Exploration of PRDE Trader Differential Evolution Over Aggregated Dynamic Inelastic Markets

Jordan Taylor

Department of Computer Science

University of Bristol

Bristol BS8 1UB U.K.

xz18677@bristol.ac.uk

***Abstract*—This paper explores the performance of PRDE; a one-dimensional differential evolution (DE) trading algorithm in a high-fidelity financial market simulation, and determining its profitability under varying stochastic market conditions and sets of evolution parameters. Experiment design took inspiration from Monte Carlo simulations, covering over 30,000 trials in stable, fluctuating, and volatile supply and demand schedules to reduce random market variables to expected values. It is shown that, for all market conditions, a statistically significant correlation exists between PRDE's profitability and its mutation rate. Secondly, this paper describes an extension to PRDE, replacing its basic DE algorithm with JADE-with-archive, and naming the new algorithm PRDJ. It is demonstrated that PRDJ's performance is similar to the expected performance of a PRDE instance with un-optimized parameters, albeit with a faster strategy convergence rate.**

Keywords— Zero-Intelligence Traders; Financial Markets; Automated Trading; Co-Evolution; Differential Evolution.

I. Introduction

Present day financial markets involve a high proportions of adaptive automated trading systems, each adapting its trading strategy in response to the distribution of competitor strategies. This is a challenging phenomenon to capture in a laboratory setting or observe in controlled markets. Agent-based modelling is a technique used to generate accurate computer-simulation models of markets and their inhabitants. This paper reports on the investigation into optimising the parameters of a trading agent guided by a process known as differential evolution (DE) to increase its profitability. [1], Known as PRDE, the trader shall be evaluated in the context of a public-domain open-source agent-based model of a contemporary financial exchange named BSE [2]. As of the date of writing this paper, PRDE has not yet been formally evaluated in a scenario against other types of trader, nor has its performance been assessed in stochastic volatile market conditions. This is important, as commodities traded in the real world do not have fixed value. Furthermore, an investigation into volatile market performance of an adaptive trading algorithm directly relates to the lucrative market of options trading where strategies have to be constantly adjusted.

Section II explains background terminology in more detail. Section III describes preliminary work prior to the investigation into PRDE. Section IV explores and discusses the variation in performance of PRDE. Section V introduces an implementation of a more advanced DE algorithm. Section VI contains the conclusion and potential future work.

II. Background

A. Differential Evolution

In the most general sense, differential evolution is a method that attempts to optimise a problem by iteratively trying to improve a candidate solution with regards to a fitness function; a given measure of quality. Part of a set of procedures known as metaheuristics, these algorithms make few assumptions about the specific problem being optimised, but do not guarantee finding a globally optimal solution. PRDE utilises a basic form of DE, known as *(DE/rand/1)* to explore its solution space. After each generation, strategy values in its candidate population are mutated in the form:

$$v_i = x_{r1} + F_{xc}(x_{r2} - x_{r3})$$

Where $\mathbf{v_i}$ is the new vector, $\boldsymbol{x_{r1}}, \boldsymbol{x_{r2}}, \boldsymbol{x_{r3}}$ are three arbitrary candidate strategies, and **F** is the mutation factor coefficient.

B. The Bristol Stock Exchange (BSE)

BSE is an open source a high-fidelity agent-based simulation of a modern financial exchange. It utilizes a continuous double auction method for buyers and sellers interfacing with a limit order book.

Multiple traders are defined already defined in BSE. Those that are mentioned in the paper are as follows:

- ZIP: Multi-parameter limited intelligence trader. Adjusts its target profit margin in response to market conditions. [3]
- SHVR: A minimally simple trader. Reads the best prices on the limit-order-book and, if no loss would occur, attempts to issue a better offer
- PRSH: A minimally simple one-dimensional stochastic hill climbing trader. It attempts to optimize its strategy value, **s**, in a manner similar to a ***k***-armed bandit. P-type trader.
- PRDE: The successor to PRSH, with the stochastic hill climbing algorithm replaced with a differential evolution system [4]. P-type trader.

Documentation on PRDE states that the strategy population size ***k*** (formally described as NP) may only be defined for ***k* >= 4** due to an otherwise insufficient search space for effective differential evolution to take place [4]. However, a higher value of ***k*** is not always more effective. The 'No Free Lunch' theorem may be invoked to describe the dilemma of exploration vs exploitation in PRDE exploring its strategies in a market session. Higher values of **k** necessitate more time to evaluate a generation during a market session – and as such the trader may forgo more optimal behaviours in order to explore its solution space.

The second parameter of interest is $\boldsymbol{F}$, PRDE's differential evolution mutation factor. Documentation states that the accepted range for a sensible mutation factor is between zero and two ($\mathbf{0 < F <= 2}$). A factor of zero would imply no evolution, whereas a factor greater than two may lead to unstable evolution and characteristically greedy behaviour – leading to premature strategy convergence.

III. Preliminary Work

Prior to investigating the performance of PRDE over varying parameters and market settings, it was crucial that the current set-up of BSE and associated scripts was tested to ensure it behaved as expected. To do this, a preliminary experiment was run with the aim of replicating the work of D. Cliff and his introductory paper on PRDE (showing its profitability over PRSH) [4]. Furthermore, it is important to authenticate the use of the key metric, actual-to-expected profit ratio, that shall be used as the crux of this paper.

A. Replicating Results of PRDE vs PRSH

To replicate the work of D. Cliff, and therefore validate the use of simulations including PRDE, a simple market simulation was defined – with the aim of measuring the cumulative profit for each class of trader over multiple trials. Figure 1 shows a 7 day experiment describing the expected cumulative profitability of each class of trader over 100 trials. The supply and demand schedule were equal; defined as a constant range between 50 and 100 with a fixed step mode. The market was implicitly inelastic. The order schedule was defined with fixed order replenishment interval of 30 seconds, as per the work of V. Smith [5]. The market was populated with 32 traders (16 buyers, 16 sellers), with 4 traders per type in each subgroup. PRDE traders were initialized with the default parameters ($\boldsymbol{k}$ = 4, $\boldsymbol{F}$ = 0.8). SHVR and ZIP traders were added as a control population due to their predictable behaviour.

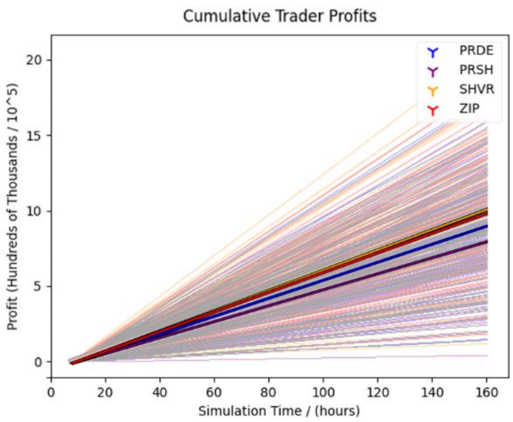

Fig. 1. Plot of the cumulative profit of multiple equally sized groups of traders for 100 trials over 7 days. Horizontal axis is time, measured in hours; vertical axis is cumulative profit for each class of trader – measured in hundreds of thousands. Highlighted lines indicate the non-linear regression of least squares for each trader calculated over all trials.

In *Figure 1*, it is evident that, on average, the cumulative profit yielded from the PRDE class of traders is greater than that of PRSH – with an experimental expected cumulative profit of $8.67\text{x}10^5$ vs $7.52\text{x}10^5$ respectively. This concurs with the findings of D.Cliff's paper showing the increased profitability of PRDE over PRSH. Furthermore, the two control trader classes, ZIP and SHVR exhibited expected behaviour – with SHVR known to outperform ZIP in most pairwise contests in the context of BSE [6].

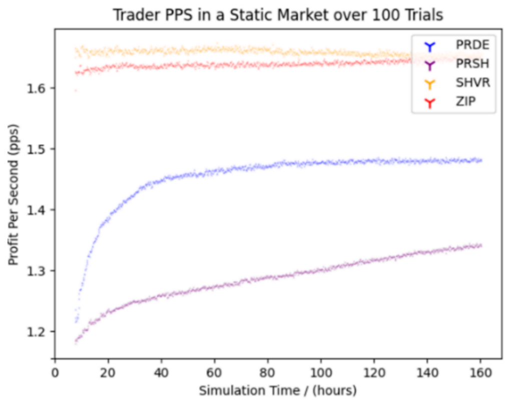

Fig. 2. Comparison of the overall profitability of the trader classes derived from the data in Figure 1, shown as profit per second over all trials. Each small dot on the graph denotes the mean profit per second for the respective trader class. for each time interval. Area under each plot is approximately equal to each trader's the expected cumulative profit..

Figure 2 shows the development of the mean profit per second metric for each class of trader over time. Once again, it is evident that the differential evolution employed in PRDE enables it to outperform PRSH and its stochastic hillclimber in terms of overall profit. Furthermore, it can be seen that PRDE converges on an optimal strategy not only at a greater rate than PRSH, but appears to have successfully achieved it, given that the gradient of the profit-per-second curve of PRDE is non-linear and tends towards zero. As such, the length of all subsequent experiments shall be increased to 14 days (336 hours) and P-type trader strategy evaluation times halved from 2 hours to a single hour (7200 seconds to 3600 seconds) to allow for more time to converge on an optimum strategy. Despite this, it can be said with confidence that the current set-up of BSE and behaviour of the strategy-optimising P-type algorithms are correct and further experiments may proceed.

B. Dynamic Market Definition

The purpose of this paper is to present a robust and comprehensive evaluation of the profitability of the PRDE trading algorithm over a domain of its differential evolution parameters; population size (NP / k) and mutation factor (F). Part of this assessment is to evaluate PRDE's performance over a range of dynamic supply and demand schedules to determine if there exists a set of particularly desirable/undesirable market conditions with respect to its profitability.

Due to the large number of trials intended to be run, the experimental design of this paper takes inspiration from Monte Carlo Simulations. Since BSE is a model of a simplified limit order book, in which the existence of a commodity is transitory (i.e., it no longer exists after being bought or sold), the efficient market hypothesis can be embraced with few confounding assumptions.

As such, a geometric Brownian-motion based diffusion model with superimposed jump component may be used to model independent inelastic supply and demand schedules. Specifically, the Weiner Process [7] of the Merton model's random walk was employed in the simulation [8], with additional jump as defined below:

$$ln(S_T) = ln(S) + \int_0^t (r - \frac{\sigma^2}{2} - \lambda(\mathrm{m} + \frac{v^2}{2})dt + \int_0^t \sigma dW(t) + J \quad (1)$$

Where the leftmost terms describe the random walk and J describes the superimposed jump component:

$$J = \sum_{i=1}^{N_t} (Q_i - 1) \quad (2)$$

In which ***N(t)*** a Poisson process with expected probability of k jumps occurring over its lifetime:

$$P(N(t) = k) = \frac{(\lambda t)^k e^{-\lambda t}}{k!} \quad (3)$$

And where $\boldsymbol{Q_i}$ is a log-normal distributed random variable defined as:

$$Q_i = \frac{1}{i\sigma\sqrt{2\pi}} e^{(-\frac{(\ln i - \mu)^2}{2\sigma^2})} \quad (4)$$

With ***S*** as the random walk's start value, $\boldsymbol{\sigma}$ as the walk expected standard deviation as per the Weiner process, ***r*** as the 'commodity safety factor', ***m*** and $\boldsymbol{v^2}$ as the mean and variance of the jump size probability distribution respectively, $\boldsymbol{\lambda}$ as the expected jumps for each process over a lifetime, and ***t*** is defined as the relative commodity lifetime as a proportion of itself (defaulted to 1). Variable ***m*** is defaulted to zero to make the chance of a positive jump equal to that of a negative jump.

With these parameters, three classes of dynamic market were defined with the following parameters:

- Stable Market: S = 100, σ = 0.07, r = 0.01, λ = 0, m = 0, v^2 = 0, t = 1
- Fluctuating Market: S = 100, σ = 0.3, r = 0.1, λ = 0, m = 0, v^2 = 0, t = 1
- Volatile Market: S = 100, σ = 0.3, r = 0.8, λ = 5, m = 0, v^2 = 0.25, t = 1.

Where the supply and demand schedules are defined as a pair of randomly generated walks given by (1). Visualisations of typical results can be seen in *figure 3.*

The parameters for each set of random walks have been chosen in such a way that the expected final commodity price is equal to within 1.5% of the starting commodity price given a large enough set.

Volatile markets are designed with the aim of constantly challenging the abstract fitness function of PRDE.

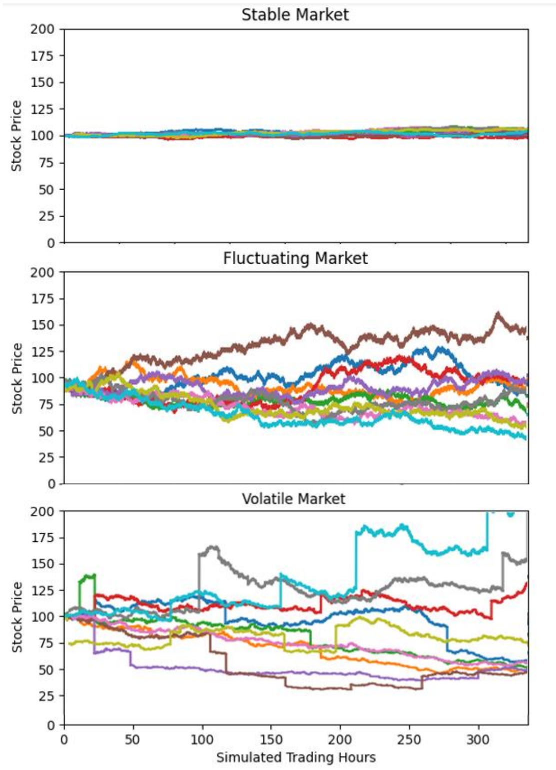

Fig. 3. Visualisations of the typical path a jump-diffusion random walk may take under three defined market conditions. Each plot shows a set of ten example supply/demand schedules.

C. Aggregated Market Performance

The investigation into the (***k***, ***F***) parameters of PRDE shall take place over stable, fluctuating, and volatile market conditions - all with unique and unknown supply and demand schedules. Consequently, overall profit and profit-per-second are no longer suitable metrics. The total amount of profit available in each market session is unknown. A new metric is required. As a result, outcomes over multiple market sessions shall be described using an actual-to-expected profit ratio.

Actual-to-expected profit ratio is defined retroactively after a market session has run to completion. It measures the overall performance of all traders of the same type for a given session. At the end of a market session, the sum of the mean profit per trader for all types is calculated. This value is used to divide the mean profit for each trader type. The resulting fraction is multiplied by the number of unique trader types in the session. This is described simply by:

$$Profit\ Ratio\ (p_0) = \frac{np_0}{\sum_{i=0}^{n} p_i} \quad (5)$$

Where ***n*** is the number of unique trader types in the market session and $\boldsymbol{p_i}$ is the mean profit for a class of traders. The result is a set of values clustered about 1. A result of 1 indicates the class of trader performed as if it yielded a share of the overall market profit proportional to its share of the market population. In other words, expected profit is based on the idea of performance if all traders had the same behaviour. Results greater than 1 indicate a trader type outperformed expectations, whereas results less than 1 indicate it was outperformed by the rest of the population.

Actual-to-expected profit can be calculated for any point in time after the market session has finished, determining the profit ratio with the current mean trader type profit as the numerator. This is demonstrated in *figure 4.* The advantage of using this metric is that it is agnostic to the conditions of the market. The amount of potential profit, supply and demand step-modes, as well as order schedule interval and time-mode have no bearing on the ratio. This is ideal as it will allow for results to be aggregated across multiple stochastic market conditions.

A second preliminary experiment was designed with the motivation of incorporating the new dynamic market conditions. Similar to the first preliminary experiment, a fixed step-mode was employed alongside a fixed order replenishment schedule of 30 seconds. It is known that different ratios of traders can have a significant effect on their performance, and by extension the average profit yielded for their population. For example, a single ZIP trader may do exceedingly well in a market of PRSH traders, whereas a ZIP majority set may collectively yield a profit ratio closer to 1 [9]. As a result, each market condition is to be trialled 5 times for all permutations of traders summing to 8 buyers and sellers per side for a total of 510 trials of length 14 days, excluding permutation sets that include populations of zero. The supply and demand schedules were independently randomly generated as per the examples in *Figure 3*, with ranges of ±30 at each point on the respective random walks. As such, each trial takes place in an inelastic market.

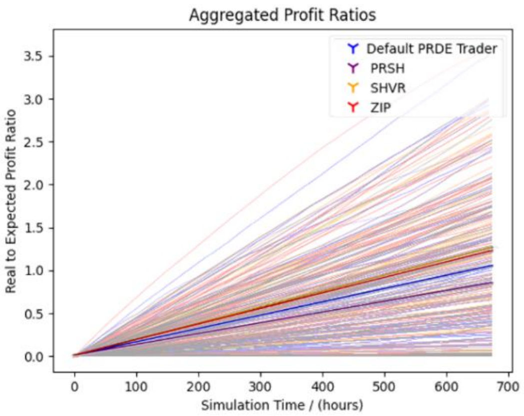

Fig. 4. Plot of the actual-expected profit ratio over two weeks for all traders, encompassing all trials, trader permutations, and market conditions. The non-linear regression of least squares for each trader type is highlighted.

There is motivation for covering all permutations. Firstly, for a sufficiently large dataset, incorporating the full range of trader permutations increases the robustness of the experiment by removing biases that may occur with fixed distributions of traders. As an aside, outlier results hold a value of entropy within the context of information theory. In this case, an expected value of self-information may be estimated from a single profit-ratio result [10]. If a profit ratio is outlying, it is more likely it was produced by a permutation trial in which it was the only member of its population. This is evidenced by the fact that the standard deviation of PRDE profit ratios at market closure for all permutations with a population of a single PRDE trader is 0.838, compared to 0.528 for trials with a population of four PRDE traders – despite the fact that the total amount of PRDE traders in single member population trials is greater than in the four-member population trials.

Comparing the cumulative profit results in *figure 1* and aggregated profit ratios in *figure 4,* the regression lines (an estimation of expected trader performance) are not only both approximately straight – but rank in the same order with a similar distribution of final expected results at market closure. As a result, it can be said with confidence that actual-to-expect profit ratio is a valid metric to be used in the investigation into the (***k***, ***F***) parameters of PRDE.

IV. Exploring PRDE Initialisation Parameters

Exploring the initialisation parameters of PRDE is the focal point of this paper. Each ***(k, F)*** parameter pair must be trialled, and a result yielded. This will be achieved using the process described in part C of the Preliminary Work section.

A. PRDE Parameter Investigation

As stated in the description of PRDE in part C of the Background section, the reasonable range of parameter ***F*** is that of a positive value less than or equal to two, and ***k*** may be defined as greater than four. As such, the investigation shall be run for all integer values of ***k*** between four and nine (inclusive) and for values of ***F*** between zero and two over ten equally sized intervals. ***F*** values of zero shall be investigated as a control measure.

Treating PRDE(***k***, ***F***) as a function, the range of actual-expected profit over the defined domain shall be tabulated for all market conditions, all permutations of buyers and sellers summing to 8v8, with 5 trials per setting and with unique supply and demand schedules generated for each trial of simulated length 14 days – pitting PRDE against varying populations of PRSH, SHVR, and ZIP.

Due to the stochastic market supply schedules, similarly to a Monte Carlo simulation, there is a motivation for a large number of trials. As the size of the dataset increases, the noise inherent in random jump diffusion schedules is minimised and converges to an expected range. The resulting dataset shall encompass a total of 30,600 fortnight long trials, totalling a simulated 428,400 days of market sessions.

TABLE I. Range of Aggregated PRDE Performance as a Function of (k, F) over all Market Conditions

Aggregated Markets		K 4	5	6	7	8	9	μ	σ
	0	0.860	0.917	0.915	0.922	0.927	0.917	0.910	0.627
	0.2	0.955	0.880	0.896	0.852	0.948	0.873	0.901	0.659
	0.4	0.916	0.982	0.909	0.896	0.932	0.864	0.916	0.599
	0.6	0.945	0.862	0.966	**1.038**	0.921	0.909	0.941	0.616
	0.8	0.878	0.905	0.994	0.983	0.932	0.943	0.939	0.615
F	**1.0**	0.982	**1.008**	**1.007**	0.925	0.984	0.994	0.983	0.650
	1.2	0.913	0.943	0.967	0.949	**1.057**	**1.020**	0.976	0.685
	1.4	0.956	0.926	**1.040**	0.976	**1.021**	0.94	0.976	0.621
	1.6	0.950	0.834	0.881	0.956	0.972	**1.011**	0.934	0.650
	1.8	**1.049**	**1.037**	**1.020**	0.954	**1.047**	**1.114**	**1.037**	0.650
	2.0	**1.103**	**1.000**	0.932	**1.022**	**1.136**	0.998	**1.032**	0.639
μ		0.956	0.964	0.957	0.952	0.958	0.962	***0.958***	
σ		0.656	0.636	0.654	0.622	0.652	0.642		***0.644***

Rs (F) = 0.806, p (2-tailed) = 0.00271. (significant)

Rs (k) = 0.2, p (2-tailed) = 0.704 (not significant)

a. Rs (F) denotes the Spearman's rank correlation coefficient between parameter F corresponding profit ratio.

b. Rs (k) denotes the Spearman's rank correlation coefficient between parameter k and corresponding profit ratio.

Fig. 5. Table 1 describes the mean actual to expected profit ratio for each pair of PRDE initialisation parameters (k, F) averaged over all market types; with each result averaged over all market conditions and all equal permutations of buyer/seller trader classes . Pairs of parameters yielding a ratio greater than 1 have their results recorded in bold. The mean profit-ratio for each ***F*** value is displayed in the vertical ***μ*** column, as well as the standard deviation. The mean profit-ratio for each value of ***k*** is displayed in the

horizontal σ row. Standard deviation σ is shown in the respective adjacent columns and rows. Mean profit-ratio of the overall dataset is displayed at the intersection of the μ column and row. It should be noted that these values are derived from the entire dataset, not the average values collated in the table.

TABLE II. RANGE OF AGGREGATED PRDE PERFORMANCE AS A FUNCTION OF (K, F) OVER STABLE MARKET CONDITIONS

Stable Market		K						μ	σ
		4	5	6	7	8	9		
F	0	0.870	0.982	**1.011**	0.874	0.914	0.957	0.926	0.663
	0.2	**1.001**	0.938	0.919	0.930	1.013	0.915	0.944	0.713
	0.4	**1.009**	**1.02**	0.979	0.894	0.869	0.776	0.916	0.528
	0.6	0.915	0.789	0.970	**1.120**	0.870	0.909	0.921	0.689
	0.8	0.954	0.852	0.983	0.977	0.815	0.878	0.900	0.615
	1.0	**1.035**	**1.037**	0.843	0.992	0.979	**1.002**	0.973	0.657
	1.2	0.929	**1.034**	**1.015**	**1.05**	**1.001**	**1.036**	**1.003**	0.727
	1.4	**1.050**	**1.005**	**1.078**	**1.015**	**1.040**	0.981	**1.020**	0.720
	1.6	**1.022**	0.802	0.854	0.959	0.938	**1.110**	**0.939**	0.738
	1.8	**1.135**	**1.03**	0.986	0.938	**1.144**	**1.125**	**1.051**	0.646
	2.0	**1.132**	**1.068**	0.847	**1.117**	**1.143**	0.983	**1.040**	0.619
μ		0.956	0.960	0.957	0.988	0.974	0.970	***0.968***	
σ		0.683	0.675	0.674	0.681	0.683	0.675		***0.680***

Rs (F) = 0.709, p (2-tailed) = 0.0146 (significant)

Rs (k) = 0.714 p (2-tailed) = 0.111 (not significant)

Fig. 6. Table 2 derives a subset of the data described in Table 1. It differs due to only describing the mean PRDE profit ratios yielded in stable market conditions – ie with stochastic yet minimally varying supply/demand schedules similar to the stable example plots in *figure 5*.

TABLE III. RANGE OF AGGREGATED PRDE PERFORMANCE AS A FUNCTION OF (K, F) OVER FLUCTUATING MARKET CONDITIONS

Fluctuating Market		K						μ	σ
		4	5	6	7	8	9		
F	0	0.861	0.934	0.938	**1.023**	0.943	0.909	0.925	0.621
	0.2	0.868	0.852	0.913	0.786	0.952	0.898	0.868	0.677
	0.4	0.832	0.926	0.907	0.985	0.954	0.874	0.903	0.634
	0.6	0.895	0.992	0.939	**1.052**	0.896	**1.000**	0.953	0.587
	0.8	0.878	**1.040**	0.994	0.979	0.986	**1.069**	0.981	0.616
	1.0	0.912	**1.033**	**1.060**	**1.014**	**1.002**	**1.030**	0.998	0.668
	1.2	0.948	0.938	**1.033**	0.862	**1.147**	1.028	0.983	0.727
	1.4	0.922	0.876	**1.134**	**1.088**	**1.062**	0.921	0.991	0.585
	1.6	0.842	0.781	0.883	0.995	**1.041**	0.935	0.903	0.606
	1.8	**1.015**	**1.021**	**1.124**	**1.022**	**1.145**	**1.030**	**1.050**	0.637
	2.0	**1.057**	0.983	0.955	**1.044**	**1.029**	**1.038**	**1.007**	0.671
μ		0.912	0.943	0.989	0.987	**1.014**	0.976	***0.960***	
σ		0.636	0.629	0.675	0.604	0.629	0.644		***0.637***

Rs (F) = 0.702, p (2-tailed) = 0.0161 (significant)

Rs (k) = 0.602, p (2-tailed) = 0.208 (not significant)

Fig. 7. Table 3 describes the mean profit ratios yielded by PRDE in fluctuating market conditions with supply and demand scedules similar to fluctuating example plots in *figure 3*.

TABLE IV. RANGE OF AGGREGATED PRDE PERFORMANCE AS A FUNCTION OF (K, F) OVER VOLATILE MARKET CONDITIONS

Volatile Market		K						μ	σ
		4	5	6	7	8	9		
F	0	0.849	0.884	0.792	0.897	0.925	0.915	0.877	0.594
	0.2	0.995	0.901	0.859	0.868	0.865	0.835	0.887	0.577
	0.4	0.897	1.049	0.841	0.839	0.973	0.972	0.929	0.614
	0.6	**1.036**	0.854	0.991	0.969	0.998	0.848	0.949	0.557
	0.8	0.807	0.872	**1.004**	**1.024**	0.995	0.917	0.937	0.601
	1.0	0.999	**1.005**	**1.117**	0.800	0.972	0.979	0.978	0.625
	1.2	0.861	0.906	0.854	0.966	**1.024**	1.025	0.939	0.592
	1.4	0.895	0.947	0.906	0.855	0.959	0.947	0.918	0.543
	1.6	0.987	0.970	0.907	0.945	0.936	**1.018**	0.961	0.581
	1.8	0.996	**1.111**	0.950	0.932	0.849	**1.219**	**1.009**	0.650
	2.0	**1.119**	0.999	0.995	0.935	**1.237**	**1.004**	**1.048**	0.625
μ		0.949	0.954	0.929	0.912	0.976	0.971	***0.948***	
σ		0.647	0.602	0.611	0.576	0.640	0.605		***0.614***

Rs (F) = 0.791, p (2-tailed) = 0.00375 (significant)

Rs (k) = 0.429, p (2-tailed) = 0.397 (not significant)

Fig. 8. Table 4 describes the mean profit ratios yielded by PRDE in volatile market conditions with supply and demand scedules similar to volatile example plots in *figure 3*.

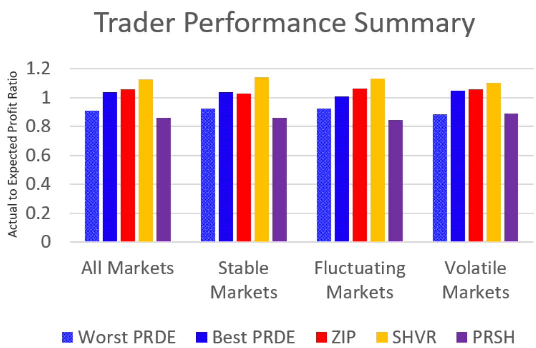

Fig. 9. Mean performance of each trading algorithm over all trials, grouped by market conditions, measured in actual-expect profit ratio. Best and worst PRDE performances are derived from a subset of the dataset where the trader was initialised with favourable/unfavourable parameters.

B. *PRDE Parameter Investigation Discussion*

Looking at the mean profit ratio column μ in Tables I to IV, there is an evident positive correlation between actual-to-expected profit ratio achieved by PRDE and the value of its mutation factor ***F*,** true for all market conditions. This can be stated with statistical confidence granted by a Spearman's Rank Correlation – a test of the monotonicity of a function. For each market condition, the whole set of profit ratios was paired with their corresponding ***F*** parameter. The correlation coefficient and two-tailed p value are displayed below each respective table (***Rs (F)***). All ***F*** **–** profit-ratio relationships are deemed to be statistically significant.

Subsequently, the Spearman's Rank Correlation procedure was conducted for each value of ***k*** and profit ratio across all market conditions, and despite a positive correlation coefficient being achieved for all data sets over the range of investigated ***k*** values, none of the relationships were deemed to be statistically significant, with two-tailed p-values greater than the standard confidence threshold 0.05.

For both ***k*** and ***F*,** there appeared to be no statistically significant relationship between the standard deviation of profit ratios and each set of parameters, with all standard deviations of results clustered about 0.644, with the greatest 0.738 (*Figure* 6) and smallest 0.527 (*Figure 8)* in opposing market settings. The consistency of standard deviation across all market settings in the dataset gives confidence to the number of trials that were run.

One of the aims of investigating multiple market settings was to determine what difference there was, if any, between the performance of PRDE over stable, fluctuating, and volatile markets. *Figure 9* displays the comparison of the mean performance of trading agents populating each market setting. From this graph, it is evident that in this transitory commodity market, SHVR continues to dominate; concurring with the findings of D. Snashall [9] - although it is worth noting that this is the first time SHVR's domination has been demonstrated with stochastic jump-diffusion defined supply and demand schedules.

In all non-stable markets ZIP continues to outperform even the strongest instances of over all settings. However, in one case where mutation factor ***F*** = 1.8, PRDE edges a 1.06% advantage in performance across its set of trials over ZIP. One hypothesis for this performance could be that, in stable market conditions, PRDE is more effectively able to search its solution space when its fitness function (profit-per-second in the current market condition) is not constantly being challenged. This has supporting evidence in the fact that, overall, PRDE was more profitable in stable markets than fluctuating and volatile settings (mean returns of 0.968, 0.960, and 0.948 respectively (*Figures 6, 7, 8)*) where variation in performance is ~2% from best to worst. However, this cannot be stated with statistical confidence. Drawing from the datasets of the markets with the greatest difference in mean profit-ratios (stable vs volatile), the two distributions of 10,200 values underwent an unpaired paired t-test; showing a difference in means but an insignificant p-value of 0.017.

Despite this, *figure 9* shows that even the worst performing PRDE where ***F*** > 0 outperforms PRSH, thus showing that differential evolution is a more effective optimization method than stochastic hill climbing in the context of P-type trader strategy.

C. Replication of Findings

In the previous section, it was shown how the mutation factor ***F*** is a statistically significant parameter to the performance of PRDE across stable, fluctuating, and volatile markets. The previous section also determined that ***k*** was not a statistically significant and as such future PRDE instances in future experiments shall be instantiated with population size ***k*** = 4. As a result, looking at aggregated market Table I (*Figure 5)*, the best and performing pairs of parameters can be defined as (***k*** = 4**,** ***F*** = 1.8) and (***k*** = 4**,** ***F*** = 0.2) respectively.

A second experiment was designed with the motivation test the validity of the previous section's findings. Two variations of PRDE were defined to describe the strongest and weakest sets of parameters. These variations, alongside the default version of PRDE, were tested in isolation under the same conditions as in the second preliminary experiment (Section 3, part C, *Figure 4*); i.e. over all permutations of traders and market conditions, for 14 days with 5 trials for each setting. The actual-to-expect profit ratios over time were calculated and extracted for each PRDE configuration.

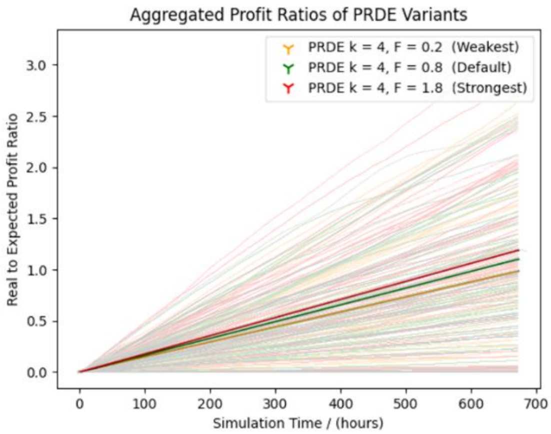

Fig. 10. Actual-expected profit ratios of each PRDE variant over time for all trials, aggregated over all market conditions and trader permutations. The non-linear regression of least squares for each PRDE variant is highlighted.

The results depicted in Figure 10 serve to validate the experimental findings. As predicted, the experimentally strongest PRDE configuration has a greater expected profit-ratio than both the default and weakest variants. The expected profit ratio of each variant at market closure is in order of their respected mutation factors, ***F*.**

V. CREATION OF PRDJ – THE JADE TRADER

As currently configured, PRDE utilizes the *'DE/rand/1'* differential evolution algorithm - a relatively simple algorithm in comparison to other examples in its family. It is proposed that a more advanced form of differential evolution may improve the profitability of PRDE.

A. JADE – A Differential Evolution Algorithm with Archive

JADE is an alternative differential evolution algorithm. As of the development of JADE, there has been no method developed solely based on a greedy DE variant which utilizes the information of the best-so-far solution in the mutation operation [11]. Despite the fact that greedy DEs are usually less reliable and tend to encounter problems such as premature convergence, JADE's *'DE/current-to-p-best'* strategy attempts to generalize *'DE/current-to-best'* to utilize the information of multiple good solutions. An advantage of JADE is that it employs the strategy mutation operation to also produce a new mutation factor ***F*** for each member of its population**.** As such, the members of a population do not need to be specifically instantiated with an initial ***F*** parameter, although an initial value may be assigned. Instead, populations employing a JADE (*'DE/current-to-p-best')* strategy are explicitly paramaterised with strategy population size: **NP / *k*,** proportion of top strategies threshold **p** and crossover update constant **c**; where **0 <** ***p*** **< 0.2** and **0 <** ***c*** **<0.5.**

Additionally, there exists a version of JADE with an optional archive [12]. Recently explored inferior solutions, when compared to the current population, provide additional information about the promising progress direction. These are stored in a set denoted as ***A***; the set of archived inferior solution, where the size of ***A*** is limited to ***k***.

B. PRDJ – Replacing Evolution Algorithm of PRDE

With JADE with archive having been introduced, this paper presents a variant of PRDE with a new DE algorithm driving its evolution. PRDJ utilizes the same strategy range as PRDE and PRSH, but uses JADE with archive to drive its evolution.

One of the challenges of implementing a more complex DE algorithm to a trader was in abstracting the fitness function. Similarly to PRDE, a PRDJ trader evaluates all strategies in its population each generation – with the fitness function being the average yielded profit-per-second over the strategy evaluation period. With respect to the control loop of BSE, each PRDJ trader now holds a record of the set of strategies used in the last generation. At the end of each generation, the evaluation period starts. Each new strategy is compared to the previous; the strategy with the highest yielded profit-per-second is kept. The strat-value and mutation factor are mutated for the next generation. Like PRDE, if the standard deviation of a population's strategy values fall below a threshold, a random individual is mutated.

An advantage of PRDJ over PRDE is that it doesn't require explicit initialization of its mutation factor ***F*;** the significant factor in section IV's investigation. As such, part C of this section shall lay out the design for an experiment to observe the change in mutation factor ***F*** and strategy value **s** over time, as well visualize PRDJs profit-per-second convergence characteristics in comparison to PRDE.

C. Discussion and Results of PRDJ

An experiment similar to that of the first preliminary experiment in Section III, part A, *Figure w*as defined. The supply and demand schedule were equal; defined as a constant range between 50 and 100 with a fixed step mode. The market was implicitly inelastic. The order schedule was defined with fixed order replenishment interval of 30 seconds. The market was populated with 40 traders (20 buyers, 20 sellers), with 4 traders per type in each subgroup – comprising PRDJ, ZIP, SHVR, and two variants of PRDE. The two PRDE variants were defined with the strongest and weakest parameters as found in Section IV, part b; (***k*** = 4, ***F*** = 1.8) and (***k*** = 4, ***F*** = 0.2) respectively. The set of PRDJ traders were initialized with crossover update constant ***c*** = 0.1, top strategy selection threshold ***p*** = 0.2, and population ***k*** = 4. The simulation was repeated for 100 trials over a duration of 14 days. A profit-per-second was plotted in *Figure 11*.

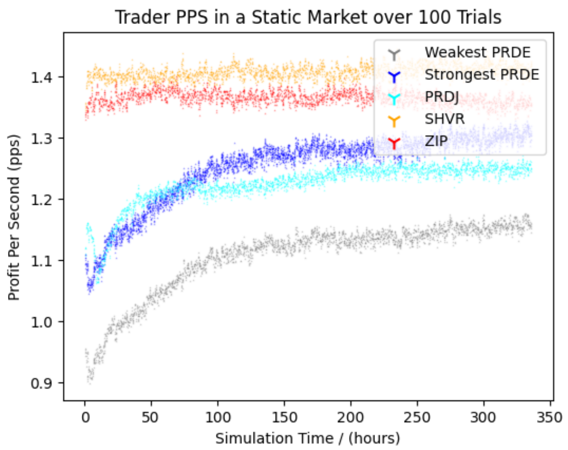

Fig. 11. Comparison of the average profitability over time PRDJ in comparison to strong and weak PRDE variants, SHVR, and ZIP, measured in profit-per-second.

Looking at *Figure 11*, PRDJ's performance in a static market can be defined as similar to PRDE; yielding slightly worse performance than the strongest PRDE variant and better performance than the weakest variant. It is worth noting that, in the context of BSE, PRDJ still loses out to the zero-intelligence traders of SHVR and ZIP.

The greediness of JADE's *'DE/current-to-p-best'* is evident in the gradient of the PRDJ plot in *Figure 11*; at its greatest point being steeper than all other plots. This is indicative of the fast convergence rate of greedy algorithms.

JADE's correctness of implantation into the PRDJ trader can be demonstrated by plotting heat-maps of the strategy value **s** and mutation factor ***F*** over time for all trials.

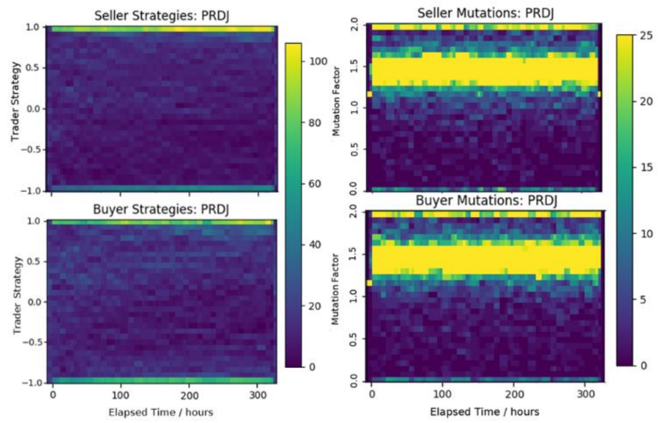

Fig. 12. Side by side heat-maps of PRDJ trader strategy values and mutation factors over time in a static market over 100 trials, separated into buyers 400 and 400 sellers. Horizontal time bins describe the current best strategy for each PRDJ trader at the start of each generation. Heatmap squares are coloured with respect to the scale to the right of each plot, with intensity denoting population.

The strategy heat-maps displayed in *Figure 12* bear a resemblance to the probability density heat-maps found in D. Cliff's inaugural PRDE paper [4]; with the majority of buyer and seller strategies converging to opposite extremes of the range throughout the market session – indicative of the expected behaviour of landscape traversing generational algorithms.

Both mutation heatmaps demonstrate behaviour expected of JADE. The mutation factor, ***F***, of most buyers and sellers converges to a normally distributed range about 1.45.

It would be expected that *'DE/current-to-p-best'*, a more advanced DE algorithm than *'DE/rand/1'*, would lead to PRDJ yielding increased profitability over its PRDE counterpart. However, the lower than expected performance of PRDJ may be due to several implementation oversights.

Firstly, the greediness of JADE may be reducing its potential profitability, especially when combined with a non-optimal set of parameters. This greediness in choosing a strategy is exemplified in the bimodality of the buyer and seller distributions clustered around the extreme values in *Figure 12*. Trading agents may find themselves drawn towards a local optimum in a basin of attraction – a consequence of premature convergence.

Secondly, the evolution of the mutation factor has an oversight in its implementation. It is evident in *Figure 12*, that a not-insignificant proportion of the population have mutation factors that converge to zero. As a result, these traders lose their capability to evolve. In converged populations, normal procedure is to intervene and mutate a random individual. However, the intervention technique introduced in PRDE fails to update the variable mutation factor ***F***, leading to stagnation.

Finally, a third factor may be the size of the strategy population in PRDJ. For sufficiently small populations, *'DE/current-to-p-best'* performs identically to the greedier *'DE/current-to-best'*, as the number of strategies selected as a basis for evolution rounds to one – hastening premature convergence. It is also hypothesized that a low population size would reduce performance in volatile markets. As such, a final experiment was designed to investigate whether ***k*** is a significant factor in the performance of PRDJ over a constantly changing fitness landscape.

Three variants of PRDJ were defined with parameters ***k*** = 3, ***k*** = 8, and ***k***= 12. Each variant was tested in isolation against PRDE and PRSH traders in a dynamic market, across all permutations of traders summing to 8 buyers and 8 sellers – trials of 14 days with five repeats per setting. PRDE traders were initialized with the default mutation factor, ***F*** = 0.8.

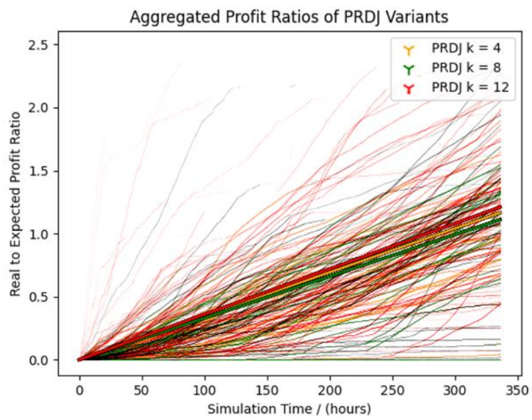

Fig. 13. PRDJ trader actual-to-expected profit ratio over time in volatile market conditions. Market sessions were populated with PRDJ, PRDE, and PRSH traders.

Figure 13 clearly demonstrates a performance advantage of PRDJ variants where ***k*** = 12 over smaller ***k*** values. There is a statistically significant difference between the sets of results of the variants where ***k*** = 12 and ***k*** = 4, in which an unpaired t-test resulted in a two tailed p value of 0.0192 indicating a significant difference between the two distributions. Despite this, the dataset was not large enough to claim the same with statistical significance for either comparison including ***k*** = 8.

The consequences of this finding imply that PRDJ may be able to yield more profit in volatile markets as the parameter ***k*** increases This implies PRDJ variants that maintain a suitably sized population of strategies are more resistant to market shocks with regards to their ability to yield a consistent profit. This may be due to the fact that, with a larger strategy population, the trader can quickly switch to a more optimal strategy already in its population, rather than relying on evolving a new strategy. However, the dataset from this experiment is insufficiently large to claim such as fact with statistical significance. Furthermore, these results suggest that the selected range of ***k*** for the PRDE parameter investigation may have been insufficient.

VI. Conclusion and Further Work

This paper set out to investigate two sets of premises. The first was to investigate the optimal pair of (***k, F***) parameters for the PRDE trading algorithm, and to determine PRDE's performance differences over stochastic stable, fluctuating, and volatile market conditions. The second premise was to introduce and evaluate an implementation of the differential evolution algorithm JADE-with-archive *('DE/current-to-p-best')* to drive the behaviour of a trader in BSE. At the end of this paper, both premises have been investigated.

The most optimal mutation factor for PRDE was determined to be ***F*** **=** 1.8, after an investigation which revealed a statistically significant positive correlation between mutation rate and profitability, ranging between ***F*** **=** 0 and ***F*** = 2.0. The bearing of strategy population ***k*** was found to be statistically insignificant with respect to profitability, after being investigated from ***k*** **=** 4 and ***k*** = 9. In a repeated set of trials, PRDE traders with higher mutation rates were found to consistently outperform those with lower rates, where the strategy population size was controlled at ***k*** = 4. Across both ranges of parameters, no statistically significant differences in PRDE's profitability were determined between stable, fluctuating, and volatile markets.

The creation of PRDJ; the JADE driven DE trader can be described as a success. Its implementation was shown with evidence to be correct. Although it did not outperform the strongest variants of PRDE, its profitability was greater than that of the default and weakest variants. This is expected behaviour for an unoptimized implementation of a 1-dimensional DE guided trader.

The results in this paper present various lines of future research. One such line would be to repeat this work, or produce a similar investigation, looking at the profitability of PRDE over a greater range of strategy population values. It has been shown in the experiment described in *Figure 13* that values for ***k*** greater than what was initially investigated may provide a statistically significant improvement in expected performance in volatile markets. This is supported by the fact that a positive correlation was found between profitability and size of ***k*** for all market conditions, but with insufficient statistical significance to formally make the claim. It is possible that an expanded set of trials could enable the claim to be made.

A second line of work would potentially be to investigate parameter sweep of strategy population ***k***, top population threshold ***p*****,** and crossover update constant ***c*** of PRDJ; comparing its performance in varying market conditions, with the aim of optimizing PRDJ's parameters. Furthermore, examining the performance impact of introducing a mutation factor intervention in converged populations could be considered.

Thirdly, an investigation into increasing the dimensionality of PRDJ's strategy is a possibility since the implementation of JADE supports the evolution of multiple strategy dimensions.

References

[1] R.Storn and K.Price, Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces, 1997

[2] D. Cliff, Minimal Simulation of a Limit-Order-Book Stock Exchange, 2018

[3] D. Cliff, Minimal-intelligence agents for bargaining behaviours in market based environments, 1997

[4] D. Cliff, Metapopulation Differential Co-Evolution of Trading Strategies in a Model Financial Market, 2022

[5] V.L. Smith, An Experimental Study of Competitive Market Behavior, 1962

[6] D.Cliff, Methods Matter: A Trading Agent with No Intelligence Routinely Outperforms AI-Based Traders, 2020

[7] T. Szabados, An elementary introduction to the Wiener process and stochastic integrals.

[8] R. C. Merton, Option pricing when underlying stock returns are discontinuous, 1976

[9] D. Snanshall, D. Cliff, Adaptive-Aggressive Traders Don't Dominate, 2019

[10] C. E. Shannon, A Mathematical Theory of Communication (pages 13-19) 1948

[11] J.Zhang, A. Sanderson, JADE: Self-adaptive differential evolution with fast and reliable convergence performance, 2007

[12] J.Zhang, A. Sanderson, JADE: Adaptive Differential Evolution With Optional External Archive, 2009