35th European Modeling & Simulation Symposium
20th International Multidisciplinary Modeling & Simulation Multiconference

2724-0029 © 2023 The Authors.
doi: 10.46354/i3m.2023.emss.043

Co-Evolution Causes Instability: Differential Evolution of ZIP Automated Traders in a Simulated Financial Market

Dave Cliff[1]

[1]Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

Email address: csdtc@bristol.ac.uk

Abstract

This paper reports results and analysis from simulation experiments in which a population of adaptive automated traders compete with one another and co-evolve their trading behaviors on a high-fidelity model of a contemporary electronic financial exchange. The automated traders are all using the adaptive Zero Intelligence Plus (ZIP) trading strategy, and the novel contribution of this paper is that here each ZIP trader has been extended to continuously use *Differential Evolution* (DE) to try to find the most profitable trading behavior in the current market conditions. The simulation experiments are structured in such a way that convergence to a stable steady-state is expected, but the results reveal that the continuous co-evolutionary interactions among traders give rise to unpredictable long-term instabilities in the traders' individual strategies. The instabilities seen in the new results presented here are qualitatively the same as those seen previously in simpler models of co-evolutionary markets involving less sophisticated adaptation mechanisms operating on less sophisticated trader-agents Thus, the results presented here add weight to the hypothesis that co-evolutionary markets are inherently unstable in strategy space, and hence that the long-term strategy instabilities seen in simpler simulation models of co-evolutionary markets are not mere artefacts of those models' simplifying assumptions. The Python source-code used in these experiments is being made freely available on GitHub, for other researchers to replicate and extend the results presented here.

Keywords: Financial Markets; Automated Trading; Differential Evolution; Algorithmic Trading; Financial Exchanges; Co-evolution.

1. Introduction

Over the past twenty years, there has been a significant transformation in major financial markets worldwide due to the widespread adoption of sophisticated automated trading systems, commonly referred to as "algorithmic traders" or "algos" or "robot traders". These automated systems have largely taken the place of human traders when it comes to actually executing trades. A simple reason for this shift is that robot traders can react to market fluctuations at speeds beyond human capabilities, while also processing massive amounts of data to inform their responses. The roots of this technology revolution lie in a groundbreaking paper by Das et al. (2001), the authors of which were all researchers from IBM T. J. Watson Research Labs, published at the *International Joint Conference on Artificial Intelligence* (IJCAI). This paper marked the first-ever demonstration, conducted through carefully controlled laboratory-style experiments, that basic algoritmic trading systems consistently outperformed human traders.

The IBM IJCAI paper received significant global media attention; it examined the interactions between human traders and two robot-trader strategies known as *GD* and *ZIP*. GD was developed by Gjerstad and Dickhaut at the University of Minnesota in 1998 (see Gjerstad and Dickhaut (1998)); while ZIP (Zero Intelligence Plus) was created in 1996 by Cliff at Hewlett-Packard's main European Research Labs in Bristol, England (see Cliff (1997)). The experiments, conducted on a simulated electronic financial exchange, revealed that both GD and ZIP consistently yielded higher profits compared to human traders. These findings have been replicated by other researchers, as evidenced in studies reported by De Luca and Cliff (2011); De Luca et al. (2011); and Cartlidge and Cliff (2013).

In the twenty-two years since the publication of the IBM IJCAI paper, adaptive automated trading technology has been widely embraced by trading entities such as investment banks and fund-management companies that previously relied on human traders. It has now become customary on major financial exchanges for both the buyer and the seller in a transaction to be algorithmic trading systems. One of the primary motivations for adopting this technology has been the significant cost-saving achieved by transitioning from human to robot traders, considering the traditionally high salaries of human traders. However, as the presence of robots increased in the markets, concerns emerged, particularly regarding the emergence of *High-Frequency Trading* (HFT). HFT involves automated trading activities occurring within extremely short timeframes, sometimes lasting only a few seconds or less, where traders buy and sell tradeable assets like stocks or currencies for quick profits. For valuable insights and critical evaluations of the rise of robot traders and HFT, as well as their impact on market fairness and stability, the reader is referred to books such as those by Arnuk and Saluzzi (2012); Patterson (2013); and Bodek and Dolgopolov (2015).

Market participants, whether they are individual traders or trading-entity corporations, along with the regulatory authorities responsible for overseeing markets, share a clear interest in ensuring that markets exhibit fairness, stability, and efficiency. However, evaluating the degree of fairness, stability, or efficiency in any given market is challenging due to the practical impossibility of conducting controlled experiments on a major financial exchange while it is actively operating. The cost and feasibility of conducting A/B testing to compare alternative structures or operational approaches for a live exchange are not realistically attainable. Although some insights can be gleaned by comparing different exchanges or *alternative trading systems* (ATSs) – which are independent trading venues resembling exchanges but without the full regulatory requirements of securities exchanges – tthere is a growing interest in instead performing comparative experiments using accurate simulation models. These models serve to explore the performance of specific types of exchanges and allow experimenters to comprehend the boundaries within which the exchange can safely operate. The aim is to understand when the market operating on the exchange veers into unfairness, instability, or inefficiency.

One notable early demonstration of the effectiveness of simulation modeling in real-world financial exchanges was reported (several years after the fact) in the book by Darley and Outkin (2007). This study involved the development of an agent-based model (ABM) for the NASDAQ exchange, which was used to anticipate the consequences of transitioning from fractional dollar prices to decimal prices (i.e., a change known as *decimalization*). Darley and Outkin's ABM model made several specific predictions regarding the impacts of this change. Subsequently, when NASDAQ actually implemented decimalization, all of the model's predictions proved to be accurate, except for one long-term prediction that couldn't be verified as true or false at the time the simulation was documented in Darley and Outkin (2007). Other instances of ABM simulations in finance and economics, tracing back to the seminal *Santa Fe Stock Market* work of Arthur et al. (1996) can be found in the surveys by Hommes and LeBaron (2018) and Chen (2018).

To create useful ABM simulations of contemporary markets, it is essential to accurately replicate relevant internal operations of the exchange and also to faithfully model the market participants, i.e. the traders in the market. The task of modeling the behavior of individual human traders or trading institutions may *prima facie* appear to be highly challenging. However, significant progress has been made over the past 30 years by employing remarkably minimal models known as "zero-intelligence" (ZI) trading strategies. The effectiveness of these strategies was established in a seminal paper published in a top economics journal by Gode and Sunder (1993), where markets populated by a trading strategy called "ZI-Constrained" (ZIC) demonstrated market dynamics that closely resembled the dynamics of directly comparable markets populated by human traders.

When prompted to provide a price for a transaction, a ZIC trader responds with a randomly generated value drawn from a uniform distribution. The range of this distribution is bound by the trader's current *limit price*, which represents the maximum price a buyer is willing to pay or the minimum price a seller is willing to accept. For ZIC buyers, the lower bound of the uniform distribution aligns with the minimum bid price allowed by the exchange (referred to as the exchange's *tick-size*, commonly one cent/penny). On the other hand, for ZIC sellers, the upper bound of the distribution is an arbitrary system maximum price, representing the highest price that can be quoted on the exchange.

The research conducted by Gode and Sunder (1993) demonstrated that markets populated by ZIC traders exhibited the same level of *allocative efficiency* as comparable markets populated by human traders. Allocative efficiency is a technical measure that evaluates the economic efficiency of a market. Since the introduction of ZIC three decades ago, ABMs of financial markets incorporating ZI traders have been extensively studied in economics, finance, and management research papers. For comprehensive reviews of these studies, see Farmer et al. (2005); Ladley (2012); and Axtell and Farmer (2018).

Although ZIC traders have proven to be remarkably useful in modelling markets, they are sufficiently simple that they clearly lack one characteristic of human traders: the ability to learn or adapt to changing market conditions. A first remedy to this was offered in the development of the Zero-Intelligence-Plus (ZIP) strategy, which adds simple machine-learning as reported in Cliff (1997), and which, as was noted above, was one of the two trading strategies that the IBM IJCAI 2001 paper showed to consistently outperform human traders (i.e., to be "super-human" in that specific sense).

Very recently, an extension of ZIC was reported by Cliff (2023b) in which the probability mass function (PMF) for the distribution from which prices are randomly generated can be varied by setting a *strategy* parameter, denoted by $s \in [-1.0, +1.0] \in \mathbb{R}$. In this extension, known as *Parameterized Response Zero Intelligence* (PRZI), when $s = 0$ the response of a PRZI trader is identical to that of ZIC, but as $s \rightarrow \pm 1$, the shape of the PMF becomes more distorted away from the rectangle of ZIC's uniform distribution, thereby allowing the PRZI trader's strategy to be either more "urgent" than ZIC (i.e., more likely to issue a randomly-generated quote-price that is close to the trader's limit price, which is hence more likely to result in a transaction but at a lower level of profitability) or more "relaxed" (generating more profitable prices than ZIC, but taking longer to find a counterparty to trade with). As originally defined in Cliff (2021), PRZI traders are non-adaptive: their s-value is fixed for all time. Two subsequent papers (Cliff (2022a,b)) have shown results from markets populated entirely by PRZI traders, each of which is continuously trying to adapt its s-value in an effort to improve its profitability. Because the profitability of any one PRZI trader's s-value is dependent at least in part on the set of s-values currently being used by all the other traders in the market, such a system is inherently co-evolutionary – and this then prompts questions of whether the co-evolutionary process within the market can reliably converge onto stable and economically efficient sets of strategy-values for the PRZI traders. To discuss this in more depth, first let us note that in a market populated by N_T traders, each running PRZI and where at time t each trader i has its own strategy-value $s_i(t)$, the vector $\vec{S}(t) = (s_1(t), s_2(t), \ldots, s_{N_T}(t))^\top$ represents a point in the N_T-dimensional *phase-space* for that market, such that as the traders each change their s_i values over time, the point $\vec{S}(t)$ traces a path or trajectory through the phase space.

In the first study of co-evolutionary PRZI markets (Cliff (2022a),) each trader used a simple stochastic hill-climbing (SHC) optimizer, much like an elementary multi-armed-bandit algorithm (see e.g. Lattimore and Szepesvari (2020); Slivkins (2021)); whereas in the second study (Cliff (2022b)), the SHC optimizer was replaced and each trader instead used *Differential Evolution* (DE: see e.g. Bilal et al. (2020)). In both cases, SHC and DE, a key notable result from the simulation studies was that the co-evolutionary process did not result in convergence to a stable equilibrium where all traders had settled on a static preferred strategy-value: instead, the system was in constant flux where a change in strategy by one trader $T1$ could trigger in reaction a change in strategy some other trader $T2$, and the change in $T2$'s strategy might prompt another trader to alter its strategy, and so on, until a consequent change in strategy by some trader Tn causes the original $T1$ trader to alter its strategy, and so the chain-reaction of strategy-changes continues. Furthermore, this constant flux in strategies, driven by each trader continuously attempting to improve its profitability, could result in the system tracing *loops* in its strategy phase-space, such that a given previous set of strategies (that was since improved upon) *recurs* and in that sense the system collectively has expended time and effort in co-evolving from some particular earlier state $S(t - \Delta_t)$ to eventually arrive at a new state $S(t)$ where $S(t) \approx S(t - \Delta_t)$, i.e. where the system has looped back on itself in strategy-space and where the transit time to complete the loop, denoted by Δ_t can be very large – potentially tens or hundreds of days of continuous round-the-clock trading and simultaneous continuous co-evolution of strategies. These recurrent states and looping paths through strategy-space were identified in Cliff (2022a,b) via the use of *Recurrence Plots* (RPs: a technique for visualizing high-dimensional complex dynamical systems, first developed by physics researchers; see e.g. Eckmann et al. (1987); Marwan et al. (2007); Webber and Marwan (2015), which are explained briefly in the Appendix to Cliff (2023b).

The conclusion drawn by Cliff (2022a,b, 2023b) was that the strategy-looping through recurring points in phase-space was evidence that that a competitive co-evolutionary process within the market, while clearly sufficient to drive constant change, was unable to guarantee even near-optimal outcomes or convergence to stable economic efficiency. However, one criticism of that argument is that the results could be a direct consequence of the choice to populate the market entirely with PRZI traders – possibly, the counterargument goes, the strategy-looping would not be seen if instead the market was populated with more sophisticated traders, such as humans or human-level robot traders. And it is that issue which is explored and resolved in this paper: here, I present first results from running co-evolutionary market simulations where all traders are instead using the ZIP strategy, known to outperform human traders, and where each trader is continuously using the DE optimization process intended to improve its profitability by searching for better settings of the five hyperparameters that govern the behavior of a ZIP trader. The key novel result in this paper is that long-term ongoing instability of strategies does indeed routinely occur in the ZIP markets, in much the same was as they did in the PRZI markets: this is an indication that competitive co-evolution in financial markets may lead to suboptimal outcomes and with traders returning to earlier strategies previously superseded by subsequent adaptive improvements, even when those traders are "super-human".

In short, the results presented here demonstrate that the co-evolutionary *Red Queen* dynamic identified by Van Valen (1973) can manifest itself in these accurately modelled financial markets, with traders constantly adapting purely to stay where they are (in terms of profitability) as the trading-strategy landscape they adapt to and operate on continuously shifts under their feet. This casts major doubts on claims that, in virtue of the competition among market participants, the market's current distribution of strategies can be expected to be "efficient" for any reasonable definition of efficiency, and bolsters the argument for thinking not in terms of the *Efficient Markets Hypothesis* (see e.g. Fama (1970)), but rather in terms of the *Adaptive Markets Hypothesis* of Lo (2004, 2019).

Section 2 of this paper presents further background details, sufficient to explain the ABM simulation experiments reported here: the text of Sections 2.1 and 2.2 is reproduced essentially verbatim from an earlier paper by Cliff (2023c) and can be skipped over by any readers already familiar with the background to this work. Having covered the relevant background, the design of the experiments conducted here to explore the effect of coevolution on strategy stability is explained in Section 3, and then key illustrative results from the comparative experiments are visualized and analyzed in Section 4. Sections 5 and 6 then offers discussion of further work and draw conclusions, respectively.

2. Background

2.1. Markets, Exchanges, and the Limit Order Book

The introduction of automated trading systems in contemporary financial markets around the world has already been discussed in sufficient depth in Section 1. Nevertheless, to fully explain the simulation experiments reported here to readers unfamiliar with financial markets, it is necessary to first introduce some standard terminology.

A *market* in any one type of tradeable asset (e.g a particular commodity or security or currency) is centered on a process by which *traders* in that asset can interact to identify potential counter-parties for a trade and then agree a price that both sides of the deal consider to be acceptable. While some important markets are organised around decentralized networks of traders interacting with each other via phone calls or computer messaging (which are known as *Over-The-Counter* or OTC markets), very many major financial markets around the world involve the traders interacting via platform offered by a centralized *exchange*. Such an exchange will typically accept *quotes* (or *orders*) from potential buyers and potential sellers – a buyer's order will name a quantity and a per-unit *bid-price* (the maximum price the trader is prepared to buy for), while a seller's order will name a quantity and a per-unit *ask-price* (the minimum price this trader is prepared to sell for). When a trader's order is received at the exchange, it will be processed by the exchange's *matching engine* to see whether any earlier orders received at the exchange and not yet matched with a counterparty are compatible with the newly-arrived order. If one or more earlier orders can be matched with the new order, then a transaction will be recorded, the relevant traders will be informed, and the transaction will be published on the exchange's *tape*, its publicly-viewable sequential record of trades that have taken place. However, if the new order cannot be matched with any previous orders already *resting* at the exchange, the new order will be added to the list of orders resting at the exchange, in the hope that a later-arriving order can be matched with it. Any unmatched order rests at the exchange until they are cancelled by the trader that originally sent them – the cancellation can be issued by the trader, or instead sometimes cancellation criteria are sent at the same time as the order is issued to the exchange, and the exchange's matching engine automatically cancels the order when the criteria are reached (e.g., orders may be specified to be automatically cancelled after a certain time-period has elapsed). After a new order is added to the exchange's list of resting orders for some asset, a summary of that list is made publicly available, by the exchange posting an updated version of that asset's *Limit Order Book* (LOB).

The LOB summarises and anonymizes the lists of currently resting bid orders and ask orders: the LOB can be thought of as a tabular data-structure, divided into two "sides" – the bid-side and the ask-side. Both sides are themselves tabular, with a nonempty row for each unique price at which a currently resting order exists, and with the data item at that row being the total quantity bid or offered at that price. The two sides of the LOB are each then sorted into order of descending goodness of price, i.e. bids are ordered highest-to-lowest, and asks are ordered lowest-to-highest, so that at the "top of the book" traders can see the current best bid price (and the quantity available at that bid price) and the current best ask price (and the quantity available at that ask price). The difference between the best bid price and the best ask price at any one time is known as the *spread*, and the arithmetic mean of the best bid and ask prices (i.e., the mid-point between the two) is known as the *mid-price*, the usual single-value summary of the market's current price. If a new order arrives at the exchange that is a bid with a price higher than the best ask, or an ask with a price lower than the best bid, that order is said to have *crossed the spread* and the exchange's matching engine then pairs the new order with the order(s) at the top of the book that have just been crossed, removing those orders from the top of the book and recording details of the transaction.

Whole books have been written on the dynamics of LOB-based financial exchanges (see, e.g., Abergel et al. (2016)). Traders interacting with a LOB-based exchange may seek to analyse the whole LOB (see e.g. Cont et al. (2021)), or perhaps instead will only look at the top of the book. In the simulation experiments reported in this paper, a long-established open-source detailed and accurate simulation of a LOB-based financial exchange, written in Python, has been used: this is the *BSE* simulator, described in more detail in Cliff (2012, 2018). BSE has been used as the underlying simulation platform for a number of research publications by various authors over recent years.

Simulating an ABM of a contemporary financial market running on a LOB-based exchange requires not only an accurate model of the exchange's internal mechanisms, but also gives rise to a manifest need for plausibly realistic models of the traders active in that market. Such model traders have been developed and validated by reference to work in *Experimental Economics*, discussed next.

2.2. Experimental Economics

One of the recipients of the 2002 Nobel Prize in Economics was Vernon Smith, in recognition for his part in pioneering and developing the research field known as *Experimental Economics*. In this field, researchers create carefully-designed laboratory experiments in which human subjects play the role of traders interacting via some market mechanism, or *auction*, and the experimenter can control the *supply and demand schedule* (SDS), which defines the market's supply and demand curves. As will be familiar to anyone with a high-school level of economics education, the supply and demand curves relate prices of bids/asks to the quantities available at those prices, and the intersection of the two curves gives the market's *equilibrium price* (denoted here by P_0) and *equilibrium quantity* (Q_0).

Technically, economists use the word *auction* to refer to the mechanism by which buyers and sellers interact to discover and agree a mutually acceptable price for a transaction, and there are many types of auction. For instance, the sole seller of a fine-art painting such as a Picasso will often be matched with a sole buyer via a process known as an *ascending-bid first-price auction*, often referred to colloquially as an *English Auction*; people who buy and sell on eBay.com are familiar with a slightly different mechanism, a *second-price sealed-bid auction*, because the buyer is whoever bid the most, but the price actually paid is whatever the second-highest bid-price was; in the Netherlands, sellers of tulip flower-bulbs start by announcing a high ask-price, and gradually lower it until a buyer announces they are willing to purchase at that price – this *Dutch Auction* is technically a *descending-ask first-price auction*; and so on. The one type of auction that has attracted the most attention from experimental economists is called the *Continuous Double Auction* (CDA), a simultaneous superposition of the English and Dutch auction mechanisms, in which any buyer is free to announce a bid-price at any time, and any seller is similarly free to announce an ask-price. One of the primary reasons why the CDA has attracted so much attention is that it is the mechanism implemented by LOB-based financial exchanges around the world.

It is beyond the scope of this paper to survey or summarise all research in experimental economics, but readers are referred to Smith (1962, 2000) for samples of Smith's seminal work in this field. For the purposes of explaining the experiments reported in this paper, it is sufficient to introduce the kind of supply and demand curves used in Smith's very first set of reported experiments, which are illustrated in Figure 1. This shows the supply and demand curves for a market in which there are 11 buyers and 11 sellers, where each trader has been assigned the right to buy/sell exactly one unit of the market's tradeable asset,

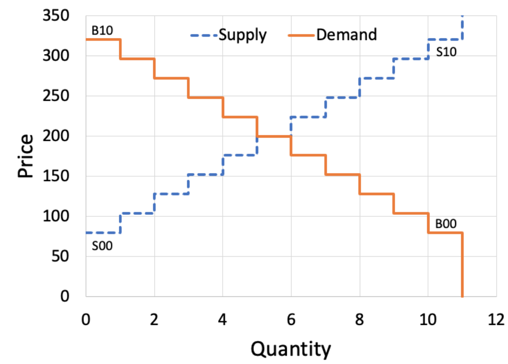

Figure 1. Supply and demand curves for an experimental economics session, of the style pioneered by Nobel-laureate Vernon Smith, in which there are 11 buyers and 11 sellers, each with an assigned *limit price* for a single unit of the tradeable asset. The schedule of seller limit prices determines the supply curve and similarly the schedule of assigned buyer limit prices determines the demand curve. The two stepped curves intersect at the *equilibrium point* (EP), giving an *equilibrium quantity* (denoted Q_0) of 5 and an *equilibrium price* (denoted P_0) of 200. Traders with assigned limit-prices to the left of the EP are referred to as having *intramarginal* limit prices, while traders whose limit-prices lie to the right of the EP are *extramarginal*. The maximally intramarginal traders are seller *S00* and buyer *B10*, while the maximally extramarginal traders are the seller *S10* and the buyer *B00*.

and where each buyer and seller have been assigned a *limit price*: the price beyond which a buyer cannot pay, and the price below which a seller cannot sell. Traders in the market are free to try to get a better price than their limit price: e.g a seller with a limit of $2.00 could agree a transaction at a price of $2.50 with a willing buyer, and walk away with a *profit* of $0.50; or a buyer with an assigned limit price of $2.00 could agree a transaction at a price of $1.50 and then that buyer would also make a profit of $0.50; and, in both cases, we could say that the trader is quoting a price that differs from their limit price by a *margin* (in this example, the margin is 25% in both cases: the seller creates an ask-price by *increasing* the the baseline limit price of $2.00 by 25%; the buyer creates a bid-price by *decreasing* the $2.00 limit-price price by 25%). One of the key findings of Smith's early experimental work (such as Smith (1962)) was that competition within the CDA among small numbers of traders, such as the 11 buyers and 11 sellers populating the market of Figure 1, was sufficient for the time-series of transaction prices in the market to converge on the theoretical equilibrium price P_0 given by the intersection of the supply and demand curves – prior to Smith publishing his results, it had been widely believed that much larger number of traders would be required for a market to show reliable and stable equilibration behavior.

Research in experimental economics by Smith and others over the 1960s, 70s, and 80s produced a large body of results illustrating the behavior of human traders in the CDA and various other auction mechanismss, via carefully structured and controlled laboratory experiments. Initially, the experiments involved humans interacting in minimal versions of the face-to-face open outcry trading pits that had formed the nexus of major national financial exchanges for more than a hundred years. But over time the experimental economists, like the real market practitioners, transitioned to basing their market interactions on networks of *trader terminals*, PCs that could be used to display market information and to input orders to a central exchange computer. Once the experiments were taking place on computer networks, it became possible to replce a human trader sat at a trader terminal with a software system doing the same job, i.e. receiving inputs of market data and producing outputs of orders sent to the exhange. This then enabled the groundbreaking work of Gode and Sunder (1993) whose seminal paper on zero-intelligence automated traders, introducing ZIC, was discussed in Section 1. Again, for reasons of brevity, no more will be said here about ZI traders in general, but to understand the experiments reported in Section 4 we do need to briefly explain more details of the ZIP trading strategy, which was demonstrated in 2001 by the IBM team to consistently outperform human traders (and is hence, in that limited sense, "super-human"), and which is extended in the experiments reported here.

2.3. ZIPDE: ZIP with Differential Evolution

For the purposes of this paper, it is sufficient to note that, as defined in Cliff (1997), any one ZIP trader is, at any one time, designated either as a buyer or a seller. ZIP is an *adaptive* trading strategy: its internal mechanisms implement a simple (and hence computationally inexpensive) form of machine learning (specifically, the *Widrow-Hoff Delta Rule with momentum*, which is also at the heart of back-propagation-based deep learning neural networks). A ZIP trader with index i maintains a time-varying *margin*, denoted by $\mu_i(t)$. ZIP trader i's initial margin value $|\mu_i(0)| \in [0,1] \in \mathbb{R}$ ($\mu_i(0) \leq 0$ for buyers, and ≥ 0 for sellers) is one of the five hyperparameters that determine the nature of its adaptivity while it is trading: the other four are a *learning rate* denoted by $\beta_i \in [0,1] \in \mathbb{R}$; a *momentum* (damping) term denoted by $\gamma_i \in [0,1] \in \mathbb{R}$; and two real-valued constants c_{a_i} and c_{r_i}, each also within the range $[0,1]$ (but expected to be much closer to zero than to one), that are used in calculating that trader's *target price*, i.e. the price it is aiming to get for the currently assigned trade. Thus any one ZIP trader i's adaptive trading behavior is set by the five-dimensional vector $\vec{V}_i = (|\mu_i(0)|, \beta_i, \gamma_i, c_{a_i}, c_{r_i})^\top \in [0,1]^5 \in \mathbb{R}^5$ and hence each ZIP trader's $\vec{V}_i$ can be thought of as a point in the five-dimensional phase-space for that individual trader. It is important to note that in the limit case of $\beta_i = 0$ no adaptation takes place at all in trader i and hence the other four hyperparameters then play no causal part in that trader's behavior; i.e., for the full set of hyperparameters to be potentially shaping the ZIP trader's activity, β_i must be greater than zero.

As was first demonstrated a more than 20 years ago (see e.g. Cliff (2001, 2009)) the setting of these hyperparameters for each ZIP trader in a market affects the dynamics of that market, it is possible to use evolutionary computation methods, such as an appropriate genetic algorithm (GA), to find vectors of hyperparameter values that optimize particular measures of market dynamics, such as the root mean square deviation of transaction prices from the theoretical competitive equilibrium price – a measure introduced by Smith Smith (1962) and popular ever since. However, the GA-based studies of Cliff (2001, 2009) concentrated on finding a *single* vector of hyperparameter values that could be cloned across *all* ZIP traders in the market, to give some desired overall market dynamics. Such an approach makes sense in applications such as market-based control (see e.g. Clearwater (1996)) but is not a realistic prospect as a model of actual real-world financial markets, in which each individual trading entity is privately attempting to maximise its own profitability, and is unlikely to ever share with competitors its own local information about the most profitable hyperparameter vectors, because such information constitutes much of that trader's competitive advantage.

To better model the dynamics of real financial markets, in the experiments reported here each ZIP trader in the market operates an evolutionary optimizer, continuously attempting to improve its profitability by altering one or more of its five hyperparameter values: the point in 5-d phase-space is now time-varying, and hence properly denoted by $\vec{V}_i(t) = (|\mu_i(0)(t)|, \beta_i(t), \gamma_i(t), c_{a_i}(t), c_{r_i}(t))^\top$. The evolutionary optimization method used here is *differential evolution* (DE: see e.g. Storn and Price (1997); Price et al. (2005); and Bilal et al. (2020)), which has repeatedly been demonstrated to be a very successful approach on a wide range of challenging problems across multiple application areas, and is commonly described as a leading-edge evolutionary optimization technique.

To distinguish this version of ZIP from its non-optimizing predecessor, we'll call this ZIPDE (ZIP with Differential Evolution, pronounced *zip-dee*). Each ZIPDE trader maintains its own private *population* of k different individual hyperparameter vectors, denoted $\vec{V}_{i,1}, \vec{V}_{i,2}, \ldots, \vec{V}_{i,k}$, and evaluates each of those k vectors sequentially, trading with each for a specified period of time $T_{\text{eval},i}$ and noting the profitability scored over that period: thus, after $kT_{\text{eval},i}$ seconds all k of the $\vec{V}_{i,j}$ hyperparameter vectors have been evaluated. Note that in the DE research literature, the population size which is here denoted by k (for compatibility with the research literature on k-armed bandits such as Gittins et al. (2011); Slivkins (2021)) is conventionally denoted by *NP* for "number in population".

In brief, DE as implemented here in each ZIPDE trader, operates iteratively, endlessly looping through a short sequence of actions: first, four distinct (non-identical) individuals are selected from the trader's population, which will be denoted here by $\vec{V}_a, \vec{V}_b, \vec{V}_c$ and $\vec{V}_d$; then a new vector denoted here as $\vec{V}_e$ is created, via an *adaptive step* of: $\vec{V}_e = \vec{V}_b + F.\vec{V}_\Delta$, where F is a DE hyperparameter known as the *differential weight*, with $F \in [0,2] \in \mathbb{R}$; and where $\vec{V}_\Delta$

is the *difference vector*: $\vec{V}_\Delta = \vec{V}_c - \vec{V}_d$. After this, a new candidate vector $\vec{V}_{\text{new}}$ is created by working along the length of the vector and setting each element at index i in $\vec{V}_{\text{new}}$, denoted by $\vec{V}_{\text{new},i}$, to be a random choice of either the element $\vec{V}_{e,i}$ with probability CR or the element $\vec{V}_{a,i}$ with probability $(1 - CR)$, where CR is the *crossover probability* hyperparameter, $CR \in [0, 1] \in \mathbb{R}$. Finally, there is the *selection step* in which the fitness (here, profitability) of $\vec{V}_{\text{new}}$ is evaluated and its fitness is greater than or equal to the fitness of $\vec{V}_a$ then $\vec{V}_a$ is deleted from the population and replaced by $\vec{V}_{\text{new}}$; but if the fitness of $\vec{V}_{\text{new}}$ is less than that of $\vec{V}_a$ then $\vec{V}_{\text{new}}$ is discarded and $\vec{V}_a$ remains untouched in the population. After the comparison of fitness values for $\vec{V}_{\text{new}}$ and $\vec{V}_a$ and any consequent adjustment to the popuation is done, the DE loop is over and it iterates back to randomly selecting another set of four distinct vectors from the population, and continues on iterating around this cycle forever.

As discussed previously in Cliff (2022b), it seems that almost the entire DE literature has been devoted to working on stationary optimization problems in which the underlying fitness landscape is fixed, constant, and hence in which it is desirable for the evolutionary process to converge on a local (and preferably global) optimum solution, and then stay just there forever more. Admittedly there are lots of engineering problems for which the assumption of this kind of fitness-landscape stationarity is not a problem, but for co-evolutionary systems such as those studied here, it is totally wrong. In the specific instance of DE as used here, as one trader's local population converges on a particular solution vector, the length of the difference vectors $\vec{V}_\Delta$ shortens to be very small, at which point the DE adaptive step reduces to $\vec{V}_e = \vec{V}_b + \epsilon$ for some small noise-level ϵ that is the residual near-zero value of $F.\vec{V}_\Delta$, and after a few iterations of that, the entire population becomes so heavily converged that all $\vec{V}_\Delta$ values are ≈ 0.0 and so differential evolution grinds to a halt.

To counter this, in the experiments described here, I added an anti-convergence mechanism which detects when a trader's local population is too converged, and then selects one individual vector from the population at random and adds mutations (zero-mean Gaussian noise) to each element of the vector. Very often, the mutated individual will have lower fitness than the un-mutated original individual and so will later be eliminated from the local population in the DE selection step; but every now and again, the mutations bring variation that does confer added fitness, possibly because the fitness landscape has altered, due to ongoing adaptations and changes in behavior of other traders within the market.

The measure used here to detect convergence is straightforward. For a trader with a local population $\mathcal{P}$ of k vectors s.t. $\mathcal{P} = \{v_1, v_2, \dots, v_k\}$, with $|v_i| = D \; \forall i$ (here, D=5), at each index j within the vector – i.e., at each *locus* on the vector, denoted as $v_{i,j}$ for locus j on individual i – first calculate the population locus mean as $\mu_j = \frac{1}{k}\sum_{i=1}^{k} v_{i,j}$ and then the population locus standard deviation as $\sigma_j = (\frac{1}{k}\sum_{i=1}^{k}(v_{i,j} - \mu_j)^2)^{\frac{1}{2}}$. Next, calculate the population locus coefficient of variation $CoV_j = \sigma_j/\mu_j$, and finally (because the CoV_j values are ratio measures) calculate the population geometric mean CoV for locus j, denoted here by GMC_j, as: $GMC_j = \left(\prod_{i=1}^{k} CoV_j\right)^{\frac{1}{k}}$ and then if any $GM_j < \epsilon$ for some small threshold level of mean CoV (e.g. for 2.5%, use ϵ=0.025), declare the population to be converged and add a mutation, as described above.

3. Experiment Design

As the key question being explored here is whether the co-evolutionary process is the cause of the observed strategic instabilities, the experiments reported here were deliberately designed to make it as easy as possible for the evolutionary process to settle onto a stable equilibrium, by configuring the BSE simulator to eliminate all other sources of variation. **This means that the experiments reported here are deliberately unrealistic in comparison to real markets**, but by clamping down on *all* sources of variation found in real markets we can directly observe any effect that the coevolutionary dynamic has on the stability of strategies in the simulated market – something that is manifestly impossible in a real market scenario – and it is intuitively obvious that introducing more sources of variation, to make the experiment scenarios closer to those of real-world markets, would only serve to introduce unwelcome confounding factors.

To this end, the assignments issued to the traders were all drawn from the single, fixed, symmetric supply and demand schedule (SDS) for $N_T = 22$ traders (11 buyers and 11 sellers) as illustrated in Figure 1: this means that the market's equilibrium price P_0 and quantity Q_0 were *constant* for the duration of each experiment, and that any one trader i was issued with the same limit price λ_i in every single assignment over the whole experiment (but different traders had differing fixed values of λ_i). By convention, the expected profitability of each trader is determined by the difference between their limit price(s) and P_0, and hence the expected profitability is zero for the 50% of traders in the market that are issued with extramarginal limit-prices. And so, while some extramarginal traders might get lucky and make a profitable trade on an occasional assignment (in just the same way as some intramarginal traders might occasionally suffer bad luck and fail to make a profitable trade for a particular assignment), in the long run the extramarginal traders get no useful fitness/PPS feedback, and are ignored in the discussion that follows. For brevity, in Section 4 of this paper we focus only the results for a single trader, the maximally intramarginal seller, denoted by S00 in the SDS of Figure 1: the choice of buyer or seller here is abitrary, and without loss of generality; the maximally intramargin trader is of interest because they are most likely to succeed in transacting every assignment issued to them, and so get the richest stream of PPS fitness signals from the market – put simply, if any trader is going to evolve to a stable optimum strategy, it is a maximally intramarginal one.

In all experiments reported here, BSE was configured to allow continuous round-the-clock trading for a single very long market session, lasting 365 days. BSE approximates continuous real-time via discrete time-slicing at a resolution of $1/N_T$ per timestep, such that each trader reacts to market events on average once per second. Fresh assignments were issued to the traders via a random process that averaged 15 seconds between assignments, so the expected number of assignments per trader is four per minute. As the maximally intramarginal traders would be expected to execute a transaction on almost every assignment they are issued with, the total number of transactions that *S*00 can be expected to execute in any one 365-day experiment is very close to $365 \times 24 \times 60 \times 4 = 2,102,400$.

For each trader *i*, the strategy evaluation time for the duration of the experiment was set to be some number of seconds draw from a uniform distribution between two and three hours (i.e., $T_{\text{eval},i} = \mathcal{U}(7200, 10800)$) – this is to avoid the kind of synchronisation artefacts highlighted by Huberman and Glance (1993) when all individuals in a simulation update at the exact same time.

For ease of comparison with prior work using differential evolution to co-evolve simpler trading strategies in BSE, I used the simplest form of DE (known in the literature as *DE/rand/1*) with *NP*=*k*=4, *F*=0.6, and *CR*=0.9.

With the constant equilibrium P_0 and Q_0 given by the unchanging SDS, the ZIP algorithm turns out to be somewhat over-specified for such an unrealistically unchanging environment: rather than spend any time at all *learning* what margin is most appropriate to the given circumstances (which will, of necessity, involve learning from mistakes, from situations where the margin could be improved upon in future), instead the most profitable approach is to disable learning by setting $\beta_i = 0.0$ and instead relying on the initial margin $\mu(0)_i$ to *evolve* via DE to be whatever value is most profitable without being so high as to be rejected by the counterparty side – this means that the evolved ZIPDE intramarginal trader's very first quote for any assignment will be near or at the best possible price for that trader given the unchanging P_0, Q_0, and λ_i values, but will be in part dependent on the prices quoted by the other traders in the market on the basis of their limit prices. For this reason then, in Section 4 we concentrate only on the values of β, on whether they evolve to zero or not.

The outcome of any one 365-day experiment is uncertain, given the multiple points in the system where randomly-generated values are injected, and so it is necessary to run multiple repetitions of each experiment, with only the seed value of the system's random-number generator(s) being varied, to give multiple independent and identically distributed (IID) experiment outcomes from which summary statistics can then be calculated.

In Section 4 results are presented from multiple IID repetitions of two experiment designs. The first design is referred to as *Evolve-1*, because in that design only a *single* trader (the maximally intramarginal seller, *S*00) is operating ZIPDE as described above – all other traders in the market are each running the standard ZIP strategy with no DE at all, no evolution of hyperparameter values, instead using hyperparameter values that are constant for any one experiment, fixed at initialization of the trader, using the standard BSE default distributions to set those values. In the second design, referred to as *Evolve-All*, ZIPDE as described above is being used simultaneously by *every* trader in the market. Thus, in the *Evolve-1* experiments, there is *no* co-evolution and the expectation is that DE will converge the sole evolving trader's β_i value to zero, or near-zero values given the anti-convergence measure explained above, and for that convergence to be stable (i.e., once $\beta_i \approx 0$, it stays ≈ 0 for the remainder of the experiment). The non-co-evolutionary *Evolve-1* experiments provide baseline reference data for evaluating the co-evolutionary *Evolve-All* experiments: if we see stable evolution to $\beta_i \approx 0$ in *Evolve-1*, and we do not see comparable stable evolution in the *Evolve-All* experiments, then given that the only difference between Evolve-1 and *Evolve-All* is the presence or absence of co-evolution, it is reasonable to conclude that the coevolutionary dynamic has destabilised the market, leading to a failure of the population of traders to reach and maintain optimum β_i values.

4. Results

Figure 2 shows the results for *S*00's profitability (PPS: profit per second) and β_{S00} values over ten IID repetitions of the 365-day *Evolve-1* simulation experiments.

The outcome of each experiment is classified as one of *converged*, *wandered*, or *missed*, by calculating a simple moving average $\hat{\beta}(t)$ of the last 12 β values (i.e., the past 24-36 hours, depending on the trader's specific value of $T_{\text{eval},i}$) and declaring $\hat{\beta}(t)$ to be near enough to zero as to be considered optimal if $\hat{\beta}(t) < \epsilon_\beta$ for some suitably small value of ϵ_β: Figure 6 uses $\epsilon_\beta = 0.05$. The sequences of timesteps on which $\hat{\beta}(t)$ is optimal in this sense are identified, and if there are no such sequences in the β time series then the experiment's outcome is classified as *missed*; if instead there are periods across the 365 days where $\hat{\beta}(t)$ is optimal, but the value of $\hat{\beta}(t)$ is non-optimal at the end of the experiment, then the outcome classification is *wandered*; while if $\hat{\beta}(t)$ is optimal at the end of the experiment, the experiment is classified as *converged*.

As can be seen in Figure 2, in seven of the ten experiments the outcome is *converged*. In total, 50 *Evolve-1* experiments were run: of which 36 (72%) converged; 10 (20%) wandered; and 4 (8%) missed. Figure 3 shows the mean and standard deviation of β across the whole set of *n*=50 *Evolve-1* experiments: Including experiments classified as *wandered* and *failed* masks the extent to which $\hat{\beta}$ approaches zero in the *converged* experiments, but this does serve to summarise the entire population of experiments. See Cliff (2023a) for visualizations of all 50 *Evolve-1*

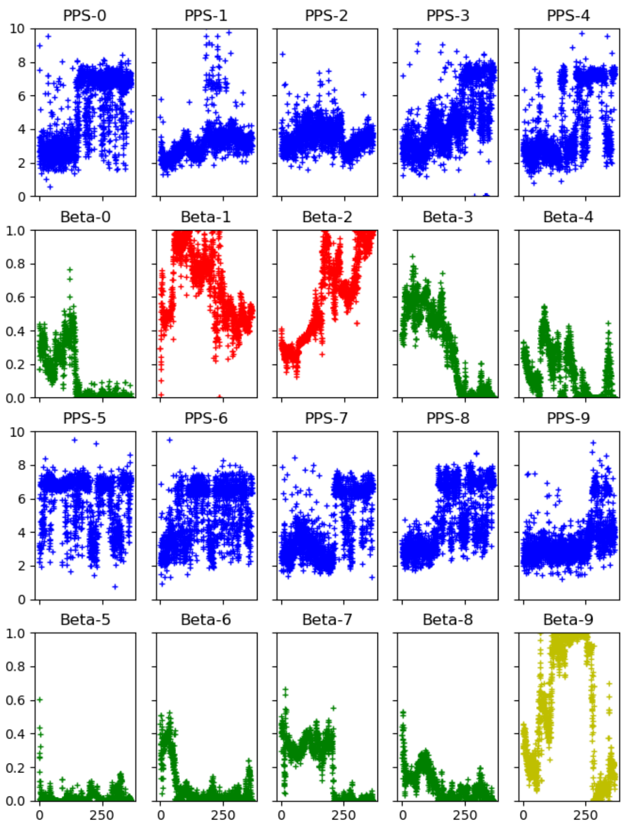

Figure 2. Results from ten IID repetitions of the 365-day ZIPDE *Evolve-1* experiment. Results from each single experiment are presented as a pair of time-series graphs for *S00*, the maximally intramarginal seller in the market: the graph labelled PPS (Profit Per Second; blue datapoint markers) vertically above the graph labelled Beta (the ZIP learning rate β). The Beta graphs are color-coded to indicate whether the outcome is classified as *converged* (green), *wandered* (yellow) or *missed* (red). In all graphs, the horizontal axis is time in days; the vertical axis is units of currency for PPS, and for Beta is the dimensionless real value $\beta \in [0.0, 1.0] \in \mathbb{R}$. Of these ten repetitions, in seven *S00*'s value of β stably converged to near-zero values, which is the optimal solution given the highly constrained nature of the *Evolve-1* experiment. See text for further discussion.

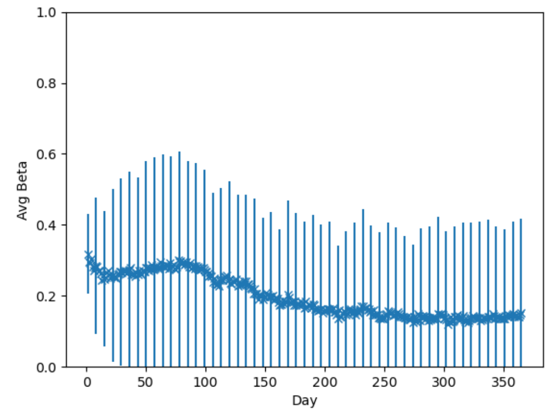

Figure 3. Average β values for *S00* in all *n*=50 IID 365-day ZIPDE *Evolve-1* experiments: this is the ten shown in Figure 6, plus results from an additional 40 repetitions shown in Cliff (2023a). Data-points are the arithmetic mean recorded at the end of each day; error-bars show plus and minus one standard deviation at the end of each week. See text for further discussion.

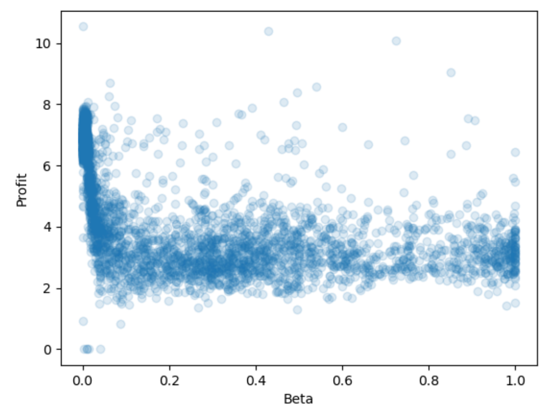

Figure 4. Scatter-plot showing the distribution of *S00*'s PPS-β pairs in the ten IID 365-day ZIPDE *Evolve-1* shown in Figure 2. See text for further discussion.

experiments conducted, and for summary time-series of only the converged results.

Visual inspection of the paired PPS and β plots in Figure 6 shows that PPS reaches its highest values when $\beta \approx 0$, as predicted. The relationship between β and PPS is further illustrated in the scatter-plot of Figure 4: peak PPS values in the range $[6.0, 8.0]$ are very tightly clustered around $\beta = 0$, and as $\beta \rightarrow 1.0$, PPS values rapidly fall to lie mainly in the range $[2.0, 4.0]$.

As Figure 5 shows, the PPS-β scatter-plot from ten *Evolve-All* experiments is qualitatively very similar to that from the ten *Evolve-1* experiments of Figure 4. Again we see a band of the highest PPS values when $\beta \approx 0$, as predicted, with a rapid fall-off to a band of lower PPS values once $\beta > 0$. But this similarity in PPS-β outcomes is only superficial: as is clear from Figure 6, the time-series of β values in all of the *Evolve-All* experiments is much more volatile, with $\beta \approx 0$ being reached in nine of the ten experiments but then wandering away to higher values such that

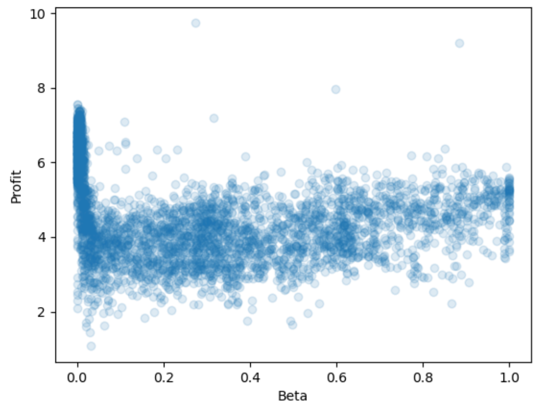

Figure 5. Scatter-plot showing the distribution of PPS-β pairs in the ten IID 365-day ZIPDE *Evolve-All* experiments as shown in Figure 6. See text for further discussion.

by the end of the 365 days, only one of the ten experiments finishes converged with $\beta \approx 0$.

Figure 7 then shows the mean and standard deviation of results from n=50 converged *Evolve-All* experiments: the ten from Figure 6 plus another 40 from experiments shown in Appendix B: of these 50, 16 (32%) converged; 18 (36%) wandered; and 16 (32%) missed. That is, switching from single-agent evolution to multi-agent co-evolution roughly halves the frequency of stable convergence.

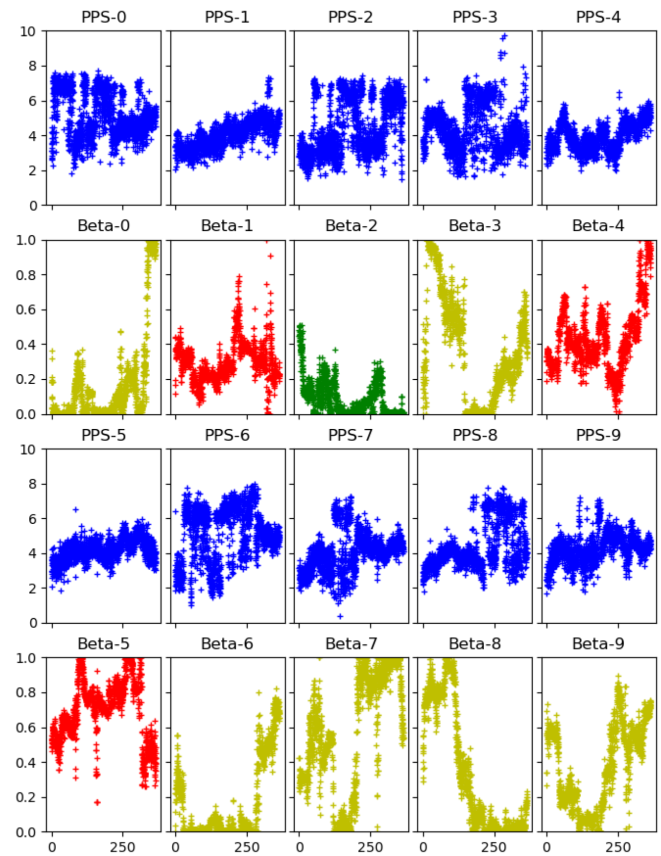

Figure 6. Results from ten IID repetitions of the 365-day ZIPDE *Evolve-All* experiment. Format is the same as for Figure 2. The optimum value of $\beta = 0.0$ is *missed* in three experiments (red datapoint markers), is reached at some point over the duration of the year in the remaining seven experiments, but six of those seven then *wandered*, failing to hold stable at $\beta = 0.0$ (yellow datapoint markers), and only one experiment *converged* and finish the year with $\beta \approx 0.0$ (green datapoint markers). All nine of the *missed* and *wandered* outcomes show wide swings in β occurring at unpredictable times over the duration of the experiment. See text for further discussion.

5. Discussion and Further work

Intuitively, a likely story for explaining for the lack of strategy-space stability seen in the co-evolutionary ZIPDE markets is that a change in the strategy of one agent a_1 somewhere within the market triggers a response, a reactive change, in the strategies of one or more other agents in $\{a_2, \ldots, a_{N_T}\}$ and this chain of actions-and-reactions eventually affects a_1 causing a fresh reaction in $a1$, and this cyclic network of causal interactions then continues forever. One avenue for further research is to develop analytical tools that can be more precise than this hand-wavy narrative: as we have total control over the simulations, we should be able to work through each timestep of the system and identify exactly when and what change in the strategy of some agent a_i then caused subsequent reactive changes in the strategies of some number of agents $a_{j \neq i}$. Untangling this what-caused-what in a network of interacting nonlinear units is a problem that is probably best addressed by use of *Granger Causality* (see e.g. Granger (1969); Shojaie and Fox (2022)), a technique developed for econometrics but which has also found great use in understanding causal patterns of activity in neural systems (see e.g. Seth et al. (2015)). There are also clear similarities with theoretical physics studies of spin-glass systems, such as the Ising-Lenz model (see e.g. Cipra (1987)), and with studies of cascades and epidemics on abstract networks (e.g. Watts (2002)). And, of course, there is a long history

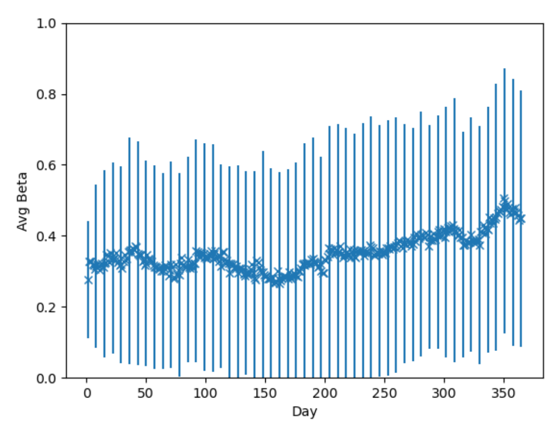

Figure 7. Average β values in n=50 IID 365-day ZIPDE *Evolve-All* experiments: this shows results from the ten experiments illustrated in Figure 6, plus an additional 40 experiments, the results from which are illustrated in Cliff (2023a). Format as for Figure 3. See text for further discussion.

of mathematical modelling of co-evolutionary processes in the theoretical biology literature (see e.g. Maynard Smith (1982); Thompson (1994); Hofbauer and Sigmund (1998)). There are many publications in these research fields, all of which can potentially be drawn upon to give better understanding the dynamics of co-evolving ZIPDE networks: there is a fair amount of work to be done here, and any progress in this direction will be reported in future papers.

Another direction for further work would be to explore whether any of the many variants of DE is for some reason more stable in the co-evolutionary context than the "plain vanilla" DE/rand/1 variant used here (for a comparative review of DE variants, see Georgioudakis and Plevris (2020)). For example, Herbert (2023) has very recently published promising preliminary results for using JADE (Zhang and Sanderson (2009)) in the context of co-evolving PRZI traders, and so an obvious next step is to re-run the experiments here using JADE instead of DE/rand/1.

6. Conclusion

This paper has reported the first ever results from extensive long-term simulations of co-evolutionary adaptive markets populated by ZIPDE traders, and has demonstrated that ZIPDE markets show qualitatively the same kind of unstable dynamics in strategy space as had been reported in previous studies where simpler trading strategies (e.g. PRZI) were co-evolving using less sophisticated evolutionary adaptation mechanisms (e.g., stochastic hill-climbing rather than DE). The simulation experiments reported here involved traders engaging in many millions of transactions over the course of each one-year continuous market session simulated at sub-second resolution, and the design of the experiments was deliberately constrained: eliminating all controllable sources of variation to give the clarity of a straightforward A/B test of evolution vs. co-evolution. The results show clearly that while simple single-agent evolution leads to stable outcomes in the scenarios studied here, co-evolution leads to instability.

Thus, the results presented here establish that even when using leading-edge evolutionary optimization techniques such as DE, operating on "super-human" trading strategies such as ZIP, the co-evolutionary Red Queen dynamic of Van Valen (1973) manifests itself, with traders constantly adapting and counter-adapting in strategy-space just to stay where they are in terms of profitability. Several avenues of further interesting research have been discussed here, and results from exploring those issues will be reported in future papers. The Python source-code used for the simulations reported here is freely available as open-source on GitHub (see Cliff (2012)), to allow other researchers to readily replicate and extend this work.

References

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

Arnuk, S. and Saluzzi, J. (2012). *Broken Markets: How High-Frequency Trading and Predatory Practices on Wall Street are Destroying Investor Confidence*. Financial Times / Prentice Hall.

Arthur, W., Holland, J., LeBaron, B., Palmer, R., and Taylor, P. (1996). Asset pricing under endogenous expectations in an artificial stock market. Technical Report 1996-12-093, The Santa Fe Institute.

Axtell, R. and Farmer, J. D. (2018). Agent-based modeling in economics and finance: Past, present, and future. Technical report.

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.

Bodek, H. and Dolgopolov, S. (2015). *The Market Structure Crisis: Electronic Stock Markets, High Frequency Trading, & Dark Pools*. Decimus.

Cartlidge, J. and Cliff, D. (2013). Evidencing the "robot phase transition" in experimental human-algorithmic markets. In Filipe, J. and Fred, A., editors, *ICAART-2013: Proceedings of the Fifth International Conference on Agents and Artificial Intelligence*, volume 1, pages 345–352.

Chen, S. H. (2018). *Agent-based computational economics: How the idea originated and where it is going*. Routledge.

Cipra, B. (1987). An introduction to the Ising Model. *The American Mathematical Monthly*, 94(10):937–959.

Clearwater, S., editor (1996). *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific.

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

Cliff, D. (2001). Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical Report 2001-99, Hewlett-Packard Laboratories.

Cliff, D. (2009). ZIP60: further explorations in the evolutionary design of trader agents and online auction-market mechanisms. *IEEE Transactions on Evolutionary Computation*, 13(1):3–18.

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. In Bruzzone, F., editor, *Proc. 30th Euro. Modeling and Simulation Symposium (EMSS2018)*, pages 194–203.

Cliff, D. (2021). *Parameterized-Response Zero-Intelligence Traders*. Manuscript, SSRN:3823317.

Cliff, D. (2022a). Co-evolutionary Dynamics in a Simulation of Interacting Financial-Market Adaptive Automated Trading Systems. In *Proc. 33rd European Modelling and Simulation Symposium (EMSS2022)*.

Cliff, D. (2022b). Metapopulation Differential Co-Evolution of Trading Strategies in a Model Financial Market. In *Proceedings of the 2022 IEEE Symposium Se-*

ries on Computational Intelligence (SSRN: 4153519), pages 1600–1609.

Cliff, D. (2023a). Co-evolution causes instability: Differential evolution of ZIP automated traders in a simulated financial market. *SSRN*, 44808214.

Cliff, D. (2023b). Parameterized-response zero-intelligence traders. *Journal of Economic Interaction and Coordination*.

Cliff, D. (2023c). Recurrence-Plot Visualisation and Quantitative Analysis of Long-Term Co-Evolutionary Dynamics in a Simulated Financial Market with ZIP Traders. In *Proc. 20th International Conference on Modeling, Simulation, and Visualisation Methods (MSV'23)*.

Cont, R., Cucuringu, M., and Zhang, C. (2021). Price impact of order flow imbalance: Multi-level, cross-asset and forecasting. *SSRN:3993561*.

Darley, V. and Outkin, A. (2007). *A NASDAQ Market Simulation: Insights on a Major Market from the Science of Complex Adaptive Systems*. World Scientific.

Das, R., Hanson, J., Kephart, J., and Tesauro, G. (2001). Agent-human interactions in the continuous double auction. In *Proc. IJCAI-2001*, pages 1169–1176.

De Luca, M. and Cliff, D. (2011). Human-agent auction interactions: Adaptive-Aggressive agents dominate. In *Proceedings IJCAI-2011*, pages 178–185.

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.

Eckmann, J.-P., Oliffson Kamphorst, S., and Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 5:973–977.

Fama, E. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *J. Finance*, 25(2):383–417.

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

Georgioudakis, M. and Plevris, V. (2020). A comparative study of differential evolution variants in constrained structural optimization. *Frontiers in Built Environment*, 6(102):1–14.

Gittins, J., Glazebrook, K., and Weber, R. (2011). *Multi-Armed Bandit Allocation Indices*. Wiley, 2 edition.

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

Gode, D. and Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *J. Political Economy*, 101(1):119–137.

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438.

Herbert, G. (2023). Differential weight and population size of PRDE traders: An analysis of their impact on market dynamics. In *Proc. 15th Int. Conf. on Agents and Artificial Intelligence (ICAART2023), also SSRN:4377728*.

Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

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

Huberman, B. and Glance, N. (1993). Evolutionary games and computer simulations. *Proc. National Academy of Sciences*, 90:7716–7718.

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

Lattimore, T. and Szepesvari, C. (2020). *Bandit Algorithms*. Cambridge University Press.

Lo, A. (2004). The adaptive markets hypothesis. *The Journal of Portfolio Management*, 30(5):15–29.

Lo, A. (2019). *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton University Press.

Marwan, N., Carmen Romano, M., Thiel, M., and Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, 438:237–329.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

Patterson, S. (2013). *Dark Pools: The Rise of AI Trading Machines and the Looming Threat to Wall Street*. Random House.

Price, K., Storn, R., and Lampinen, J. (2005). *Differential Evolution: A Practical Approach to Global Optimization*. Springer.

Seth, A. K., Barrett, A. B., and Barnett, L. (2015). Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297.

Shojaie, A. and Fox, E. B. (2022). Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9(1):289–319.

Slivkins, A. (2021). *Introduction to Multi-Armed Bandits*. Arxiv:1904.07272v6.

Smith, V. (1962). An Experimental Study of Competitive Market Behaviour. *Journal of Political Economy*, 70(2):111–137.

Smith, V., editor (2000). *Bargaining and Market Behavior: Essays in Experimental Economics*. Cambridge University Press.

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

Thompson, J. (1994). *The Coevolutionary Process*. University of Chicago Press.

Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, 1:1–30.

Watts, D. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771.

Webber, C. and Marwan, N., editors (2015). *Recurrence Quantification Analysis: Theory and Best Practice*. Springer.

Zhang, J. and Sanderson, A. (2009). JADE: Adaptive Differential Evolution with Optional External Archive. *IEEE Trans. Evolutionary Computation*, 13(5):945–958.