

Market fragmentation and market consolidation: Multiple steady states in systems of adaptive traders choosing where to trade

Aleksandra Alorić*

Scientific Computing Laboratory, Center for the Study of Complex Systems, Institute of Physics Belgrade, University of Belgrade, Pregrevica 118, 11080 Belgrade, Serbia

Peter Sollich

Institut für Theoretische Physik, Georg-August-Universität Göttingen, Friedrich-Hund-Platz 1, 37077 Göttingen, Germany and Disordered Systems Group, Department of Mathematics, King's College London, Strand, London WC2R 2LS, United Kingdom



(Received 19 February 2019; revised manuscript received 8 April 2019; published 21 June 2019)

Technological progress is leading to proliferation and diversification of trading venues, thus increasing the relevance of the long-standing question of market fragmentation versus consolidation. To address this issue quantitatively, we analyze systems of adaptive traders that choose where to trade based on their previous experience. We demonstrate that only based on aggregate parameters about trading venues, such as the demand-to-supply ratio, we can assess whether a population of traders will prefer fragmentation or specialization towards a single venue. We investigate what conditions lead to market fragmentation for populations with a long memory and analyze the stability and other properties of both fragmented and consolidated steady states. Finally, we investigate the dynamics of populations with finite memory; when this memory is long the true long-time steady states are consolidated but fragmented states are strongly metastable, dominating the behavior out to long times.

DOI: [10.1103/PhysRevE.99.062309](https://doi.org/10.1103/PhysRevE.99.062309)

I. INTRODUCTION

Whether a consolidated or a fragmented market is more beneficial to a population of traders is a long-standing debate [1–6]. In a consolidated or concentrated market, the majority of trades occurs in one (or a few) as opposed to numerous trading venues. With technological advances we have seen a proliferation of trading venues such as online marketplaces. Even more recently, alternative or dark trading venues have appeared, e.g., dark pools. These are popular not least for their lack of transparency, which makes them interesting for trading large quantities of shares without strongly influencing the price (see, e.g., [6,7]).

The emergence of collective behavior in systems of autonomous agents is a research topic that has seen widespread interest among physicists in the past couple of decades. The main reason for this is the recognition that statistical physics techniques, which contributed to the understanding of macroscopic phenomena arising in large systems of interacting microscopic entities, can be applied to a range of biological, economic, and social systems. A large body of work exists in the physics literature on collective effects in socioeconomic systems [8,9], e.g., mass movement of people [10,11], herd behavior of traders [12], and voting patterns [13,14]. One of the most prominent examples is the minority game, which continues to attract interest due to its simplicity and its ability to reproduce at least “stylized” facts about financial markets [15,16]; extensions of the model also predict interesting grouping phenomena when multiple assets are

available to agents [17]. In a similar vein, in this paper, we investigate whether fragmentation and consolidation can arise solely as a consequence of interactions at the level of the agents, combined with individual adaptation.

Some studies of stylized models of market competitions already exist, often pointing out the emergence of monopolies whereby the majority of trades occurs in one trading venue. Pagano [18] argued that when markets are identical (in terms of their transaction costs), risk averse traders will concentrate in a single market. On the other hand, when there is asymmetry, fragmentation might arise with traders being clustered based on the sizes of their desired transactions. Chowdhry and Nanda [2] reached the same conclusion in a system with asymmetrically informed traders and a general number of markets.

Ellison *et al.* [19] and Shi *et al.* [20] also studied competition among markets and the conditions under which such competition can lead either to monopolies or to coexistence of multiple markets. The authors named two significant effects in a competition of double auction trading venues. One of them is the positive size effect, i.e., agents prefer to trade in a market where there are already many traders of the opposite type. As an example, sellers like trading at markets where there are many buyers as this gives them a wider choice of offers. The authors of Refs. [19,20] also suggested the existence of a negative size effect in a double auction market: Agents will prefer being in the minority group of traders more often, with, e.g., buyers benefiting from trading at a market where there are not many other buyers (see, e.g., [21]). Ellison *et al.* [19] pointed out that such negative size effects can enable the coexistence of many markets. On the other hand, Shi *et al.* [20] investigated which of the two effects is stronger and found that

*aleksandra.aloric@gmail.com

due to more substantial positive effects, a monopoly will be the favored end state in many situations. The authors of [20] argued that market coexistence remains a possibility when there is strong market differentiation, especially for markets that have different pricing policies: One market might charge a fixed participation fee while another might take a profit fee. A common feature of the studies mentioned above is that some form of Nash equilibrium analysis was used, assuming perfect information about the activity of all traders and maximization of an underlying utility function for the agents.

The increased proportion of trades that take place in dark trading venues—15% of the U.S. share market volume was traded in dark pools in 2013 [22]—suggests increased fragmentation at least between traditional and dark trading. Calling for more research on reasons behind market fragmentation, Gomber *et al.* [23] suggested the heterogeneity of traders and their needs as one of the main drivers of market fragmentation. However, studies of similar effects such as the emergence of market loyalty in fish markets [24], herding [12], and grouping of agents in multiresource minority games [17] show that fragmentationlike phenomena may also be emergent. Nonetheless, existing models for such emergent fragmentation often assume a considerable amount of structure, e.g., in the connectivity among agents, the information available about the actions of other players, the rules of interaction via the market mechanism, the asymmetry between buyers and sellers, etc. In contrast, we study here a model in which initially homogeneous agents adapt only to their private information and show that even in such a system both market fragmentation and consolidation can occur depending on global system parameters.

Based on observations from the CAT tournament [25], where the spontaneous emergence of long-lived market loyalties was seen in complicated systems of adaptive markets and traders, we hypothesize that the reason for fragmentation may not lie in the intricacies of different market mechanisms or trading strategies. Instead, we conjecture that fragmentation is a collective phenomenon arising as a consequence of the continuous adaptation of the individual agents to an evolving system. To test this hypothesis, we developed a stylized model of double auctions and adaptive traders [26,27] that does indeed predict emergent fragmentation under minimal assumptions on the complexity of market and trading mechanisms. The model also shows market consolidation under some circumstances. Our focus in this study is to pin down under what conditions fragmentation and consolidation occur and what relative benefits they bring for the traders. As we will see, the behavior of the model is remarkably rich in spite of its simplicity, with multiple steady states coexisting in the limit of long agent memory. For finite memory length, this can lead to the existence of long-lived metastable states that dominate before the true steady state is reached eventually.

We start with a short description of the model [27] in Sec. II and then proceed to the large memory limit analysis of small systems with $N = 2$ and 4 agents in Sec. III. These can be thought of as two- and four-player games. They are convenient as we can easily track each trader's adaptation. At the same time they already reveal qualitative phenomena related to those we find later in large systems, in particular

coordination at the same market (for $N = 2$) and onset of fragmentation via pairwise coordination (for $N = 4$). Moving on to the large population limit ($N \rightarrow \infty$), we then first analyze a population with homogeneous buying preferences in Sec. IV A. We develop the relevant mathematical framework and techniques of analysis here and then generalize the results to systems with separate buyer and seller agent types. Finally, we study the system dynamics in some detail to go beyond the steady-state analysis in Sec. V.

Overall we follow a typical statistical physics philosophy in using a model that reduces the underlying market choice dynamics to its key ingredients, allowing us to obtain detailed insights into the origins of the resulting collective behavior. The analysis also relies significantly on statistical physics concepts and methods: We focus mostly on the thermodynamic limit of large agent populations, where we exploit the fact that the behavior of N interacting agents for $N \rightarrow \infty$ can be captured by the dynamics of a single agent subject to self-consistently determined population-level order parameters. The main outcome from this physical point of view is the emergence of multiple nontrivial steady states in the large interacting nonequilibrium systems that we study.

II. MODEL

Here we summarize basic assumptions and properties of the model introduced in [26,27], which is the foundation for the analysis in this paper.

Learning. In the model, agents choose among the available markets once in every trading period and submit their order to the chosen market. A key assumption is that agents base their decision of where to trade on their previous experience at the different markets. Agents rely on the following reinforcement rule, which is based on the experience-weighted attraction rule [28,29] but neglects knowledge about the other markets (via so-called fictitious payoffs):

$$A_m^i(n+1) = \begin{cases} (1-r)A_m^i(n) + rS_m^i(n) & \text{for } m \text{ chosen in round } n \\ (1-r)A_m^i(n) & \text{otherwise.} \end{cases} \quad (1)$$

Here $A_m^i(n+1)$ is agent i 's attraction to market m at trading period $n+1$ given the agent's score or return $S_m^i(n)$ obtained in the previous trading period (discussed below) and the previous attraction $A_m^i(n)$. To understand the role of r , one can write down the resulting general expression for the attraction at trading round n :

$$A_m^i(n) = \sum_{j=0}^{n-1} r(1-r)^{n-j} \delta_{m^i(j),m} S_m^i(j) + (1-r)^n A_m^i(0),$$

where the Kronecker δ restricts updates to rounds where the agent's chosen market $m^i(j)$ is the one (m) being considered. The factor $r(1-r)^{n-j}$ in this expression is a weight that decays exponentially into the past, becoming small once $n-j$ is of order $1/r$. Thus each agent effectively averages scores over a sliding window into the past of length approximately equal to $1/r$, so $1/r$ can be thought of as setting the length of the agents' memory.

To choose a market at each trading round, an agent translates the learned attractions into probabilities of choosing each markets, using the multinomial logit or softmax function

$$P_m^i(n) = \frac{\exp[\beta A_m^i(n)]}{\sum_{m'} \exp[\beta A_{m'}^i(n)]}. \quad (2)$$

This aspect of the model is also in line with the experience-weighted attraction literature [28,29]; β is the intensity of choice and regulates how strongly the agents bias their preferences towards actions with high attractions. For $\beta \rightarrow \infty$ the agents choose the option with the highest attraction, while for $\beta \rightarrow 0$ they choose randomly and with equal probabilities among all options.

We study agents whose choice of the type of trading order (to buy or to sell) is not adaptive but rather set by a fixed buying preferences p_B^i . This assumption simplifies the analysis while still allowing both consolidation and fragmentation behavior as shown previously [27].

Trading strategies. Agents do not have sophisticated trading strategies in our model and are essentially zero-intelligence traders [30–32]. Their orders to buy (bid) or sell (ask) a single unit of the underlying good at a certain price are independent of previous returns or other information. We assume specifically that bids b and asks a are normally distributed as $a \sim \mathcal{N}(\mu_a, \sigma_a^2)$ and $b \sim \mathcal{N}(\mu_b, \sigma_b^2)$, where we fix $\mu_b > \mu_a$ as in [27]. After each round of trading, each agent receives a score, reflecting their payoff in the trade. This depends on the global trading price set by a chosen market m as well as the order the agent has submitted. The scores of agents who do trade are assigned as in previous studies [30,33]: buyers value paying less than they offered (b) and so their score is $S = b - \pi$. Sellers value trading for more than their ask a and so $S = \pi - a$ is a reasonable model for their payoff; in both cases π is the trading price.

Market mechanism. In the spirit of keeping the model as simple as possible we consider double auction markets in discrete time, counted as before in trading rounds. In every round the global trading price is set by the market: Once all orders have arrived, these are used to determine the average bid $\langle b \rangle$ and average ask $\langle a \rangle$ and set the price

$$\pi = \langle a \rangle + \theta(\langle b \rangle - \langle a \rangle), \quad (3)$$

where θ fixes the price closer to the average bid ($\theta > 0.5$) or the average ask ($\theta < 0.5$), as in [26]. This parameter thus represents the bias of the market towards buyers or sellers. Once the trading price has been set, all bids below this price, and all asks above it, are marked as invalid orders as they cannot be executed at the current trading price. The remaining orders are executed by randomly pairing buyers and sellers. Excess buyers or sellers, i.e., those that cannot be paired, receive zero score, as do the agents who submitted invalid orders.

Note that traders are not informed about the market biases, or the market mechanism in general. The only information they have at their disposal to adapt their market preferences is their personal score.

III. FINITE N

A. Two traders: Coordination

To understand collective effects in trading systems, we first build up some intuition by looking at a very simple model with only one buyer and one seller. The traders have a choice between two markets with different biases. As the system consists of only two agents and two markets, fragmentation (or segregation as introduced previously [27]), in which a population will split into distinctive groups favoring one option, is not feasible. However, we can investigate if long-lasting loyalty to a single market emerges, which can signal market consolidation.

To make trading possible the two agents effectively need to coordinate, i.e., to submit orders to the same market. This can lead to one of the agents earning less than they could have done at the other market. One question of interest concerns the conditions under which the agents prefer random decisions of who will be a winner or loser in this manner, as opposed to settling in these roles over longer periods of time. Thus we will focus on the existence of coordination of traders and investigate for which parameter settings agents develop strong preferences for the same market. Intriguingly, this two-player analysis ends up being largely similar to the work by Hanaki *et al.* [34], where a two-agent case was likewise studied as a first step to understanding collective effects. (In [34] these concerned specialization behavior of agents searching for parking spots.)

For the $N = 2$ analysis it is convenient to label the two players as $i = \pm 1$ and similarly for the two markets. We use the following specific parameter settings.

(i) Of the two players, player $i = 1$ always buys while player $i = -1$ always sells ($p_B^1 = 1$ and $p_B^{-1} = 0$).

(ii) Bids and asks are deterministic, i.e., $b \sim \mathcal{N}(\mu_b, 0)$ and $a \sim \mathcal{N}(\mu_a, 0)$, with their difference being fixed to $\mu_b - \mu_a = 1$.

(iii) The trading price at each market is set as defined in [26], $\pi_m = \langle a \rangle + \theta_m(\langle b \rangle - \langle a \rangle)$.

(iv) We assume that the market biases are symmetric, $(\theta_1, \theta_{-1}) = (\theta, 1 - \theta)$, where $\theta \in [0, 0.5]$.

The simplification over our previous work [26,27] of making bids and asks deterministic allows us to focus solely on the coordination of the market choices and does not change the behavior of the system qualitatively. The deterministic order prices then also make the trading prices deterministic: $\pi_m = \mu_a + \theta_m(\mu_b - \mu_a) = \mu_a + \theta_m$.

We can summarize the attraction update rule (1) as

$$A_m^i(n+1) = (1-r)A_m^i(n) + rS_m^i(n),$$

with the convention that $S_m^i(n) = 0$ if market m was not chosen by agent i in round n . This generalized score is fully determined by the market choice of the opposite player

$$S_m^i(n) = \delta_{m^{(-)i}(n), m} \delta_{m^{-i}(n), m} \Sigma_m^i, \quad (4)$$

where $m^{(-)i}(n)$ denotes the market of choice of the (co-)player $(-i)$ during trade n and

$$\Sigma_m^i = \begin{cases} \mu_b - \pi_m = 1 - \theta_m, & i = 1 \\ \pi_m - \mu_a = \theta_m, & i = -1 \end{cases} \quad (5)$$

encodes the relevant nonzero score values that depend on the type of market and agent. The logit assignment (2) by which

agents choose a market m simplifies for $N = 2$ to

$$P_m^i(n) = \frac{1}{1 + \exp[-\beta m \Delta^i(n)]} = \sigma_\beta(m \Delta^i(n)),$$

where $\sigma_\beta(z) = [1 + \exp(-\beta z)]^{-1}$ is the logistic sigmoid. The choice probabilities do not depend on the attractions to the two markets individually but only on their difference $\Delta^i = A_1^i - A_{-1}^i$. The latter is updated as

$$\begin{aligned} \Delta^i(n+1) &= A_1^i(n+1) - A_{-1}^i(n+1) \\ &= r S_1^i(n) + (1-r) A_1^i(n) \\ &\quad - [r S_{-1}^i(n) + (1-r) A_{-1}^i(n)]. \end{aligned}$$

The stochastic variable $\Delta^i(n+1)$ thus depends on the choices the agents make in trading round n , $m^i(n)$ and $m^{-i}(n)$, which are drawn from distributions that depend on $\Delta^i(n)$ and $\Delta^{-i}(n)$. This situation simplifies in the long memory limit $r \rightarrow 0$, where the attraction differences change sufficiently slowly to average out stochastic fluctuations. One can then effectively replace $\delta_{m^i(n),1}$ by its expected value $\sigma_\beta(\Delta^i(n))$ (and similarly for $-i$ and other market choices) in the score (4). This gives

$$\begin{aligned} \Delta^i(n+1) &= r [\sigma_\beta(\Delta^i(n)) \sigma_\beta(\Delta^{-i}(n)) \Sigma_1^i \\ &\quad - \sigma_\beta(-\Delta^i(n)) \sigma_\beta(-\Delta^{-i}(n)) \Sigma_{-1}^i] \\ &\quad + (1-r) \Delta^i(n), \end{aligned}$$

which can be further simplified into

$$\begin{aligned} \frac{\Delta^i(n+1) - \Delta^i(n)}{r} \\ &= -\Delta^i(n) + [\sigma_\beta(\Delta^i(n)) \sigma_\beta(\Delta^{-i}(n)) \Sigma_1^i \\ &\quad - \sigma_\beta(-\Delta^i(n)) \sigma_\beta(-\Delta^{-i}(n)) \Sigma_{-1}^i]. \end{aligned}$$

The finite difference on the left-hand side becomes a derivative in the limit of small r if we switch to the rescaled time $t = nr$, for which a unit time interval corresponds to $1/r$ trading periods:

$$\begin{aligned} \partial_t \Delta^i(t) &= -\Delta^i(t) + [\sigma_\beta(\Delta^i(t)) \sigma_\beta(\Delta^{-i}(t)) \Sigma_1^i \\ &\quad - \sigma_\beta(-\Delta^i(t)) \sigma_\beta(-\Delta^{-i}(t)) \Sigma_{-1}^i]. \end{aligned}$$

A convenient change in variables that simplifies this pair of differential equations is $\Delta^1(t) = \xi(t) + \rho(t)$ and $\Delta^{-1}(t) = \xi(t) - \rho(t)$, which after some algebra and exploiting the market symmetry gives

$$\begin{aligned} \partial_t \xi(t) &= -\xi(t) + \frac{1}{2} \frac{\sinh[\beta \xi(t)]}{\cosh[\beta \xi(t)] + \cosh[\beta \rho(t)]}, \\ \partial_t \rho(t) &= -\rho(t) + \frac{1-2\theta}{2} \frac{\cosh[\beta \xi(t)]}{\cosh[\beta \xi(t)] + \cosh[\beta \rho(t)]}. \end{aligned} \quad (6)$$

Note that $\xi = (\Delta^1 + \Delta^{-1})/2$ describes the average of the attraction differences of the two agents, while $\rho = (\Delta^1 - \Delta^{-1})/2$ captures the deviation between them.

To understand the dynamics, we first consider its fixed points, which need to satisfy

$$\begin{aligned} \xi^* &= \frac{1}{2} \frac{\sinh(\beta \xi^*)}{\cosh(\beta \xi^*) + \cosh(\beta \rho^*)}, \\ \rho^* &= \frac{1-2\theta}{2} \frac{\cosh(\beta \xi^*)}{\cosh(\beta \xi^*) + \cosh(\beta \rho^*)}. \end{aligned} \quad (7)$$




FIG. 1. Two-trader dynamics: flow diagrams (6) for (a) the intensity of choice $\beta = 2$, with a unique fixed point where agents decide largely randomly, and (b) $\beta = 6$, with two new fixed points indicating where coordinated states appear. For the market bias used, $\theta = 0.3$, the critical intensity of choice where coordinated states emerge is $\beta_c = 4.16$.

The first of these equations is always satisfied if $\xi^* = 0$, and in that case the equation for ρ^* has a unique solution whose sign depends on the sign of $1 - 2\theta$. When market 1 is favorable towards buyers ($\theta < 0.5$), ρ^* will be positive. As $\Delta^{\pm 1} = \pm \rho^*$, this can be interpreted as a state where buyers and sellers learn which market is good for them and thus have preferences for opposite markets. (Here Δ^1 is positive, meaning that player 1, the buyer, prefers market 1, which is good for buyers.) As we will see shortly, this solution is only stable for low intensities of choice where the agents' market choice dynamics remains largely random. The intuition for the appearance of an instability with increasing β is that, if agents were to follow through fully on their attractions towards opposite markets, they would never get to trade.

The stability of the solution ($\xi^* = 0, \rho^*$) can be studied by linearizing the dynamical equations (6), resulting in the stability criterion

$$\frac{\beta}{2} \frac{1}{1 + \cosh(\beta \rho^*)} \leq 1.$$

Expressed in the original variables Δ^i , the solution with $\Delta^{1*} + \Delta^{-1*} = 0$ is stable as long as

$$\frac{\beta}{2} \frac{1}{1 + \cosh[\beta(\Delta^{1*} - \Delta^{-1*})/2]} \leq 1. \quad (8)$$

This stability condition is exactly the same as in Ref. [34] because the learning dynamics we follow is essentially the same and differs only in the details of the deterministic returns.

We illustrate in Fig. 1 that for low intensities of choice, where the stability criterion (8) is satisfied, the fixed point discussed so far is the only one. At higher β the criterion is violated and two new stable fixed points appear. Here the agents' attraction differences are of the same sign, i.e., they prefer going to the same market. This happens even though market 1 favors buyers while market -1 favors sellers.

At first sight it may seem puzzling that for high intensity of choice, one of the agents decides to settle for less in persistently choosing the market where the agent will be awarded lower scores. However, this pattern of behavior in fact maximizes the number of trades that take place. In the




FIG. 2. Two traders: the (θ, β) phase diagram and returns. (a) Coordination and indecisiveness regions for different intensity of choice and market biases (β and θ). (b) Returns for different β in a system with two fair markets $\theta = 0.5$. (c) Returns for different β for market bias $\theta = 0.3$. At the critical $\beta_c = 4.16$, the average return of the two agents in the coordinated state is higher than it would be in the continuation of the low- β fixed point (yellow dashed line), but one of the agents needs to settle for less.

low- β regime, all four pairings of market choices are equally probable, $(m^1, m^{-1}) \in \{(1, 1), (1, -1), (-1, 1), (-1, -1)\}$, but only the first and the last enable trading. On the other hand, in the high- β regime, when both agents persistently choose the same market, they always get to trade, although one of the traders always receives a lower return. For the market parameters used in Fig. 1, $(\theta_1, \theta_{-1}) = (0.3, 0.7)$, the agent who settles for a lower score then receives a score of 0.3, while the other one obtains 0.7. This has to be compared to the average payoff at low β , which by averaging over the four market choice pairings is seen to be $\frac{1}{4}(\theta_1 + \theta_{-1}) = \frac{1}{4}$. Hence both agents clearly earn more in the coordinated regime than by choosing randomly.

We can find the domain of parameters θ and β where the agents will coordinate [Fig. 2(a)] by starting from the regime of agents choosing largely randomly and tracking where the stability condition (8) is first violated as β is increased. As in the case of large populations [27], we observe that β_c increases with increased market difference or bias. The symmetry breaking between markets that coordination requires is therefore not driven by the difference between the markets. In fact, the coordination threshold is lowest for a system with two identical markets ($\theta = 0.5$). One can rationalize this by saying that the agents are happiest to coordinate at one of the markets in this limit as neither needs to settle for less. We show average returns for this setup—a pair of traders choosing between two unbiased markets—as a function of β in Fig. 2(b). One observes the expected average score of $1/4$ for low β ; as β is increased, the agents effectively realize that coordination at a single market enables more trades and consequently higher average returns.

In Fig. 2(c) we show analogous results for the case of two biased markets ($\theta = 0.3$). We plot the individual agents' payoffs and their average in the state where they coordinate at one market and compare this to the payoff in the largely random low- β state. As a reference we also plot the continuation of the latter to larger β , where it is unstable. It is notable that returns decrease with β on this branch: The more the agents act on their preference for opposite markets, the less often they manage to meet at the same market. This results in more and more trading rounds where both receive a return of zero, dragging down average returns.

By contrast, in the coordinated state the average return increases with β , i.e., as the agents make more and more definite choices. Interestingly, Fig. 2(c) shows that this increase in the average return is accompanied by a growing difference between the returns of the individual agents. These payoff differences can occur in our model because agents are unaware of the opposite player's return, making decisions only on the basis of their own scores. Borrowing terminology from the large system limit [27], we will refer to the agent with the higher return as return driven and the other as volume driven. It is notable in Fig. 2(c) that there is a range of β where the volume-driven agent receives an average return that is lower than not only that of the return-oriented agent, but also the hypothetical return both agents would achieve in the (unstable) uncoordinated state; this regime grows as the markets become more biased.

Intuitively, the return-driven player develops a strong preference for the market where the player can earn more. The other agent will occasionally try the other market, but typically not get to trade there. As this results in a zero return, the player is better off persisting with the coordinated choice, which offers a low but at least nonzero return.

The two-agent systems studied so far can be mapped to two-player games: the symmetric pure coordination game when the markets are unbiased and the battle of the sexes when markets are symmetrically biased. For these games it is known that the two coordinated states correspond to pure Nash equilibria (see, e.g., [35]). In the symmetric pure coordination game, both of these are envy-free (i.e., both agents earn the same), but not so in the battle of the sexes; this is consistent with the differences we saw between unbiased and biased markets, and the Nash equilibria correspond to the $\beta \rightarrow \infty$ limit of the coordinated states. There are also mixed Nash equilibria. These correspond to the continuation to $\beta \rightarrow \infty$ of our uncoordinated state for the symmetric pure coordination game, but not otherwise. A full correspondence to Nash equilibria could be obtained by modifying the learning rule so that the attractions to markets that were not chosen are kept unchanged. This can be interpreted as fictitious play and is discussed in more detail in [36].

The results described above can be generalized to a pair of traders who do not have strict buyer and seller roles but instead decide to buy with some probability. We assume symmetric preferences for buying, $p_B^1 = 1 - p_B^{-1} = p_B$. For a trade to occur, agents now need to be at the same market and need to submit opposite (buy and sell) orders. As the buying preferences p_B^i are fixed, this only changes Σ_m^i from (5) to

$$\Sigma_m^i = p_B^i(1 - p_B^{-i})(1 - \theta_m) + (1 - p_B^i)p_B^{-i}\theta_m. \quad (9)$$




FIG. 3. Two traders: coordination threshold as a function of θ and p_B . Note that the threshold is finite for all system parameters and increases the more similar the agents become, i.e., as p_B decreases towards 0.5.

To see this, note that agent i receives a buyer payoff $1 - \theta_m$ when the agent assumes the role of a buyer (with prob p_B^i) while the opposite player acts as seller; agent i also receives a seller payoff in the opposite configuration. Repeating the calculation above, one then finds that the fixed point conditions (7) both acquire a factor of $p_B^2 + (1 - p_B)^2$ on the right-hand side, while the stability condition (8) is multiplied by the same factor. Figure 3 shows contours of the resulting critical β_c for coordination. We note that the coordination threshold increases as p_B approaches $1/2$: Agents without strong buy and sell preferences need higher intensities of choice to benefit from the coordinated state. This makes sense because agents with p_B closer to $1/2$ derive a lower benefit from coordinating at a market: As they need to assume buyer and seller roles, trades at the same market happen only with some probability, specifically $p_B^2 + (1 - p_B)^2$ in our setting, which approaches $1/2$ for $p_B \rightarrow 1/2$.

B. Four traders: Onset of fragmentation

The two-player system studied above already exhibited an interesting collective phenomenon: coordination at a market to enable more trades, sometimes even to the detriment of an individual agent. Turning to fragmentation, where otherwise homogeneous agents nonetheless learn to adopt different policies, the minimal system size where we can expect a similar effect is $N = 4$. We first study two identical buyers and two identical sellers, choosing agents with deterministic buy and sell behavior ($p_B^i = 0$ or 1) for simplicity. A system with four agents is small enough so that we can still easily write down deterministic equations for the evolution of market attractions, but large enough for the first signals of fragmented states to appear as agents can split across the markets in pairs. We consider again symmetrically biased markets, $\theta_1 = 1 - \theta_{-1} = \theta$. As before, the market choice behavior of each agent is determined by their market attraction difference $\Delta^{g,i}$.

Here the index g denotes the agent group (buyers or sellers), while i labels agents within each group. For small r the attraction differences again obey deterministic time evolution equations that can be derived by following the reasoning in the preceding section. The only difference lies in the calculation of the return $S_m^{g,i}(n)$ at a chosen market, which now depends on the choices made by all other players:

$$S_m^{g,i}(n) = \delta_{m^{g,i}(n),m} \left\{ \frac{\Sigma_m^g}{2} \delta_{m^{g,-i}(n),m} (\delta_{m^{-g,i}(n),m} + \delta_{m^{-g,-i}(n),m}) + \Sigma_m^g (1 - \delta_{m^{g,-i}(n),m}) (\delta_{m^{-g,i}(n),m} + \delta_{m^{-g,-i}(n),m}) - \delta_{m^{-g,i}(n),m} \delta_{m^{-g,-i}(n),m}) \right\},$$

In this expression, Σ_m^g denotes the deterministic part of the return, which only depends on the chosen market m and the agent type g , by analogy with the two-player case in Eq. (9). The Kronecker δ symbols ensure that other agents are present at the same market m . The first term describes the situation where both agents of the same type go to a single market m ; the return is then zero if no agents of the opposite type are at the same market, Σ_m^g if there are two, and $\Sigma_m^g/2$ if there is only one (as our chosen agent then only has probability $1/2$ of being allowed to trade). On the other hand, when the second player of the same type is not at the same market, the player receives the full return if there is at least one trader of the opposite group present. This is described by the second term. The deterministic equations for $r \rightarrow 0$ then take the form

$$\partial_t \Delta^{g,i}(t) = -\Delta^{g,i}(t) + \sum_{m=-1}^1 m S_m^{g,i}(t),$$

where $S_m^{g,i}(t)$ has the meaning of returns averaged over a long time window so that the Kronecker δ 's in $S_m^{g,i}(n)$ are replaced with their expected values, exactly as in the derivation for two players. We solve these equations numerically and find that for low and intermediate intensity of choice β the behavior is analogous to that for $N = 2$, showing a transition from a single uncoordinated fixed point to two coordinated fixed points as β increases; throughout this range the agents within each group have identical market attractions. The novel feature of the $N = 4$ system is that, when β is increased yet further, four new stable states appear. We call these fragmented as each group of agents now “fragments” into two individuals with distinct, and essentially opposite, market preferences. Both markets are then populated by a pair of traders, one from each group. The fragmented fixed points appear in pairs (stable and unstable fixed point) and for high enough value of β unstable fragmented fixed points become partially fragmented, e.g., only one group splits across the markets, while the other group specialize for one market. As these fixed points are not stable we do not show this transition line in Fig. 4.

In Fig. 4 we show the two critical β lines (the coordination and the fragmentation threshold) as a function of the market bias θ , for the above scenario of four players with strict buy and sell roles. The coordination line is very close to the one for two players, which is included for comparison. Both the coordination and fragmentation lines follow the same trend, with the threshold in β increasing as θ departs from 0.5.

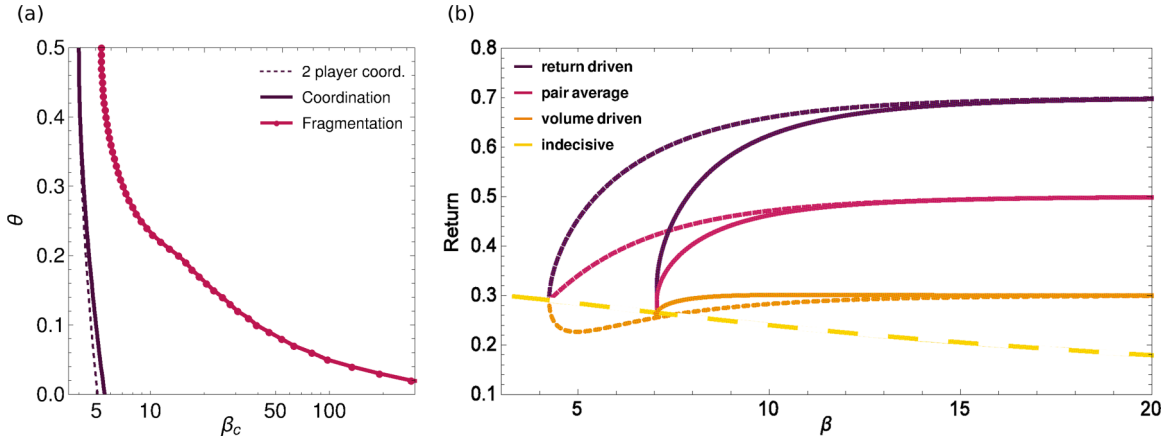


FIG. 4. Four agents (two buyers and two sellers): phase diagram and returns. (a) Phase diagram showing steady states as a function of intensity of choice β and market bias θ . Coordinated steady states exist to the right of the dark violet, solid line and fragmented steady states to the right of the pink solid line with markers. (b) Returns against intensity of choice β for all agents and separately for return- and volume-driven agents; market biases are $(\theta_1, \theta_{-1}) = (0.3, 0.7)$. Dashed lines show coordinated states and solid lines fragmented states. The yellow dashed line shows the average return in the uncoordinated steady state, continued into the instability region at high β .

Figure 4(b) shows return lines for different intensities of choice β : Dashed lines correspond to coordinated states, while solid lines are averages in the fragmented state. Note that the difference in returns in the coordinate state is between groups of agents, with all agents in a group either return driven or volume driven. In the fragmented states there is one return-driven and one volume-driven agent in each group, on the other hand. We note that in the large- β limit the returns achieved in coordinated and fragmented states become identical. This is because with either pattern of market choices, if these choices are made deterministically, then all agents are guaranteed to be able to trade. For finite β , returns in the coordinated state are generally higher than for fragmentation.

Note that the four fragmented states arise because in each agent group there are two ways to assign the two agents to the two markets. For N agents, one therefore expects $\{(N/2)! / [(N/4)!]^2\}^2$ such states. This number grows very rapidly with N , while the number of coordinated states remains at 2.

Finally, as in the analysis of the two-agent system, we can generalize the results by allowing agents to assume the role of buyer with some group-dependent probability $p_B^{(g)}$. We again take these probabilities as symmetric between groups, $p_B^1 = 1 - p_B^{-1} = p_B$. The deterministic part Σ_m^g of the agents' returns is then modified in a manner directly analogous to Eq. (9), and one can determine the effect on the existence of the various steady states. Figure 5 shows the results for the symmetric markets $(\theta_1, \theta_{-1}) = (0.3, 0.7)$ and symmetric groups. As in the system with only two agents, when the traders' preferences for buying are similar ($p_B \approx 1/2$) they have a weaker incentive to coordinate, resulting in a higher coordination threshold for β (for the sake of clearer visualization we use $1/\beta$ on the y axis). The same behavior is seen also for the fragmentation threshold.

We indicate in Fig. 5 also the regime where a further type of fixed point exists: partially fragmented states. In these states there is a single agent whose market preference is the opposite

of that of the other three players, so only one agent group is fragmented. These states evolve for high enough β out of unstable fragmented states, which themselves appear in pairs with the stable fragmented states at the onset of fragmentation. As we will see below, partially fragmented states exist in the large population limit too, though in a limited region of parameter space. In the small system here they are unstable. Intuitively this is likely to be due to the smaller number of trades: In the large- β limit of a partially fragmented fixed point, there will be at most one trade per trading period (only two out of the three agents going to one market will be able to trade), while both fragmented and coordinated states lead to two trades.




FIG. 5. Four traders: phase diagram when buy and sell roles are probabilistic. Coordination takes place below the dark violet (solid) line and fragmentation below the pink line (solid line with circles). These regions shrink when the difference in the buy and sell preferences of the agents decreases ($p_B \rightarrow 0.5$), similarly to the trend in the two-player system (dashed line). Below the orange line partially fragmented fixed points exist, where one of the four agents has a preference for the opposite market.

IV. LARGE POPULATION LIMIT

Population with a fixed buying preference

In studying systems with a small number of agents we have already encountered a rich phenomenology: coordination of agents at a single market, pairwise fragmentation across two markets, and even some mixed states where one group fragments while the other specializes in trading at a single market. We now complement and generalize these results by investigating the possible types of steady state in the large population limit. We start with a simple setting, a population in which all agents have identical preferences for buying $p_B^i = p_B \forall i$. The assumption of population homogeneity is a strong one, but these traders still undergo fragmentation for a broad range of parameters, while the system is simpler to analyze. Thus it is a useful prelude to the analysis of a population consisting of two or more groups with distinct buying preferences.

To describe the steady states of such an initially homogeneous agent population we follow the distribution of attraction differences ($\Delta^i = A_1^i - A_2^i$) across the population [27]. The state of each market m enters via the probability with which buy (B) and sell (S) orders are executed successfully [27,37]:

$$\begin{aligned} T_{Bm} &= \min\left(1, \frac{Q_{Sm}}{Q_{Bm}D_m}\right), \\ T_{Sm} &= \min\left(1, \frac{Q_{Bm}D_m}{Q_{Sm}}\right). \end{aligned} \quad (10)$$

Here the factors $Q_{B,S,m}$ are the probabilities for submitted buy or sell orders to be valid, i.e., on the right side of the trading price (the explicit expressions are given in Appendix A). Note that, whereas for small systems we simplified to deterministic order prices, we return here to the full model where bids b and asks a are stochastic and the trading price is calculated as in Eq. (3). [As explained in Sec. II, we choose Gaussian distributions for bids and asks, $a \sim \mathcal{N}(\mu_a, \sigma_a^2)$ and $b \sim \mathcal{N}(\mu_b, \sigma_b^2)$; for numerical evaluations we set $\mu_b - \mu_a = 1$ and $\sigma_a = \sigma_b = 1$.] The D_m are demand-to-supply ratios, defined as the number of buyers over the number of sellers at market m . For small r , the attraction difference distribution evolves according to a Fokker-Planck equation

$$\begin{aligned} \partial_t P(\Delta|p_B, T_\gamma) &= -\partial_\Delta [M_1(\Delta|p_B, T_\gamma)P(\Delta|p_B, T_\gamma)] \\ &+ \frac{r}{2} \partial_\Delta^2 [M_2(\Delta|p_B, T_\gamma)P(\Delta|p_B, T_\gamma)], \end{aligned} \quad (11)$$

where the drift M_1 and diffusion M_2 both depend on the buying preference of the agents and on the four trading probabilities T_γ . [We use $\gamma = (\tau, m)$ as the generic label for a combination of order type $\tau = B, S$ and market choice m .] The drift term is (see Appendix A for details and for the explicit expression of the return distribution from which $\langle S_\gamma \rangle$ is calculated)

$$M_1(\Delta|p_B, T_\gamma) = \sum_{m=-1}^1 \sum_{\tau \in \{B,S\}} m p_\tau T_{\tau m} \langle S_{\tau m} \rangle \sigma_\beta(m\Delta) - \Delta, \quad (12)$$

where the sum runs over markets m and order types τ and we use the convention $p_S = 1 - p_B$. The strength of the diffusion




FIG. 6. Critical intensities of choice as a function of market biases for (a) an indecisive population ($p_B = 0.5$) and (b) a population of decisive buyers ($p_B = 0.8$).

term is

$$\begin{aligned} M_2(\Delta|p_B, T_\gamma) &= \Delta^2 + \sum_{m=-1}^1 \sum_{\tau \in \{B,S\}} [p_\tau T_{\tau m} (\langle S_{\tau m}^2 \rangle \\ &- 2m\Delta \langle S_{\tau m} \rangle)] \sigma_\beta(m\Delta). \end{aligned} \quad (13)$$

The steady state of the Fokker-Planck equation is (see, e.g., [38])

$$P(\Delta|p_B, T_\gamma) \propto \frac{1}{M_2(\Delta|p_B, T_\gamma)} \exp\left(-\frac{f(\Delta)}{r}\right), \quad (14)$$

where

$$f(\Delta) = -2 \int_0^\Delta d\Delta' \frac{M_1(\Delta'|p_B, D_m)}{M_2(\Delta'|p_B, D_m)} \quad (15)$$

plays a role analogous to a free energy in thermodynamics. When $f(\Delta)$ has a single minimum, $P(\Delta)$ will approach a narrow peak at this location for $r \rightarrow 0$ and we have an unfragmented state. Otherwise, as many peaks as there are local minima in $f(\Delta)$ will appear, corresponding to a fragmented state: Each peak represents a subgroup of agents following a distinct market choice strategy.

Note that in the Fokker-Planck description, the market order parameters D_m that determine the trading probabilities T_γ have to be calculated self-consistently from $P(\Delta)$ [27,37]. The same self-consistency condition then also needs to hold at a steady state. Initially we will treat the order parameters as fixed exogenously however. Such a situation could arise if, for example, our agents are just a very small fraction of the overall trading cohort, with the latter fixing the demand-to-supply ratio.

Fragmentation for $r \rightarrow 0$

In Fig. 6 we show how the threshold value of the intensity of choice depends on the market biases (θ_1, θ_{-1}) for different agent populations, one indecisive ($p_B = 0.5$) and one made up of decisive buyers ($p_B = 0.8$); for this calculation we set the order parameters to their endogenous value following the self-consistent procedure outlined in (19). We see that for every pair of market biases there is a finite threshold β_c above which fragmentation sets in. When agents are indecisive with regard to buying and selling, the region where fragmentation occurs is greatest when markets are identical or symmetrically biased. For




FIG. 7. Steady-state attraction difference distributions of decisive buyers ($p_B = 0.8$). We compare steady states at $\beta = 20$ for (a) two unbiased markets $(\theta_1, \theta_{-1}) = (0.5, 0.5)$ and (b) two symmetrically biased markets $(\theta_1, \theta_{-1}) = (0.3, 0.7)$, for $r = 0.1$ (dashed dark violet line) and 0.05 (solid pink line). The distributions are strongly and weakly fragmented, respectively: on the right, the relative height of the lower peak decreases as r is reduced.

decisive buyers ($p_B = 0.8$), on the other hand, the fragmentation threshold is lowest when the markets are identical. Intermediate values of p_B provide a smooth interpolation between these two situations.

To understand more closely the properties of the fragmented states, we show in Fig. 7 the steady-state distributions of traders with $p_B = 0.8$ when faced with the choice between two unbiased or two symmetrically biased markets. To understand the trend with r , we show the distributions for two different values of r in each case. As expected from (14), the peak width decreases as approximately \sqrt{r} with decreasing r , but in Fig. 7(b) we see that the relative peak heights also depend on r . In fact, if the peaks are located at attraction differences Δ_1 and Δ_2 , then the peak height ratio can be written as

$$\frac{P(\Delta_1 | p_B, T_\gamma)}{P(\Delta_2 | p_B, T_\gamma)} = \frac{M_2(\Delta_2 | p_B, T_\gamma)}{M_2(\Delta_1 | p_B, T_\gamma)} \exp\left(-\frac{f(\Delta_2) - f(\Delta_1)}{r}\right). \quad (16)$$

This ratio can stay finite for $r \rightarrow 0$ only when

$$f(\Delta_1) = f(\Delta_2), \quad (17)$$

and we call this situation strong fragmentation as it survives even in the $r \rightarrow 0$ limit. This is the situation in Fig. 7(a). If the free energies at the two peaks are unequal, on the other hand, one continues to have two peaks in $P(\Delta)$ for any nonzero r but the height of one peak decreases (exponentially in $1/r$) as r goes to zero. We call this behavior, which is illustrated in Fig. 7(b), weak fragmentation because the lower peak may become unobservably small for low r ; in the strict limit $r \rightarrow 0$, the distribution $P(\Delta)$ becomes unimodal again. The strong-weak distinction as defined applies literally only to this $r \rightarrow 0$ limit; at nonzero r it becomes a crossover between fragmented states where the emergent subgroups have roughly even (strong) or very different (weak) sizes. At the weakly fragmented state most of the trades happen at a single market (increasingly so as r decreases); we relate this state to market consolidation, and thus the question of fragmentation versus

consolidation becomes a question of strong versus weak fragmentation in our setup.

Now that we have a method for finding steady states and classifying them, we return to the space of market order parameters and investigate where fragmentation occurs. In Fig. 8 we show where weakly (colored regions) and strongly (solid lines inside these regions) fragmented states appear, at a fixed intensity of choice $\beta = 8.5$. We compare again indecisive ($p_B = 0.5$) and decisive buyers ($p_B = 0.8$), for three different market setups. We first note that the weak fragmentation region encompasses a very wide range of market conditions (order parameters D_m) for indecisive buyers, but shrinks significantly when the agents have stronger preferences for buying. Looking at the dependence on market setup, an obvious feature is that for two unbiased markets [shown in Fig. 8(c)], equal demand-to-supply ratios ($D_1 = D_{-1}$ line) produce strong fragmentation for both types of agents. This makes sense as the markets are then identical both in their setup $\theta_1 = \theta_{-1}$ and in the prevailing market conditions, making it easy for groups of agents with opposite market preferences to coexist. For the indecisive agents who will act as buyers or sellers with equal probability, the same situation arises when the markets have exactly opposite demand-to-supply ratios ($D_1 = 1/D_{-1}$) and therefore still offer them identical average returns.

With increasing market biases [Figs. 8(a) and 8(b)] the picture obtained for two unbiased markets changes largely smoothly, though note that for decisive buyers (top row) the two crossing lines of strong fragmentation detach into two separate lines [Fig. 8(b), top], with one eventually disappearing out of range.

Market order parameters. So far we have looked at fragmentation behavior driven by exogenously set market conditions (demand-to-supply ratios). We now return to our model as originally set out, where only the adaptive agents we describe trade at the two markets. This leads to the following question: Will a population endogenously create market conditions needed for its fragmentation?

For the case of traders with homogeneous buying preferences p_B this question can be answered relatively straightforwardly. If the steady-state distribution of attraction differences is $P(\Delta | p_B, T_\gamma)$, then the fractions of the whole population buying and selling at market m are

$$N_{Bm} = p_B \int d\Delta P(\Delta | p_B, T_\gamma) \sigma_\beta(m\Delta),$$

$$N_{Sm} = (1 - p_B) \int d\Delta P(\Delta | p_B, T_\gamma) \sigma_\beta(m\Delta). \quad (18)$$

The demand-to-supply ratio then does not in fact depend on the market preference distribution

$$D_m = \frac{N_{Bm}}{N_{Sm}} = \frac{p_B}{1 - p_B} \quad (19)$$

and is fully determined by p_B . In the space of market order parameters in Fig. 8, this endogenous set of market conditions is marked with a black dot. We see that, at high enough β , the population of indecisive buyers fragments strongly when the markets are unbiased or symmetrically biased, and one can check that these results hold independently of the specific market biases used in the figure. Decisive buyers, on the

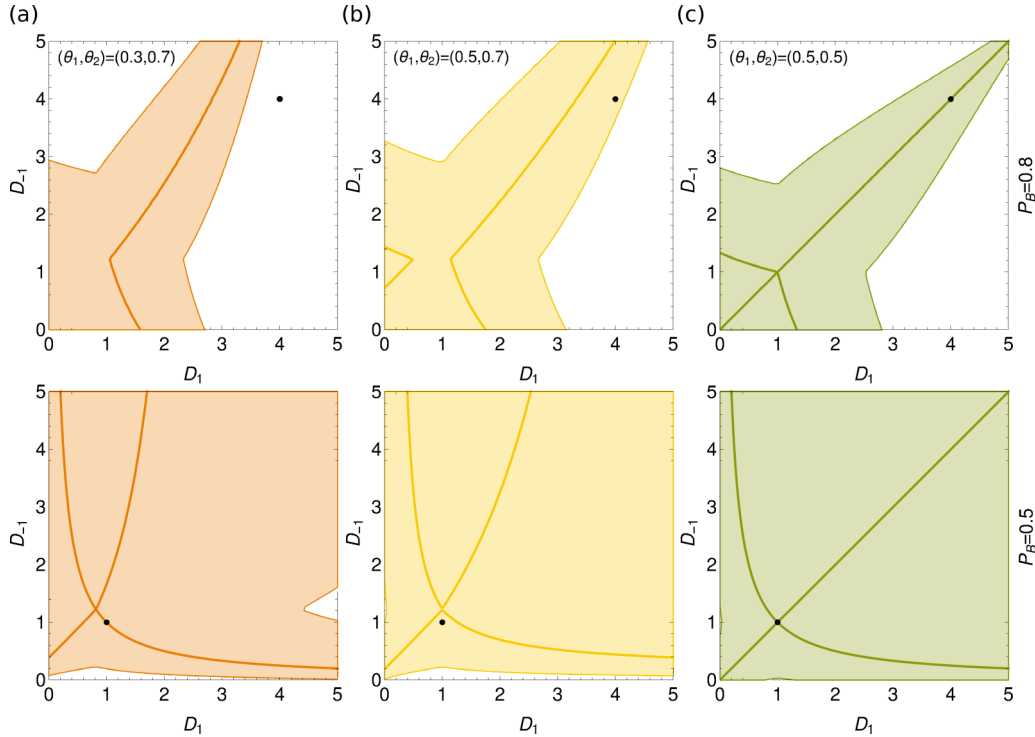


FIG. 8. Single population: steady-state types in the space of market order parameters (D_1, D_{-1}) for $\beta = 8.5$. Shown on top is the population of decisive buyers ($p_B = 0.8$) and on bottom the indecisive population ($p_B = 0.5$) for (a) symmetrically biased markets $(\theta_1, \theta_{-1}) = (0.3, 0.7)$, (b) one unbiased and one biased market $(\theta_1, \theta_{-1}) = (0.5, 0.7)$, and (c) two unbiased markets $(\theta_1, \theta_{-1}) = (0.5, 0.5)$. Colored regions indicate where weakly fragmented states exist (for $r \rightarrow 0$). Solid lines inside these regions indicate strongly fragmented states.

other hand, fragment strongly only if the markets are equal (the figure shows only $\theta = 0.5$, but the statement is true for general θ). Otherwise weak fragmentation occurs, although [see Fig. 8(c), top] when the markets are very different at a β above that used in Fig. 8.

With these insights, it is worth revisiting Fig. 6. It shows the existence of the fragmentation threshold β_c for all market biases, and we recall that this threshold is defined as the point where $P(\Delta)$ first acquires two peaks. From what we have seen above, we now understand that for most combinations of market biases, the steady state one finds above β_c is a weakly fragmented one. The exceptions are the dark lines in Fig. 6, which indicate equal ($\theta_1 = \theta_{-1}$) or symmetrically biased ($\theta_1 = 1 - \theta_{-1}$) markets.

V. TWO-GROUP POPULATION

So far in our analysis of the large size limit of a homogeneous population of traders with buying preference p_B , we have shown how for any given pair of market order parameters (D_1, D_{-1}) we can determine the population steady state. We identified three possible types of steady states: unfragmented (U), weakly fragmented (W), and strongly fragmented (S). We now generalize the investigation to populations of agents consisting of groups with different buying preferences. We demonstrate the approach for the case of two groups of the same size, but the principles are general and can be extended to larger numbers of groups or different group sizes. We denote a steady state of a population consisting of two groups by

a pair of letters (X, X') . Here $X, X' \in \{U, W, S\}$ indicates the type of steady state for each group, producing nine different types of population steady states.

We can now find, in the space of market order parameters (D_1, D_{-1}) , the domains of different state types as we did in Fig. 8. We can use the figure directly to read off the steady states of a population of two groups with buying preferences $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.5)$. For example, when the market order parameters are $(D_1, D_{-1}) = (5, 5)$ (the top right corner of all the diagrams), the steady state of the two-group population is (U,W) when the markets are symmetrically biased or biased and unbiased (Fig. 8 left and center) and (S,S) when both markets are unbiased (right diagrams). This simple analysis can be extended to any number of groups because, for market order parameters that are fixed exogenously, the groups are independent.

Our primary interest, however, lies in the case of endogenous market conditions where the agents we model capture the entire trading population and thus define their own market order parameters. In this case, we need to find the steady states self-consistently. We have previously described a procedure for doing this, for nonzero r [27]: Starting from some initial market order parameters (D_1, D_{-1}) , one calculates the steady states and updates (D_1, D_{-1}) iteratively, converging eventually to a self-consistent set of order parameters. Here we aim to get a complete picture of all possible steady states, independently of initial conditions. To do this, we start from the update equation for the market order parameters from the iterative approach. These are simply the definitions of

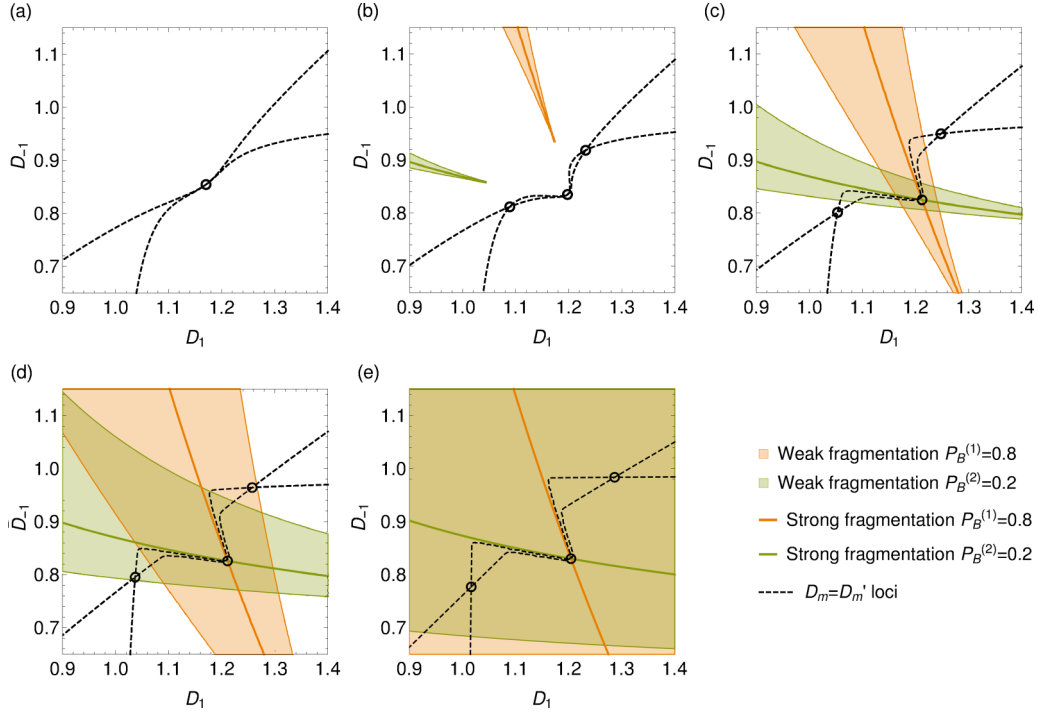


FIG. 9. Steady states of traders with decisive buy and sell preferences: order parameter diagrams. For a two-group system with $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$, each diagram shows the order parameter self-consistency lines (dashed black) and for each group the weak fragmentation region and the strong fragmentation line; $r = 0.001$ throughout. (a) Single (U,U) solution with $\beta = 1/0.31$, (b) three (U,U) solutions with $\beta = 1/0.29$, (c) one (S,S) and two (U,U) solutions with $\beta = 1/0.265$, (d) (U,W), (S,S), and (W,U) solutions with $\beta = 1/0.245$, and (e) (W,W), (S,S), and (W,W) solutions with $\beta = 1/0.2$. We use the following abbreviations for the steady state of each group: U, unfragmented; W, weakly fragmented; and S, strongly fragmented.

the market order parameters (18) and (19) extended to two groups:

$$D'_m = \frac{N_{Bm}^{(1)} + N_{Bm}^{(2)}}{N_{Sm}^{(1)} + N_{Sm}^{(2)}} = \frac{p_B^{(1)} \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(1)}, T_\gamma) + p_B^{(2)} \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(2)}, T_\gamma)}{(1 - p_B^{(1)}) \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(1)}, T_\gamma) + (1 - p_B^{(2)}) \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(2)}, T_\gamma)}. \quad (20)$$

We can now define, in the market order parameter space, the two loci where $D'_1 = D_1$ and $D'_{-1} = D_{-1}$, respectively, meaning that one of the order parameters is already self-consistent. The intersection of these loci (two lines, for our case of two markets) then gives us all the self-consistent sets of market order parameters. To distinguish weak and strong fragmentation, the limit $r \rightarrow 0$ ought to be taken. To avoid numerical issues we use here instead a small nonzero r to determine the attraction difference distributions $P(\Delta|p_B^{(g)}, T_\gamma)$ from which the D'_m are calculated. In most of what follows we focus on symmetric market setups ($\theta_1 = 1 - \theta_{-1}$) and symmetric agent buying preferences ($p_B^{(1)} = 1 - p_B^{(2)} = p_B$). To avoid having too many parameters to vary, we will fix the market biases to the default values $(\theta_1, \theta_{-1}) = (0.3, 0.7)$ unless stated otherwise.

A. Transitions in populations of decisive and indecisive traders

As shown in previous sections, the intensity of choice β is a crucial parameter determining whether the steady state in a system is fragmented or consolidated. Here we build upon this analysis by investigating how the nature of steady states changes as β is increased. We start this section with

examples of steady states of a population with decisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$ as well as one with largely indecisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$. We then generalize these results to a full phase diagram for the $r \rightarrow 0$ limit, giving the number and type of steady states as a function of the intensity of choice β and the buying preference p_B .

1. Decisive traders

In Fig. 9 we show, for a series of different β , the market order parameter space (D_1, D_{-1}) with the weak fragmentation region and the strong fragmentation line marked for both groups of a population with $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$. The order parameter self-consistency lines are also shown.

In Fig. 9(a) we show the low- β regime ($\beta = 1/0.31$), just before the onset of fragmentation. Note that for this β , the steady states of both groups are unfragmented across the entire range of market order parameters shown. The unique intersection of the $D'_m = D_m$ loci identifies a single steady state of type (U,U). Figure 9(b) shows a just slightly increased $\beta = 1/0.29$ where most market order parameters settings still give unfragmented states but there are now three intersections of the self-consistency loci, giving as many (U,U) steady

states. In the steady state that is a continuation of the low- β solution the agents show only mild preferences among the markets, with buyers slightly preferring the market that gives higher returns for buyers and similarly for sellers. The other two unfragmented solutions correspond to coordination at one of the two markets so that overall the situation is similar to the one we saw for $N = 2$ and $N = 4$.

Increasing β further [Fig. 9(c)], one crosses the threshold ($\beta_c \approx 1/0.28$ here) where one of the unfragmented solutions first fragments; the continuation of the low- β state is now in the strongly fragmented domain of both groups, while the other two steady states remain unfragmented. Note that since the weak fragmentation regions surround the strong fragmentation lines for both agent groups, there must in fact be a narrow range of slightly-lower- β values where the fragmentation is weak: The low- β solution must change from (U,U) through (W,W) to (S,S) as β increases.

In Figs. 9(d) and 9(e) one observes that with increasing β the fragmentation regions keep growing. This results in the two unfragmented (U,U) solutions changing first into (U,W) and finally (W,W).

We note that inferences about stability from diagrams like Fig. 9 are in general unwarranted; for example, the initial pitchfork bifurcation from one to three (U,U) states in Figs. 9(a) and 9(b) does not necessarily imply that the middle solution is unstable. It would be unstable under repeated updating from D_m to D'_m . However, Eq. (20) shows that this is not the real dynamics but would correspond to a scenario where the dynamics of the order parameters is slowed down artificially so that agents always have time to equilibrate their attraction difference distributions $P(\Delta|p_B^{(g)}, T_\gamma)$ to the current order parameter values.

We highlight one further feature of Fig. 9: For small r as used in the figure, the order parameter self-consistency lines tend to follow segments of the strong segregation lines before emerging on either side into a weak segregation region. This can be understood by noting that the self-consistency line for D_1 , for example, is the zero contour of the function $D'_1(D_1, D_{-1}) - D_1$ in the order parameter plane. This function varies steeply as a strong segregation line is crossed, developing discontinuities that look like cliff edges for $r \rightarrow 0$. A contour line that hits such a cliff must follow the line of the cliff before returning to the smooth parts of the landscape, which is the effect we see in Fig. 9.

The cliff edges themselves arise because on a segregation line, the free energy function $f(\Delta)$ in Eq. (16) has two minima of equal height. A small change of $O(r)$ in D_1 or D_{-1} will cause similar small changes in the height of these minima, but from Eq. (16) this is enough to cause the weight ratio between the two peaks in $P(\Delta)$ to shift by a factor of order unity. Changes larger than this will transfer all weight from one peak to the other and correspondingly modify D'_1 by a finite amount. For $r \rightarrow 0$ the required order parameter changes become infinitesimal, leading to the cliff edge structure of $D'_1 - D_1$ and analogously $D'_{-1} - D_{-1}$.

2. Indecisive traders

We now compare the results above with those for a population consisting of two agent groups with only weak

preferences for buying and selling, $(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$. The motivation for this comes from the fact that agents with only mild buy and sell preferences should develop weaker preferences for markets that offer higher returns for buyers or sellers. They will also not be penalized much if only a single group populates a market, as such an arrangement will still sustain a large number of trades.

In Fig. 10 we observe several differences compared to the situation in Fig. 9 for decisive traders, most notably with regard to the number of solutions. Specifically, as can just be discerned from Fig. 10, on crossing the fragmentation threshold four new states appear. These states are also different in nature: They are partially fragmented in the sense that one group of agents is strongly fragmented and thus retains a bimodal distribution of attraction differences for $r \rightarrow 0$ while the other is either weakly fragmented or unfragmented. We have seen a similar state in the systems with four agents, although there it was unstable because it reduced the number of possible trades. In the large population limit, having one fragmented and one unfragmented group of agents still leaves many possibilities for trading, especially for indecisive agents where roughly half of each group of agents will probabilistically assume the role of buyer or seller in each trading round. On the general grounds discussed above, the appearance of the partially fragmented (U,S) states is expected to proceed via (U,W) states, though again the β range where the latter appear is numerically small.

When the intensity of choice β is increased beyond that in Fig. 10(a), the low- β solution transitions from (U,U) to (W,W) [Fig. 10(b)] and eventually (S,S) [Fig. 10(d)], i.e., both agent groups fragment first weakly and then strongly. Comparing the partially fragmented solutions in Figs. 10(b) and 10(c), we see that they change from (U,S) to (W,S); finally two of them merge with the uncoordinated (W,W) state into an (S,S) state. The other two partially fragmented states eventually transition into (W,W) states; as in the case of decisive traders, these represent coordination of the agents at a single market.

B. The (β, p_B) phase diagram

We have observed both market consolidation and fragmentation when a population is faced with a choice of two symmetric markets, depending on the different choice of system parameters (p_B, β) . We next vary these parameters systematically to construct a detailed phase diagram and study the regions where one finds the various states that we described above. The size of these regions then also gives an indication of how typical the different scenarios are. We continue to focus on symmetric markets with $(\theta_1, \theta_{-1}) = (0.3, 0.7)$ but note that calculations for other (symmetric) market settings give qualitatively similar results. In Fig. 11 we show the phase diagram in the space of intensity of choice β and group preference for buying $p_B \equiv p_B^{(1)}$. This diagram is the large population analog of the diagram for four agents ($N = 4$) shown in Fig. 5. There we had identified regions with states that are unfragmented and indecisive (low β), unfragmented and coordinated, fragmented, and partially fragmented. Broadly, these types of states persist in the large population limit, but they have additional structure that makes for a richer phase diagram.

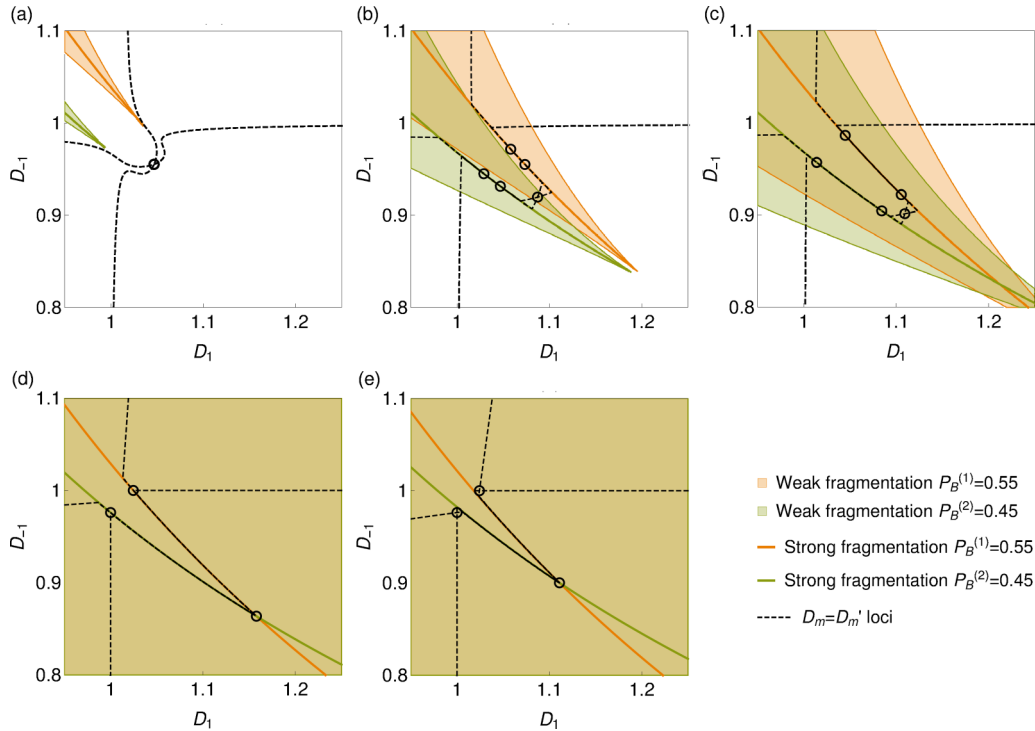


FIG. 10. Steady states of largely indecisive traders: order parameter diagrams. We show the behavior of largely indecisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$ for different intensities of choice β in the large memory limit, evaluated numerically for $r = 0.00001$: (a) $\beta = 1/0.31$, one unfragmented solution (U,U); (b) $\beta = 1/0.285$, one weakly fragmented (W,W) and four partially fragmented (U,S) states; (c) $\beta = 1/0.27$ one weakly (W,W) and four partially fragmented (W,S) states; (d) $\beta = 1/0.1$, one strongly fragmented (S,S) and two partially fragmented (W,S) states; and (e) $\beta = 1/0.05$, one strongly fragmented (S,S) and two weakly fragmented states (W,W).

To help visualize the structure of the phase diagram, we show an additional version as an inset that has been distorted to preserve the topology but make even small regions in the phase diagram visible. Also, to avoid having too many separate regions we do not distinguish in the diagram between unfragmented and weakly fragmented states, which both have distributions of market preferences that become unimodal for $r \rightarrow 0$. We label such states collectively V to separate them from strongly fragmented states with their bimodal market preference distributions. Two vertical dashed lines mark the two scenarios of decisive and indecisive traders studied above (see Figs. 9 and 10).

We now look in more detail at the structure of Fig. 11. Crossing any line in the phase diagram changes either the number of population solutions or the nature of the steady state for one or both agent groups. We note that, due to the symmetry of the system we consider, many of the changes for the two groups happen simultaneously. In the inset, regions of the parameter space are laid out according to the number of solutions: five solutions on the left and three on the right, with a single solution in the small- β region at the top.

The dark violet line in Fig. 11 is the line where the multiplicity of states changes from 1 to 3 or 5. Looking at the order parameter, self-consistency lines shows that this transition takes place via a pitchfork bifurcation in the former case and two symmetric saddle-node bifurcations in the latter. The dark violet line is an analog of the line shown in the same color in the phase diagram (Fig. 5) of the system with $N = 4$ players. The region of multiple solutions has grown for large

N , but the inverse critical intensity of choice $1/\beta_c$ is still an increasing function of p_B .

As in Fig. 5, the pink line with circles in Fig. 11 marks the appearance of a steady state where both groups are strongly fragmented. We observe that the critical intensity of choice where this happens diverges ($1/\beta \rightarrow 0$) as $p_B \rightarrow 0.5$, i.e., the region of strong fragmentation shrinks as the difference between the groups' buying preferences diminishes. Further lines in the phase diagram show where the solution multiplicity changes directly from 3 to 5 (yellow line) and where partial fragmentation occurs (green and orange line) as individual solutions transition from (V,V) to (V,S). Note that in the large population limit such partially fragmented states appear only for populations with moderate preferences for buying, in contrast to the system with $N = 4$ agents (Fig. 5) where they exist for all p_B .

We mark one further line (dashed pink) in the main graph of Fig. 11, showing the transition within the small- β (V,V) solution from the unfragmented (U,U) to the weakly fragmented (W,W) state. With this we make an explicit connection to results reported previously (Fig. 7 in [27]) where we investigated the appearance of (weak) fragmentation with increasing intensity of choice. We note that for the system of our first case study $p_B = 0.8$ the thresholds for weak and strong fragmentation almost overlap; the region of the weakly fragmented indecisive state is very narrow for this choice of parameters and in general for p_B above approximately 0.7, while it becomes larger for indecisive traders. The pink circle on the y axis marks the end of the weak fragmentation line.




FIG. 11. Types of steady states for two symmetric agent groups, shown in (β, p_B) space. Crossing each of the lines changes either the number or the type of steady states. The inset shows the distorted but topologically equivalent diagram to show the phase diagram regions more clearly. The dark violet and yellow lines show the change in solution multiplicity, the pink solid line with circles shows the strong fragmentation, the pink dashed line shows the weak fragmentation of the uncoordinated low- β solution, the orange and green lines show the partial fragmentation. Here V denotes a unimodal distribution of agents' market preferences in the $r \rightarrow 0$ limit, i.e., an unfragmented (U) or weakly fragmented (W) steady state, and S denotes a strongly fragmented steady state.

It turns out that this is the strong fragmentation threshold of the homogeneous population with even preference for buying and selling ($p_B = 0.5$), with the change from weak to strong fragmentation caused by the additional symmetry between the two groups for this value of p_B .

Interestingly, there are two distinct regions in the phase diagram of Fig. 11 where we observe three (V,V) and two (V,S) states, i.e., three unfragmented and two partially fragmented solutions. It turns out that in the region at lower β (higher $1/\beta$) the partially fragmented solutions are coordinated, insofar as both groups of agents have an overall preference for the same market. For high β one has the opposite situation, and it is those uncoordinated (V,S) solutions together with an unfragmented (V,V) solution that then merge into a single (S,S) state as p_B is increased.

We note briefly that the various lines shown in Fig. 11 were detected by solution tracking, e.g., by carefully varying p_B and β and tracking the number and type of solutions; further details can be found in Appendix B. The tracking approach is chosen as it is numerically faster and more reliable than the finite- r procedure we used in previous figures, avoiding, e.g., the numerical noise visible in the two loci in Fig. 10. It is important to remember that the results only provide information about the existence of steady states, not their stability; the latter can be probed only using actual dynamics as discussed below. Figure 11 also relates to fixed market biases so trends with changes in these biases cannot be seen; we have checked, however, that the overall structure of the phase diagram remains intact as long as market biases are symmetric. Quantitative trends were explored in our previous work [27], where we saw that the fragmentation region shrinks as markets become increasingly different.

In summary, the diagram in Fig. 11 shows that for systems with two symmetric markets and two groups of traders with symmetric buying preferences both fragmented and coordinated (or consolidated) steady states exist across a substantial range of values for the intensity of choice β . Single-market dominance happens when the steady state is either unfragmented or weakly or partially fragmented but coordinated: The majority of trades then happens at a single market. On the other hand, markets can coexist, receiving a roughly even share of trades, when the steady state is strongly fragmented or weakly or partially fragmented but uncoordinated. In the former case both markets are visited by both groups, while in the latter case an effective market and group loyalty appears. In the following sections we analyze these different steady states further, with regard to the average population returns they produce and their stability in simulated systems with finite N and r .

C. Average population returns

The phase diagram in Fig. 11 reveals a plethora of possible steady states in the system of two markets and a large population of traders, depending on the traders' learning parameter β and their propensity to act as buyers p_B . We now investigate whether these steady states induce differences in average population returns as we saw in small systems, e.g., Figs. 2 and 4. We look at the average population return per trading round, where we count also zero returns that arise from an order being invalid or no trading partner being available.

In Fig. 12 we show average population returns for the two scenarios of decisive [$(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$, Fig. 12(a)] and indecisive [$(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$, Fig. 12(b)] traders. The

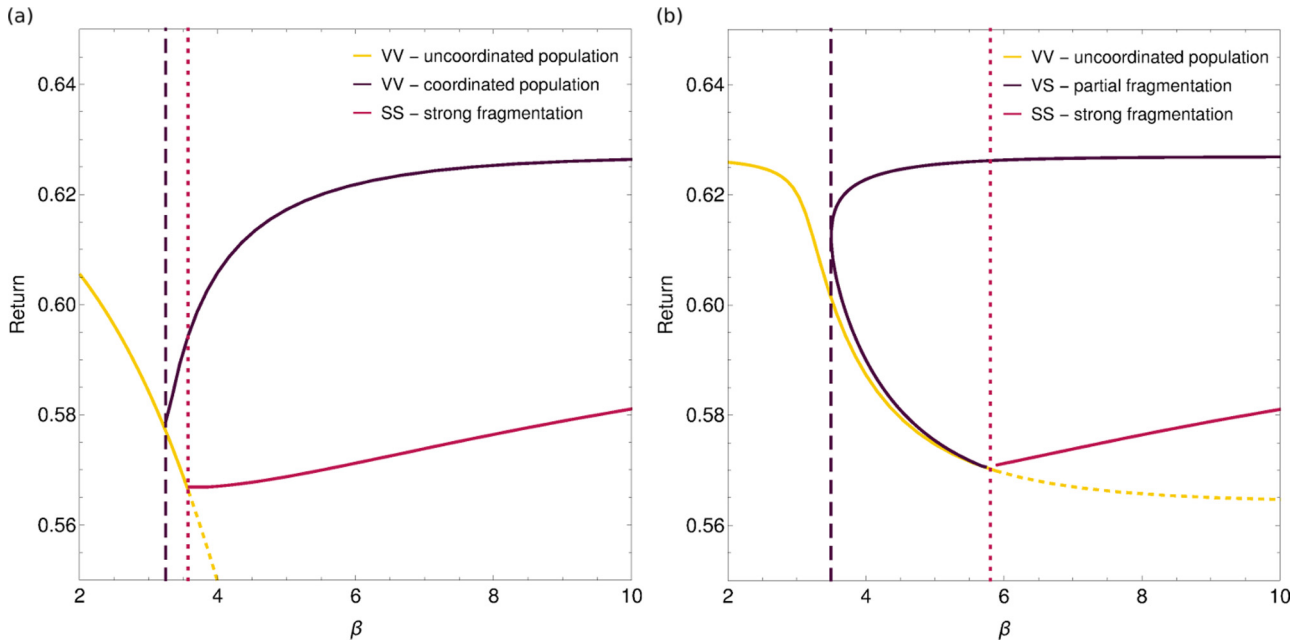


FIG. 12. Average population returns for different steady states in the $r \rightarrow 0$ limit. The yellow line shows the low- β steady state representing uncoordinated population (dashed in the regime where it is no longer a bona fide steady state). The dashed dark violet line marks the value of β where the multiplicity of solutions changes (see Fig. 11); at the dashed pink line strongly fragmented steady states first appear. (a) Decisive population $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$. The dark violet line shows the average population return for a coordinated unfragmented or weakly fragmented steady state and the pink line similarly for a strongly fragmented state. (b) Indecisive population $(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$. The dark violet line gives the average population returns for partially fragmented steady states (coordinated on top, uncoordinated on bottom) and the pink line similarly for a strongly fragmented state.

β dependences reflect the transitions between solution types we saw earlier (in Figs. 9 and 10 and the phase diagram in Fig. 11). The overall trends resemble those for finite N . First, we note that the return of the uncoordinated low- β solution (marked in yellow in Fig. 12) is the lowest among the alternatives once multiple solutions exist. Second, the coordinated states (dark violet) lead to the highest average return. Interestingly, this is not influenced by the type of fragmentation, i.e., it is true for both weakly fragmented and partially fragmented states as long as a majority of the population develops a preference for a single market. By comparison, the strongly fragmented state (pink) always leads to a lower average population return.

The differences in the returns achieved by populations of decisive and indecisive traders, respectively, are driven mainly by the fact that indecisive groups can sustain more trades without requiring the presence of other groups at a market. This is particularly visible in the higher population average return for low β ; in this range the decisive population suffers from the group-specific market preferences that tend to separate traders towards different markets and consequently result in a lower number of trades. Additionally, the continuation of the low- β solution is a viable steady state for a broader range of intensities of choice for the indecisive population. The dashed yellow line marks the region of β for which this fixed point is no longer a genuine steady state, as the free energy has multiple minima when evaluated at the order parameters calculated for this fixed point. Along this line the indecisive

population return again does not drop as far as it does in the case of a more decisive population.

In Fig. 12(b) we note the occurrence of the saddle node bifurcation in the transition of the indecisive population, with four new (V,S) solutions (which come in two pairs giving identical returns) emerging at once. The top branch corresponds to the average population returns at the coordinated partially fragmented states; for greater values of β (outside the range shown) these states smoothly transition into weakly fragmented, coordinated, states. The bottom branch relates to uncoordinated partially fragmented states that merge into the strongly fragmented (S,S) state for greater β .

Interestingly, the average population return in the high- β limit of the coordinated state also corresponds to the average population return when all traders choose randomly (i.e., $\beta = 0$). This is true because in both limits the average number of agents trading at each market is equal. Intriguingly, this means that when learning is introduced, for low intensities of choice, an agent who makes decisions based on their previous history may be worse off than an agent who plays at random. This effect disappears again only in the large- β limit of the weakly fragmented state, though note that in the latter case one group earns more than the other. Returning to the strongly fragmented state, despite indications that for a given β this is best among the states that do not distinguish between groups in the long run (see Fig. 6 of [27]), in terms of average population return this state is outperformed by random traders ($\beta = 0$).




FIG. 13. Memory length dependence of phase boundaries. The pink line (solid with circles and dotted) shows the fragmentation threshold, where at least one steady state is fragmented (weakly or strongly). The dark violet line (dark solid and dashed) shows the boundary of the region where multiple steady states exist. Solid lines represent indecisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.55, 0.45)$ and dashed lines decisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$. Multiple steady states exist for small enough r , i.e., long enough memory $1/r$. The market parameters are $(\theta_1, \theta_{-1}) = (0.3, 0.7)$.

D. Dynamics

We now ask what effect the existence of multiple steady states, as predicted by theory for infinite populations, has on the dynamics. We simulate the dynamics numerically, necessarily for finite N and with learning rate $r > 0$, i.e., for finite memory length $1/r$. In previous work we have already shown that the theory predicts the steady-state properties of finite populations quite well (see, e.g., Fig. 4 in [27]). The role of r is more important as this can shift phase boundaries [27]. (Conceptually, the precise distinction for $r \rightarrow 0$ between weakly and strongly fragmented states is also lost for $r > 0$ and becomes a crossover.)

In Fig. 13 we illustrate the r dependence of two key phase boundaries for the two populations we have mainly considered so far (decisive $p_B = 0.8$ and indecisive $p_B = 0.55$). We note that the region of multiple steady states shrinks with increasing r for both populations while the fragmentation line is only weakly r dependent. The lines are related to the lines of the same color in the (β, p_B) phase diagram in Fig. 11 and the dashed gray lines in Fig. 11 correspond to the $r \rightarrow 0$ limit of the (r, β) phase diagram in Fig. 13.

Overall, Fig. 13 tells us that we need to use reasonably small r , certainly below 0.05 for $p_B = 0.8$, to see multiple steady states in numerical simulations. As smaller r slow the dynamics, we choose in practice values of r that are as large as possible while staying well within the multiple states regime.

In Fig. 14 we show numerical data for the actual dynamics of a system of decisive traders $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$ at our standard market parameters $(\theta_1, \theta_2) = (0.3, 0.7)$, taken from a single run for a population with $N = 2000$ traders (see [37] for simulation details) using the learning rate and inverse decision strength $(r, 1/\beta) = (0.05, 0.16)$. For these parameters the phase diagram of Fig. 13 predicts the existence of three steady states, two weakly fragmented states (with the majority of both groups coordinated at the same market, $m = -1$ or $m = 1$) and a strongly fragmented state (this state was studied in [27]; see Fig. 3 there; it is the unique steady state for the larger $r = 0.1$ used in [27]). As a global summary statistic of the shape of the attraction distributions of the two groups of agents we use the Binder cumulant [39]

$$B = 1 - \frac{\langle \Delta^4 \rangle_{P(\Delta)}}{3 \langle \Delta^2 \rangle_{P(\Delta)}^2}$$

and plot this over time (see further discussion in [27,37]). Away from the strongly fragmented state the attraction




FIG. 14. Metastability of the strongly fragmented state and transition to the weakly fragmented state: dynamical evolution of a system with $N = 2000$ agents using $(r, 1/\beta) = (0.05, 0.16)$, with preferences for buying $(p_B^{(1)}, p_B^{(2)}) = (0.8, 0.2)$, and market parameters $(\theta_1, \theta_2) = (0.3, 0.7)$. (a) Evolution of Binder cumulants of the two attraction distributions of the two agent groups [buyers (green) and sellers (orange)]. Dashed lines are theoretical predictions for the strongly fragmented steady state (dark violet denotes equal for both groups) and weakly fragmented state (green and orange for the two groups). (b) Attraction distributions predicted from theory for the weakly fragmented steady state (solid line) compared to simulation data at $t = 500$ (histogram).




FIG. 15. Lifetime of strongly fragmented state for different system sizes N . (a) Binder cumulant time series (averaged over the two agent groups for compactness) along with the $N \rightarrow \infty$ theoretical prediction for the strongly fragmented state (dashed line), showing an increase of lifetime with N . $1/\beta = 0.15$ while the other parameters are as in Fig. 14. (b) Autocorrelation function of single agent attraction differences $C(t) = \langle [\Delta^i(\tau) - \bar{\Delta}(\tau)][\Delta^i(\tau + t) - \bar{\Delta}(\tau + t)] \rangle$. The single agent autocorrelation time is essentially N independent.

distributions of the two groups are not related by a symmetry, so we plot their Binder cumulants separately.

Figure 14 shows that the system quickly reaches the strongly fragmented state, with the Binder cumulants being close to the theoretically predicted value; the slight deviation can be attributed to the finite population size. The dynamics then branches off from the theoretical prediction at $t \approx 50$, showing that the strongly fragmented state is, for finite N , only metastable. The departure is led by one of the agent groups and reaches one of the theoretically expected weakly fragmented states at $t \approx 500$, as shown in Fig. 14 by the agreement of both the relevant Binder cumulants [Fig. 14(a)] and the full attraction distributions [Fig. 14(b)].

We proceed in Fig. 15(a) to analyze the lifetime of the strongly fragmented steady state in more detail. The figure displays Binder cumulant time series for different population sizes at the learning parameters $(r, 1/\beta) = (0.05, 0.15)$ and shows that the lifetime increases with system size (we have not analyzed the N dependence in detail; in the range shown it is approximately linear). We can compare this with the time correlations of the attraction difference Δ for individual agents: Fig. 15(b) graphs this correlation function, measured from the point in time when the strongly fragmented state is first reached. One sees clearly that the single agent correlation time is essentially independent of N , while the lifetime of the strongly fragmented state grows significantly with system size N . The conclusion is that strong fragmentation is a long-lived state of the population for large N , within which single agents effectively “equilibrate” by losing all memory of their initial preferences.

In Fig. 16 we move to the r dependence of the lifetime of the strongly fragmented state, showing Binder cumulants for a small system $N = 200$ for different r values at fixed $1/\beta = 0.15$. For all values of r , rapid initial convergence to the strongly fragmented state is observed. Within this state the Binder cumulants depend weakly on r (as has been noted

previously [27]), reflecting the r dependence of the attraction distributions. The lifetime of the strongly fragmented state, set by the decay of the Binder cumulant to lower values, increases with r . This is consistent with the results of Fig. 13, which showed that above some β -dependent threshold value for r the strongly fragmented state is the only steady state and thus must be stable, corresponding to an infinite lifetime. For the value $\beta = 1/0.15$ in Fig. 16, theory predicts this threshold to be $r \approx 0.055$. Numerically, we see that the strongly fragmented state has a finite lifetime up to $r = 0.07$, presumably due to finite population effects for the relatively small $N = 200$ used in the simulations presented in the figure.

We find qualitatively the same features as above also in numerical simulations of the dynamics of a system of indecisive traders, with populations first reaching a long-lived (for large N) strongly fragmented state and eventually decaying




FIG. 16. Binder cumulant time series for different learning rates r at the fixed intensity of choice $1/\beta = 0.15$. All parameters are as in Fig. 14, except for the smaller population size $N = 200$ and r as shown. The lifetime of the strongly fragmented state lifetime increases with r , eventually becoming infinite when this state is the only steady-state solution.

into a partially fragmented state. This is the behavior when such multiple steady states are predicted by our theory, i.e., for small r ; for larger r (above $r_c \approx 0.02$; see Fig. 13) strong fragmentation is the only steady state. Quantitatively, we find that where strong fragmentation is metastable its lifetimes are significantly longer than for the decisive traders, exceeding our maximal simulation times of 10^6 trading rounds ($t = 20\,000$) for the largest $r < r_c$.

We comment finally on the role of initial conditions. In the dynamical simulations shown so far we used for these $P(\Delta|p_B) = \delta(\Delta)$, corresponding to the reasonable assumption that the agents have no initial preference for either market. We also explored Gaussian initial distributions for the attraction differences, $P(\Delta|p_B) = \mathcal{N}(\mu, \sigma^2)$. Where there is only a single steady state we then find, as expected, that this state is reached irrespective of the chosen initial condition. On the other hand, where the theory predicts multiple steady states, the initial conditions do matter. We observe that the metastable strongly fragmented state continues to be reached whenever the mean initial attraction difference $|\mu|$ is small enough, irrespective of the standard deviation σ . As $|\mu|$ is increased we see that the dynamics “misses” the metastable strongly fragmented state and rapidly moves to a final weakly or partially fragmented state. This is consistent with the intuition that these states break the symmetry between markets, and hence are favored when the population already starts off with an overall initial preference for one of the markets.

VI. DISCUSSION AND CONCLUSIONS

In this paper, our aim was to investigate the existence of coordinated and fragmented steady states in a system of agents choosing adaptively between two markets. We focused primarily on the long memory limit, where the transition to fragmentation is sharp. We first studied two traders who learn how to coordinate at a market and maximize their average return even though one of them will necessarily earn less. Moving to a four-player system, we observed fragmentation in addition to coordination. Interestingly, we found that coordinated and fragmented states lead to the same average population return for high intensity of choice β , in spite of the presence of two different types of agents (buyers and sellers). In the coordinated state one of the agent types will always earn less, while in the fragmented state both types have the same average, but one agent from each group is less satisfied. Thus, at the fragmented state, average returns do not discriminate between types of agents.

We then introduced a general method for determining the type and number of steady states in the limit of large populations with long memory. This can be done in our setup with only a single order parameter per market. After a preliminary analysis for exogenously determined order parameters, we saw that in the general case a self-consistency criterion determines the order parameters in the steady state. Analyzing a quantity analogous to a free energy then allows one to say whether a population (or one of its groups) is fragmented and whether this fragmentation is strong or weak.

Already for small system sizes we noticed that the agents’ preference for buying p_B is an important system parameter. Not only does it influence the critical intensity of choice β on

p_B for the onset of fragmentation, but for $N \geq 4$ it also qualitatively affects the nature of the steady states. This remains true also for the $N \rightarrow \infty$ limit, where we find a rich variety of steady states in the (β, p_B) diagram, in spite of the simplified nature of our models for markets and traders. These include market coexistence, where both markets attract both types of traders (S,S) and where market–trader specialization occurred (W,W) (uncoordinated weakly fragmented state for moderately indecisive traders); single market dominance (W,W) (coordinated weakly fragmented states); market indifference (U,U) (e.g., for low β); and general vs specialized markets [e.g., (U,S), where a single market attracts both groups of agents while the other can be viewed as specializing towards only one group]. Interestingly, all these different steady states arise without imposing any heterogeneity onto the agents (in contrast to assumptions elsewhere [23]) and fragmentation is the preferred state even when the markets have identical properties (contrary to views expressed in [2,18]).

To interpret our results for the prevalence of fragmentation more broadly we can draw on the work of Cheung and Friedman [40], who used evidence from behavioral game theory to suggest that values of β are consistent across games but increase in more informative environments. The authors also argued that a parameter closely analogous to r increases with the trustworthiness of information in the system. Bearing in mind the results shown in Fig. 13, where for large r and large β the only steady state is the fragmented one, this suggests that more informative environments, or ones where information is more trustworthy because of, e.g., stability over long timescales, might naturally lead to fragmented states. The prevalence of the strongly fragmented state is clear also from Fig. 11, which shows that this state exists for all populations with groups symmetrically biased towards buying and selling, respectively.

One of the nontrivial predictions of our theory is the existence of partially fragmented states, where one group of agents (e.g., those who have a preference for buying) fragments while the other (where agents prefer to sell) does not. We saw that the region in the phase diagram where such states appear increases with N for indecisive traders and shrinks for decisive traders (compare Fig. 5 for $N = 4$ and Fig. 11 for $N \rightarrow \infty$).

We studied also the average population returns achieved by agents in the various steady states. For large populations we saw that the coordinated weakly fragmented steady state leads to the highest population average returns, even though one agent group earns less in that state. We also noticed that such steady states, which essentially represent coordination at a single market when $r \rightarrow 0$, lead to the same average payoff for large β as for random agents ($\beta = 0$). This is because coordination at a single market, just like random market choice, leads to the same number of buyers and sellers at a single market and thus the same number of successful trades and average returns. Interestingly, this shows that weak learning (finite β) leads to lower returns, e.g. not choosing the strictly best trading venue (in terms of returns) can be worse for an agent than random guessing. This behavior is rather similar to the J-curve effect studied in [41,42] where, in the context of trading agents with different information levels, moderately informed agents earn less from higher

informed agents but also from uninformed, randomly trading, agents.

Finally we investigated, by means of numerical simulations, how the theoretically predicted steady states appear in the dynamics of finite agent populations. If the agents start as “blank canvasses” (without initial market preference), we found that the adaptation process always leads to the strongly fragmented state first. This state is metastable, with a lifetime that grows large with population size, and the system eventually settles into one of the weakly fragmented states. This remains true even if there is scatter in the agents’ initial preferences, while a systematic initial bias towards one of the markets can cause the dynamics to miss the metastable strongly fragmented state. To put this result into more intuitive terms, two markets that enter into competition to attract on average indifferent traders will always exhibit a period of coexistence in a strongly fragmented state (and if $r > r_c$ this coexistence will last indefinitely), whereas if the population is not indifferent initially then a market monopoly will arise much more quickly.

$$\begin{aligned}
 P(S|m, \mathcal{B}) &= Q_{\mathcal{B}m} T_{\mathcal{B}m} \frac{1}{Q_{\mathcal{B}m} \sigma_b \sqrt{2\pi}} \exp\left(-\frac{[S - (\mu_b - \pi_m)]^2}{2\sigma_b^2}\right) \theta(S) + \delta(S)(1 - Q_{\mathcal{B}m} T_{\mathcal{B}m}), \\
 P(S|m, \mathcal{S}) &= \underbrace{\frac{Q_{\mathcal{S}m} T_{\mathcal{S}m}}{Q_{\mathcal{S}m} \sigma_a \sqrt{2\pi}} \exp\left(-\frac{[S - (\pi_m - \mu_a)]^2}{2\sigma_a^2}\right)}_{\text{agent trades}} \theta(S) + \delta(S) \underbrace{(1 - Q_{\mathcal{S}m} T_{\mathcal{S}m})}_{\text{agent does not trade}}.
 \end{aligned} \tag{A1}$$

(Note that in statements of these distributions in previous publications [27], μ_a and μ_b were omitted due a typographical error.) When agents have fixed buying preferences p_B , their return distribution is then dependent only on the chosen market m :

$$P(S|m) = p_B P(S|m, \mathcal{B}) + (1 - p_B) P(S|m, \mathcal{S}).$$

The probabilities that an order is valid Q_γ are given by

$$\begin{aligned}
 Q_{\mathcal{B}m} &= \frac{1}{\sigma_b \sqrt{2\pi}} \int_{\pi_m}^{\infty} db \exp\left(-\frac{(b - \mu_b)^2}{2\sigma_b^2}\right), \\
 Q_{\mathcal{S}m} &= \frac{1}{\sigma_a \sqrt{2\pi}} \int_{-\infty}^{\pi_m} da \exp\left(-\frac{(a - \mu_a)^2}{2\sigma_a^2}\right)
 \end{aligned}$$

and can be expressed in terms of error functions [37].

The transition kernel between two states Δ and Δ' of an agent with buying preference p_B is

$$\begin{aligned}
 K(\Delta'|\Delta, p_B) &= \int dS \sum_{m=-1}^1 [p_B P(S|m, \mathcal{B}) \\
 &\quad + (1 - p_B) P(S|m, \mathcal{S})] P(m|\Delta) \\
 &\quad \times \delta(\Delta' - mrS - (1 - r)\Delta). \tag{A2}
 \end{aligned}$$

The resulting drift and diffusion terms for small r are discussed in detail in [37]; here (Fig. 17) we provide plots corresponding to Eqs. (12) and (13), evaluated at three different sets of market order parameters for illustration. We consider the value $\beta = 1/0.265$ for the intensity of choice, in order to match Fig. 9(c). The three sets of market order parameters all

ACKNOWLEDGMENTS

The authors are grateful to P. McBurney and R. Nicole for useful discussions. P.S. acknowledges the stimulating research environment provided by the EPSRC Centre for Doctoral Training in Cross-Disciplinary Approaches to Non-Equilibrium Systems (Grant No. EP/L015854/1). A.A. acknowledges support from the Ministry of Education, Science, and Technological Development of the Republic of Serbia under Project No. ON171017.

APPENDIX A: DETAILS OF THE FOKKER-PLANCK DESCRIPTION

In this Appendix we provide some of the explicit expressions appearing in the Fokker-Planck description of our market choice model. As per definitions of bid and ask distributions and score assignments, defined in the discussion of trading strategies of Sec. II, the return distributions for an agent choosing a market m and an order type \mathcal{B} or \mathcal{S} are

lie on a horizontal line ($D_{-1} = -1$), while D_1 is changed so that the order parameters lie in the unfragmented, the weakly fragmented, or the strongly fragmented region, respectively. Plots in Fig. 17(a) illustrate market conditions leading to an unfragmented distribution; there is a unique solution of $M_1(\Delta|p_B, T_\gamma) = 0$, corresponding to the unique free energy minimum [calculated from Eq. (15) and shown in the bottom row of the figure]. Both are marked by a circle. Figure 17(b) illustrates the weakly fragmented case, where there are three zeros of the drift term (two stable fixed points and one unstable one), corresponding to two minima of the free energy; as the minima are at different heights, the resulting (steady-state) distribution of Δ will become concentrated around the lowest minimum for $r \rightarrow 0$, as discussed in the main text. Finally, the case shown in Fig. 17(c) has two equal minima of the free energy and thus represents a strongly fragmented scenario. Note that the diffusion term $M_2(\Delta)$ is in all three cases of order unity and does not affect the number of free energy minima; it only makes a quantitative contribution to the free energy and hence to $P(\Delta|p_B, T_\gamma)$.

APPENDIX B: ALGORITHMIC REMARKS

The method of finding all steady-state solutions by identifying loci of self-consistent market order parameters is the best way to exhaust market order parameter space and thus to find all the solutions for the finite r . By identifying the domains where these solutions lie, we can also fully characterize the solution at nonzero r , obtaining information




FIG. 17. Drift $M_1(\Delta)$ (top), diffusion $M_2(\Delta)$ (middle), and free energy $f(\Delta)$ (bottom) functions for a subgroup with preference for buying $p_B = 0.8$. Function plots illustrate three qualitatively different conditions for the following pairs of market order parameters: (a) $(D_1, D_{-1}) = (1, 1)$, unfragmented region; (b) $(D_1, D_{-1}) = (1.1, 1)$, weakly fragmented region; and (c) $(D_1, D_{-1}) = (1.15, 1)$, strongly fragmented region. Market biases are set to the standard values $(\theta_1, \theta_{-1}) = (0.7, 0.3)$. All functions are evaluated at the intensity of choice $\beta = 1/0.265$.

about the limit $r \rightarrow 0$ by extrapolation. However, this method is numerically demanding as for every point in order parameter space we need to find a steady-state distribution (its normalization usually takes most of the processing time) and recalculate the corresponding order parameters. Checking what corrections arise for $r \rightarrow 0$ takes additional time. We describe numerically less demanding alternatives below.

Population with homogeneous market preferences. We have seen that, depending on system parameters, the attraction distribution for a group of agents can be unimodal in the $r \rightarrow 0$ limit (U and W states). These states represent a population where the market preferences within the group are homogeneous. This realization offers a straightforward way to find all the states of this type for any system parameter. The demand-to-supply order parameters simplify to

$$D_m = \frac{p_B^{(1)} \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(1)}) + p_B^{(2)} \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(2)})}{(1 - p_B^{(1)}) \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(1)}) + (1 - p_B^{(2)}) \int d\Delta \sigma_\beta(m\Delta) P(\Delta|p_B^{(2)})} = \frac{p_B^{(1)} \sigma_\beta(m\Delta^{(1)}) + p_B^{(2)} \sigma_\beta(m\Delta^{(2)})}{(1 - p_B^{(1)}) \sigma_\beta(m\Delta^{(1)}) + (1 - p_B^{(2)}) \sigma_\beta(m\Delta^{(2)})}, \quad (\text{B1})$$

where in the right hand side we have used $\langle \sigma_\beta(\Delta) \rangle = \sigma_\beta(\langle \Delta \rangle)$, a relation that is exact in the $r \rightarrow 0$ limit where the steady-state distribution is a δ distribution centered at $\Delta^{(g)}$. To identify these peak positions we find the zeros of the first jump moment M_1 as defined in Eq. (12), taking into account the dependence of D_m on the attraction difference $\Delta^{(g)}$ in each group. This means that when searching for a steady state in which both groups of traders have homogeneous market preferences, we need to solve the peak position equations for the two groups simultaneously:

$$\begin{aligned} M_1^{(1)}(\Delta^{(1)} | p_B^{(1)}, D_m(\Delta^{(1)}, \Delta^{(2)})) &= 0, \\ M_1^{(2)}(\Delta^{(2)} | p_B^{(2)}, D_m(\Delta^{(1)}, \Delta^{(2)})) &= 0. \end{aligned} \quad (\text{B2})$$

Every solution $(\Delta^{(1)*}, \Delta^{(2)*})$ found in this way needs to be checked for consistency with the initial assumption of homogeneous market preferences, i.e., the market order parameters corresponding to every solution pair need to belong to the unfragmented or weakly fragmented solution domain. This is done by calculating the corresponding order parameters D_m from Eq. (B1) and finding the free energy corresponding to these order parameters. If the global free energy minimum is centered at $\Delta^{(g)*}$ the solution is consistent with our initial assumption and we have found a (homogeneous) population steady state. Depending on the signs of Δ^* , we classify such steady states further as either coordinated for $\Delta^{(1)*} \Delta^{(2)*} > 0$ or uncoordinated for $\Delta^{(1)*} \Delta^{(2)*} < 0$. For any finite intensity of choice β , a single agent can of course choose another market even if the state is categorized as coordinated at market 1, but the categorization is exact for the $\beta \rightarrow \infty$ limit.

$$D_m = \frac{N_{Bm}}{N_{Sm}},$$

$$D_m = \frac{p_B^{(1)}[\omega^{(1)}\sigma_\beta(m\Delta_1^{(1)}) + (1 - \omega^{(1)})\sigma_\beta(m\Delta_2^{(1)})] + p_B^{(2)}[\omega^{(2)}\sigma_\beta(m\Delta_1^{(2)}) + (1 - \omega^{(2)})\sigma_\beta(m\Delta_2^{(2)})]}{(1 - p_B^{(1)})[\omega^{(1)}\sigma_\beta(m\Delta_1^{(1)}) + (1 - \omega^{(1)})\sigma_\beta(m\Delta_2^{(1)})] + (1 - p_B^{(2)})[\omega^{(2)}\sigma_\beta(m\Delta_1^{(2)}) + (1 - \omega^{(2)})\sigma_\beta(m\Delta_2^{(2)})]}. \quad (\text{B3})$$

If there are weights $\omega^{(g)} \in [0, 1]$ corresponding to the intersection point (D_1^*, D_2^*) , then the strongly fragmented state exists. These states leave both markets equally active and as we showed in the discussion in Sec. V they entail benefits for the population as a whole, not favoring any of the symmetric groups.

Partially fragmented states. Finally, we outline a procedure to find a population steady state that is a combination of a bimodal (S) state in one group and a unimodal (U or W) state in the other, for $r \rightarrow 0$. A starting point for this search can be obtained by solving the homogeneous population equations (B2). When one of the groups is consistent with the homogeneous population assumption while the other is not, we can investigate whether the strongly fragmented solution for this other population exists. To find these states, we assume that the group that is inconsistent with a given

In the second case study ($p_B = 0.55$), the continuation of the low- β fixed point is a solution we can consistently find by this method for a wide range of intensities of choice, much wider than when the groups have more pronounced buy and sell preferences. Crossing the dark violet line in the phase diagram (Fig. 11), the new fixed points that arise all turn out to be inconsistent with the homogeneous population assumption until very high intensities of choice. This is why we need to employ different techniques to find the other solutions presented in Fig. 10. Only when the intensity of choice is increased further do partially fragmented states cease to exist, and solutions consistent with the homogeneous population assumption return.

Strongly cofragmented state (S,S). To find if these states exist we apply a procedure based on a Maxwell construction argument outlined in Sec. IV A and in [37] for a population consisting of a single group. For each group we define a locus in the space of order parameters (D_1, D_2) for which the strong fragmentation condition (8) is satisfied. If there is an intersection (D_1^*, D_2^*) between the two loci there are market demand-to-supply ratios in which both groups favor a strongly fragmented state. We finally need to confirm that the two order parameters can be created if only the two fragmented groups trade on the markets. If we assume the strongly fragmented distributions are of the form

$$P(\Delta | p_B^{(g)}) = \omega^{(g)}\delta(\Delta - \Delta_1^{(g)}) + (1 - \omega^{(g)})\delta(\Delta - \Delta_2^{(g)}),$$

then the corresponding order parameters are

homogeneous population solution is in the fragmented state. Thus possible order parameters for this state are on the locus defined by the Maxwell construction. For every pair (D_1, D_2) from the fragmented state locus we investigate the free energy of the second group (whether it is unfragmented or weakly fragmented). We find the peak position and represent the attraction distribution as a unimodal distribution centered at the (global) free energy minimum. We only need to examine whether by peak weight redistribution of the strongly fragmented group we can retrieve the initial order parameters (D_1, D_2) . When this is possible, the partially fragmented state exists. In the example shown in Fig. 10, due to mild buy and sell preferences, when one of the groups is fragmented there are two unfragmented options for the second group, corresponding to specialization to either of the two markets.

[1] H. Mendelson, *J. Financ. Quant. Anal.* **22**, 189 (1987).

[2] B. Chowdhry and V. Nanda, *Rev. Financ. Stud.* **4**, 483 (1991).

[3] A. Madhavan, *Rev. Financ. Stud.* **8**, 579 (1995).

[4] B. Biais, L. Glosten, and C. Spatt, *J. Financ. Mark.* **8**, 217 (2005).

[5] P. Bennett and L. Wei, *J. Financ. Mark.* **9**, 49 (2006).

- [6] H. Degryse, F. de Jong, and V. van Kervel, *Rev. Financ.* **19**, 1587 (2015).
- [7] S. Buti, B. Rindi, and I. M. Werner, *J. Financ. Econ.* **124**, 244 (2017).
- [8] C. Castellano, S. Fortunato, and V. Loreto, *Rev. Mod. Phys.* **81**, 591 (2009).
- [9] J.-P. Bouchaud, *Europhys. News* **50**, 24 (2019).
- [10] D. Helbing and P. Molnar, *Phys. Rev. E* **51**, 4282 (1995).
- [11] M. Moussaïd, D. Helbing, and G. Theraulaz, *Proc. Natl. Acad. Sci. USA* **108**, 6884 (2011).
- [12] G. Tedeschi, G. Iori, and M. Gallegati, *J. Econ. Behav. Organ.* **81**, 82 (2012).
- [13] S. Galam, Y. Gefen, and Y. Shapir, *J. Math. Sociol.* **9**, 1 (1982).
- [14] J. Fernández-Gracia, K. Suchecki, J. J. Ramasco, M. San Miguel, and V. M. Eguíluz, *Phys. Rev. Lett.* **112**, 158701 (2014).
- [15] D. Challet, M. Marsili, and Y.-C. Zhang, *Physica A* **294**, 514 (2001).
- [16] D. Challet and M. Marsili, *Phys. Rev. E* **68**, 036132 (2003).
- [17] Z.-G. Huang, J.-Q. Zhang, J.-Q. Dong, L. Huang, and Y.-C. Lai, *Sci. Rep.* **2**, 703 (2012).
- [18] M. Pagano, *Q. J. Econ.* **104**, 255 (1989).
- [19] G. Ellison, D. Fudenberg, and M. Möbius, *J. Eur. Econ. Assoc.* **2**, 30 (2004).
- [20] B. Shi, E. H. Gerding, P. Vytelingum, and N. R. Jennings, *Autonomous Agents Multi-Agent Syst.* **26**, 245 (2013).
- [21] B. Caillaud and B. Jullien, *RAND J. Econ.* **34**, 309 (2003).
- [22] G. Shorter and R. S. Miller, Dark pools in equity trading: Policy concerns and recent developments, Congressional Research Service Report No. R43739, Library of Congress, Washington, DC (2014).
- [23] P. Gomber, S. Sagade, E. Theissen, M. C. Weber, and C. Westheide, *J. Econ. Surv.* **31**, 792 (2017).
- [24] A. P. Kirman and N. J. Vriend, in *Interaction and Market Structure*, edited by D. Gatti, M. Gallegati, and A. Kirman, Lecture Notes in Economics and Mathematical Systems Vol. 484 (Springer, Berlin, 2000), pp. 33–56.
- [25] K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps, Overview of CAT: A market design competition, University of Liverpool, Department of Computer Science Report No. ULCS-09-005, 2009 (unpublished), Version 2.0.
- [26] A. Alorić, P. Sollich, and P. McBurney, in *Advances in Artificial Economics*, edited by F. Amblard, F. J. Miguel, A. Blanchet, and B. Gaudou, Lecture Notes in Economics and Mathematical Systems Vol. 676 (Springer International, Cham, 2015), pp. 79–90.
- [27] A. Alorić, P. Sollich, P. McBurney, and T. Galla, *PLoS One* **11**, e0154606 (2016).
- [28] C. Camerer and T. H. Ho, *Econometrica* **67**, 827 (1999).
- [29] T. H. Ho, C. Camerer, and J.-K. Chong, *J. Econ. Theory* **133**, 177 (2007).
- [30] D. K. Gode and S. Sunder, *J. Polit. Econ.* **101**, 119 (1993).
- [31] J. Duffy, *Handb. Comput. Econ.* **2**, 949 (2006).
- [32] D. Ladley, *Knowl. Eng. Rev.* **27**, 273 (2012).
- [33] M. Anufriev, J. Arifovic, J. Ledyard, and V. Panchenko, *J. Evol. Econ.* **23**, 539 (2013).
- [34] N. Hanaki, A. Kirman, and M. Marsili, *J. Econ. Behav. Organ.* **77**, 382 (2011).
- [35] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning about a Highly Connected World* (Cambridge University Press, Cambridge, 2010).
- [36] R. Nicole and P. Sollich, *PLoS One* **13**, e0196577 (2018).
- [37] A. Alorić, Spontaneous segregation of adaptive agents in auctions, Ph.D. thesis, King's College London, 2017.
- [38] N. G. van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier Science, Amsterdam, 1992).
- [39] K. Binder, *Phys. Rev. Lett.* **47**, 693 (1981).
- [40] Y.-W. Cheung and D. Friedman, *Games Econ. Behav.* **19**, 46 (1997).
- [41] J. Huber, *J. Econ. Dyn. Control* **31**, 2536 (2007).
- [42] B. Toth, E. Scalas, J. Huber, and M. Kirchler, *Eur. Phys. J. B* **55**, 115 (2007).