Microfoundations for diffusion price processes

Mikko S. Pakkanen*

February 18, 2010

Abstract

We study microeconomic foundations of diffusion processes as models of stock price dynamics. To this end, we develop a microscopic model of a stock market with finitely many heterogeneous economic agents, who trade in continuous time, giving rise to an endogeneous pure-jump process describing the evolution of stock prices over time. When the number of agents in the market is large, we show that the price process can be approximated by a diffusion, with price-dependent drift and volatility coefficients that are determined by small excess demands and trading volume in the microscopic model. We extend the microscopic model further by allowing for non-market interactions between agents, to model herd behavior in the market. In this case, price dynamics can be approximated by a process with stochastic volatility. Finally, we demonstrate how heavy-tailed stock returns emerge when agents have a strong tendency towards herd behavior.

Keywords: Stock prices; microfoundations; diffusion processes; stochastic volatility; heavy tails
2000 Mathematics Subject Classification: 91B26; 60F17; 60J75
JEL Classification: C65; D53; G12

1 Introduction

In mathematical finance, macroscopic price dynamics of a risky financial asset are often modeled by some given, exogeneous *diffusion process* $(S_t)_{t\in[0,T]}$, a strong Markov process defined by stochastic differential equation

$$\mathrm{d}S_t = \mu(S_t)S_t\mathrm{d}t + \sigma(S_t)S_t\mathrm{d}W_t,$$

*Department of Mathematics and Statistics, University of Helsinki, P.O. Box 68, FI-00014 Helsingin yliopisto, Finland. E-mail: `msp@iki.fi`

driven by a standard Brownian motion $(W_t)_{t\in[0,T]}$. To cater for various various stylized facts (e.g. heavy-tailed return distributions and intermittency) observed in financial time series, these one-dimensional diffusion models have been extended further by allowing the volatility coefficient σ to depend on some additional diffusion process, leading to *stochastic volatility* models. While these diffusion models are a cornerstone of continuous-time financial modeling, not many attempts have been made to find *microfoundations* for them—that is, to derive diffusions from "first principles," viz. microeconomic considerations concerning trading mechanisms and behavior of individual economic agents in the market. This is the case especially with stochastic volatility models.

The paucity of such studies is not necessarily surprising, since already the question of what a *derivation* entails, is not obvious. Namely, if we consider a financial market from a microeconomic perspective, we observe that

1. asset prices move *endogeneously* according to actions of agents,
2. for an individual agent, continuous, *dynamic trading* is impossible,
3. the number of agents is *finite*,

and consequently, asset price dynamics follow a (piecewise-constant) *pure-jump process*. Thus, to develop a microscopic market model which gives rise to a *continuous* price process, or a diffusion, in particular, we should be able to gloss over at least one of these three features. From a microscopic point of view, the first one is obviously crucial. However, the significance of the remaining two might be lesser: while *perfect* dynamic trading is impossible, an agent can still trade with a very high frequency, and moreover, while there are only finitely many agents, for practical purposes the number of them may be very large.

The purpose of this paper is to explore the latter case, specifically to develop a model of a market with finitely many agents, in which the resulting price process can be *approximated* by a diffusion or a stochastic volatility model, provided that the number of agents is large enough. The precise mathematical justification of the approximation is given as a functional limit theorem, which ensures the convergence of prices to such a process, when we pass to the *large-market* limit with infinitely many agents.

1.1 Existing approaches

Previously, Föllmer and Schweizer [10], and Horst [15] have studied discrete-time models, in which asset prices are the outcome of a sequence of temporary equilibria determined by random excess demand functions of agents in the market. In their models, the asset price process is a first-order stochastic recurrence in a random environment, which converges, assuming certain conditions on the random environment, to an Ornstein–Uhlenbeck process in a random environment. Frey and Stremme [11] take a somewhat similar approach, with a model they use to study effects of demand generated by hedging strategies—in the limit, they obtain a diffusion price process with time- and price-dependent volatility, as a consequence of hedging.

The market mechanisms of the aforementioned models mimic that of a *batch market*, where excess demands of *several* agents are matched *simultaneously*. However, as pointed out in Bayraktar et al. [2, p. 646], most modern financial markets operate in a *continuous* fashion, with buy and sell orders arriving *asynchronously* in continuous time. In such a situation, there appears to be no invisible hand that could both compute and enforce *Walrasian* equilibrium prices matching excess demands of several agents simultaneously—even for a moment. To gain a better understanding of the price dynamics arising in these *continuous markets*, Bayraktar et al. [2], and Horst and Rothe [16] have proposed models based on ideas from *queueing theory*. Pointing to a parallel between buy and sell orders arriving to the market and customers arriving to a queueing system, they derive pure-jump processes describing price dynamics when the arriving orders are unit-sized and executed immediately by a market maker, leading to constant-sized price jumps. These price processes are similar to the processes measuring the number of customers in a state-dependent queueing network. Using approximation methods developed for such queueing networks, the authors prove that when the number of agents in the market tends to infinity, the price processes converge to solutions to (deterministic) differential equations—or *fluid limits* in the parlance of queueing theory. Further, they show that rescaled price fluctuations around these fluid limits converge to various stochastic processes, such as fractional diffusions and solutions to stochastic delay differential equations, depending on the behavioral assumptions made about agents.

1.2 Motivation and main results

As the models of Bayraktar et al. [2], and Horst and Rothe [16] seem to be primarily geared towards understanding microfoundations of *long-range dependence* of prices and macroscopic effects of financial strategies that have *delayed* dependencies on past prices, respectively, their stochastic large-market limits of the price processes are somewhat specific, covering only certain special cases of (Markovian) diffusions. Moreover, with these models, the limiting price processes exhibit deterministic volatility, while an emergence of stochastic volatility is only hinted at (see [2, Section 4.1]). This leads to a natural question, whether more general diffusion models, and in particular, ones with stochastic volatility, could be obtained as large-market limits of the price processes in similar models of continuous markets. Of course, at the same time, one would like to determine how these limiting diffusions depend on the behavioral characteristics of the agents in the market.

To study these questions, we develop in this paper a microscopic model of a stock market where finitely many agents trade shares of a stock asynchronously in continuous time. While the model is similar to the ones in [2, 16], there are also some dissimilarities. The most noteworthy of them are that we allow for more general heterogeneity among agents—going beyond boundedly many agent types—and that the price impacts of individual trades are larger asymptotically. These are crucial features that give rise to non-deterministic and unbounded volatility in the large-market limit. It is, however, worth pointing out that compared to the model of this paper, in [16] agents' trades are allowed to depend on past prices in a more sophisticated way, and in [2] the distributions of times elapsed between trades are more general.

In Section 2, we formulate the first version our microscopic market model. At this stage, the agents in the market interact with each other solely through trading. For this model, we prove that the resulting price process can be approximated by a general one-dimensional diffusion, the *drift* and *volatility coefficients* of which are connected to *small excess demands* and *trading volume* in the microscopic model. As an example, how this diffusion approximation can be applied, we consider a more specific model, in which agents consists of *fundamentalists* and *noise traders*, similarly to the models of Föllmer and Schweizer [10], and Horst [15]. Additionally, to demonstrate that the diffusion approximation indeed applies to settings with *completely* heterogeneous agents, we allow in this example some of the behavioral rules of the agents to originate from a *random environment*.

The second, extended version of the microscopic market model is considered in Section 3. In this model, agents are allowed to have *non-market* interactions, i.e. interactions besides trading. To be specific, agents' behavior is influenced by an *opinion index*, a composite of agents' individual opinions of the attractiveness of the stock, to model *herd behavior* in the market. Then, extending the earlier diffusion approximation, we show that in the large-market limit, the price of the stock follows a diffusion with *stochastic volatility*, the volatility process of which is a function of the limit of the fluctuations of the opinion index. Since the volatility process is possibly unbounded, the price process may exhibit *heavy-tailed* fluctuations. We show that when the behavior of agents depends very strongly on the opinion index. the log returns generated by the price process in the limit can indeed be heavy-tailed, in the sense that they have infinite second moments, and in some cases approximately *power-law tails*. This result complements some earlier studies (see e.g. [5]) that have identified herd behavior as a potential mechanism behind heavy tails. We conclude with numerical simulations of log returns, evaluating how well the power-law tails approximate the tails of log returns in the large-market limit and in markets with finitely many agents.

2 Microfoundations for one-dimensional diffusions

2.1 Basic microscopic market model

We consider a simple model of a continuous market where a single stock is traded. The market is formed by a finite number of ordinary agents, indexed by the set $\mathbb{A}_n := \{1, \ldots, n\}$, who trade the stock, and a *market maker*, or a *specialist*, who handles the trading. Similarly to [2, 16], we assume that all trading occurs through the market maker—that is, no direct transactions between ordinary agents take place. Essentially, this setup is what Garman [13] calls a *dealership market*. The agents we consider are *boundedly rational*, and their behavior is determined by "minimal intelligence"-type of simple randomized rules, instead of utility maximation based on intertemporal preferences or game theoretic considerations. This methodology is frequently used in both theoretical and computational models of *market microstructure*, see e.g. [2, 9] for motivation. Agents' actions are allowed to depend on previous prices through the price in the most recent trade. Economically, this simplifying assumption of "Markovian feedback" means that the agents are proponents of the weak form of the *efficient markets hypothesis*, while mathematically, the assumption allows us to use techniques developed for *Markov processes*, leaving the model rather tractable.

While a finite time horizon would be a natural choice for this model, which essentially attempts to describe *short-term* stock price dynamics,[1] we extend the time interval to be $[0, \infty)$, purely out of mathematical convenience. In any finite time interval, only a finite number of trades occur, hence only countably many trades occur ultimately. For each $k \in \mathbb{Z}_+$ ($\mathbb{Z}_+ := \{1, 2, \ldots\}$) we log the k-th trade by triple (T_k, A_k, P_k), where T_k is the time when the trade occurs, A_k is the agent

[1] Long-term price dynamics might be affected by institutional changes, evolution of trading strategies through learning, dramatic news, or other "structural breaks" which this model does not attempt to capture.

who trades with the market maker at that time, and P_k is the *logarithmic price per share*—which we henceforth simply call the *price*—in this trade. Moreover, $\mathscr{G}_k := \sigma\{T_i, A_i, P_i : i \leq k\}$ represents the market history up to the k-th trade. The behavior of the agents is characterized by pairs (λ_a, e^n_a), $a \in \mathbb{A}_n$ of functions, the role of which is described in the following.

1. Trading intensity functions λ_a, $a \in \mathbb{A}_n$. The continuous, bounded *trading intensity function* $\lambda_a : \mathbb{R} \to \mathbb{R}_+$ ($\mathbb{R}_+ := (0, \infty)$) quantifies agent a's propensity to trade the asset, given the price in the most recent trade. Namely, the probability that a certain agent is the one who trades next, given the latest price, is proportional to her trading propensity. Specifically, we set

$$\mathbf{P}(A_k = a | \mathscr{G}_{k-1}) := \frac{\lambda_a(P_{k-1})}{\lambda_{\mathbb{A}_n}(P_{k-1})}, \quad a \in \mathbb{A}_n, \tag{2.1}$$

where the normalizing factor is the *aggregate trading intensity*

$$\lambda_{\mathbb{A}_n} := \sum_{a=1}^{n} \lambda_a.$$

To ensure that (2.1) makes sense, we assume that regardless of the latest price, there are always agents who are willing to trade, that is, $\lambda_{\mathbb{A}_n}(x) > 0$ for all $x \in \mathbb{R}$. Moreover, we assume that, given the market history, the k-th *intratrade time* $\tau_k := T_k - T_{k-1}$ is exponentially distributed with rate parameter $\lambda_{\mathbb{A}_n}(P_{k-1})$, that is

$$\mathbf{P}(\tau_k \in [0, t] | \mathscr{G}_{k-1}) = 1 - e^{-\lambda_{\mathbb{A}_n}(P_{k-1})t}, \quad t \geq 0. \tag{2.2}$$

Informally, this means that, given the market history, the expected number of trades occuring in a short time interval equals approximately the length of the interval multiplied by the aggregate trading intensity suggested by the latest price.

Remark 2.3. More precisely, once technical Assumption 2.8 (to be formulated below) is in force, one can show that given $\mathscr{G}_{k-1}$, the expected number of trades taking place between T_{k-1} and $T_{k-1} + h$ is $\lambda_{\mathbb{A}_n}(P_{k-1})h + o(h)$ when $h \downarrow 0$.

Remark 2.4. The condition (2.2), which is necessary to ensure that the resulting price process is Markovian, is equivalent to the perhaps more intuitive assumption that intratrade times are *memoryless*, in the sense that

$$\mathbf{P}(\tau_k > t + h | \tau_k > h, \mathscr{G}_{k-1}) = \mathbf{P}(\tau_k > t | \mathscr{G}_{k-1}), \quad t, h \geq 0.$$

This can be interpreted so that any recent absence of trades neither increases nor decreases the probability that a trade occurs in the future.

2. Excess demand functions e_a^n, $a \in \mathbb{A}_n$. The Borel measurable *excess demand function* $e_a^n : \mathbb{R}^2 \to \mathbb{R}$ quantifies *how much* agent a would buy or sell, were she to trade. More precisely, this works so that at time T_k, agent A_k contacts the market maker and places a market order for

$$e_a^n(P_{k-1}, \xi_k)$$

shares of the stock, where positive (resp. negative) amounts indicate buy (resp. sell) orders. In addition to the price of the most recent trade, we thus allow the amount to depend on exogeneous randomness, via the random variable ξ_k, which we call a *signal*. The signals $(\xi_k)_{k=1}^{\infty}$ can be considered e.g. as new information concerning macroeconomic variables which are relevant to the valuation of the stock, or alternatively, they can simply be interpreted as unforeseen factors affecting the behavior of the agents. In any case, we assume that ξ_k is independent of $\mathscr{G}_{k-1}$ and A_k. We might want to allow ξ_k to depend on A_k, or possibly even on P_{k-1}, but this is not really necessary, since such dependence can readily be "factored" into excess demand functions.

The market maker carries a sufficient inventory of shares (or alternatively, is allowed to *short*), and she executes the market order instantaneously, setting the price using a *pricing rule* $r_n : \mathbb{R}^2 \to \mathbb{R}$, a Borel measurable function. Namely, the price P_k in this trade is given by

$$P_k = r_n(e_{A_k}^n(P_{k-1}, \xi_k), P_{k-1}). \tag{2.5}$$

Thus, we allow the market maker to use the price of the latest trade as a reference price. As a deterministic function, the pricing rule r_n is known to agents. Consequently, the agents know also the *quotes* $r_n(q, P_{k-1})$, $q \in \mathbb{R}$, and they may use the quotes to decide their excess demands $e_a^n(P_{k-1}, \xi_k)$, $a \in \mathbb{A}_n$. We could allow the pricing rule to depend on some source of randomness (like the excess demand functions) e.g. to model variable liquidity, but for simplicity, we refrain from doing so.

Additionally, we introduce an initial price P_0 with some cumulative distribution function F_{P_0} (which does not depend on n, for simplicity), and fix $T_0 := 0$. The evolution of the price over time is now conveniently described by the càdlàg process

$$X_t^n := \sum_{k=0}^{\infty} P_k \mathbf{1}_{[T_k, T_{k+1})}(t), \quad t \geq 0. \tag{2.6}$$

We complete the formulation of the model by introducing two technical assumptions that ensure that the model can be given a probabilistic construction so that $(X_t^n)_{t \in [0,\infty)}$ becomes a time-homogeneous Markov process.

Assumption 2.7 (Signals)**.** The random variables $(\xi_k)_{k=1}^{\infty}$ are i.i.d. according to a cumulative distribution function F_ξ.

Assuming independence of the signals is natural since they are thought convey unforeseen exogeneous factors. Requiring that the signals are identically distributed is slightly more restrictive—however, since agents respond to the signals heterogeneously, effectively this assumption only rules out time-dependent signals.

Additionally, we need to be a bit more elaborate with the distributional specification (2.2) and fix the joint distribution of prices $(P_k)_{k=0}^{\infty}$ and intratrade times $(\tau_k)_{k=1}^{\infty}$ as follows.

Assumption 2.8 (Intratrade times)**.** The intratrade times $(\tau_k)_{k=1}^{\infty}$ are given by

$$\tau_k := \frac{\gamma_k}{\lambda_{\mathbb{A}_n}(P_{k-1})}, \quad k \in \mathbb{Z}_+, \tag{2.9}$$

where the $(\gamma_k)_{k=1}^{\infty}$ are i.i.d. random variables independent of $(P_k)_{k=0}^{\infty}$, such that γ_1 is exponentially distributed with rate parameter 1.

Using now some well-known characterizations of Markov chains and pure-jump Markov processes (see e.g. Proposition 8.6 and Theorem 12.18 of [20]), it is straightforward show that a Markovian construction of the model exists.

Lemma 2.10 (Markovianity)**.** *If the preceding assumptions hold, then there exists a probability space* $(\Omega, \mathscr{F}, \mathbf{P})$ *that carries the model, such that* $(X_t^n)_{t\in[0,\infty)}$ *is a time-homogeneous pure-jump Markov process with rate kernel*

$$K_n(x, \mathrm{d}y) = \lambda_{\mathbb{A}_n}(x) k_n(x, \mathrm{d}y), \tag{2.11}$$

where the kernel $k_n(x, \mathrm{d}y)$ *is a regular version of the conditional distribution* $\mathbf{P}(P_1 - P_0 \in \mathrm{d}y | P_0 = x)$.

2.2 Large-market diffusion approximation for the price process

We will now discuss the convergence of the price process $(X_t^n)_{t\in[0,\infty)}$ of the stock to a one-dimensional diffusion, as the number of agents tends to infinity, that is $n \to \infty$. For our purposes, the most useful notion of convergence of price processes is the weak convergence of their probability distributions. To this end, we regard price processes as random elements of the space $D[0, \infty)$ of real-valued càdlàg functions on $[0, \infty)$, equipped with the usual Skorohod topology (see e.g. [3, 8, 18] for more details). In what follows, we will implicitly assume that $(X_t^n)_{t\in[0,\infty)}$ is a Markov process, as per Lemma 2.10, carried by probability space $(\Omega, \mathscr{F}, \mathbf{P})$.

Regarding notation, all asymptotic statements are for $n \to \infty$, unless stated otherwise, "u.o.c." is an abbreviation for "uniformly on compact sets", $\mathbf{E}$ denotes expectation with respect to $\mathbf{P}$, and $\xrightarrow{\mathscr{D}}$ (resp. $\stackrel{\mathscr{D}}{=}$) denotes convergence (resp. equality) in distribution.

The model formulated in the preceding section is clearly too general to be always able to admit a diffusion approximation—indeed we can easily find counterexamples with price movements that are nowhere near "diffusive" when passing to the large-market limit. However, since the model is Markovian, it suffices to ensure that, firstly, the limit exists, and secondly, that no *jumps* appear in the limit. In what follows, we shall see that the model does admit a diffusion approximation if it conforms to the following postulates:

P^1 *Small price impacts.* Impacts of inviduals trades on the price are small.

P^2 *Market is cleared.* Supply and demand do not diverge excessively.

P^3 *Steady flow of trades.* The flow of trades does not contain "bursts" or very large, erratic trades.

We shall elaborate on these postulates, and give a precise mathematical formulation of what they entail in the model, below. If some of these postulates are violated, a limit might still exist, but there could be jumps. Any careful analysis of this case is, however, beyond the scope of this paper.

Of course, alternatively we might try to be much more specific with the formulation of the model, by assigning certain functional forms to excess demand and trading intensity functions. However, it seems preferable to be "non-parametric" at this stage, since it allows us to identify the key quantities whose asymptotic behavior is crucial for the diffusion approximation, and further, it accommodates the possibility that these functions might be (completely) random. A more concrete example is deferred to Section 2.3.

Let us first consider P^1. Note that unlike perfectly competitive markets, the centralized *dealership* market, at least in the simple form used in this paper and in [2, 16], has *a priori* no intrinsic mechanism that would make price impacts of individual trades small when the number of agents is large. For this reason, it is necessary to assume that the market maker uses a pricing rule that gives rise to this effect asymptotically.

Assumption 2.12 (P^1, Pricing rule)**.** The pricing rule r_n of the market maker is given by

$$r_n(q,x) = x + \frac{\alpha}{\sqrt{n}} q + u_n(q,x), \quad q,x \in \mathbb{R}, \tag{2.13}$$

where $\alpha > 0$, and u_n, $n \in \mathbb{Z}_+$ are Borel measurable functions such that for any $\delta > 0$ there exists positive constants C_δ^n, $n \in \mathbb{Z}_+$ such that $C_\delta^n = o(n^{-1})$, and

$$\sup_{|x|\leq\delta} |u_n(q,x)| \leq C_\delta^n |q|$$

for all $q \in \mathbb{R}$, $n \in \mathbb{Z}_+$.

Essentially, Assumption 2.12 means that the market maker has a pricing rule which is in the large market setting nearly *affine* (allowing for certain deviations from this through u_n) with a very small slope, corresponding to the idea that individual trades have only small impacts. We choose the affine functional form mostly for the sake of tractability, but affine pricing rules (also known as *linear price impact functions*) can be seen to be economically motivated by the desire the rule out some forms of price manipulation, see e.g. [17].

Remark 2.14. We have chosen the price impact of an individual trade to be proportional to $1/\sqrt{n}$, while in the models in [2, 16], a unit-sized trade changes the price of the stock by $1/n$. While these choices of scaling produce the same qualitative effect (P^1), the resulting limits are rather distinct from each other. As we shall see, our choice of scaling leads to limits that typically exhibit *non-deterministic* volatility, as opposed to the stochastic limits in [2, 16], the deterministic volatility of which depends at most on the corresponding fluid limits (see, however, Section 4.1 of [2]). Ultimately, the choice of scaling depends on what one wants to model—it seems that $1/\sqrt{n}$ is suited to the study of short-term fluctuations and volatility, whereas $1/n$ is perhaps more appropriate in studies of long-term behavior.

Contrary to *equilibrium* models that usually admit equilibrium prices that balance supply and demand, in the *disequilibrium* model we are studying, postulate P^2 need not hold, even approximately. Asymptotically, bought shares may outnumber sold shares (or vice versa) so excessively that the scaling in Assumption 2.12 is defeated. To be more precise, let us first note that conventional concepts of aggregate supply and demand are not meaningful because trades occur asyncronously, one by one, and after each trade, agents' trading intensities and excess demands change. However, by taking expectations, a natural surrogate concept of *expected aggregate excess demand* per capita at each price level $x \in \mathbb{R}$ can be defined by

$$z_n(x) := \frac{1}{n} \sum_{a=1}^{n} \lambda_a(x) \mathbf{E}[e_a^n(x, \xi_1)].$$

The expected aggregate excess demand of all agents is thus nz_n. Since the price impact of one traded share is by Assumption 2.12 proportional to $1/\sqrt{n}$, this suggests that to ensure finiteness of prices in the limit, z_n should be (at most) proportional to $1/\sqrt{n}$.

Assumption 2.15 (P^2, Excess demands)**.** The function z_n is well-defined (and finite) for all $n \in \mathbb{Z}_+$, and there exists a continuous function $z : \mathbb{R} \to \mathbb{R}$ such that

$$z_n = \frac{z}{\sqrt{n}} + o(n^{-1/2}) \quad \text{u.o.c.}$$

While Assumption 2.15 implies that bought shares offset sold shares on average, since we have not imposed any uniform bounds on the excess demand and trading intensity functions of the agents, there is still a possibility of short, but excessive bursts of trades, or individual large trades that might prevent any finite limit for the price. This necessitates P^3, which we formulate by demanding that excess demands of the agents are (locally) uniformly square integrable and that trading intensity per capita does not explode asymptotically on compact sets. Note that when the excess demand and trading intensity function of the agents are uniformly bounded, these requirements are indeed met.

Assumption 2.16 (P^3, No explosive trading)**.** For every $\delta > 0$,

1. $\left\{e_a^n(x, \xi_1)^2 : |x| \leq \delta, a \in \mathbb{A}_n, n \in \mathbb{Z}_+\right\}$ is uniformly integrable, and
2. $\limsup_{n\to\infty} \sup_{|x|\leq\delta} \left|\frac{\lambda_{\mathbb{A}_n}(x)}{n}\right| < \infty.$

Finally, in addition to the postulates P^1, P^2, and P^3, which we have now formulated precisely, we need to measure *trading volume* in the market. Analogously to expected aggregate excess demand, we may define a similar concept of *expected trading volume* per capita at each price level $x \in \mathbb{R}$ by

$$v_n(x) := \left(\frac{1}{n}\sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[e_a^n(x, \xi_1)^2]\right)^{1/2}.$$

This is, in effect, a *root mean square* of the expected trades, however weighted by trading intensities. We assume that the functions v_n, $n \in \mathbb{Z}_+$ have a limit.

Assumption 2.17 (Trading volume)**.** There exists a continuous function $v : \mathbb{R} \to \mathbb{R}_+$ such that

$$v_n \xrightarrow{\text{u.o.c.}} v. \tag{2.18}$$

These assumptions have left us with three asymptotic, large-market quantities: α, $z(\cdot)$, and $v(\cdot)$ measuring price impact, small excess demands, and trading volume, respectively. We can now state the diffusion approximation in terms of them as follows.

Theorem 2.19 (Diffusion approximation)**.** *Suppose that the preceding assumptions hold. If z and v are locally Lipschitz and of linear growth, then*

$$(X_t^n)_{t\in[0,\infty)} \xrightarrow{\mathscr{D}} (X_t)_{t\in[0,\infty)} \quad \text{in } D[0,\infty),$$

where $(X_t)_{t\in[0,\infty)}$ is the unique strong solution to the stochastic differential equation

$$\mathrm{d}X_t = \alpha z(X_t)\mathrm{d}t + \alpha v(X_t)\mathrm{d}W_t, \quad X_0 = \zeta, \tag{2.20}$$

where the process $(W_t)_{t\in[0,\infty)}$ is a standard Brownian motion and ζ is distributed according to F_{P_0} and independent of $(W_t)_{t\in[0,\infty)}$.

Proof. See Appendix A.1. □

2.3 Concrete example: a fundamentalist–noise trader model

Since the model we have studied so far is rather general in the sense that it does not fix any concrete form for its diffusion approximation, it is instructive consider how Theorem 2.19 applies in a more concrete setting.

Let us consider a model with agents similar to ones in the models studied by Föllmer and Schweizer [10], and Horst [15]. Namely, suppose that some of the agents in $\mathbb{A}_n$ are *fundamentalists*, and the remaining agents are *noise traders*. Fundamentalists' behavior is characterized by their belief that the price of the stock should eventually settle to some *fundamental level*, e.g. the present value of expected dividends to be paid to the shareholders. For simplicity, let us assume that this level is time-invariant and agreed upon by the fundamentalists, and denote it by $F \in \mathbb{R}$. Let us set $f_a := 1$ if $a \in \mathbb{A}_n$ is a fundamentalist, and $f_a := 0$ if not. Further, let $k_n := |\{a \in \mathbb{A}_n : f_a = 1\}|$ denote the number of fundamentalists in $\mathbb{A}_n$, and write $\mathbb{A}_\infty := \cup_{n\in\mathbb{Z}_+}\mathbb{A}_n$. It is natural to expect that the proportion k_n/n of fundamentalists converges to a limit as $n \to \infty$. We shall denote the limit by $\phi \in [0,1]$.

We assume that each agent's excess demand function is a sum of a systematic demand and a purely random demand. The systematic demand accommodates fundamentalists' key behavioral trait that if the price in the latest trade has been smaller (resp. greater) than F, they consider the stock cheap (resp. expensive) and hence are willing to buy (resp. sell) shares. Noise trader's systematic demand equals zero identically, as their excess demands shall be purely random. To define the excess demand functions, let $w : \mathbb{R} \to \mathbb{R}$ be an increasing continuous function, such that $w(0) = 0$. We set

$$e_a^n(x,s) := \begin{cases} n^{-1/2}w(F-x) + s, & f_a = 1, \\ s, & f_a = 0. \end{cases}$$

Thus, random demands arise from the signals $(\xi_n)_{n=1}^{\infty}$, regarding which we assume $\mathbf{E}[\xi_1] = 0$ and $\sigma_\xi^2 := \mathbf{E}[\xi_1^2] < \infty$. We have scaled fundamentalists' systematic demand by $n^{-1/2}$ to ensure that it does not generate, on aggregate, too large excess demands. As for the market maker, we assume that she uses the simple affine pricing rule

$$r_n(q,x) := x + \frac{\alpha}{\sqrt{n}} q, \tag{2.21}$$

where $\alpha > 0$. (See the discussion below Assumption 2.12 regarding affine pricing rules.)

As for trading intensity functions, we assume that agents have chosen them *at random* before trading has begun. To make this precise, we assume that λ_a, $a \in \mathbb{A}_\infty$ are independent random functions carried by a complete probability space $(\tilde{\Omega}, \mathscr{G}, \mathbf{Q})$, separate from $(\Omega, \mathscr{F}, \mathbf{P})$. That is, for any $a \in \mathbb{A}_\infty$, λ_a is a measurable map $\tilde{\Omega} \to C_b(\mathbb{R}, \mathbb{R}_+)$, where $C_b(\mathbb{R}, \mathbb{R}_+)$ is the space of bounded continuous functions $\mathbb{R} \to \mathbb{R}_+$, equipped with the topology of uniform convergence on compact sets, and the corresponding Borel σ-algebra. For technical reasons, we need to assume some integrability, namely that for all $a \in \mathbb{A}_\infty$ and $\delta > 0$,

$$\mathbf{E}^{\mathbf{Q}}\Big[\sup_{|x|\le\delta} \lambda_a(x)\Big] < \infty. \tag{2.22}$$

Moreover, we assume that the intensity functions are identically distributed within the groups of fundamentalists and noise traders. Finally, we define the expected trading intensities among fundamentalists and noise traders at each price level $x \in \mathbb{R}$ by

$$\overline{\lambda}_i(x) := \mathbf{E}^{\mathbf{Q}}[\lambda_a(x)], \quad f_a = i,\ i \in \{0,1\}.$$

Now that the trading intensity functions are random, the price process $(X_t^n)_{t\in[0,\infty)}$ lives on the product space $(\Omega \times \tilde{\Omega}, \mathscr{F} \otimes \mathscr{G}, \mathbf{P} \otimes \mathbf{Q})$.[2] Then, typically $(X_t^n)_{t\in[0,\infty)}$ is not a Markov process, since the intensity functions of the agents are initially unknown (and may remain so) to the observer of the price process, but a Markov process in a *random environment* given by the realization $\tilde{\omega} \in \tilde{\Omega}$.

Nevertheless, when passing to the large-market limit, the price process converges to a (Markovian) diffusion.

Proposition 2.23 (Diffusion approximation)**.** *Suppose that the model is as described in this subsection and that Assumptions 2.8 and 2.7 hold. If the functions*

$$x \mapsto \overline{\lambda}_1(x) w(F-x), \quad x \mapsto \overline{\lambda}(x)^{1/2} := [\phi\overline{\lambda}_1(x) + (1-\phi)\overline{\lambda}_2(x)]^{1/2}$$

[2] The process $(X_t^n)_{t\in[0,\infty)}$ can be constructed so that it depends on the intensity functions in a measurable way, so no problems will arise in this respect.

are locally Lipschitz and of linear growth, then

$$(X_t^n)_{t\in[0,\infty)} \xrightarrow{\mathscr{D}} (X_t)_{t\in[0,\infty)} \quad \textit{in } D[0,\infty),$$

where $(X_t)_{t\in[0,\infty)}$ *is the unique strong solution to the stochastic differential equation*

$$\mathrm{d}X_t = \alpha\phi\overline{\lambda}_1(X_t)w(F - X_t)\mathrm{d}t + \alpha\sigma_\xi\overline{\lambda}(X_t)^{1/2}\mathrm{d}W_t, \quad X_0 = \zeta, \tag{2.24}$$

where the process $(W_t)_{t\in[0,\infty)}$ *is a standard Brownian motion and* ζ *is distributed according to* F_{P_0} *and independent of* $(W_t)_{t\in[0,\infty)}$.

Proof. See Appendix A.2. □

Remark 2.25. The limiting diffusion (2.24) is drawn towards the fundamental level F at a rate which depends on the mean trading intensity function $\overline{\lambda}_1$ and systematic excess demand w of the fundamentalists. This is somewhat analogous to the approximations obtained in [10, 15], in which the same rate is modulated by a random environment describing "investor sentiment." However, in our model randomness from the random environment is suppressed in the limit because of averaging, as opposed to the diffusive scaling of random environments in the aforementioned papers.

3 Microfoundations for stochastic volatility and heavy tails

In the model we have considered so far, the only form of *interaction* between agents is *trading* via the market maker. In this section, we consider an extended model, with agents that can also have interactions *beyond the market*. The purpose of this extension is to show how *stochastic volatility* and, in some cases, *heavy-tailed returns* appear as consequences of these interactions.

3.1 Extended microscopic market model with non-market interactions and herding

Let us consider the setup of the fundamentalist–noise trader model above, but with only noise traders present in the market. Unlike before, in addition to prices in preceding trades, we allow agents' behavior to depend on *opinions* of the stock. Each agent has an opinion on the attractiveness of the stock, and the opinion may change over time. When an agent finds the stock attractive, it simply means that she would like to engage in trading of the stock—without distinguishing between tendencies to buy and sell. Conversely, when she finds the stock unattractive, she is just indifferent to it. To model *herd behavior* (see e.g. [1] for further discussion on this behavioral trait), we assume, however, that when trading, agents choose to follow a general opinion on attractiveness, ignoring their individual opinions. Specifically, we define an *opinion index* that will be a composite of agents' individual opinions on the attractiveness of the stock, and agents' trading will be influenced by the current value of the index. Since this interaction structure is the main point of this extension, we shall eschew the generality present in the model of Section 2, and formulate the model to be otherwise as simple as possible.

We assume that agents' opinions are *dichotomous*—each agent considers the stock either attractive or unattractive. At each time point, an agent can choose to either trade the stock or change her opinion of the stock—we shall call these collectively *agent actions*. For each $k \in \mathbb{Z}_+$, let us denote by T_k the time of the occurence of the k-th agent action, and by B_k an indicator assuming value one, should the action be a trade. For each $a \in \mathbb{A}_n$, let m_k^a denote agent a's opinion at time T_k. In the case of a positive opinion, m_k^a assumes value 1 and in the case of a negative opinion, it assumes value -1. We define the opinion index at time T_k, denoted by M_k, to be

$$M_k := \frac{1}{\sqrt{n}} \sum_{a=1}^{n} m_k^a. \tag{3.1}$$

(Here we use the scaling $1/\sqrt{n}$, instead of the more natural $1/n$ because of technical requirements. However, this seems innocuous, as the point made about scaling in Remark 2.14 applies also here.) The market history up to the k-th action is defined to be $\mathscr{H}_k := \sigma\{T_i, B_i, A_i, P_i, M_i : i \leq k\}$.

We equip each agent $a \in \mathbb{A}_n$ with a trading intensity function $\lambda_a : \mathbb{R}^2 \to \mathbb{R}_+$ that quantifies her propensity to trade, given the latest price and the latest value of the opinion index. Further, we assume that agents' propensity to change their opinions can be quantified by a common *constant*[3] intensity $\overline{\mu} > 0$. Using these intensities, we define the conditional probabilities that agent $a \in \mathbb{A}_n$ performs the k-th agent action, and that she trades ($B_k = 1$) or changes her opinion ($B_k = 0$)

[3] Extending the model and the corresponding diffusion approximation (Proposition 3.8, below) to allow for a price-dependent $\overline{\mu}$ would be straightforward. However, for the Markovianity of the model, it is crucial that $\overline{\mu}$ is common to all agents.

to be

$$\mathbf{P}(A_k = a, B_k = i | \mathscr{H}_{k-1}) := \begin{cases} \dfrac{\lambda_a(P_{k-1}, M_{k-1})}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}}, & i = 1, \\ \dfrac{\overline{\mu}}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}}, & i = 0, \end{cases} \tag{3.2}$$

where $\lambda_{\mathbb{A}_n} := \sum_{a=1}^{n} \lambda_a$ and $\mu_{\mathbb{A}_n} := n\overline{\mu}$. Analogously to (2.2), given the market history, the *inter-action time* $\tau_k := T_k - T_{k-1}$ is taken to be exponentially distributed with rate parameter $\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}$, that is

$$\mathbf{P}(\tau_k \in [0, t] | \mathscr{H}_{k-1}) = 1 - e^{-(\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n})}, \quad t \geq 0.$$

Since the opinion index changes only when some of the agents changes her own opinion, we can deduce dynamics of the index from the probabilities (3.2). Suppose that $B_k = 0$. Then, if agent A_k's current opinion is 1 (resp. -1), the opinion index decreases (resp. increases) by $2/\sqrt{n}$. Hence, we have[4]

$$\begin{aligned} \mathbf{P}\left(M_k - M_{k-1} = -\frac{2}{\sqrt{n}} \middle| \mathscr{H}_{k-1} \right) &= \mathbf{P}\big(B_k = 0, m_{k-1}^{A_k} = 1\big) \\ &= \frac{\frac{1}{2}\left(1 + \frac{M_{k-1}}{\sqrt{n}}\right) \mu_{\mathbb{A}_n}}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}} \\ \mathbf{P}\left(M_k - M_{k-1} = \frac{2}{\sqrt{n}} \middle| \mathscr{H}_{k-1} \right) &= \mathbf{P}\big(B_k = 0, m_{k-1}^{A_k} = -1\big) \\ &= \frac{\frac{1}{2}\left(1 - \frac{M_{k-1}}{\sqrt{n}}\right) \mu_{\mathbb{A}_n}}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}} \end{aligned} \tag{3.3}$$

We assume that the initial opinions m_0^a, $a \in \mathbb{A}_n$ are i.i.d. and unbiased, in the sense that $\mathbf{P}(m_0^1 = 1) = \mathbf{P}(m_0^1 = -1) = 1/2$.

As mentioned before, the agents are noise traders, as in the model of the preceding section. Hence, whenever an agent trades, she trades a random amount that is independent of the market history. For simplicity, we assume that trades are unit-sized and that buying and selling occur with equal probabilities. Consequently, the market maker either raises or lowers the price by $\alpha/\sqrt{n}$ with equal

[4] To be fully rigorous, in this step we need to assume that (3.2) holds even when $\mathscr{H}_{k-1}$ is enlarged with the σ-algebra generated by the past individual opinions m_l^a, $a \in \mathbb{A}_n$, $l \leq k-1$.

probabilities

$$\begin{aligned}\mathbf{P}\left(P_k - P_{k-1} = \frac{\alpha}{\sqrt{n}}\middle|\mathscr{H}_{k-1}\right) &= \mathbf{P}\left(P_k - P_{k-1} = -\frac{\alpha}{\sqrt{n}}\middle|\mathscr{H}_{k-1}\right) \\ &= \frac{1}{2}\mathbf{P}(B_k = 1|\mathscr{H}_{k-1}) \\ &= \frac{\frac{1}{2}\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1})}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}}.\end{aligned} \tag{3.4}$$

The initial price P_0 does not have very significant role in what follows, so we assume that $P_0 = p_0 \in \mathbb{R}$ (a constant).

In the diffusion approximation we will derive below, the key càdlàg processes describing the temporal evolution of this model are $(X_t^n)_{t\in[0,\infty)}$ for the price, defined as before by (2.6), and

$$V_t^n := \sum_{k=1}^{\infty} M_k \mathbf{1}_{[T_k, T_{k-1})}(t), \quad t \geq 0,$$

for the opinion index. We replace the earlier technical Assumption 2.8 by the following.

Assumption 3.5 (Inter-action times)**.** The inter-action times $(\tau_k)_{k=1}^{\infty}$ are given by

$$\tau_k := \frac{\gamma_k}{\lambda_{\mathbb{A}_n}(P_{k-1}, M_{k-1}) + \mu_{\mathbb{A}_n}}, \quad k \in \mathbb{Z}_+, \tag{3.6}$$

where the $(\gamma_k)_{k=1}^{\infty}$ are i.i.d. random variables independent of $(P_k, M_k)_{k=0}^{\infty}$, such that γ_1 is exponentially distributed with rate parameter 1.

Analogously to the original model, one can show that the extended model can be given a probabilistic construction such that $(X_t^n, V_t^n)_{t\in[0,\infty)}$ becomes a time-homogeneous pure-jump Markov process.

Lemma 3.7 (Markovianity)**.** *If Assumption 3.5 holds, then there exists a probability space* $(\Omega, \mathscr{F}, \mathbf{P})$ *which carries the model, such that* $(X_t^n, V_t^n)_{t\in[0,\infty)}$ *is a time-homogeneous pure-jump Markov process with rate kernel*

$$K_n(x, v, \mathrm{d}y, \mathrm{d}w) = \big(\lambda_{\mathbb{A}_n}(x, v) + \mu_{\mathbb{A}_n}\big) k_n(x, v, \mathrm{d}y, \mathrm{d}w),$$

where the kernel $k_n(x, v, \mathrm{d}y, \mathrm{d}w)$ *is a regular version of the conditional distribution* $\mathbf{P}(P_1 - P_0 \in \mathrm{d}y, M_1 - M_0 \in \mathrm{d}w | P_0 = x, M_0 = v)$.

3.2 Large-market diffusion approximation for the price and opinion index processes

We are now ready to derive a large-market diffusion approximation for the price and the opinion index in the extended model we have just defined. In the limit, the opinion index follows a stationary Ornstein–Uhlenbeck process and the price follows a diffusion process with stochastic volatility that depends on the Ornstein–Uhlenbeck limit of the opinion index.

Before stating the approximation result, let $D_{\mathbb{R}^2}[0,\infty)$ denote the space of $\mathbb{R}^2$-valued càdlàg functions on $[0,\infty)$, again equipped with the corresponding Skorohod topology.

Proposition 3.8 (Diffusion approximation). *Suppose that the model is as described in this section and Assumption 3.5 holds. If*

$$\frac{\lambda_{\mathbb{A}_n}}{n} \xrightarrow{\text{u.o.c.}} \overline{\lambda}, \tag{3.9}$$

where $\overline{\lambda} : \mathbb{R}^2 \to \mathbb{R}_+$ *is a continuous function such that* $(x,v) \mapsto \overline{\lambda}(x,v)^{1/2}$ *is Lipschitz, then*

$$(X_t^n, V_t^n)_{t\in[0,\infty)} \xrightarrow{\mathscr{D}} (X_t, V_t)_{t\in[0,\infty)} \quad \textit{in } D_{\mathbb{R}^2}[0,\infty), \tag{3.10}$$

where $(X_t, V_t)_{t\in[0,\infty)}$ *is the unique strong solution to the stochastic differential equations*

$$\begin{cases} \mathrm{d}X_t = \alpha\overline{\lambda}(X_t, V_t)^{1/2}\mathrm{d}W_t, & X_0 = p_0, \\ \mathrm{d}V_t = -2\overline{\mu}V_t\mathrm{d}t + 2\overline{\mu}^{1/2}\mathrm{d}B_t, & V_0 = \eta, \end{cases} \tag{3.11}$$

where the process $(W_t, B_t)_{t\in[0,\infty)}$ *is a standard planar Brownian motion and* η *is a standard Gaussian random variable independent of* $(W_t, B_t)_{t\in[0,\infty)}$.

Proof. See Appendix A.3. □

Remark 3.12. The assumption that $(x,v) \mapsto \overline{\lambda}(x,v)^{1/2}$ is Lipschitz is not needed whenever we know *a priori* that (3.11) has a unique strong solution.

3.3 Emergence of heavy-tailed log returns when herding is strong

Let us consider a particular setting, where agents' trading increases rapidly when the opinion index assumes high values, that is, there is a strong tendency towards herd behavior. For simplicity, suppose that agents' trading intensities depend only on the opinion index, and that they are given by some continuous increasing function $\varphi : \mathbb{R} \to \mathbb{R}_+$, common to all agents. However, the agents are not entirely homogeneous: their abilities in high-frequency trading, that is, their maximal trading intensities differ. Moreover, we hypothesize that in a large market, there are agents that are able to trade with very high frequencies. In accordance to these principles, we assume that agent a's maximal trading capability is simply given by her index a, and that her trading intensity function is

$$\lambda_a(x,v) = \varphi(v) \wedge a, \quad (x,v) \in \mathbb{R}^2.$$

Thus, the highest possible trading intensity among the agents in $\mathbb{A}_n$ is n, which tends to infinity as the number agents tend to infinity.

With trading intensity functions as above, we find that condition (3.9) in Theorem 3.8 holds with

$$\overline{\lambda}(x,v) = \varphi(v), \quad (x,v) \in \mathbb{R}^2.$$

In this case, the question of existence and uniqueness of solutions to (3.11) is trivial, so in the large-market limit the dynamics of the price are given (see Remark 3.12) by the stochastic differential equation

$$\mathrm{d}X_t = \alpha\varphi(V_t)^{1/2}\mathrm{d}W_t, \quad X_0 = p_0. \tag{3.13}$$

Recall that we have used the convention that *prices* are always *logarithmic*, so increments $X_{t+\Delta t} - X_t$, $\Delta t > 0$ are actually *log returns* of the *non-logarithmic* price process e^{X_t}, $t \geq 0$. We shall now show that if the asymptotic trading intensity per capita as a function of the opinion index, i.e. the function φ, grows fast enough, quantified by condition (3.15) below, the log returns are heavy-tailed in the sense that they have infinite second moments.

Proposition 3.14 (Infinite second moment)**.** *Let $(X_t)_{t\in[0,\infty)}$ be as in (3.13), and suppose that φ satisfies*

$$\liminf_{v\to\infty} \frac{\varphi(v)}{\exp(\frac{1}{2}v^2)} > 0. \tag{3.15}$$

Then, for all $t \geq 0$ and $\Delta t > 0$, we have $\mathbf{E}[(X_{t+\Delta t} - X_t)^2] = \infty$.

Proof. See Appendix A.4. □

Remark 3.16. Heavy-tailed behavior is an *emergent* phenomenon in this model. It appears only in the large-market limit—for finite n, the log return $X^n_{t+\Delta t} - X^n_t$ is *light-tailed*. To see this, note that in (3.6),

$$\tau_k \geq \frac{\gamma_k}{n(n+\overline{\mu})} =: \tilde{\tau}_k, \quad k \in \mathbb{Z}_+.$$

Clearly, $N_t := \max\{k : \tilde{\tau}_1 + \cdots + \tilde{\tau}_k \leq t\}$, $t \geq 0$ is a homogeneous Poisson process with intensity $n(n + \overline{\mu})$, and we have

$$|X^n_{t+\Delta t} - X^n_t| \leq \frac{\alpha}{\sqrt{n}} N_{t+\Delta t}.$$

Since $N_{t+\Delta t}$ is Poisson-distributed, and hence light-tailed in the sense that its moment generating function is finite, it follows that also $X^n_{t+\Delta t} - X^n_t$ is light-tailed.

While $X^n_{t+\Delta t} - X^n_t$ is light-tailed in the strict mathematical sense, for large n, its tail may, however, resemble a heavy tail up to a cut-off. Then, it might be difficult to distinguish its tails from *bona fide* heavy tails using finitely many observations. This effect is demonstrated by numerical simulations in the following subsection.

Among criteria of heavy-tailedness, infiniteness of second moment is perhaps the one with the most far-reaching consequences in practice. It precludes the use of e.g. quadratic hedging, mean–variance portfolio selection, or statistical methods that require finite variance. That said, there is no clear consensus of what should be regarded as heavy tails. Namely, "heavy tail" is sometimes used as a synonym for *power-law* tails, which exhibit *polynomial* decay of tail probabilities (as opposed to *exponential* decay of tail probabilities e.g. in the Gaussian distribution). This criterion is not fully compatible with the one based on the infinite second moment—some power-law tails lead to infinite second moments, some do not. However, power-law tails have drawn much attention in recent empirical and theoretical studies of stock price dynamics. In particular, empirical evidence suggests that log returns of actively traded stocks over short time intervals could be modeled by power-law distributions (see e.g. [4, 12]).

Hence, it is worthwhile to study whether the log returns in the model have power-law tails. To this end, let us first recall the notion of *regularly varying* random variable, which defines a wide and extensively-studied class of random variables with power-law tails.

Definition 3.17. Random variable Z is said to be *regularly varying* with index $\theta > 0$, if

$$\mathbf{P}(Z > z) \sim pL(z)z^{-\theta}, \quad \mathbf{P}(Z \leq -z) \sim qL(z)z^{-\theta}, \quad z \to \infty,$$

where $p + q = 1$ and L is a *slowly varying* function, that is, for all $c > 0$, we have $\lim_{z\to\infty} L(cz)/L(z) = 1$.

In the definition, the index θ is connected to the heaviness of the tail, while the slowly varying function L is merely a small perturbation. Namely, using the representation theorem for slowly varying functions (e.g. Theorem A3.2 of [7]) one can show that for any $\varepsilon > 0$, there exists $z_\varepsilon > 0$ such that

$$z^{-\varepsilon} \leq L(z) \leq z^{\varepsilon}, \quad z \geq z_\varepsilon.$$

Using these bounds, it is straightforward to deduce a fundamental property of the index θ, namely, that $\mathbf{E}[|Z|^r] = \infty$ for all $r \in (\theta, \infty)$, and $\mathbf{E}[|Z|^r] < \infty$ for all $r \in [0, \theta)$. Note also that if the random variable Z is multiplied by a non-zero constant, the resulting random variable is regularly varying with the same index.

Some preliminary numerical simulations (some more thorough simulations are reported in the next subsection) suggested that when the trading intensity φ grows like the function $v \mapsto \exp(\gamma v^2)$, where $\gamma > 0$, then the log-return $X_{t+\Delta t} - X_t$ appears to follow a power-law distribution. This observation is explained by the following result, which establishes that then $X_{t+\Delta t} - X_t$, with a suitable rescaling, converges in distribution to a regularly varying random variable.

Proposition 3.18 (Power-law tails)**.** *Let $(X_t)_{t\in[0,\infty)}$ be as in (3.13), and suppose that for some $v_0 > 0$, $\varphi(v) = C\exp(\gamma v^2)$ for all $v \geq v_0$, where $C, \gamma > 0$. Then, for all $t \geq 0$,*

$$\frac{1}{\sqrt{\Delta t}}(X_{t+\Delta t} - X_t) \xrightarrow[\Delta t \to 0+]{\mathscr{D}} Z, \tag{3.19}$$

where Z is a symmetrically distributed regularly varying random variable with index $1/\gamma$. Specifically,

$$\mathbf{P}(|Z| > z) \sim \frac{1}{\pi}\alpha^{1/\gamma}C^{1/(2\gamma)}\Gamma\left(\frac{1}{2\gamma} + \frac{1}{2}\right)2^{1/(2\gamma)-1}\left(\frac{\gamma}{\log z}\right)^{1/2} z^{-1/\gamma}, \quad z \to \infty, \tag{3.20}$$

where Γ denotes the Gamma function.

Proof. See Appendix A.5. □

It seems plausible that this result holds also in a more general class of trading intensity functions φ, but unfortunately it seems somewhat difficult to pinpoint any neat general criterion characterizing such functions. That said, the parametric form of φ was also beneficial, in the sense that it allowed us to compute the tail estimate (3.20), which is of interest in connection with the simulation study below.

Remark 3.21. Propositions 3.14 and 3.18 are consistent when it comes to finiteness of second moments. Namely, when φ is as in Proposition 3.18 and $\gamma > 1/2$, then we have $\mathbf{E}[Z^2] = \infty$ and $\mathbf{E}[(X_{t+\Delta t} - X_t)^2] = \infty$. Conversely, when $\gamma < 1/2$, we have $\mathbf{E}[Z^2] < \infty$ and $\mathbf{E}[(X_{t+\Delta t} - X_t)^2] < \infty$ (this can be checked adapting the method used to prove Proposition 3.14).

Remark 3.22. While empirical studies suggest that log returns over short time intervals obey power-law distributions, in longer time scales, observed log returns appear to have distributions close to Gaussian, corresponding to the phenomenon called *aggregational Gaussianity* (see e.g. [4]). This phenomenon can be observed also in our model. Namely, if φ satisfies

$$\int_0^\infty \frac{\varphi(v)\mathrm{d}v}{\exp(\frac{1}{2}v^2)} < \infty$$

(and φ does *not* need to have the exponential form, as in Proposition 3.14), then we have

$$\frac{1}{\sqrt{\Delta t}}(X_{t+\Delta t} - X_t) \xrightarrow[\Delta t \to \infty]{\mathscr{D}} Z',$$

where Z' is a Gaussian random variable with mean 0 and variance $\alpha^2 \mathbf{E}[\varphi(V_0)]$. This convergence follows from the ergodicity of the opinion index process $(V_t)_{t\in[0,\infty)}$ by simple arguments using characteristic functions.

3.4 Numerical simulations of log returns

Intuitively, Proposition 3.18 says that for small $\Delta t > 0$ we may approximate tail probabilities of the log return $X_{t+\Delta t} - X_t$ by

$$\mathbf{P}(|X_{t+\Delta t} - X_t| > x) \approx \mathbf{P}\Big(|Z| > x/\sqrt{\Delta t}\Big) \approx \tilde{L}\Big(x/\sqrt{\Delta t}\Big)\Big(x/\sqrt{\Delta t}\Big)^{-1/\gamma}, \quad (3.23)$$

where $\tilde{L}$ is the slowly varying part of the tail estimate (3.20). To judge the fitness of this approximation for some fixed Δt's, we performed some numerical simulations of the log return $X_{t+\Delta t} - X_t$.

The method used to simulate $X_{t+\Delta t} - X_t$ is described in Appendix B.1. In the simulations, the values of parameters of the model were $\alpha := 1$, $\overline{\mu} = 1$, and $\varphi(v) := \exp\big(v^2 \mathrm{sgn}(v)\big)$, $v \in \mathbb{R}$. Four interval lengths, $\Delta t := 1, 0.1, 0.01, 0.001$, were considered. Moreover, $t := 0$, which entails no loss of generality, since by the stationarity of $(V_t)_{t\in[0,\infty)}$, the log return $X_{t+\Delta t} - X_t$ has the same distribution for any $t \geq 0$. For each interval length, three runs of simulations of $X_{\Delta t} - X_0$ were performed using $k = 1000, 10000, 100000$ discretization steps (to assess the robustness of the results to changes in the level of discretization). The number of observations produced in each run was 100000.

The empirical cumulative distributions of the observations produced in these simulation runs, along with the corresponding approximations given by (3.23) are plotted in Figure 1. The empirical cumulative distribution functions with different levels of discretization seem rather consistent, apart from the extreme parts of the plots, where sampling noise affects the shapes of these functions. When $\Delta t = 1$, we find that the approximation underestimates the tail probabilities of the log return. Indeed, the plotted empirical cumulative distribution functions seem to decay slower than the approximation suggesting that the log return has heavier tails than the approximation Z. When $\Delta t = 0.1$, the same is true, albeit to smaller extent. However, when Δt is decreased to 0.01 or 0.001, the approximations fit rather nicely to the observed distributions.

Encouraged by the good fit of the approximations when $\Delta t = 0.01, 0.001$, we run some further simulations, to study how well the it applies to the tail of $X^n_{\Delta t} - X^n_0$, the corresponding log return in a market with finitely many agents. Of course, since $X^n_{\Delta t} - X^n_0$ is light-tailed (as pointed out in Remark 3.16), we expected this approximation to be valid only for a bounded range of tail probabilities.

The log return $X^n_{\Delta t} - X^n_0$ was simulated using the method described in Appendix B.2. The simulation runs covered the interval lengths $\Delta t = 0.01, 0.001$, with $n = 500, 5000, 50000$ agents in the market. The values of the other parameters were the same as before, and again, 100000 observations were produced in each run. The empirical cumulative distributions of these observations and the corresponding approximations given by (3.23) are plotted in Figure 2. According to these plots, the approximation follows rather closely the empirical tail probabilities up to a cut-off (which increases as number of agents increases), after which the approximation provides significant overestimates of the tail probabilities, as expected. However, for a (hypothetical) risk manager, who observes these log returns and does not know the true distribution of them, it might require too large a sample to infer reliably where this cut-off lies. In her situation, using the approximation—or in practice fitting a power-law distribution to the observed data—would thus appear to be preferable, as it provides conservative estimates of tail probabilities, erring on the right side from the perspective of risk management.

4 Conclusions

We have proposed a potential microfoundation for diffusion price processes, as large-market approximations of price processes in a microscopic market model with heterogeneous agents trading a stock asynchronously in continuous time.

We have first explored microfoundations for one-dimensional diffusions based on a simpler version of the microscopic market model. With this model, we have found that the price process can be approximated by a one-dimensional diffusion with drift and diffusion coefficients that are characterized in an economically intuitive way, respectively, by small excess demands and trading volume in the microscopic model. As an illustration, we applied this result to a more concrete model with fundamentalists and noise traders, in which the stock price follows a Markov process in a random environment, determined by random trading intensity functions. When passing to the large-market limit, we obtained an Ornstein–Uhlenbeck-type process as an approximation of the price process.

Additionally, to explore potential microfoundations for stochastic volatility and heavy-tailed log returns, we have studied an extended version of the microscopic market model, in which agents are allowed to also have interactions besides trading, through an opinion index—to model herd behavior in the market. In the case of this model, we obtained a two-dimensional large-market diffusion approximation, jointly for the price and the opinion index. In the limit, the price follows a process with stochastic volatility that depends, in addition to the price itself, on the Ornstein–Uhlenbeck limit of the opinion index process. Moreover, we have demonstrated that when agents have strong tendency towards herd behavior, the limiting price process exhibits heavy-tailed log returns with infinite second moments, and in some cases, power-law tails.

Acknowledgements

I would like to thank an assistant editor and an anonymous referee for insightful comments. Moreover, I thank Esa Nummelin and Tommi Sottinen for discussions and comments on several earlier drafts of this paper, and Lasse Leskelä for valuable discussions concerning regular variation. Additionally, I gratefully acknowledge financial support from the Academy of Finland (through project 116747) and from the Finnish Cultural Foundation.

Appendix A: Proofs

A.1 Proof of Theorem 2.19

We will apply Theorem IX.4.21 of [18], and thus this proof is merely a verification that the required conditions are satisfied. First, let us note that since the functions λ_a, $a \in \mathbb{A}_n$ are bounded, the kernel $K_n(x, \mathrm{d}y)$ is finite. Next, define auxiliary functions $b_n : \mathbb{R} \to \mathbb{R}$ and $\tilde{c}_n : \mathbb{R} \to \mathbb{R}_+$ by

$$b_n(x) := \sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[\overline{r}_n(e_a^n(x,\xi_1),x)],$$
$$\tilde{c}_n(x) := \sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[\overline{r}_n(e_a^n(x,\xi_1),x)^2],$$

where $\overline{r}_n(q,x) := r_n(q,x) - x = \alpha q/\sqrt{n} + u_n(q,x)$. Using the representation (2.5) and the disintegration theorem (e.g. Theorem 6.4 of [20]) we obtain

$$\int k_n(x,\mathrm{d}y)y = \frac{b_n(x)}{\lambda_{\mathbb{A}_n}(x)}, \quad \int k_n(x,\mathrm{d}y)y^2 = \frac{\tilde{c}_n(x)}{\lambda_{\mathbb{A}_n}(x)}. \tag{A.1}$$

Now,

$$b_n(x) = \alpha\sqrt{n}z_n(x) + \sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[u_n(e_a^n(x,\xi_1),x)],$$

where for any $\delta > 0$ and $|x| \le \delta$,

$$\begin{aligned}\left|\sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[u_n(e_a^n(x,\xi_1),x)]\right| &\le nC_\delta^n \sup_{|x'|\le\delta}\left|\frac{\lambda_{\mathbb{A}_n}(x')}{n}\right| \\ &\quad\times \sup_{|x'|\le\delta, a\in\mathbb{A}_n, n\in\mathbb{Z}_+} \mathbf{E}[|e_a^n(x',\xi_1)|].\end{aligned} \tag{A.2}$$

The right hand side converges to zero, since we assumed that $M_\delta^n = o(n^{-1})$, and since quantities $\sup_{|x'|\le\delta}\left|\frac{\lambda_{\mathbb{A}_n}(x')}{n}\right|$, $n \in \mathbb{Z}_+$ are bounded and

$$\sup_{|x'|\le\delta, a\in\mathbb{A}_n, n\in\mathbb{Z}_+} \mathbf{E}[|e_a^n(x',\xi_1)|] < \infty$$

by uniform integrability (Assumption 2.16). Similarly,

$$\begin{aligned}\tilde{c}_n(x) = \alpha^2 v_n(x)^2 &+ \sum_{a=1}^{n} \lambda_a(x)\frac{\mathbf{E}[2\alpha e_a^n(x,\xi_1)u_n(e_a^n(x,\xi_1),x)]}{\sqrt{n}} \\ &+ \sum_{a=1}^{n} \lambda_a(x)\mathbf{E}[u_n(e_a^n(x,\xi_1),x)^2],\end{aligned}$$

where the latter two terms are asymptotically negligible (u.o.c.) by estimates analogous to (A.2). Hence, by Assumptions 2.15 and 2.17, we have

$$\int K_n(\cdot,\mathrm{d}y)y = b_n(\cdot) \xrightarrow{\text{u.o.c.}} \alpha z(\cdot), \quad \int K_n(\cdot,\mathrm{d}y)y^2 = \tilde{c}_n(\cdot) \xrightarrow{\text{u.o.c.}} \alpha^2 v(\cdot)^2.$$

Next, let $\varepsilon > 0$ and $\delta > 0$ be fixed. Again, by using the representation (2.5) and the disintegration theorem (e.g. Theorem 6.4 of [20]) we obtain

$$\sup_{|x|\le\delta} \int K_n(x, \mathrm{d}y) y^2 \mathbf{1}_{\{|y|>\varepsilon\}} = l_n(\varepsilon, \delta),$$

where

$$l_n(\varepsilon, \delta) := \sup_{|x|\le\delta} \sum_{a=1}^{n} \lambda_a(x) \mathbf{E}\left[\overline{r}_n(e_a^n(x, \xi_1), x)^2 \mathbf{1}_{\{|\overline{r}_n(e_a^n(x,\xi_1),x)|>\varepsilon\}}\right]. \tag{A.3}$$

Clearly, by Assumptions 2.12 and 2.16,

$$\{|\overline{r}(e_a^n(x, \xi_1), x)| > \varepsilon\} \subset \{|e_a^n(x, \xi_1)| > \tilde{\varepsilon}_n\}$$

for all $|x| \le \delta$, where $\tilde{\varepsilon}_n := \sqrt{n}\varepsilon/(\alpha + \sqrt{n}C_\delta^n) \to \infty$. Hence, we have the bound

$$\begin{aligned} l_n(\varepsilon, \delta) \le \sup_{|x|\le\delta} \left|\frac{\lambda_{\mathbb{A}_n}(x)}{n}\right| \Bigg(&\sup_{|x|\le\delta, a\in\mathbb{A}_n, n\in\mathbb{Z}_+} \alpha^2 \mathbf{E}\left[e_a^n(x, \xi)^2 \mathbf{1}_{\{|e_n(x,\xi_1)|>\tilde{\varepsilon}_n\}}\right] \\ &+ \sup_{|x|\le\delta, a\in\mathbb{A}_n, n\in\mathbb{Z}_+} n(C_\delta^n)^2 \mathbf{E}[e_a^n(x, \xi_1)^2] \Bigg), \end{aligned}$$

which converges to zero, by uniform integrability and the assumption that $C_\delta^n = o(n^{-1})$.

Finally, the assumptions about z and v imply that the stochastic differential equation (2.20) has a unique strong solution (see e.g. Theorem III.2.32 of [18]). Hence, the associated martingale problem is well-posed. Moreover, the measurability condition IX.4.3(ii) in [18, p. 555] follows from standard results (e.g. Theorem 21.10 of [20]). Hence the conditions of Theorem IX.4.21 of [18] are fulfilled and the assertion follows. □

A.2 Proof of Proposition 2.23

We proceed in two steps. Initially we prove so-called *quenched* convergence, that is, for $\mathbf{Q}$-a.e. realization $\tilde{\omega} \in \tilde{\Omega}$ of the random environment, we have

$$(X_t^n(\cdot, \tilde{\omega}))_{t\in[0,\infty)} \xrightarrow{\mathscr{D}} (X_t)_{t\in[0,\infty)} \quad \text{in } D[0, \infty). \tag{A.4}$$

Then, using this intermediate result, we argue that also the asserted *annealed* convergence follows.

Step 1: Quenched convergence. First, we study convergence trading intensities per capita. This requires some elementary facts from the theory of *Bochner* *integration* of Banach-valued maps, to be found e.g. in the monograph [6]. Let us note that for any $k \in \mathbb{Z}_+$, restrictions $\lambda_a|_{[-k,k]}$, $a \in \mathbb{A}_\infty$ are random elements of the separable Banach space $C[-k,k]$ equipped with the usual sup-norm, and by (2.22) Bochner integrable (see e.g. Theorem II.2.2 of [6]). Hence, by the strong law of large numbers for Banach-valued random variables (e.g. Theorem III.1.1 of [14]), we have

$$\frac{1}{k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 1}} \lambda_a|_{[-k,k]} \to I_{1,k} \quad \mathbf{Q}\text{-a.s.},$$

where $I_{1,k} \in C[-k,k]$ is the Bochner integral $\int_{\tilde{\Omega}} \lambda_a|_{[-k,k]}(\cdot)(\tilde{\omega})\mathbf{Q}(\mathrm{d}\tilde{\omega})$ for some $a \in \mathbb{A}_\infty$ such that $f_a = 1$. Similarly,

$$\frac{1}{n - k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 0}} \lambda_a|_{[-k,k]} \to I_{0,k} \quad \mathbf{Q}\text{-a.s.}$$

where $I_{0,k} := \int_{\tilde{\Omega}} \lambda_a|_{[-k,k]}(\cdot)(\tilde{\omega})\mathbf{Q}(\mathrm{d}\tilde{\omega})$ for some $a \in \mathbb{A}_\infty$ such that $f_a = 0$. Since taking an evaluation $f \mapsto f(x)$, $x \in [-k,k]$, commutes with Bochner integration, which follows e.g. from Theorem II.2.6 of [6], and since the Bochner integral of a real-valued function coincides with the Lebesgue integral, we have $I_{0,k}(x) = \overline{\lambda}_0(x)$ and $I_{1,k}(x) = \overline{\lambda}_1(x)$ for all $x \in [-k,k]$. Hence, by taking intersections over all $k \in \mathbb{Z}_+$, we obtain $\mathbf{Q}$-a.s.

$$\frac{1}{k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 1}} \lambda_a \xrightarrow{\text{u.o.c.}} \overline{\lambda}_1, \quad \text{and} \quad \frac{1}{n - k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 0}} \lambda_a \xrightarrow{\text{u.o.c.}} \overline{\lambda}_0, \tag{A.5}$$

and consequently,

$$\frac{\lambda_{\mathbb{A}_n}}{n} = \frac{k_n}{n} \times \frac{1}{k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 1}} \lambda_a + \left(1 - \frac{k_n}{n}\right) \times \frac{1}{n - k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 0}} \lambda_a \xrightarrow{\text{u.o.c.}} \overline{\lambda} \quad \mathbf{Q}\text{-a.s.} \tag{A.6}$$

Second, observe that for all $a \in \mathbb{A}_\infty$, and $\delta > 0$, we have the bound

$$e_a^n(x, \xi_1)^2 \leq 2\left(\sup_{|x| \leq \delta} w(F - x)^2 + \xi_1^2\right),$$

so together with (A.6), this implies that Assumption 2.16 holds $\mathbf{Q}$-almost surely. Further, a straightforward computation yields

$$\sqrt{n} z_n(x) = \frac{k_n}{n} \times \frac{1}{k_n} \sum_{\substack{a \in \mathbb{A}_n \\ f_a = 1}} \lambda_a(x) w(F - x),$$

which, by (A.5), implies that Assumption 2.15 holds **Q**-a.s. with $z(x) = \phi\overline{\lambda}_1(x)w(F-x)$. Similarly,

$$v_n(x)^2 = \frac{1}{n}\Bigg[\sum_{\substack{a\in\mathbb{A}_n\\ f_a=1}} \lambda_a(x)\Big(\frac{w(F-x)^2}{n} + \sigma_\xi^2\Big) + \sum_{\substack{a\in\mathbb{A}_n\\ f_a=0}} \lambda_a(x)\sigma_\xi^2\Bigg],$$

which, for any $\delta > 0$ and $|x| \le \delta$, implies the bounds

$$\frac{\lambda_{\mathbb{A}_n}(x)}{n}\sigma_\xi^2 \le v_n(x)^2 \le \frac{\lambda_{\mathbb{A}_n}(x)}{n}\Big(\sigma_\xi^2 + \frac{1}{n}\sup_{|x|\le\delta} w(F-x)^2\Big).$$

Hence, by (A.6), Assumption 2.17 holds **Q**-a.s. with $v(x) = \sigma_\xi\overline{\lambda}(x)^{1/2}$.

To conclude the proof of convergence (A.4), we note that pricing rule (2.21) obviously conforms to Assumption 2.12, and the local Lipschitz and growth conditions for z and v have been assumed *a priori*, so (A.4) follows now from Theorem 2.19.

Step 2: Annealed convergence. Let $f : D[0,\infty) \to \mathbb{R}$ be bounded and continuous. Using Fubini's theorem we obtain the estimate

$$\begin{aligned}&\Big|\int_{\Omega\times\tilde{\Omega}} f\big(X^n_\cdot(\omega,\tilde{\omega})\big)(\mathbf{P}\otimes\mathbf{Q})(\mathrm{d}\omega,\mathrm{d}\tilde{\omega}) - \mathbf{E}[f(X_\cdot)]\Big| \\ &\le \int_{\tilde{\Omega}}\Big|\int_{\Omega} f\big(X^n_\cdot(\omega,\tilde{\omega})\big)\mathbf{P}(\mathrm{d}\omega) - \mathbf{E}[f(X_\cdot)]\Big|\,\mathbf{Q}(\mathrm{d}\tilde{\omega}),\end{aligned}$$

from which the assertion follows by (A.4) and Lebesgue's dominated convergence theorem. □

A.3 Proof of Proposition 3.8

We use the same method as in the proof of Theorem 2.19. First, we compute the first and second moments of the bivariate distributions $k_n(x, v, \mathrm{d}y, \mathrm{d}w)$, $n \in \mathbb{Z}_+$ using (3.3) and (3.4). The first moments are

$$\begin{aligned}\int k_n(x,v,\mathrm{d}y,\mathrm{d}w)y &= 0,\\ \int k_n(x,v,\mathrm{d}y,\mathrm{d}w)w &= \frac{-2\overline{\mu}v}{\lambda_{\mathbb{A}_n}(x,v)+\mu_{\mathbb{A}_n}}.\end{aligned} \tag{A.7}$$

Further, the second moments are

$$\begin{aligned}\int k_n(x,v,\mathrm{d}y,\mathrm{d}w)y^2 &= \frac{\alpha^2 n^{-1}\lambda_{\mathbb{A}_n}(x,v)}{\lambda_{\mathbb{A}_n}(x,v)+\mu_{\mathbb{A}_n}},\\ \int k_n(x,v,\mathrm{d}y,\mathrm{d}w)w^2 &= \frac{4\mu_{\mathbb{A}_n}}{\lambda_{\mathbb{A}_n}(x,v)+\mu_{\mathbb{A}_n}},\end{aligned} \tag{A.8}$$

and since the price and the opinion index cannot change simultaneously,

$$\int k_n(x, v, \mathrm{d}y, \mathrm{d}w) yw = 0. \tag{A.9}$$

Using (A.7), (A.8), and (A.9), we see that the kernels $K_n(x, v, \mathrm{d}y, \mathrm{d}w)$, $n \in \mathbb{Z}_+$ satisfy

$$\int K_n(x, v, \mathrm{d}y, \mathrm{d}w)(y, w) = (0, -2\overline{\mu}v), \quad (x, v) \in \mathbb{R}^2,$$

and

$$\sup_{\|(x,v)\| \leq \delta} \left\| \int K_n(x, v, \mathrm{d}y, \mathrm{d}w) \begin{bmatrix} y^2 & yw \\ wy & w^2 \end{bmatrix} - \begin{bmatrix} \alpha^2 \overline{\lambda}(x, v) & 0 \\ 0 & 4\overline{\mu} \end{bmatrix} \right\|_{\mathbb{R}^{2\times 2}} \to 0,$$

for any $\delta > 0$ by (3.9). Additionally, we note that for all $(x, v) \in \mathbb{R}^2$, the support of the measure $K_n(x, v, \mathrm{d}y, \mathrm{d}w)$ is contained in the rectangle $R_n := [-\alpha/\sqrt{n}, \alpha/\sqrt{n}] \times [-2/\sqrt{n}, 2/\sqrt{n}]$, so for any $\varepsilon > 0$ and $\delta > 0$, we have

$$\sup_{\|(x,v)\| \leq \delta} \int K_n(x, v, \mathrm{d}y, \mathrm{d}w) \|(y, w)\|^2 \mathbf{1}_{\{\|(y,w)\| > \varepsilon\}} \to 0,$$

since $R_n \cap \{\|(y, w)\| > \varepsilon\} = \varnothing$ for all sufficiently large n.

Finally, since $V_0^n \xrightarrow{\mathscr{D}} \eta$ by the central limit theorem, and since the Lipschitz assumption about $\overline{\lambda}$ implies that the stochastic differential equations (3.11) have a unique strong solution, similarly to the proof of Theorem 2.19, the assertion follows now from Theorem IX.4.21 of [18]. □

A.4 Proof of Proposition 3.14

Using characteristic functions and complex Doléans exponentials (see e.g. [20, p. 351]) one deduces that

$$X_{t+\Delta t} - X_t \stackrel{\mathscr{D}}{=} \left(\int_t^{t+\Delta t} \alpha^2 \varphi(V_s) \mathrm{d}s \right)^{1/2} \xi, \tag{A.10}$$

where ξ is a standard Gaussian variable independent of the process $(V_t)_{t\in[0,\infty)}$. (To be fully rigorous, one must first replace the upper bound $t + \Delta t$ by the the stopping time $t + \tau_k$, where $\tau_k := \inf\{u > 0 : \int_t^{u+t} \alpha^2 \varphi(V_s)\mathrm{d}s > k\}$, prove the identity in this case, and then pass to the limit $k \to \infty$.) Using independence and Fubini's theorem, we obtain from (A.10)

$$\mathbf{E}[(X_{t+\Delta t} - X_t)^2] = \mathbf{E}\left[\int_t^{t+\Delta t} \alpha^2 \varphi(V_s) \mathrm{d}s \right] = \int_t^{t+\Delta t} \alpha^2 \mathbf{E}[\varphi(V_s)] \mathrm{d}s.$$

The assertion follows now, since for any $s \in [0, \infty)$,

$$\mathbf{E}[\varphi(V_s)] \geq \frac{1}{\sqrt{2\pi}} \int_0^\infty \frac{\varphi(v)}{\exp(\frac{1}{2}v^2)} \mathrm{d}v = \infty$$

by the growth condition (3.15). □

A.5 Proof of Proposition 3.18

We assume that $C = 1$, as the extension to the general case is straightforward. First, by (A.10) we have

$$\frac{1}{\sqrt{\Delta t}}(X_{t+\Delta t} - X_t) \overset{\mathscr{D}}{=} \left(\frac{1}{\Delta t} \int_t^{t+\Delta t} \alpha^2 \varphi(V_s) \mathrm{d}s \right)^{1/2} \xi,$$

where ξ is a standard Gaussian variable independent of the process $(V_t)_{t\in[0,\infty)}$. Since φ is a continuous function and $(V_t)_{t\in[0,\infty)}$ is a continuous process,

$$\left(\frac{1}{\Delta t} \int_t^{t+\Delta t} \alpha^2 \varphi(V_s) \mathrm{d}s \right)^{1/2} \xi \xrightarrow[\Delta t \to 0+]{\text{a.s.}} \alpha \varphi(V_t)^{1/2} \xi =: Z.$$

The random variable Z is symmetrically distributed, since $(V_t)_{t\in[0,\infty)}$ and ξ are independent and since ξ is symmetrically distributed.

Next, we show that $\varphi(V_t)^{1/2}$ is a regularly varying random variable. For this purpose, we note that for any $x > \varphi(v_0)$,

$$\mathbf{P}(\varphi(V_t)^{1/2} > x) = \mathbf{P}(V_t > (2\gamma^{-1} \log x)^{1/2}).$$

Since the Gaussian cumulative distribution function Φ satisfies $1 - \Phi(x) \sim \Phi'(x)/x$, $x \to \infty$ (see e.g. Example 3.3.29 of [7]), we have

$$\begin{aligned} \mathbf{P}(V_t > (2\gamma^{-1} \log x)^{1/2}) &\sim \frac{\frac{1}{\sqrt{2\pi}} \exp\big(-\gamma^{-1} \log x\big)}{(2\gamma^{-1} \log x)^{1/2}} \\ &= \left(\frac{\gamma}{4\pi \log x} \right)^{1/2} x^{-1/\gamma}, \quad x \to \infty. \end{aligned} \tag{A.11}$$

Denote now $L(x) := \big(\gamma/(4\pi \log x)\big)^{1/2}$, $x > 1$. For any $c > 0$, we have

$$\lim_{x\to\infty} \frac{L(cx)}{L(x)} = \lim_{x\to\infty} \left(\frac{\log x}{\log c + \log x} \right)^{1/2} = 1,$$

so L is a slowly varying function. Hence, $\varphi(V_t)^{1/2}$ is a regularly varying random variable with index $1/\gamma$.

Finally, by Breiman's lemma [19, Lemma 4.2], we have

$$\mathbf{P}(|Z| > x) \sim \mathbf{E}[|\alpha W_1|^{1/\gamma}]\mathbf{P}(\varphi(V_t)^{1/2} > x), \quad x \to \infty,$$

so by symmetry of the distribution, we find that Z is a regularly varying random variable with index $1/\gamma$. Expression (3.20) can now be found by computing (integrating by substitution)

$$\mathbf{E}[|W_1|^{1/\gamma}] = \frac{2^{(\gamma+1)/(2\gamma)}}{\sqrt{2\pi}}\Gamma\left(\frac{\gamma+1}{2\gamma}\right)$$

and plugging in (A.11). □

Appendix B: Technical details on the numerical simulations of log returns

The implementations of the following simulation schemes were written in C programming language, and they were run inside R statistical computing environment [22], making use of the built-in *Mersenne twister* random number generator of R. The source codes of the implementations are available on request.

B.1 Simulation of log returns in the large-market limit

To simulate the log return $X_{\Delta t} - X_0$, we used the following scheme, based on the distributional equality (A.10). First, for a chosen number, $k \in \mathbb{Z}_+$, of discretization steps, an *exact* skeleton $V_{\frac{i\Delta t}{k}}$, $i = 0, 1, 2, \ldots, k$ of the Ornstein–Uhlenbeck process $(V_u)_{u\in[0,\Delta t]}$ was simulated using its Gaussian transition density (see e.g. [21, p. 218]). Using this skeleton, a Riemann sum approximation of the integral in (A.10) was computed. Finally, the square root of this approximation was multiplied by a simulated value drawn from the standard Gaussian distribution to produce the actual observation.

B.2 Simulation of log returns in markets with finitely many agents

In the simulation of the log return $X^n_{\Delta t} - X^n_0$, it was necessary to eschew methods that require pathwise simulation of $(X^n_u)_{u\in[0,\Delta t]}$, which would be computationally very costly when the number of agents is large, due to the plenitude of trades triggered by rapidly-growing trading intensities. To this end, we developed an efficient method that exploits the fact that the process $(X^n_u)_{u\in[0,\infty)}$ is equals, in distribution, a linear combination of two independent Poisson processes, run according to a "clock" given by a functional of the opinion index process $(V_u^n)_{u\in[0,\Delta t]}$. This way, it suffices to perform a pathwise simulation of $(V_u^n)_{u\in[0,\Delta t]}$ only, which is computationally less demanding since the opinion-change intensities remain constant throughout.

The simulation scheme can be described as follows. First, using (3.3) and Proposition 3.7, we find that $(V_t^n)_{t\in[0,\infty)}$ is a homogeneous pure-jump Markov process in its own right, with rate kernel

$$K_n^V(v,\mathrm{d}w) = \lambda_{\mathbb{A}_n} \underbrace{\left[\frac{1}{2}\left(1+\frac{v}{\sqrt{n}}\right)\delta_{-2/\sqrt{n}}(\mathrm{d}w) + \frac{1}{2}\left(1-\frac{v}{\sqrt{n}}\right)\delta_{2/\sqrt{n}}(\mathrm{d}w)\right]}_{=:k_n^V(v,\mathrm{d}w)}.$$

Then, to simulate a path of $(V_u^n)_{u\in[0,\Delta t]}$, by Theorem 12.17 of [20], one needs to draw sufficiently many inter-jump times from the exponential distribution with rate parameter $\mu_{\mathbb{A}_n}$, so that the sum of the inter-jump times exceeds Δt, and then sample the jumps from a Markov chain with $k_n^V(v,\mathrm{d}w)$ as its transition kernel. The key observation is that, by Theorem 6.4.1 of [8], there exists independent Poisson processes $(N_u^1)_{u\in[0,\infty)}$ and $(N_u^2)_{u\in[0,\infty)}$, independent of $(V_t^n)_{t\in[0,\infty)}$, such that

$$X_{\Delta t}^n - X_0^n \overset{\mathscr{D}}{=} \frac{\alpha}{\sqrt{n}}\left(N^1_{\frac{1}{2}\int_0^{\Delta t}\lambda_{\mathbb{A}_n}(V_s^n)\mathrm{d}s} - N^2_{\frac{1}{2}\int_0^{\Delta t}\lambda_{\mathbb{A}_n}(V_s^n)\mathrm{d}s}\right). \tag{B.1}$$

Thus, to produce the actual observation using (B.1), it remains to compute the integral $\int_0^{\Delta t}\lambda_{\mathbb{A}_n}(V_s^n)\mathrm{d}s$ using the simulated path of $(V_t^n)_{t\in[0,\Delta t]}$, and then draw two values from the Poisson distribution with parameter $\frac{1}{2}\int_0^{\Delta t}\lambda_{\mathbb{A}_n}(V_s^n)\mathrm{d}s$.

References

[1] Banerjee, A.V.: A simple model of herd behavior. Quart. J. Econ. **107**(3), 797–817 (1992).

[2] Bayraktar, E., Horst, U., Sircar, R.: Queueing theoretic approaches to financial price fluctuations. In: Birge, J.R., Linetsky, V. (eds.) Financial Engineering, pp. 637–680. Handbooks in Operations Research and Management Science, Vol. 15, Elsevier, Amsterdam (2007)

[3] Billingsley, P.: Convergence of probability measures, 2nd edn. Wiley, New York (1999)

[4] Cont, R.: Empirical properties of asset returns: stylized facts and statistical issues. Quant. Finance **1**(2), 223–236 (2001)

[5] Cont, R., Bouchaud, J.-P.: Herd behavior and aggregate fluctuations in financial markets. Macroecon. Dynam. **4**(2), 170–196 (2000)

[6] Diestel, J., Uhl, J.-J. Jr.: Vector measures. American Mathematical Society, Providence (1977)

[7] Embrechts, P., Klüppelberg, C., Mikosch, T.: Modelling extremal events. Springer, Berlin (1997)

[8] Ethier, S.N., Kurtz, T.G.: Markov processes. Wiley, New York (1986)

[9] Farmer, J.D., Patelli, P., Zovko, I.I.: The predictive power of zero intelligence in financial markets. Proc. Natl. Acad. Sci. **102**(11), 2254–2259 (2005)

[10] Föllmer, H., Schweizer, M.: A microeconomic approach to diffusion models for stock prices. Math. Finance **3**(1), 1–23 (1993)

[11] Frey, R., Stremme, A.: Market volatility and feedback effects from dynamic hedging. Math. Finance **7**(4), 351–374 (1997)

[12] Gabaix, X.: Power laws in economics and finance. Ann. Rev. Econ. **1**, 255–293 (2009)

[13] Garman, M.B.: Market microstructure. J. Finan. Econ. **3**(3), 257–275 (1976)

[14] Hoffmann-Jørgensen, J.: Probability in Banach space. In: École d'Été de Probabilités de Saint-Flour, VI-1976, pp. 1–186. Lecture Notes in Math., Vol. 598, Springer, Berlin (1977)

[15] Horst, U.: Financial price fluctuations in a stock market model with many interacting agents. Econom. Theory **25**(4), 917–932 (2005)

[16] Horst, U., Rothe, C.: Queuing, social interactions, and the microstructure of financial markets. Macroecon. Dynam. **12**(2), 211–233 (2008)

[17] Huberman, G., Stanzl, W.: Price manipulation and quasi-arbitrage. Econometrica **72**(4), 1247–1275 (2004)

[18] Jacod, J., Shiryaev, A.N.: Limit theorems for stochastic processes, 2nd edn. Springer, Berlin (2003)

[19] Jessen, A.H., Mikosch, T.: Regularly varying functions. Publ. Inst. Math. (Beograd) (N.S.) **80**(94) 171–192 (2006)

[20] Kallenberg, O.: Foundations of modern probability, 2nd edn. Springer, New York (2002)

[21] Karlin, S., Taylor, H.M.: A second course in stochastic processes. Academic Press, San Diego (1981)

[22] R Development Core Team: R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna Austria (2008) URL: http://www.R-project.org

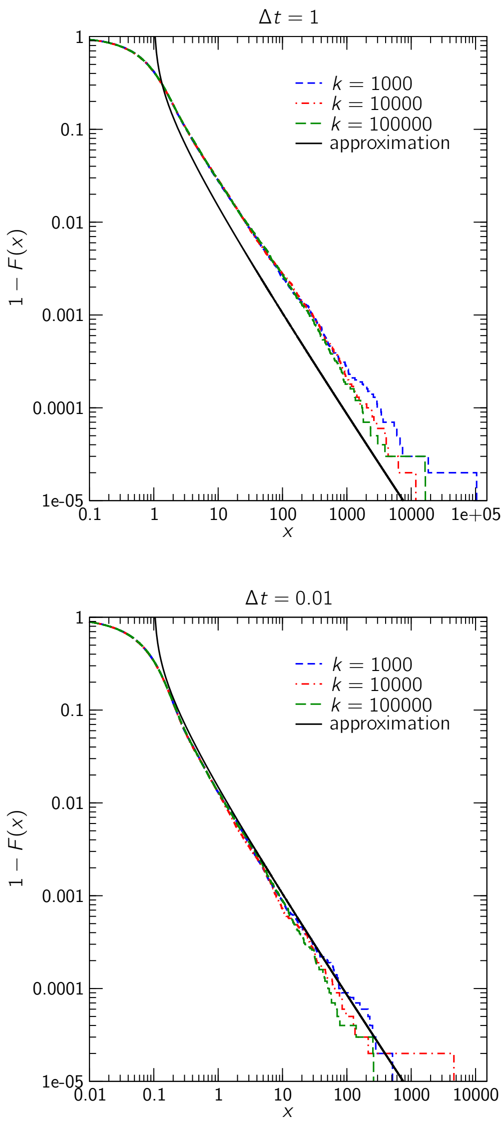

Figure 1: The approximations (solid lines) given by (3.23) compared to the empirical cumulative distribution functions (dashed/dotted lines) F of the simulated values of $|X_{\Delta t} - X_0|$, when $\alpha = 1$, $\overline{\mu} = 1$, $\varphi(v) \equiv \exp(v^2 \mathrm{sgn}(v))$, $\Delta t = 1, 0.1, 0.01, 0.001$, $k = 1000, 10000, 100000$ (discretization steps). Number of observations in each simulation run: 100000.

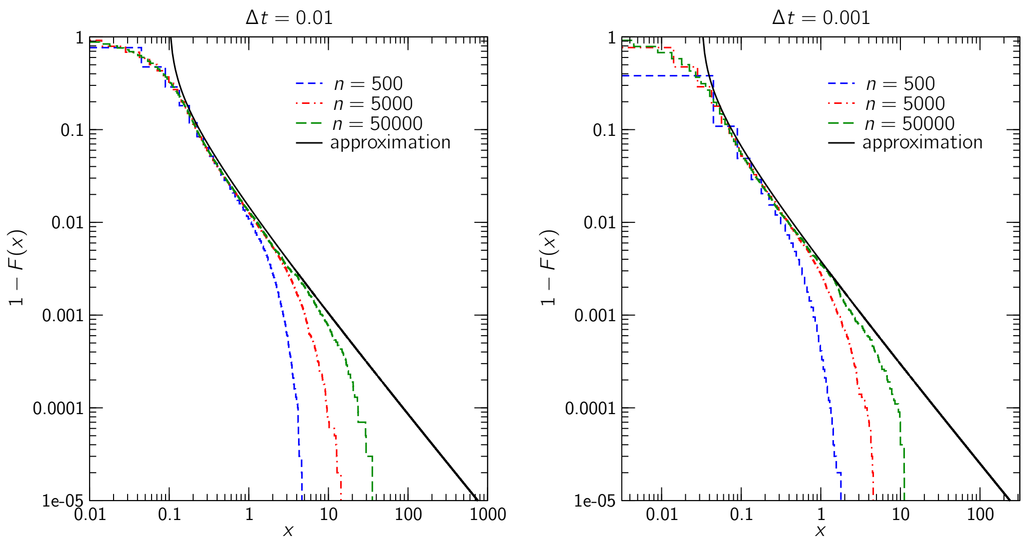

Figure 2: The approximations (solid lines) given by (3.23) compared to the empirical cumulative distribution functions (dashed/dotted lines) F of the simulated values of $|X^n_{\Delta t} - X^n_0|$, when $\alpha = 1$, $\overline{\mu} = 1$, $\varphi(v) \equiv \exp(v^2 \mathrm{sgn}(v))$, $\Delta t = 0.01, 0.001$, $n = 500, 5000, 50000$. Number of observations in each simulation run: 100000.