in some detail.

Section 6: Qualitative Arguments

These are significant claims and so many strands of argument will be advanced in defense. First we will lay out the incentives at the individual or micro level. Then macro level improvements which support these. We will close with a variety of proofs and examples to create a behavioral feel for the nature of this solution. One principle that will come up again and again is that CDM is a kind of obverse or dual of existing market structure and so reasoning and conclusions will frequently run backward from what is normally encountered when discussing marketplaces.

We begin with the speculator. By offering a parimutuel system with commissions he enjoys a positive sum game rather than the negative sum game that confronts investors currently. Secondly risk is contained to whatever level of monetary commitment chosen at the time of speculating. Finally Average return on risk assumed may be set to any level desired and so people capable of correcting the marketplace can enjoy average returns on investment which are higher than any historical example from existing markets.

Turning to the producers and consumers they enjoy a marketplace with lower total costs of transaction and an effective suspension of the need to hedge. More on that in the macro discussion.

In order to compete a CDM probably has to accept less revenue than an existing marketplace it will supplant. However as the costs of operation of a CDM are small fractions of a percent of the costs of operating anything like a traditionally designed marketplace per capita profits can be higher.

So we will now explore the origin of these benefits. Structurally producers and consumers are symmetrically hedged. That is the aggregate desire by consumers for a steady supply is

matched by the aggregate desire by producers for steady demand. This means a clearing price will not only produce trade it will produce globally maximizing trade. Consequently by setting up a marketplace which will reward individual efforts that approach this global maximum is in the collective interest of ultimate traders. With settlement of trades made disinterestedly each producer's or consumer's cost of trade can be set at a level which will be lower than the average currently encountered. This means that everybody whose trading costs, including hedging, isn't better than average will have an immediate monetary incentive to switch markets. As they leave, the zero-sum nature of the existing marketplace will reduce the rewards gained from enabling those markets leading to a kind of inverse cream skimming until the existing market has been abandoned and the CDM is the primary benchmark.

Then there are regulatory issues to consider. Speculating is now paid off by the individual contribution to market function. Speculative errors have their investments redistributed to pay off those who correct the mistake. This eliminates the competition between speculators and traders for commodity flow while actually sharpening the competition for information provision. This major change has three effects.

The first of these is leverage elimination. Since high speculative return is a structural part of the market design the need for leverage to achieve acceptable returns vanishes. Without margin trades the need to monitor the fiscal health of speculating entities is non-existent as the maximum loss they can incur is the amount they invested in their speculation. Likewise they do not need access to any of the commodity being traded because they have not and very likely will not take a trading position. Pure speculating earns much higher returns than speculative trading as it incurs no delivery or storage costs, no commissions, and any potential gains to be made from speculative trading could be profitably eliminated through price speculation. With these issues becoming non-issues the liability and expense of dealing with them is eliminated.

The second regulatory problem which is effectively gone is money laundering. Since the amount of speculative capital has been greatly reduced, from a multiple of trade volume to a fraction, and the destination of speculative capital is uncontrollable the only entity capable of directing a monetary flow for ulterior purposes is the marketplace itself. Since the marketplace operates by executing a fairly simple algorithm and publishing everything it is doing any attempt to do such a thing would be trivially determined by anyone on earth that cared to look. That level of risk/reward makes monitoring this problem effectively free.

Finally market manipulation, since CDM makes trades for delivery in order to validate and reward its published price calendar successfully inserting erroneous information only serves to lower everyone's return including then hypothetical manipulator. However since correcting erroneous information pays a massive ROI anyone aware of the manipulation can make a perfectly safe investment with a great return undoing the manipulation. Furthermore manipulating prices to damage markets is illegal. In short attempts to manipulate CDM markets will be obvious, and therefore trivially prosecutable; unprofitable, since the only practical effect of success is depression of trade; difficult, since anybody aware of the manipulation can make a great return correcting it; and finally dangerous, since these factors mean at best one can expect to lose their entire investment and at worst the criminal fines and incarceration will result in having even less money plus loss of freedom. So this too

becomes largely a non-issue.

Section 7: Gaming the system

The failsafe nature of CDM will be demonstrated by four extreme gedanken experiments. The ridiculousness of these premises and the relative harmlessness which results will be very reassuring.

The first is 'The Asteroid Strike' some global scale unpredictable catastrophe hits., CDM outcome: with no leverage nor commitments there is zero systemic risk the market readjusts to the new reality on the ground and life goes on. So lets unpack this. To begin the disaster must be huge, a train derailment or a bad storm wiping out one farmer's crops won't have any significant impact on a global price index. Second the disaster must be unanticipated, if everyone knows that a brewing war will disrupt shipping and effectively shut down global supply chains then that information will be properly reflected in the evolution of prices. Finally we turn to the lack of systemic risk. Within this system every participant's counter-party is the marketplace which operates at full reserve for all transactions, consequently market behavior won't exacerbate any crisis as there is no counter-party unwinding, de-leveraging, nor any other category of unmeetable promise in the system to cause risk.

Then there is 'The Lex Luthor Scenario' some person or small group that controls more wealth and more information than the entire rest of the marketplace exists, this person can sell you soap as food and tell you its tasty and you will continue to believe them even while you're eating it, CDM outcome: whatever they say same as any other system saddled with this kind of actor. It may seem useless to consider such outlandish ideas and to make it clear this presupposes the existence of a person or group with the abilities evinced by the DC comic book supervillain. However as touched on previously discussing market manipulation this is the minimum set of capacities required for successful manipulation of CDM markets. The resources required to derange a market are large but they pale in comparison with the resources required to get everyone's cooperation, and it literally takes everyone to not step in and fix the problem, putting this firmly in the realm of fantasy.

Which leads us to 'Bill Gates Gets an Aneurysm' a billionaire wakes up one morning and believes that copper should be more valuable than gold, CDM outcome: the copper market has a pretty bad day then it fixes itself and all of the billionaire's money is gone. Assuming a large ticket commodity trading at roughly a billion dollars a day then large scale price derangement using the design parameters put forth in this paper will be on the order of ten billion dollars. This will have to be provided as currency in the form the market trades, each marketplace must have some designated currency. Which isn't an inconsiderable challenge as the consequences will be dire. It isn't even clear that such a trade would even be legal to accept in the first place and caps could easily be calculated for price adjustments in the near term which simply eliminate the fuss and bother of losing one trading day. Also such activity would effectively destroy a day or more worth of trading in the existing markets as margin trading would allow such an actor to cause a liquidity event in the tens of billions eventually ending in their impoverishment.

Finally I will walk through an idea that CDM cannot survive 'Non-existent Clearing Prices' imagine that the supply and demand curves do not intersect in the positive realm. The lowest

price for which there is supply is higher than the highest price for which there is demand. This may seem ridiculous and certainly for any existing commodity market it is but consider a privately held company whose owner will not sell the stock at any reasonable price. In this case a CDM will persistently fail until the participants give it up for a bad job. However no market scenario can create deal flow under these conditions by definition.

Section 8: Quantitative Arguments

First capital liquidity requirements in existing markets are on the order of the value of the transacted volume of goods. The number of trades is a large multiple of the number of deliveries but margin trading allows leverage to stand in for actual cash. Still with trillions of dollars being transacted every year an incredible amount of capital need to be committed to operating the existing markets. Setting the parameter functions $V()$ and $R()$ to lower transaction costs by an order of magnitude and raise marginal expected returns on capital by an order of magnitude lowers the capital requirements of operating a competing CDM by 99%. There is nothing special about these choices of function so in setting up a CDM one is free to choose whatever will make inter-market competition most effective from the operators standpoint.

A small digression on designing price calendars. Some useful figures to keep in mind when making these decisions: 8.64×10^{10} μs per day and 0.02mm per μs . Microseconds are used as a unit of measure since existing markets are converging on microsecond transactions, however $0.02\text{mm}/\mu\text{s}$ is approximately 45 mi/hr which is actually a fairly high speed for logistical supply chain operation. So trading windows need to be chosen so that reasonable amounts of deliveries are actually possible and so that the potential loss of a trading opportunity to one of the morbid scenarios can be tolerated. These speeds give an approximate radius of 1000 miles per day so national to international markets would trade on a calendar in the day to week range. This provides a factor of 8×10^{10} to 6×10^{11} lowering of the stringency of the computational system design. This factor then constrains the lower limits of $V()$ since a CDM must meet its operational costs to continue functioning. As one can see its not much of a constraint as underbidding a competitive market's electricity budget is possible leaving the competition with no viable mechanism for defense outside of politics or warfare.

Another massive improvement CDM offers by shifting value from liquidity to information is a straightforward strategy to allow the producers and consumers to take the lion share of the speculative pool for themselves. To reiterate an earlier point the best price under the dual strictures of competition and voluntary exchange is the level where marginal cost of production makes its closest approach to marginal value of consumption. At that level the maximum number of trades take place and every trade is worth making. These margins are the margins of the producers and consumers so they have a privileged position in access to this information since its simply their own operational data. Thus bidding your own margins becomes the ultimate hedging strategy since your payoff will increase the closer to the marginal participant you actually are. To estimate how close this strategy will come to the correct price estimate we can examine the conditions of the game. Each trading period can be updated until it is active on the same schedule as trading that means with trading windows ranging from day to week lengths and business forward planning running from 2-6 quarters we can expect roughly 25-500 update opportunities until the price needs to be acted upon. Since each set of speculations becomes public once it is integrated into the system each

speculator can update their priors based on the common understanding of everyone's opinions. That make the sample size effectively exponential $O(\text{speculator pool}^{\text{trading windows}})$ which even for modest group sizes should home in on effective clearing prices essentially all of the time.

The last property I'd like to highlight is the early adopter windfall. A CDM maintains and updates $V()$ as it operates to track what level of reward can be expected for information received. Starting a CDM in competition with an existing marketplace is possible with extremely low market share 0.01-0.1% will frequently be quite reasonable. This means that as the CDM takes market share from the competition $V()$ will increase by a factor of 1000-10,000. Since speculative costs are sunk speculating during the initial phase of operation on trading periods which fall after CDM replacement of its competition will payoff with this higher $V()$ creating powerful incentives for advocacy of a CDM once it goes into operation.

Section 9: Logical Arguments

We will begin with the simplest and most devastating proof of CDM efficacy to whit:

Prop 1. Markets work by extracting and aggregating price information from human sources by offering attractive marginal returns to capital.

Prop 2. CDM extracts the same type of information, though arguably less restricting itself to a useful subset and offers more attractive marginal returns to do so.

Conclusion. Believing CDM doesn't work, and in fact believing that it doesn't work better, requires either believing that existing markets do not work, and thus losing access to any counter-arguments, or providing a novel, and by novel I mean supernatural, explanation for market function.

Being non-constructive this isn't the most satisfying demonstration and indeed it is possible that markets don't function the way that we think that they do or that humanity has some unknown cognitive deficit that makes CDM fail and we haven't encountered it before because no one has ever tried this before. So turning to a constructive demonstration beginning with some preliminaries. The first concept we need is price function space. The price function for a given good is the function of clearing price over time. While access to the true price function is impossible all possible price functions can be found within the space of functions. The differential economic value of two candidate price functions can be calculated by operating an economy with each of them and comparing the results. This is impossible, thankfully marginal analysis tells us that the ideal price function will match supply to demand and it is possible to test a candidate price function for that behavior. Since trade is valuable enough to generate commissions it is possible to have a reward which is proportional to the success of a candidate price function. Since price functions are functions, embedded in the space of functions, integration is a natural metric. Since the present is more valuable than the future, as only the present can be acted within, weighting the integral with a suitably decaying exponential function will represent this and allow for any level of aggregate return one would wish for the resulting system. Setting the commission level allows for the definition of transaction overhead and potential reward. So all that is required to create a functioning marketplace is the function of transacted volumes. By empirically updating this function based on the volumes being transacted by the now mutating candidate price function all of the requirements of a functioning marketplace have been met assuming one speculator, one producer, and one consumer. All that is then required for a CDM are settlement strategies for

each player category to allow this to operate as an evolutionary game.

Settling trades is a vast area. In CDM buyers and sellers register amounts they wish to trade but there is no further discriminator among them. In the event of mismatch between supply and demand (which I expect, in very small amounts, to be essentially constant) some mechanism for allocating the shortfall must be implemented and whatever it is it needs to be clearly articulated beforehand and execution needs to be accountable. I'm proposing a system of randomized order bottom up filling or to put it inefficiently: take the larger side randomly order them and then traverse the list incrementing so long as the entry you are on has capacity until the smaller side's aggregate is reached. This is based on the idea of diminishing marginal returns on order fulfillment thus minimizing total marginal costs of misses. The big positives are that small and even average size players essentially never miss an order, which is a lot of customers, and that encourages competitive markets which makes the market work better and increases individual dependence on the market. The downside is that it provides an incentive to split into multiple accounts to take advantage. Trading under false pretenses isn't exactly legal and having a market reg is pretty trivial but awareness and enforcement are issues at scale.

However that isn't close to the only way to do matching, so I'm going to go over some less than ideal options and look at why they may see the light of day anyway. One interesting possibility is "The Buyers Market" or "The Sellers Market". Basically if a monopoly or cartel controls one side of trade already they may be less than enthused at the idea of a market incentive to break them up. If entry isn't possible without their consent then control over settlement could get things off the ground with say designated actors having their full order fill first always. This could allow monopolistic squeeze tactics to operate within normal market operations. This doesn't have to be sinister, for example a country that exported a cash crop heavily could start a CDM with mercantilist leanings and then if that were the first market it could grow to become the global benchmark. The country would enjoy a small but significant advantage for going first and the rest of the world would have a better marketplace.

Another natural but dangerous idea I want to mention is pro rata, the proportional distribution of resources so if 90% as much as desired is available everybody gets 90% of their order filled. This sounds attractive but its actually quite dangerous because it encourages false signaling to the market. When the market does price miss it is vital to have an accurate gauge of how bad the error is. With pro rata if you suspect or know that the market is favorable to your side of the trade (and therefore orders on your side will not fill) then you can over bid your needs in an attempt to get a great deal. This will then falsely signal how extreme the imbalance is causing short term problems and provide anti-competitive market incentives that could develop into long term problems. With this loophole in market behavior it becomes an advantage to be large enough to have the resources to exploit it and this market competitive strategy if it dominates productivity based strategies damages the market and the broader economy.

Speculation settling is another situation that has a vast option space. The system receives and then applies speculations to the future price calendar. This could be done in real time but that would lead to shifting information available during the decision window for buyers and sellers creating the possibility for increased returns from more active market awareness. This

is exactly against the direction I'm seeking to push things so batching this process is indicated. Once these guesses are batched there needs to be a method for deciding how to calculate the outcome. The most obvious method is time stamped serial application but since the last move is always "right" this creates an incentive to drive up operational costs which isn't really something I'm interested in rewarding.

One way to to flatten the time value across the bidding period is averaging. This creates a single update for each set of speculations but the assignation of ownership of individual information bits is challenging. Using arithmetical averages the greatest influences would be those furthest from the consensus but we want informed agreement and therefore want to reward the middle folks. We can also get into a "Beauty Contest" problem where correctly following general sentiment rather than correctly predicting the future is most valuable. Because of these issues non-individual methods of speculative effect resolution are undesirable.

So having worked out that we want to serially apply speculations but we don't what to use temporal ordering what ordering do we wish to use? There are $n!$ potential ordering strategies for n objects and the only thing we have shown is we don't want the one we have. One dangerous possibility is preference ordering where some select group through either money, political influence, market favoritism, or illicit appropriation are granted preferential settlement access. This is certainly damaging to the reputation of the market operator, though in the case of political influence perhaps unavoidable, but isn't necessarily ruinously destructive to operation of the market. Two decent options will be explored. First median ordering, for each trading period all speculations to be settled are ordered least to greatest and then this list is traversed from alternating opposite ends until each entry has been visited once. This maximizes the cost of speculation and delivers reported prices that always reflect consensus. Note that this doesn't invalidate an extreme opinion that eventually proves correct as the history of the price fluctuation will still record that wild eyed radical's thought and when mainstream opinion bends that way the reward for being first will be preserved. However by clustering the middle who by definition are closer together and by interspersing highs with lows there is a move of money towards the extremes at the expense of the middle compared to the final option. So now my implemented method which is: randomness. This sounds really bad which is why this section is almost all setup of all the pitfalls that exist in this part of the design. The nice thing about randomness is that it maximizes the incomes of the median bidders and so by incentive rather than force drives the market to desirable behavior. The nasty things about randomness is that considerable care must be taken to insure unpredictability and it sounds bad.

Conclusion

CDM is viable alternative design for commodity markets. Adoption in in the personal interest of all critical participants. The results of general adoption are more efficient and prosperous economies.