Studying Narrative Economics by Adding Continuous-Time Opinion Dynamics to an Agent-Based Model of Co-Evolutionary Adaptive Financial Markets

Arwa Bokhari[a] and Dave Cliff[b]
Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

Keywords: Narrative Economics, Opinion Dynamics, Co-Evolutionary Systems, Adaptive Markets, Financial Markets, Automated Trading, Agent-Based Computational Economics.

Abstract: In 2017 Robert Shiller, a Nobel Laureate, introduced *Narrative Economics*, an approach to explaining aspects of economies that are difficult to comprehend when analyzed using conventional methods: in light of narratives (i.e., stories) that participants in asset markets hear, believe, and tell each other, some observable economic factors, such as price dynamics of otherwise valueless digital assets, can be explained largely within the context of those narratives. As Shiller argues, it is best to explain and understand seemingly irrational and hard-to-explain behaviors, such as investing in highly volatile cryptocurrency markets, in narrative terms: people invest because they believe that it makes sense to do so, or have a heartfelt opinion about the prospects of the asset, and they share these beliefs and opinions with themselves and others in the form of narratives. In this paper, we address the question of how an agent-based modeling platform can be developed to be used for studying narrative economics. To do this, we integrate two very recently published developments. From the field of agent-based models of financial markets, we use the *PRDE* adaptive zero-intelligence trader strategy introduced by Cliff (2022), and we extend it to integrate a continuous-time real-valued nonlinear opinion dynamics model reported by Bizyaeva *et al.* (2022). In our integrated system, each trader holds an *opinion* variable whose value can be altered by interaction with other agents, modeling the influence that narratives have on an agent's opinions, and which can also be altered by observation of events in the market. Furthermore, the PRDE algorithm is modified to allow each trader's trading behavior to smoothly alter as that trader's opinion dynamically varies. Results reported for the first time here show that in our model there is a tightly coupled circular interplay between opinions and prices: changes in the distribution of opinions can affect subsequent price dynamics; and changes in price dynamics can affect the consequent distribution of opinions. Thus this paper presents a first demonstration of the reliability and effectiveness of our new agent-based modeling platform for use in studying issues in narrative economics. Python source-code for our model is being made freely available as open-source release on GitHub, to allow other researchers to replicate and extend our work.

1 INTRODUCTION

The notion of *narrative economics*, introduced and popularized by Nobel Prize winner Robert Shiller (Shiller, 2017; Shiller, 2019), involves the study of the spread and dynamics of "narratives", i.e. the stories (particularly those of human interest and emotions) told and believed by market participants about the nature of the market, and how these change over time, to understand the dynamics of economic and market systems. In particular, Shiller argues that phenomena which are difficult to explain in any other way, as such cryptocurrency price fluctuations, can best be explained by reference to the narratives in play among participants in the relevant markets.

In their 2021 prize-winning ICAART paper, Lomas & Cliff (Lomas and Cliff, 2021), noted that a narrative is simply an expression of an opinion, and used this to forge an initial link between agent-based models of financial markets, and the long-established field of research in the social sciences literature known as *Opinion Dynamics* (OD). Research in OD aims to understand how individuals in a population can influence others to change their opinions, and how they

[a] https://orcid.org/0000-0003-2987-4601
[b] https://orcid.org/0000-0003-3822-9364

Bokhari, A. and Cliff, D.
Studying Narrative Economics by Adding Continuous-Time Opinion Dynamics to an Agent-Based Model of Co-Evolutionary Adaptive Financial Markets.
DOI: 10.5220/0011797000003393
In *Proceedings of the 15th International Conference on Agents and Artificial Intelligence (ICAART 2023) - Volume 2*, pages 355-367
ISBN: 978-989-758-623-1; ISSN: 2184-433X

may change their own opinions, as a result of those influences. Lomas & Cliff presented a novel approach to integrating mathematically naive 20-year-old OD models with agent-based computational economics. In the Lomas & Cliff model, traders buy and sell assets by submitting bid and/or ask orders to an automated central financial exchange (such as a real-life stock exchange or a cryptocurrency exchange website), and showed how changes in the trader-agent's opinions could be reflected in the price dynamics of the market. While notable for its novelty, Lomas & Cliff's work could be improved upon in two major ways: first, their method for integrating OD with trader-agents has since been shown to have a vulnerability that can result in the market collectively grinding to an irreversible halt; and second they did not discuss or explore any mechanisms by which market events could alter the traders' opinions – in their model, causality was one-directional from opinions to prices. In this paper, we present our model which addresses both those issues and also relies on a less mathematically naive OD model, which enables further work in exploring the stability of the dynamics of our model, as is discussed further later in this paper.

Although the OD literature is substantial, the vast majority of it discusses models where agents hold abstract opinions that are altered only by interactions with other agents and which do not contain external references to establish whether certain opinions are true or false (for example, opinions may be held about politics or religion, on which people hold strong personal views and firmly believe that their opinions are correct, but for which there is no one objectively correct opinion). In contrast to much of the existing OD literature, Lomas & Cliff's work was notable because the opinions of their trader-agents regarding the current price of an asset could later prove to be true or false. For example, a trader may believe that the price of an asset would decrease in the near term, but in fact it increases. Therefore, in agent-based models of narrative economics traders need to be able to adjust their opinions based not only on *local* interactions with other agents but also on uncertain future *globally* observable outcomes that have not yet been resolved or revealed: this is unusual in the field of OD, with almost no prior papers addressing this issue.

As with OD research, research in agent-based modelling of economic and financial systems has been underway for several decades, and it is a sufficiently mature field that it has its own name: *Agent-Based Computational Economics* (ACE: see e.g., (Chen, 2018; Hommes and LeBaron, 2018)). In ACE models, a simulated economic or financial system is created and populated with some number of autonomous agents, where each agent follows some specific algorithm or strategy for interacting with other agents in the system. It has been repeatedly demonstrated that ACE models employing only minimally simple trader strategies can be surprisingly informative, and these so-called *zero intelligence* (ZI) trading strategies are now commonly used to illustrate and explore various economic and market phenomena. The seminal paper in establishing markets populated by ZI traders as worthwhile objects of study is (Gode and Sunder, 1993); and for reviews of the impact and effectiveness of ZI studies in ACE, see (Farmer et al., 2005; Ladley, 2012; Axtell and Farmer, 2022).

Very recently, (Cliff, 2022) proposed the *PRDE* (**P**arameterized-**R**esponse zero-intelligence with **D**ifferential **E**volution) as an instance of an efficiently adaptive ZI trading strategy, i.e. a minimal-intelligence strategy that can usefully alter its trading behavior to better fit its immediate market conditions, with the aim of maximizing its individual profitability at all times. Cliff's PRDE is an extension of his earlier *PRZI* (**P**arameterized-**R**esponse **Z**ero-**I**ntelligence) trading strategy (Cliff, 2021) which was introduced to address a problem identified in the original Lomas & Cliff work: the modified ZI trading strategies reported by Lomas & Cliff did integrate the trader's opinions into their trading behavior, but could result in situations where the opinionated traders all settled to a state in which no further transactions could take place in the market, so the market ceased to show any activity at all: for full details of this, see the discussion in (Cliff, 2021). PRZI and then PRDE were introduced as remedies to this problem in Lomas & Cliff's 2021 paper, but to the best of our knowledge our paper here is the first to integrate PRDE with a leading-edge OD model.

While the Lomas & Cliff paper is definitely novel, it is somewhat disappointing that the OD models they explored adding to existing ZI trader strategies were old, extremely simple, and mathematically naive. Specifically, Lomas & Cliff explored the integration of three OD models to two ZI strategies. The two ZI strategies were the *Zero-Intelligence Constrained* (ZIC) of (Gode and Sunder, 1993) and the *Near-Zero Intelligence* (NZI) of (Duffy and Ünver, 2006): when Lomas & Cliff added OD models to these two strategies, they named the extended strategies OZIC and ONZI, respectively. The three OD models were the *Bounded Confidence* (BC) model of (Hegselmann and Krause, 2002), the *Relative Agreement* (RA) model described in (Deffuant et al., 2002; Meadows and Cliff, 2012), and the *Relative Disagreement* (RD) model of (Meadows and Cliff, 2013), the most recent of which is over a decade old. All three of BC, RA, and RD are simple abstract discrete-time models where time advances according to $t \leftarrow t+1$ with no clear links from the integer-valued t to continuous real-world time – that is, in these simple old OD models there is no use of differential equations with stated time-constants, which would in principle allow the whole arsenal of continuous-time dynamical systems modelling and analysis tools and techniques to be deployed, such as Lyapunov stability analysis To address this issue, in our paper here we switch away from the simplistic discrete-time OD models used by Lomas & Cliff, and use instead a contemporary continuous-time nonlinear differential-equation based OD model recently reported by Bizyaeva, Franci, & Leonard (2022) which we refer to hereafter as the *BFL* model.

The novel results in our paper concern the overall dynamics of the integration of the BFL OD model with PRDE traders. For consistency with the naming convention introduced by Lomas & Cliff, we refer to opinionated-PRDE traders as OPRDE, and then to distinguish our use of BFL from any later attempts to create opinionated PRDE traders using different OD models, we refer to our model as OPRDE-BFL.

We show here that our ACE-style market simulations populated by OPRDE-BFL traders can exhibit situations in which changes in market opinion alter subsequent price dynamics (as was demonstrated by Lomas & Cliff with OZIC and ONZI) and also that changes in prices can plausibly alter the subsequent distribution of opinions in the population of traders (as was *not* demonstrated by Lomas & Cliff).

In Section 2 we give further details of the background to our work, concentrating on explaining the PRDE trader strategy and the BFL opinion dynamics model. In Section 3 we illustrates how the BFL model can be customized for PRDE integration. In Section 3.3 we explain how BFL is integrated into PRDE traders. In Section 4 we describe the details of our experiments and discuss the results. In Section 5 we discuss future directions and the implications of our findings.

2 BACKGROUND

2.1 Simulated Financial Markets

The study of simulation models of financial markets has been known for a long time as a means of exploring the fine-grained dynamics of various form of market. Such models very often require populating a market mechanism with a number of trader-agents: autonomous entities that work independently within the framework of the particular market mechanism being simulated. As per the market simulation literature reviewed here, usually each trader in the market is assigned a role, either that of a buyer or a seller, and there is a systematic way in which each buyer (seller) is assigned an order to buy (sell) a certain quantity of an arbitrary abstract commodity in which the market trades, as well as a private (secret to the trader) *limit price*, which is the maximum (minimum) unit-price at which they can buy (sell) the item. The difference between the transaction price and the trader's limit price is the *profit* (sometimes referred to as *utility* or *surplus*). Over the 30 years or so that ACE has been the topic of active development, a few specific trader-agent algorithms are notable for their longevity in the literature: SNPR (Rust et al., 1992); ZIC (Gode and Sunder, 1993); ZIP (Cliff, 1997); GD (Gjerstad and Dickhaut, 1998); MGD (Tesauro and Das, 2001); GDX (Tesauro and Bredin, 2002); HBL (Gjerstad et al., 2003); and AA(Vytelingum et al., 2008). The seminal ZIC developed by Gode and Sunder (1993) is highly stochastic however it shows surprisingly human-like market dynamics. In a landmark paper by IBM researchers (Das et al., 2001), GD and ZIP were the first to be demonstrated to consistently show superior performance to human traders (see also: (De Luca and Cliff, 2011b; De Luca and Cliff, 2011a; De Luca et al., 2011)), and the IBM result is widely cited as initiating the rise of algorithmic trading in real-world financial markets. All of the these strategies, except for SNPR and ZIC, use some form of machine learning (ML) or artificial intelligence (AI) to modify their responses over time, better-adapting their trading behavior to the market conditions they find themselves in, and the details of these algorithms have often been published in major AI/ML conferences and journals.

Recently (Cliff, 2021) proposed the PRZI trading algorithm: PRZI traders are firmly in the ZI tradition but they each have a scalar real-valued *strategy* parameter $s \in [-1,+1] \in \mathbb{R}$ which governs their response to market events: when a PRZI trader has $s=0$ its trading behavior is identical to the seminal ZIC of (Gode and Sunder, 1993), but when $s>0$ it becomes more "urgent", quoting prices that are more likely to find a counterparty and lead to a transaction, but for which the expected profitability of the transaction will be reduced relative to the prices quoted when $s=0$; and similary when $s<0$ the PRZI trader is more "relaxed" quoting prices that are more profitable if they do lead to a transaction, but less likely to lead to a transaction than the prices quoted when $s=0$.

At the extremes, when a PRZI trader i has $s_i = -1.0$ its trading strategy is equivalent to the maximally relaxed *Shaver* (abbreviated SHVR) strategy proposed in (Cliff, 2012; Cliff, 2018); and when $s_i = +1.0$ it is acting as the maximally urgent *Give-away* (GVWY) strategy, also described in (Cliff, 2012; Cliff, 2018).

As thus defined, an individual PRZI trader is non-adaptive: it is assigned an s-value at creation, and keeps that same s value for all of its lifetime. However, in subsequent work PRZI traders have been extended to adapt their s-values dynamically, as market conditions change, attempting to always increase or maintain their profitability as market circumstances alter. This allows populations of adaptive PRZI traders to be used as a tool for simulation modelling of contemporary real-world financial markets, in which all traders are simultaneously adapting their trading strategies, each seeking to maximize their own profitability, while burdened by the complexity and uncertainty of adapting to a market environment where every other trader is simultaneously adapting, continuously adjusting its strategy in real-time.

The most recent and currently most efficient adaptive PRZI trader was described in (Cliff, 2022): *PRzi with Differential Evolution* (PRDE), which makes use of Differential Evolution (DE: see e.g. (Storn and Price, 1997; Bilal et al., 2020)) as its optimization strategy. Each PRDE trader maintains a population of size $k \geq 4$ candidate s-values, and iterates over an infinite loop in which on each iteration it evaluates each of the k candidates in turn, and then uses a basic DE process to create a new candidate s-value, which is also evaluated: if that new candidate is better than one of the four in the original population, it replaces that one; if not, it is discarded; and the loop iterates again. Evaluation of any one candidate s-value involves the PRDE trader operating in the market using only that s-value for some period of time, and then calculating the profit-per-unit-time for that s-value as its "fitness" score in the DE process. For all our experiments reported here, we used the Python PRDE reference implementation published on Github as part of the freely-available *BSE* platform for agent-based modelling of contemporary financial markets, available as (Cliff, 2012).

In all of our experiments reported here we used BSE in its default configuration, where it allows for the definition of some number of buyer-agents N_B, and some number of seller-agents N_S. Each buyer (seller) is periodically issued with assignments to buy (sell) a unit of the exchange's tradeable asset at a price no higher (lower) than that trader's given private limit price, and to find a willing seller (buyer) as a trading counte-party via interacting within a continuous double auction (CDA) running on a centralised financial exchange that operates a Limit Order Book (LOB), by submitting bid (ask) orders to the exchange. The CDA is an auction mechanism in which any buyer can submit a bid order to the exchange at any time, and any seller can submit a sell order at any time, and the exchange continuously runs a *matching engine* to pair up buyers and sellers whose orders are compatible – e.g. if a seller S1 quotes an ask of \$100 for a unit of the asset and a buyer B1 then quotes a bid price of more than \$100, the exchange matches B1 and S1 as counter-parties to a transaction and S1 then sells to B1 for \$100, because that was the earlier-quoted of the two prices. However, whenever a trader quotes a price that cannot be matched with a counter-party, that quote "rests" at the exchange and is entered on the LOB; with the exchange publishing an updated LOB to all market participants every time the LOB changes. The published LOB shows a summary of the array of all unmatched buyer orders resting at the exchange (on the *bid-side* of the LOB) and all unmatched seller orders (on the *ask-side*) of the LOB, with the two sides of the LOB sorted in price-order from best (highest bid, lowest ask) to worst (lowest bid, highest ask). All major financial exchanges around the world for stocks/shares, currencies, commodities, and digital assets run LOB-based CDAs, so in this respect BSE is an excellent model of real-world exchanges. For further details of CDAs and LOBs, see e.g. (Friedman and Rust, 1992; Gould et al., 2013; Abergel et al., 2016). Limit prices on the trader's assignments were drawn from a pair of supply and demand curves that we specified, allowing us to control the equilibrium price and quantity in each market session.

3 OPRDE: OPINIONATED PRDE

In this section we describe our rationale for the specific set of changes we introduced to give *Opinionated PRDE* (OPRDE) traders, which fully enable the study of narrative economics via agent-based models.

3.1 Opinions About What?

First, we need to be specific about what our trader-agents hold opinions *about*. In the first instance, we are working with agents whose opinions at time t concern the qualitative shift in price $P(t)$ of a tradeable asset at a particular future point in time ω, denoted by T_ω, relative to the price at some earlier point in time T_0 when the opinion was first formed. Each trader i's individual opinion at time t is a real number $x_i(t) \in [-1.0, +1.0] \in \mathbb{R}$, which we interpret *qualitatively* as mapping onto one of two opinion-states such that $x_i(t) \in \{-1,+1\}$; where $x_i(t) = +1$ implies that i's opinion is that the price will significantly *rise*; and correspondingly that $x_i(t) = -1$ implies that i's opinion is that the price will significantly *fall*. Here the notion of "significant" rises or falls in price is defined in terms of whether the absolute percentage change in price is above some threshold percentage Δ_p; that is, a price change is significant if:

$$\left|\frac{P(T_\omega)-P(T_0)}{P(T_0)}\right| > \Delta_p$$

This definition, with each trader i having its own individual values for $T_{0,i}$, $T_{\omega,i}$, and $\Delta_{p,i}$, allows the model be able to capture situations in which different traders form their opinions at different times (i.e., the $T_{0,i}$ values may vary with i), have differing time-horizons over which they judge the rise/fall/unchanged nature of any difference in price (i.e., the $T_{\omega,i}$ values may vary with i), and have different views of what counts as a significant change in price (i.e., the $\Delta_{p,i}$ values may vary with i, too). However, for simplicity, in the experiments reported here we study only situations in which all traders have the same values of these three variables: that is, in our experiments here, all traders form their opinions at the same time, when the market session opens; all traders assess the price-change at the same time, when the market session closes; and all traders have the same definition of "significant", here $\Delta_{p,i} = 0.2, \forall i$. This means that our traders will exhibit opinion dynamics both *within* any one market session, and also *across* a sequence of successive market sessions. In this paper we concentrate only on the within-session dynamics, studying only a single market session at a time. In later papers we will explore opinion dynamics in our system when it runs for continuous sequences of multiple successive market sessions.

3.2 The BFL Nonlinear OD Model

Here we present our condensed summary of (Bizyaeva et al., 2022): for full details, readers are of course referred to the original paper.

Consider a network of $N_T = N_B + N_S$ trading agents each forming a scalar opinion $x_1,...,x_{N_T} \in \mathbb{R}$ about the price of a tradeable asset, and $M \geq 1$ communication sources (such as mass media), offering static opinions $c_1,...,c_M \in \mathbb{R}$ about the same asset. Consider x_i the real-valued opinion variable of trader i, where negative values indicate a belief that prices are about to decline, while positive values indicate an increase, and the same is true for the value of the opinions from the communication sources, for instance c_k denotes the opinion from communication source k regarding the price of the asset. Let $X = (x_1,...,x_{N_T})$ be the *opinion state* of the agent network, a special case being the neutral state at the origin $X = 0$ which would be the case if all agents held no firm opinion about the directionality of any near-term price movements. Agent i is unopinionated if its opinion state is small, i.e., $\|x_i\| \leqslant \vartheta$, for a fixed threshold $\vartheta \simeq 0$; agent i is opinionated if $\|x_i\| \geq \vartheta$. Agents can agree and disagree: when two agents have the same qualitative opinion state (e.g., they both favor the same option), they agree. When they have qualitatively different opinions, they disagree.

In the BFL model proposed by (Bizyaeva et al., 2022), the opinions of agents are assumed to evolve nonlinearly in continuous time, and we have tailored the BFL model to best suit it to the financial market context in which we are using it. Our application consists of the following parameters and property specifications:

• **Saturated Opinion Exchanges.** Almost every signaling network, whether natural or artificial, exhibits saturated nonlinearities due to the limits of action and sensing. In the case of an agent opinion dynamics network, the saturation of interactions between agents efficiently bounds the attraction between their opinions, thus, overcoming the paradox of linear weighted average models. (Bizyaeva et al., 2022).

• **Local Opinion.** A trader's local opinion is obtained as a linear weighted sum of the opinions expressed by all the other traders in the network. The weighted average is based on an adjacency matrix denoted here by $A = [a_{ij}] \in \mathbb{R}^{N_T} \times \mathbb{R}^{N_T}$. It is important to note that there are two types of network interactions: self-reinforcing interactions, weighted by a_{ii}; and neighbour interactions, weighted by a_{ij}, where the sign of the adjacency weight determines the sign of the network interaction, for all existing links between traders in the traders' social network satisfying $a_{ij} \neq 0$, with a_{ij} weighing the extent to which agents are influenced by their neighbors.

• **Opinions From Communication Sources.** The impact of communication sources, such as the mass media, on a trader's opinion, is captured by a weighted linear sum of the static opinions of these sources. This sum is weighted by the weight matrix $B = [b_{ik}] \in \mathbb{R}^{N_T} \times \mathbb{R}^M$ which identifies to what extent each trader can be influenced by a given communication source.

• **Attention.** It is important to note that traders' attention or susceptibility to exchange opinions may vary. In our model, each trader i's attention parameter u_i is linked to i's profit, meaning that when a trader loses profitability, its attention increases, otherwise it remains the same; thus if the agent loses it will be more alert to its neighbors and communication sources. Let π_t represents the trader's profit at time t and π_{t-1} is the trader's profit at a previous period $t-1$, then the trader's attention is updated as follows:

$$u_{t+1}\begin{cases} u_t+\xi & \pi_t-\pi_{t-1}<0 \\ u_t & \pi_t-\pi_{t-1}\geq 0 \end{cases} \tag{1}$$

where ξ is an incremental value that gets increased as the market session approaches the end.

• **Resistance.** A trader's resistance parameter d_i represents its willingness to change its opinion.

• **Inputs.** Our agents have exogenous inputs: two input parameters g_i and e_i are introduced for each agent. The first input parameter g_i represents the collective opinion of the market, as summarised by public data shown on a central exchange's LOB, such as the market's current *mid-price* or *micro-price* (both of which are defined below), which is analogous to the *global opinion* in the work of (Guzelyte and Cliff, 2022). The second input parameter e_i represents some indication of the likelihood of a specific opinion proving to be true at its stated T_x – what (Guzelyte and Cliff, 2022) referred to as the "event opinion". In principle, supply and demand data can be used to predict whether prices are likely to rise or fall in the short term. In the event that there is an excess of supply relative to demand, the price of the asset will likely fall; and *vice versa*. A simple way for ZI traders to estimate supply/demand imbalance, as introduced in (Church and Cliff, 2019), is by taking the difference between the current market *mid-price*, denoted here by $p_m(t)=(p^*_{bid}(t)+q^*_{ask}(t))/2$, and the current market *micro-price* (as defined by (Cartea et al., 2015)), denoted here by $p_\mu(t)$, where

$$p_\mu(t)=\frac{p^*_{ask}(t)q^*_{bid}(t)+p^*_{bid}(t)q^*_{ask}(t)}{q^*_{bid}(t)+q^*_{ask}(t)} \tag{2}$$

in which $p^*_{ask}(t)$ is the price of the best ask at time t (i.e., the price at the top of the ask side of the exchange's LOB); $p^*_{bid}(t)$ is the price of the best bid at time t (i.e. the price at the top of the bid side of the LOB); $q^*_{bid}(t)$ is the total quantity available at $p^*_{bid}(t)$; and $q^*_{ask}(t)$ is the total quantity available at $p^*_{ask}(t)$. In the case of zero supply/demand imbalance at the top of the LOB (i.e., Equation 2 reduces to the equation for the market midprice, and therefore the difference between the two prices, denoted by $\Delta_m(t)=p_\mu(t)-p_m(t)$ is zero). Accordingly, a positive imbalance (i.e. $\Delta_m(t)>>0$) indicates that the subsequent transaction prices are likely to increase, and a negative imbalance (i.e. $\Delta_m(t)<<0$) indicates that subsequent transaction prices are likely to fall.

As per (Bizyaeva et al., 2022), the networked opinion dynamics $\dot{x_i}$ are then summarised as:

$$\frac{dx_i}{dt}=-d_ix_i+u_i\left(S_1\left(s_A\right)+S_2\left(s_B\right)\right)+\left(g_i+e_i\right) \tag{3}$$

where $s_A=\sum_{j=1}^{N_T}A_{ij}x_j$ and $s_B=\sum_{k=1}^{M}B_{ik}c_k$.

The evolution of agent i's opinion is determined by four terms. These terms are the linear damping term, the saturated network interaction term, the saturated communication sources influence term, and the environmental signal term. $S_1,S_2:\mathbb{R}\rightarrow\mathbb{R}$ are bounded saturation functions satisfying $S(0)=0,\ S(0)'=1,\ S(0)''=0,\ S(0)'''\neq 0$ with an odd symmetry $S(-y)=-S(y)$; S_1 saturates the network interactions, S_2 saturates the communication sources influence; the damping coefficient $d_i>0$ represents a trader's reluctance to form a strong belief since it drives the values $\dot{x_i}$ to the neutral point, which implies that higher d_i indicates that the trader will be more resistant to forming an opinion, however, the parameter $u_i\geq 0$ indicates how attentive agents are to their social interactions, which affects the degree to which their opinions are socially influenced. Environment inputs e_i and g_i represent "opinions" (or, more accurately, opinion-influencing factors) derived from the environment, which are independent of the other agents' opinions.

3.3 Adding BFL to PRDE

In this study, we provide PRDE traders with a real-valued opinion variable, which means that opinionated PRDE buyers will behave differently from opinionated PRDE sellers in the same market environment. For instance, when opinions indicate that prices are on the rise, OPRDE buyers will respond with a heightened sense of urgency as a hybrid of GVWY and ZIC. As a result, they will quote prices based on strategy values influenced by their opinions; conversely, OPRDE sellers will act more relaxed and quote prices based on strategy values that are a hybrid of ZIC and SHVR. On the other hand, when opinions indicate that asset prices are falling, the same reasoning applies *mutatis mutandis*.

By this point, it should be obvious that we require some function that maps from trader i's opinion variable x_i to its PRDE strategy s_i, i.e. $s_i=F_i(x_i)$. Considering the simplest case, since both s_i and $x_i\in[-1,+1]\in\mathbb{R}$, it is possible for the mapping to be the identity function, or its opposite, depending upon whether i is a buyer or a seller. A buyer's simplest F_i is identity function: $F_i(+1)=+1,\ F_i(-1)=-1$, whereas a seller's simplest F_i is negative identity function: $F_i(+1)=-1,\ F_i(-1)=+1$. It should be noted, however, that this fairly rapidly shifts the trader's strategy to the extremes (either SHVR or GVWY) as $|x|\rightarrow 1$, therefore, this may not always be the most effective approach: because at the extremes, buyers usually lose. Consider, for example, a buyer playing SHVR. It will be moving away from it minimum price to it limit price, which means it ends up quoting high quote prices. In contrast, a seller playing SHVR moves from it maximum price while still making high quote prices, and generates more profit as a result. Additionally, when a buyer plays GVWY, it loses since it limit price is it maximum price. At the extremes, a buyer only has a chance of winning when the seller plays GVWY. It is due to the fact that the seller will maintain it limit price as it minimum price. This results in the deal being in favor of the buyer, even if the buyer is bidding higher than it limit price. Furthermore, it is important to note that if the strategies are always at -1, 0, and +1, the PRZI family of traders will be unable to utilize the full potential of the infinite number of strategies that PRZI offers. Thus, it is clear that we need a nonlinear mapping from an opinion to a strategy that will not push the strategies to the extreme edges $\{-1,+1\}$. In such a manner that when prices fall, a buyer acts as a hybrid with a high probability of quoting low prices; and when prices rise, however, the seller acts as a hybrid, with a high probability of quoting high prices. We have found that the sine function usefully maps a trader's opinion to its strategy: it has the virtue of simplicity, and the fact that it provides values for a strategy that are close to those for the identity function. This is necessary to determine the appropriate quote prices. Nevertheless, when the sine function has a value of +1 or 1, we clip it to keep it below those extreme values.

Recall that each PRDE trader maintains a private local population of potential strategy-values, of population-size $k \geq 4$, which for trader i can be represented by $s_{i,1}, s_{i,2}, ..., s_{i,k}$. Since PRDE traders use just a single real scalar value to specify their bargaining behavior, every individual in trader i's local DE population is just a single value. Consequently, the conventional DE concept of crossover (i.e., selecting alleles from two parents, one allele per dimension of the genomes) is not applicable here: PRDE constructs a genome entirely by operating on the base vector. In its current configuration, PRDE applies the basic "vanilla" DE/rand/1 where, after evaluating a particular strategy $s_{i,x}$, three other s-values are randomly selected from the population: $s_{i,a}, s_{i,b}$, and $s_{i,c}$ where $x \neq a \neq b \neq c$, and therefore a new candidate strategy $s_{i,y}$ is created s.t. $s_{i,y} = max(min(s_{i,a} + F_i(s_{i,b} - s_{i,c}, +1), -1)$ where F_i represents the trader's differential weight coefficient (in the experiments reported here, $F_i = 0.8; \forall i$), with the nested min and max functions keeping the candidate strategy value within the range [-1.0, +1.0]. In OPRDE we introduce the trader's opinion $s_{i,o}$ as a new candidate strategy. Then the fitness of $s_{i,y}$ and $s_{i,o}$ are evaluated and the best strategy replaces $s_{i,x}$, otherwise, it is discarded; and then the next strategy $s_{i,x+1}$ is evaluated.

3.4 Implementation Details

The experiments reported in Secction 4 use BSE to simulate a financial market for a single abstract tradeable commodity, where $N_B = N_S = 30$ and hence $N_T = 60$; and where each trader is running OPRDE with $k = 5$. Throughout the process, buyers are given a maximum purchase price of \$140 per unit, while sellers are restricted to a minimum sale price of \$60 per unit. A schedule such as this provides what economists refer to as perfect elasticity of supply and demand, and it is widely used in experimental economics research (such as (Smith, 1965)): this ensures that every seller can find a buyer that could act as a counterparty, and vice versa. In other words, no traders would be given extra-marginal prices that would limit their ability to find counterparties. As a result of a transaction between two traders, both their cash assignments and their stock assignments are depleted to render them inactive, and they each wait for a random period of up to five seconds before they are re-assigned fresh cash or stock, enabling them to rejoin the market as active traders once again.

BSE simulates continuous time using a discrete time-slicing approach using a temporal step-size of $\Delta t = 1/N_T$, i.e. 0.0167sec for $N_T = 60$, as a result, each trader can interact with the market at least once a second. Our experiments here simulate 300 days of continuous round-the-clock 24×7 tradings: the trading takes place on sub-second timescales, but the co-evolutionary dynamics play out over much longer periods. The profit per second (PPS) of a strategy s at time t is calculated by summing all profits generated over the time period $[t - \Delta_E, t]$ and dividing that accumulated profit by the evaluation period Δ_E: in our experiments here we use Δ_E=7200sec (i.e., 2 hours); given there are $k = 5$ strategies, any one trader takes 10 simulated hours to evaluate all its candidate strategies. As stated earlier, every transaction's profit is determined by the difference between the price agreed by the buyer and seller and their individual limit prices: e.g., if a transaction takes place at a price of \$90, then the buyer's profit on the transaction is \$50 (denoted by π_B) and the seller's profit is \$30 (denoted by π_S) because in these experiments all buyers had a limit price of \$140, and all sellers had a limit price of \$60.

For the BFL opinion dynamics model, for the illustrative outputs shown in this paper we used the following parameter values: $u_i = 0.5$; $d_i = 1, \forall i =$

$1, \ldots, N_T$; the adjacency matrix $A = [a_{ij}] = [1] \forall i \forall j$; and for these preliminary experiments the communication sources are switched off: exploring the interplay between the agents' opinions and the communicatins sources is something that we intend to explore in depth in a later paper.

4 RESULTS

This section evaluates OPRDE as an extension of PRDE by giving illustrative answers to the following three questions: (1) How does the OD of OPRDE affect market prices? (2) How can the distribution of prices affect the opinions of traders? And finally (3) How do markets populated wholly by OPRDE traders compare to those populated only by PRDE traders? [1]

4.1 Opinions Affecting Prices

Our first evaluation question is addressed by examining the results from extreme opinion distributions, such as those in which all traders hold extremely positive opinions, those in which all hold extremely negative opinions, and those in which the opinion distribution shifts over time.

Figure 1 illustrates how market prices change when traders' exogenously imposed opinions start out as negative for the first half of the period , and then shift to positive toward the end of the 25-day experiment. In both phases, transaction prices were highly correlated with opinions, demonstrating a causal link between traders' opinions and their price-quoting decisions. Figure 2 then shows the corresponding changes in traders' profit per second (PPS) values: in the first five days, with all traders holding negative opinions, sellers trade urgently and hence enter into transactions that yield less profit for them while the buyers trade with less urgency (i.e., are more relaxed) and are able to hold out for more profitable prices; and in the second period of five days the fortunes of the buyers and sellers reverse as the opinions do.

This causal effect of opinions on prices in our opinionated-agent-based model of a financial market functionally replicates what was first demonstrated by (Lomas and Cliff, 2021). In the next section, we demonstrate functionality in our system *beyond* what was demonstrated by Lomas & Cliff: that is, causality in the opposite direction, where changes in the distribution of prices in the market go on to affect the distribution of opinions among the population of traders.

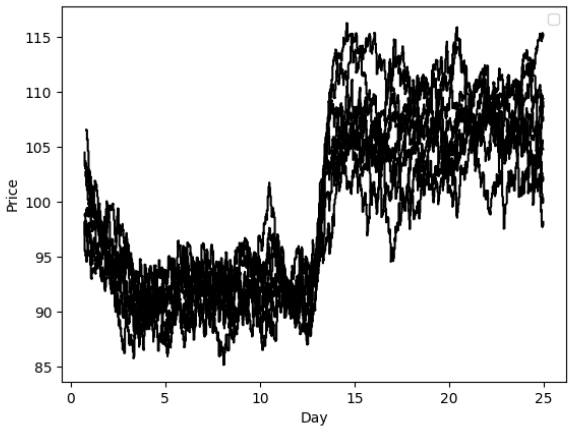

Figure 1: Plot of transaction prices from 10 IID experiments each involving all-OPRDE market over a 25-day period in which opinions initially start out as extremely negative for the first half of the period, and are then positive for the remaining of the period. See text for further discussion.

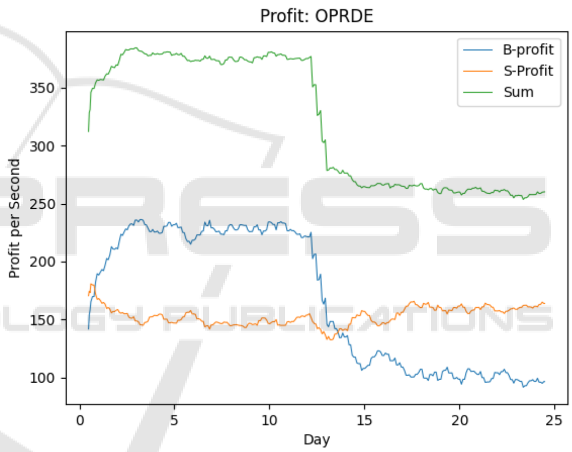

Figure 2: Profitability-per-second (PPS) plot for one of the 25-day experiments from Figure 1: the traces show aggregate PPS for all buyers; aggregate PPS for all sellers; and the sum of the two aggregates, i.e. total PPS extracted by the whole population of traders. See text for further discussion.

4.2 Prices Affecting Opinions

Among traders in financial markets, the phrase *market impact* is commonly used to refer to situations where the price of an asset quoted by potential counterparties to a large trade (where "large" means a trading volume sufficiently large to shift the supply or demand curve for that asset) moves in such a way that the trader trying to execute the large trade gets a worse price than was quoted on the exchange at the time that trader issued the order; and crucially this shift in

[1] For clarity, data from single experiments are presented here; for data from additional experiments which demonstrates that the results presented here are typical (rather than cherry-picked special cases), see (Bokhari and Cliff, 2023).

price occurs *before any transaction has actually taken place*. For example if a real-world stock-trader issues a bid-quote to buy a single share of IBM, the price that trader is quoted by potential sellers of IBM will be very close to whatever the current best ask price is on the exchange's LOB for IBM; but if the trader instead issues a bid-quote for one million IBM shares, this sudden revelation of excess demand for IBM stock means that potential sellers of IBM are all likely to alter the prices they quote, shifting upward, to reflect the rise in IBM's share-price that the sellers anticipate occurring as an immediate consequence of the newly increased demand for that stock. Similarly, the arrival of a large ask order will prompt potential buyers to instantly revise their bid-prices down – and, in both cases, the price change happens before any transaction takes place.

In Section 3, we introduced $\Delta_m(t)$ as a way of measuring supply/demand imbalance at the top of the LOB and so we add $\Delta_m(t)$ to the opinion dynamics model as the environment factor e_i in Equation 3. In consequence, OPRDE traders alter their behavior in response to anticipated changes in price – that is, they will be sensitive to market impact, because their opinions will be affected by the distribution of prices (and quantities) on the LOB. For example, when prices are likely to rise, OPRDE buyers should feel an increased sense of urgency and OPRDE sellers should feel a decreasing sense of urgency. In contrast, when prices are likely to fall, OPRDE sellers should increase their urgency, whereas OPRDE buyers should reduce their urgency.

To test this, here we show results from one long-term experiment in which the market is suddenly flooded with a large number of sell orders all priced at \$60 during the period from day 30 to day 60, after which the excess sell orders are abruptly removed: this step-change in excess supply in the market causes an imbalance at the top of the LOB, resulting in a negative value of $\Delta_m(t)$, which indicates an expected near-term decline in the asset price. Figure 3 shows buyers' strategies at this point clustering around -0.5, which means buyers are being relaxed and acting as a hybrid of ZIC and SHVR; on the other hand, Figure 4 shows sellers' strategies over the same period clustering around $+0.5$ i.e. acting urgently as a hybrid of ZIC and GVWY. As can be seen from Figure 5, the change in strategies is reflected in the quoted prices, leading to a substantial decline in transaction prices during the period of excess supply. The same effect is shown in the PPS time series shown in Figure 6: sellers have suffered significant losses, while buyers have made huge profits during the period of excess supply. This demonstrates that market dynamics have

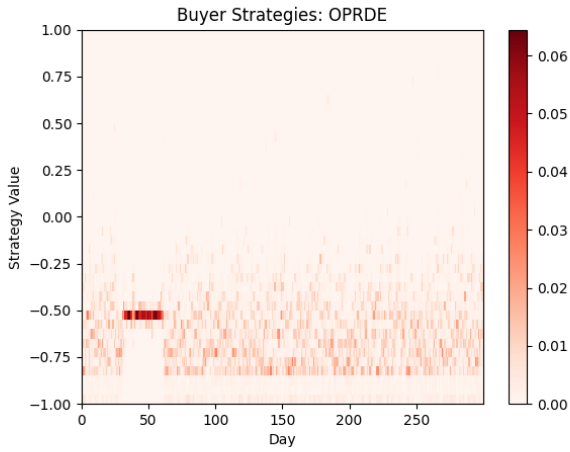

Figure 3: Heat-map of individual strategy-values for the population of 30 OPRDE buyers in a 300-day experiment where severe excess supply is suddenly introduced to the market on Day 30 and suddenly removed on Day 60. Horizontal axis is day-number; vertical axis is strategy; shading intensity shows proportion of traders in the population with that strategy (darker shading signifies more traders).

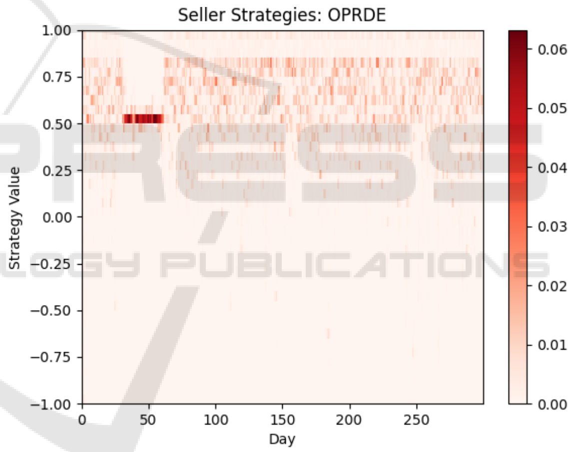

Figure 4: Heat-map of individual strategy-values for the population of 30 OPRDE sellers in the experiment for the same experiment as Figure 3.

an impact on opinion dynamics. Thus, we are able to affirmatively answer our second evaluation question.

A simplified version of OPRDE is presented here. However, Equation 2 has the weakness of being sensitive only to imbalances at the top of the LOB (the measure is not sensitive to imbalances at deeper levels of the LOB, thus being quite fragile). Multi-level order flow imbalance (MLOFI) is an alternative metric that can be used to measure imbalance as proposed by (Cont et al., 2021), which takes into account multiple levels of the LOB when determining LOB supply/demand imbalance.

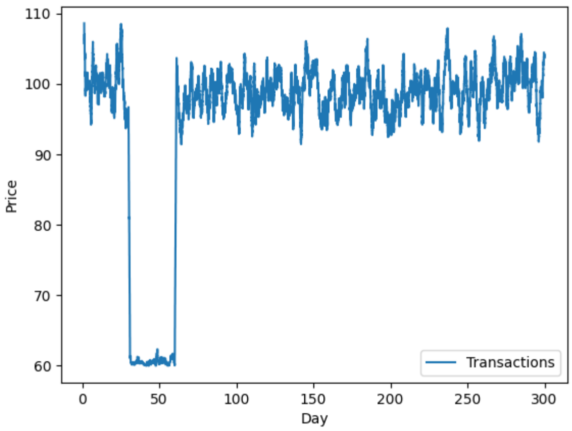

Figure 5: Transaction prices plot for the same experiment as Figure 3.

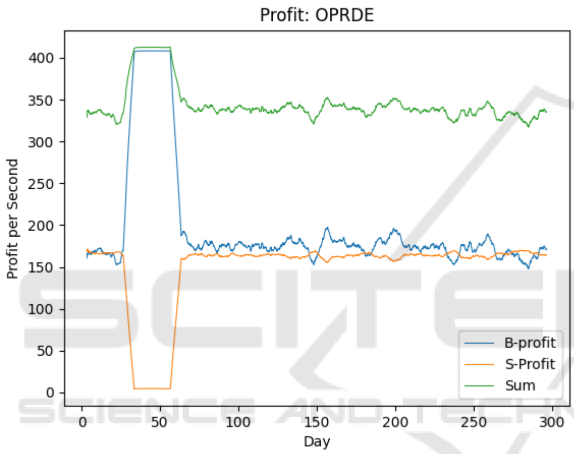

Figure 6: A profitability plot of one 300-day experiment populated exclusively with OPRDE traders (from the same experiment as Figure 5).

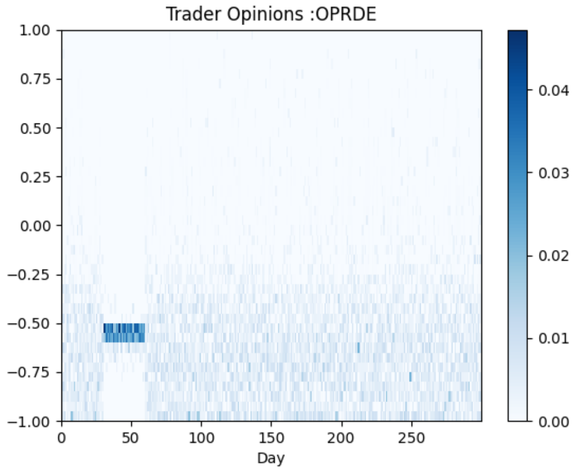

Figure 7: Heatmap demonstrating the impact of introduced excess supply on the distribution of opinions among traders (from the same experiment as Figure 5): vertical axis is opinion-value; darker shading significes more traders having that opinion.

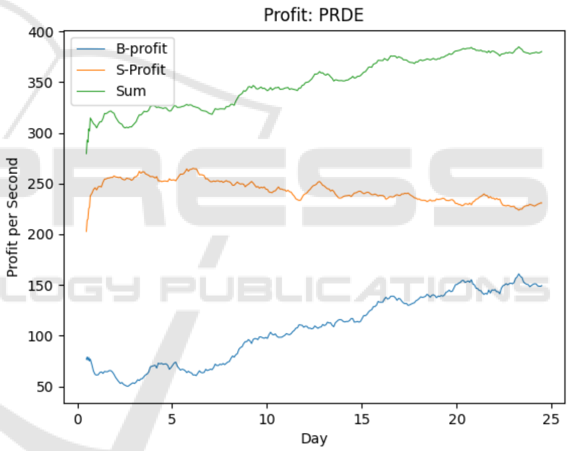

Figure 8: Plot of profitability data from one 25-day experiment in a market populated entirely by PRDE traders.

4.3 OPRDE Compared to PRDE

To illustrate the co-evolutionary dynamics under the BFL OD model at the level of individual traders' strategies, Figure 9 and 10 and Figure 12 and 13 show heatmaps illustrating the elite s-values of the 30 buyers and sellers in 25-day experiment of a market populated by PRDE and OPRDE, respectively, for which the corresponding PPS is shown in Figure 8 and 11. Both experiments assigned initial strategy values at (t=0) randomly from a uniform distribution over the range $[-1.0, +1.0]$. Upon visual inspection, it is clear that OPRDE traders are moving in a more diverse strategy space than PRDE traders.

Figure 11 illustrates a PPS plot for a typical experiment on a market populated by OPRDE traders. Although, the total profit extracted (i.e. $\pi_T(t)$)) by OPRDE traders is less than of PRDE traders Figure 8, it is clear that OPRDE sellers and buyers are both making profits interchangeably and no side is constantly getting ripped off as is the case with PRDE traders, as PRDE buyers are constantly losing, which may indicate that OPRDE traders are adapting faster than PRDE traders since they are opinionated.

5 DISCUSSION & CONCLUSIONS

The purpose of this work is to provide a platform for the experimental exploration of agent-based models of narrative economics: this paper is the first to report on our model; and, having introduced the rationale for the model and the details of its mechanisms, we now expect to follow this with other papers that ex-

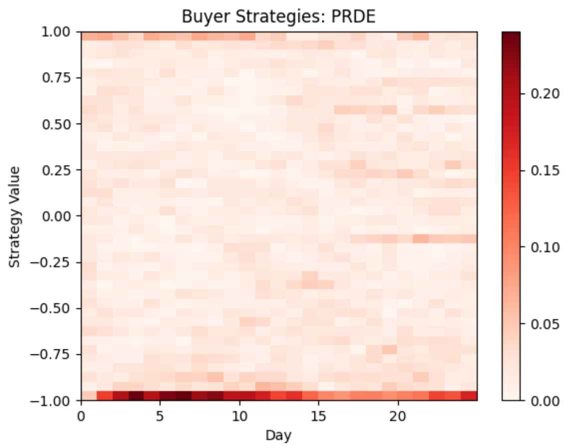

Figure 9: Heatmap of PRDE buyers strategy distribution from the same experiment as Figure 8.

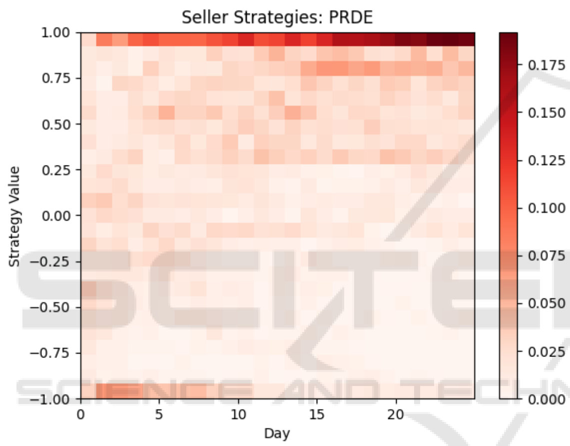

Figure 10: Heatmap of PRDE sellers strategy distribution from the same experiment as Figure 8.

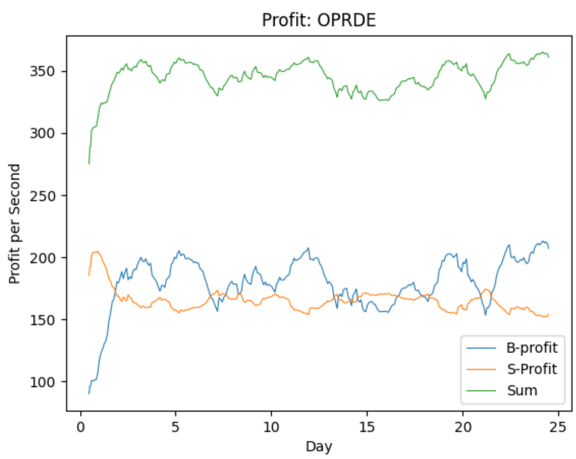

Figure 11: Plot of profitability data from one 25-day experiment in a market populated entirely by OPRDE traders.

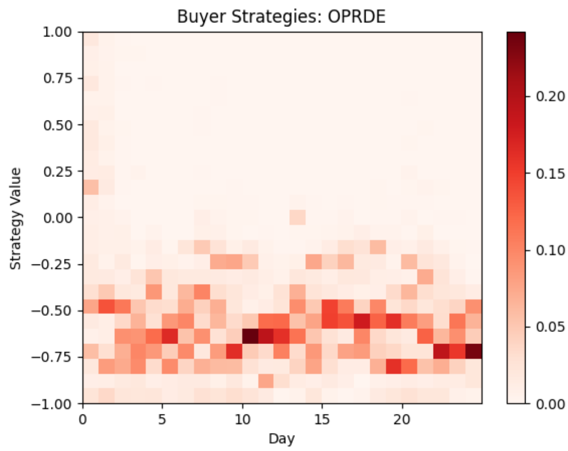

Figure 12: Heatmap of OPRDE buyers strategy distribution from the same experiment as Figure 11.

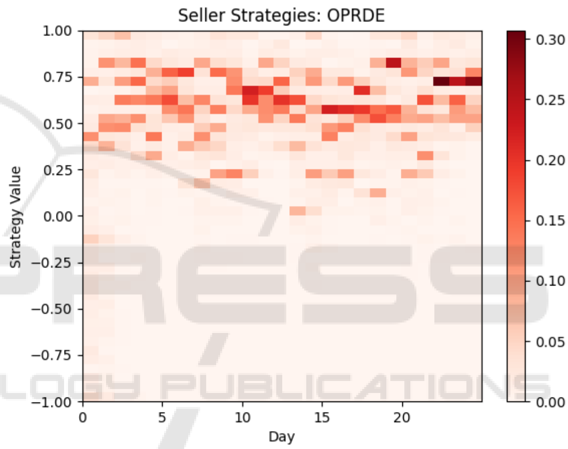

Figure 13: Heatmap of OPRDE sellers strategy distribution from the same experiment as Figure 11.

plore various aspects of the space of market systems in which narrative and opinions can move prices, and where prices can can alter opinions. Shiller argues, in his seminal proposal for work on narrative economics, that empirical research should be conducted by gathering data on the stories that individuals tell each other regarding economic matters, which influence their perceptions of future economic events and are themselves significant factors in economic dynamics. The work we have described here offers an alternative experimental approach that complements that proposed by Shiller: experimentalists now have access to agent-based simulations using our platform, allowing them to gain a deeper understanding of how opinions, expressions of those opinions, and subsequent economic outcomes interact dynamically. To aid other researchers in replicating and extending our results, we will make the source-code for our system freely available as an open-source repository on GitHub:[2] we look forward to seeing what uses other people put our system to.

REFERENCES

Abergel, F., Anane, M., Chakraboti, A., Jedidi, A., and Toke, I. (2016). *Limit Order Books*. Cambridge University Press.

Axtell, R. L. and Farmer, J. D. (2022). Agent-based modeling in economics and finance: past, present, and future. *Journal of Economic Literature*.

Bilal, Pant, M., Zaheer, H., Garcia-Hernandez, L., and Abraham, A. (2020). Differential evolution: A review of more than two decades of research. *Engineering Applications of Artificial Intelligence*, 90:103479.

Bizyaeva, A., Franci, A., and Leonard, N. E. (2022). Nonlinear opinion dynamics with tunable sensitivity. *IEEE Transactions on Automatic Control*.

Bokhari, A. and Cliff, D. (2023). Studying narrative economics by adding continuous-time opinion dynamics to an agent-based model of co-evolutionary adaptive financial markets. *SSRN*, 4316574.

Cartea, Á., Jaimungal, S., and Penalva, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.

Chen, S. H. (2018). *Agent-Based Computational Economics*. Routledge.

Church, G. and Cliff, D. (2019). A simulator for studying automated block trading on a coupled dark/lit financial exchange with reputation tracking. In *Proc. EMSS2019*, pages 284–293. CAL-TEK Srl.

Cliff, D. (1997). Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical Report HPL-97-91, HP Labs Technical Report.

Cliff, D. (2012). *Bristol Stock Exchange: open-source financial exchange simulator*. https://github.com/davecliff/BristolStockExchange.

Cliff, D. (2018). BSE: A minimal simulation of a limit-order-book stock exchange. *ArXiv:1809.06027*.

Cliff, D. (2021). Parameterised-response zero-intelligence traders. *arXiv preprint arXiv:2103.11341*.

Cliff, D. (2022). Metapopulation differential co-evolution of trading strategies in a model financial market. *Proc. IEEE SSCI Symposium on Differential Evolution*.

Cont, R., Cucuringu, M., and Zhang, C. (2021). Price impact of order flow imbalance: Multi-level, cross-sectional and forecasting. *ArXiv:2112.13213*.

Das, R., Hanson, J. E., Kephart, J. O., and Tesauro, G. (2001). Agent-human interactions in the CDA. In *Proc. IJCAI2001*, pages 1169–1178.

De Luca, M. and Cliff, D. (2011a). Agent-human interactions in the CDA, redux. *Proceedings ICAART-2011*.

De Luca, M. and Cliff, D. (2011b). Human-agent auction interactions: Adaptive-aggressive agents dominate. In *Proceedings IJCAI 2011*.

[2]https://github.com/NarrativeEconomics/OD

De Luca, M., Szostek, C., Cartlidge, J., and Cliff, D. (2011). Studies of interaction between human traders and algorithmic trading systems. Technical report, UK Government Office for Science, London.

Deffuant, G., Neau, D., and Amblard, F. (2002). How can extremism prevail? A study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation*, 5(4):1.

Duffy, J. and Ünver, M. U. (2006). Asset Price Bubbles and Crashes with Near-Zero-Intelligence Traders. *Economic Theory*, 27:537–563.

Farmer, J. D., Patelli, P., and Zovko, I. (2005). The Predictive Power of Zero Intelligence in Financial Markets. *PNAS*, 102(6):2254–2259.

Friedman, D. and Rust, J., editors (1992). *The Double Auction Market*. Addison-Wesley.

Gjerstad, S. and Dickhaut, J. (1998). Price formation in double auctions. *Games & Econ. Behav.*, 22(1):1–29.

Gjerstad, S. et al. (2003). The impact of pace in double auction bargaining. *University of Arizona, Department of Economics Working Paper*, pages 03–03.

Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.

Gould, M., Porter., M., Williams, S., McDonald, M., Fenn, D., and Howison, S. (2013). Limit order books. *Quantitative Finance*, 13(11):1709–1742.

Guzelyte, R. and Cliff, D. (2022). Narrative economics of the racetrack: An agent-based model of opinion dynamics in in-play betting on a sports betting exchange. In Rocha, A.-P., Steels, L., and van den Herik, J., editors, *Proceedings ICAART2022*, pages 225–236.

Hegselmann, G. and Krause, U. (2002). Opinion dynamics and bounded confidence: models, analysis and simulation. *J. Art. Societies & Social Simulation*, 5(3):2.

Hommes, C. and LeBaron, B., editors (2018). *Computational Economics: Heterogeneous Agent Modeling*. North-Holland.

Ladley, D. (2012). Zero Intelligence in Economics and Finance. *The Knowledge Engineering Review*, 27(2):273–286.

Lomas, K. and Cliff, D. (2021). Exploring narrative economics: An agent-based-modeling platform that integrates automated traders with opinion dynamics. In *Proceedings ICAART2021*.

Meadows, M. and Cliff, D. (2012). Reexamining the Relative Agreement Model of Opinion Dynamics. *Journal of Artificial Societies and Social Simulation*, 15(4):4.

Meadows, M. and Cliff, D. (2013). The Relative Disagreement Model of Opinion Dynamics: Where Do Extremists Come From? *7th International Workshop on Self-Organizing Systems (IWSOS)*, pages 66–77.

Rust, J., Miller, J., and Palmer, R. (1992). Behavior of trading automata in a computerized double auction market. In Friedman, D. and Rust, J., editors, *The Double Auction Market*, pages 155–198. Addison-Wesley.

Shiller, R. (2017). Narrative economics. *American economic review*, 107(4):967–1004.

Shiller, R. (2019). *Narrative Economics: How Stories Go Viral & Drive Major Economic Events*. Princeton.

Smith, V. L. (1965). Experimental auction markets and the walrasian hypothesis. *Journal of Political Economy*, 73(4):387–393.

Storn, R. and Price, K. (1997). Differential evolution: A simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11:341–359.

Tesauro, G. and Bredin, J. (2002). Sequential strategic bidding in auctions using dynamic programming. In *Proceedings AAMAS2002*.

Tesauro, G. and Das, R. (2001). High-performance bidding agents for the continuous double auction. In *Proc. 3rd ACM Conference on E-Commerce*, pages 206–209.

Vytelingum, P., Cliff, D., and Jennings, N. R. (2008). Strategic bidding in continuous double auctions. *Artificial Intelligence*, 172(14):1700–1729.