Estimated Dynamic Equilibrium Model: Supply and Demand as a Sample Path of a Stochastic Process

Mikhail L. Arbuzov
Independent Researcher
Mike.arbuzov54@gmail.com

Sisong Bei
Independent Researcher
qurining@gmail.com

Alexey Shvets
Palo Alto Networks
ashvets@paloaltonetworks.com

Abstract

We introduce the *Estimated Dynamic Equilibrium Model* (EDEM), an agent-based framework that treats supply and demand as a coupled stochastic process driven by heterogeneous, noisy agent valuations. The model's primary technical contribution is the identification of a generative mechanism for persistent disequilibrium: when market-clearing prices are sequentially sampled from the upper tail of noisy bid distributions and recycled as inputs for future valuations, expected prices drift upward despite strictly zero-mean estimation errors. We derive this order-statistic bias in closed form for i.i.d. uniform bids and use simulations to show that compounding this bias across epochs yields exponential price growth without requiring assumptions of investor optimism or irrationality. This framework extends Miller's divergence-of-opinion theory to a dynamic setting, recovering Walrasian equilibrium and Miller's static premium as limiting cases. Through controlled experiments and sensitivity analysis on a simulated real-estate neighborhood, we identify six distinct regimes—ranging from band-stability to runaway bubbles—emerging from a single agent ruleset. These results offer a potential explanation for the contradictory findings in the empirical divergence-of-opinion literature and suggest that machine-learning valuation algorithms may inadvertently amplify this inherent statistical bias.

Keywords: agent-based model; bottom-up; stochastic process; disequilibrium; divergence of opinion; real-estate market

1 Introduction

The Efficient Market Hypothesis (EMH) remains the dominant theory of price formation. It predicts when prices reflect available information but does not provide a generative microstructural account of when and why they fail to: bubbles, crashes, persistent cross-sectional return anomalies, and the stylised observation that real markets very rarely sit at the textbook supply–demand fixed point. Decades of empirical work have produced rival theories whose claims often contradict one another, because abundant and noisy financial data can be mined to support nearly any hypothesis [Moffitt, 2017]. To avoid contributing yet another empirical controversy, this paper takes the less traveled *atheoretical* route: constructing an economy from the ground up in a controlled environment and asking what behavior emerges from a small number of common-sense rules.

The central technical claim of the paper is that two structural features — max-bid clearing and per-epoch feedback of clearing prices into agent valuations — generate persistent positive drift in realised market prices even when individual agents' estimation errors are exactly zero-mean. The mechanism is order-statistic bias: each seller selects the maximum of multiple noisy bids, and the maximum of n zero-mean draws is, in expectation, strictly above their mean. Feeding the realised maximum back into the next round's anchor compounds the bias multiplicatively. No behavioural assumption about investor optimism, risk aversion, or attention is required. Bubbles, business cycles, persistent overshoots, and persistent undershoots all follow from this single mechanism interacting with heterogeneous time-on-market dynamics.

We use agent-based modelling [Wilensky, 1999] because it allows this kind of structural mechanism to emerge from a few simple, transparent rules. The setting is a real-estate market in a small neighbourhood — a market in which every participant must form a personal estimate of an asset whose true value is fundamentally uncertain. Real estate is a useful laboratory: prices are observable, transactions are costly enough that naïve arbitrage arguments fail, and divergence of opinion among buyers and sellers is the rule rather than the exception.

The central object of study is the *Estimated Dynamic Equilibrium Model* (EDEM), in which supply and demand are not static schedules but realisations of a stochastic process driven by heterogeneous, error-prone agent valuations. EDEM is a dynamic extension of Miller [1977]'s divergence-of-opinion theory, but it relaxes Miller's static one-shot setting to allow valuations, the price level, and the population of buyers and sellers to co-evolve over time.

The paper makes three contributions:

1. **An order-statistic mechanism for bubbles without biased agents** (Section 3.7): we derive in closed form, for a clean special case, that the expected winning bid exceeds the home value by $\sigma(n-1)/(n+1)$ when n bidders draw zero-mean errors with dispersion σ. Compounded across epochs, this produces the exponential price drift of Section 5.2.1.
2. **A unified framework** (Section 3) that nests the classical Walrasian equilibrium and Miller's static premium as limiting cases, and admits a family of out-of-equilibrium regimes (business cycles, persistent overshoots and undershoots, bubbles, constant transition) parameterised by estimation noise, balancer strength, and time-on-market.
3. **Eight primary controlled experiments plus a 30-cell sensitivity grid** (Section 5) implemented in the Python ABM framework Mesa [Masad and Kazil, 2015], replacing earlier Wilensky [1999] NetLogo prototypes. The Mesa codebase is open source and reproduces every figure in the paper from a clean checkout.

A practical implication is that empirical premium-vs-discount findings depend strongly on the sample's position along the underlying stochastic trajectory, which short-window studies cannot identify. Two studies that find opposite signs may both be correct characterisations of distinct epochs in the same process. A secondary implication concerns valuation algorithms: a machine-learning estimator trained on historical clearing prices inherits the bid-selection asymmetry and re-deploys it as a positively-biased point estimator, providing the market with a coordinating signal that can amplify the very bubble the algorithm was meant to price. Section 6.4 discusses this in connection with the 2021 wind-down of Zillow Offers [Zillow Group, Inc., 2021].

The remainder of the paper is organised as follows. Section 2 reviews Miller [1977] and the divergence-of-opinion debate. Section 3 presents the EDEM framework formally and derives the order-statistic drift. Section 4 describes the Mesa implementation. Section 5 reports the eight primary experiments and their results. Section 6 discusses implications. Section 7 sketches extensions and Section 8 concludes.

2 Background: Divergence of Opinion

2.1 Miller's static premium hypothesis

Miller [1977] proposed a deceptively simple model of how heterogeneous beliefs determine an asset's market price. Suppose any single investor can purchase only one share—perhaps because of limited funds—and there are N shares available. The shares will end up owned by the N investors with the highest valuations. The marginal investor, the one with the lowest valuation among the holders, sets the price. As the divergence of opinion about the asset's value widens, this marginal valuation rises, pushing the price above the mean estimate of all potential investors.

The implication is striking: when short selling is restricted, an asset's price reflects the optimism of a minority, not the average view of the market. Greater disagreement among investors thus produces a *premium*, and because eventual returns must be earned against this elevated price, future returns on high-disagreement assets should be lower. This is the "divergence of opinion premium hypothesis"—or, equivalently, the *overvaluation hypothesis* [Doukas et al., 2006].

Miller offered initial empirical support: stocks about which there is the greatest divergence of opinion at issuance experience smaller subsequent price appreciation than "seasoned" stocks over horizons of one to five years. Although Miller's model was largely ignored for decades, its narrative remains unusually compelling. As Moffitt [2017] puts it, "although Miller's model has largely been ignored, it has lost none of its relevance."

2.2 The premium–discount debate

The empirical literature that grew up around Miller's hypothesis is remarkable for its inability to converge. Doukas et al. [2006] provide a careful overview of the half-century of contested findings, of which the broad outlines are as follows.

The *discount* camp argues that divergence of opinion proxies for risk: the higher the disagreement, the riskier the asset, and therefore the lower its price relative to fundamentals [Williams, 1977, Mayshar, 1983, Varian, 1985, Merton, 1987, Epstein and Wang, 1994]. A direct consequence is that future returns on high-disagreement assets should be *higher*, not lower.

The *premium* camp—supporting Miller's original prediction—finds positive associations between dispersion in analyst earnings forecasts and contemporaneous prices, with low subsequent returns. Doukas et al. [2006] themselves report evidence consistent with the discount hypothesis after controlling for analyst-related dispersion measures, framing the apparent contradictions as artefacts of measurement rather than substantive disagreement.

Decades on, the debate is unresolved. Moffitt [2017] summarises the state of play as follows:

> "Though we have over two centuries of financial market history and various theories of price formation, today there exists no single theory acceptable to market practitioners, yet rigorous enough to satisfy market theorists. The only major theory that purports universality is the efficient market theory, which fails to explain some of the most important market phenomena, e.g. bubbles and crashes."

2.3 What is missing: dynamics

A common feature of Miller's model and its successors is that they are fundamentally *static*. They specify a one-shot relationship between the distribution of beliefs and the market-clearing price. Real markets, however, are continuous processes in which today's price feeds back into tomorrow's beliefs, and in which the population of active buyers and sellers responds to recent prices with adjustment lags.

This paper argues that the static framing is responsible for much of the empirical confusion. When supply and demand are sample paths of a stochastic process rather than fixed schedules, the same underlying mechanism— heterogeneous valuations with random error—can produce *either* a divergence-of-opinion premium *or* a discount, depending on the balance between estimation noise and the market's adjustment capacity. Section 3 formalises this claim.

Shiller [2003] reviews how behavioural finance has progressively encroached on EMH territory; the present work occupies a related niche, treating price formation itself as the locus of behavioural noise rather than treating noise as a deviation from a frictionless benchmark.

3 The EDEM Framework

This section formalises the Estimated Dynamic Equilibrium Model. We state EDEM in its general form, then identify the Dynamic Equilibrium (DE) model of Section 5.1 and Walrasian equilibrium as limiting cases. The central technical claim is that two structural features — max-bid clearing and per-epoch feedback of clearing prices into agent valuations — generate persistent positive drift in realised prices even when individual agents' estimation errors are zero-mean. Section 3.7 derives this drift in a clean order-statistic form.

3.1 Primitive objects

A market is staged on a finite torus G of homes; each home $h \in G$ has a fixed but unobservable *fair value* $v^*(h)$ and a *market value* $v_t(h) \in \mathbb{R}_{>0}$ that evolves over discrete time $t = 0, 1, 2, \ldots$. The market is populated by two agent classes:

- **Sellers** $\mathcal{S}_t \subset G$, each occupying one home and posting an ask price $a_t(s)$;
- **Buyers** $\mathcal{B}_t$, each free to walk the torus and place bids on sellers it lands with.

We let $Q_s(t) = |\mathcal{S}_t|$ and $Q_d(t) = |\mathcal{B}_t|$ denote the supply and demand quantities at time t; both are state variables, not parameters.

3.2 Heterogeneous estimation

Each agent is endowed with a personal *estimation function*. For an agent i valuing home h at time t:

$$E_{i,t}(h) \;=\; v_t(h) \cdot \big(1 + \varepsilon_{i,t,h}\big), \qquad \varepsilon_{i,t,h} \overset{\text{iid}}{\sim} F_i(\sigma_i), \tag{1}$$

where F_i is a zero-mean error distribution with agent-specific dispersion σ_i. The three-way index is essential: each *bid* that agent i places — on a different home, or even on the same home at a different tick — draws an independent realisation of the error. Without per-bid independence, the buyer-side order-statistic logic of Cond. 2 (Section 3.6) and the per-epoch update of Eq. (10) would degenerate to single-draw statistics. Concretely, $\mathbb{E}[\varepsilon_{i,t,h}] = 0$ and $\operatorname{Var}(\varepsilon_{i,t,h})$ is increasing in σ_i. The DE specialisation (Section 3.8) takes F_i as the symmetric uniform $\mathcal{U}[-\sigma_i, +\sigma_i]$; in the broader EDEM the error distribution is left unspecified beyond zero-mean and the support condition $\Pr(\varepsilon_{i,t} > -1) = 1$ that keeps Eq. (1) strictly positive.

Unit convention. Throughout the formal model, σ_i is a dimensionless dispersion parameter expressed as a *decimal*; e.g. $\sigma_i = 0.05$ corresponds to a $\pm 5\%$ maximum estimation error. The parameter tables and figure captions report σ as a percentage for readability; the conversion is implicit when these values appear in Eq. (1) and its descendants.

The dispersion parameter σ_i is the *divergence of opinion* that Miller [1977] treats statically. Allowing σ_i to vary across agents (estimation heterogeneity) and over time (rising or falling market-wide divergence of opinion) is the key generalisation that makes EDEM a dynamic theory of price formation.

3.3 Per-tick interaction

At each tick:

1. Buyers move and post bids: a buyer b that lands on a seller s's home draws a bid $\beta^{(t)}_{b,s} = E_{b,t}(h(s))$ from Eq. (1) and presents it to s. Bids accumulate; the seller tracks the best bid received.
2. Sellers process bids: when its patience timer elapses, s offers to sign with the buyer holding the best bid. The buyer either commits (a sale clears) or declines (the bid is dropped); the commitment rule is the key behavioural primitive and is discussed below in Section 3.6.

3.4 Market-clearing functional

The realised market price at time t is the output of a clearing algorithm A that consumes the full market state. Let

$$X_t \;=\; \big(\mathcal{S}_t,\, \mathcal{B}_t,\, \{v_t(h)\}_{h\in G},\, \{\beta^{(t)}_{\bullet}\},\, \{a_t(s)\}_{s\in\mathcal{S}_t},\, \{\text{patience}_t(s)\},\, \mathcal{H}_t\big) \tag{2}$$

denote the model's state at t (agent populations, home values, outstanding bids and asks, seller patience timers, and the history $\mathcal{H}_t$ of completed sales). The realised market price is

$$p_t \;=\; A(X_t;\, \theta), \tag{3}$$

where θ collects structural parameters. The reduced state variables $Q_s(t)$, $Q_d(t)$, and a cross-population dispersion summary σ are the principal — but not the only — quantities that A depends on; we use the abbreviated notation $A(Q_s(t), Q_d(t), \sigma)$ when the remaining state is held implicit. In our two specific instantiations A is realised differently:

- In DE, A is the rolling average of the last W *sale events* (Eq. (9)). Sales are discrete events and not every tick records one.
- In the speculative-market EDEM variant, A is implicit in the multiplicative update rule Eq. (10); sales are not realised, and the market price is the average home value.

3.5 Stochastic supply and demand schedules

Crucially, Q_s and Q_d in Eq. (3) are not the deterministic schedules $Q_s(p)$ and $Q_d(p)$ of textbook microeconomics. They are integer-valued random walks driven by the price increment $\Delta p_{t-1} = p_{t-1} - p_{t-2}$ via two balancer parameters C_b, C_s:

$$Q_s(t) = Q_s(t-1) + Z_t^s, \qquad \mathbb{E}\big[Z_t^s\big] = C_b \operatorname{sgn}(\Delta p_{t-1}), \tag{4}$$

$$Q_d(t) = Q_d(t-1) - Z_t^d, \qquad \mathbb{E}\big[Z_t^d\big] = C_s \operatorname{sgn}(\Delta p_{t-1}). \tag{5}$$

The integer increments Z_t^s, Z_t^d are required because agent counts are integers; they realise possibly non-integer expected counts via the deterministic-plus-Bernoulli rule of Section 4.3, with a population floor of one agent per side. With $C_b, C_s > 0$ rising prices add sellers and remove buyers (mean-reverting); with $C_b, C_s < 0$ rising prices add buyers and remove sellers (trend-following); with $C_b = C_s = 0$ the populations are frozen and the price evolves under the bid-selection asymmetry alone. The simulations in Section 5.2 use $C_s = C_b$ and report the single number C_b.

Equations (4) and (5) describe the *interior* dynamics. A population floor of one agent per side (Section 4.3) truncates the random walk whenever $\max(Q_s, Q_d)$ would otherwise dictate killing the last remaining agent on the depleted side; under that boundary the expectation equations no longer hold literally. The one-vs-many regime that emerges for strong $|C_b|$ in Run 7 (Section 5.2.2) is exactly this boundary state.

Together, Eqs. (1) and (3) to (5) define a coupled stochastic process whose realisation is a sample path. Equilibrium in the textbook sense corresponds to the steady state $\Delta p_t = \Delta Q_s(t) = \Delta Q_d(t) = 0$ — a non-generic state that the process visits only fleetingly under most parameter regimes, which is the central observation of the paper.

3.6 Bid commitment (Cond. 2)

A buyer b offered to sign at price β commits iff β is at or above the buyer's own running benchmark over its outstanding bids:

$$b \text{ commits} \iff \beta \geq \frac{1}{|\mathcal{B}_b|} \sum_{\beta' \in \mathcal{B}_b} \beta', \tag{6}$$

where $\mathcal{B}_b$ is the set of all bids b has currently outstanding (the offered β included, so $|\mathcal{B}_b| \geq 1$ always). The intuition is order-statistic selection on the buyer's side: each $\beta' \in \mathcal{B}_b$ is a noisy estimate of fair value, and the bids b has placed on different sellers reflect different draws of b's estimation error. Committing only when the offered β is at or above the buyer's running benchmark restricts b's actual purchase to the upper tail of its own bid distribution — the homes for which b's estimate happened to land on the optimistic side. This is the buyer-side mirror of the seller-side max-bid selection of Section 3.7, and it is the second of the two ingredients that produces the order- statistic drift. The prior write-up of the model in Arbuzov [2018a] stated this rule with the inequality reversed; we restore the version that matches the simulation results actually reported there (Section 4.4).

3.7 The order-statistic drift: a special-case mechanism

The bubble mechanism that Runs 6–8 exhibit is a consequence of order-statistic bias. We give a closed-form result for a tractable special case below, then connect it empirically to the actual update rule Eq. (10).

Closed form for the seller-side maximum. Consider a seller s that has received $n \geq 2$ bids $\beta_1, \ldots, \beta_n$ from buyers with iid zero-mean estimation errors $\varepsilon_i \sim \mathcal{U}[-\sigma, +\sigma]$:

$$\beta_i \;=\; v_t(h(s)) \cdot (1 + \varepsilon_i). \tag{7}$$

The seller-level winning bid is $\max_i \beta_i$; its expected ratio to the home's current value is

$$\mathbb{E}\left[\frac{\max_i \beta_i}{v_t(h(s))}\right] \;=\; 1 + \mathbb{E}\Big[\max_i \varepsilon_i\Big] \;=\; 1 + \sigma\,\frac{n-1}{n+1}, \tag{8}$$

using the standard order-statistic identity for the maximum of n iid uniform variates. The right-hand side strictly exceeds one for all $n \geq 2$ and all $\sigma > 0$; the gap grows in both n (more bidders per seller) and σ (wider divergence of opinion).

What Eq. (8) does and does not prove. Equation (8) establishes that the *seller-side maximum* has expected ratio strictly above one. The per-epoch update in Eq. (10) compounds a related but distinct quantity $\bar{r}_t$: a buyer-side mean over each buyer's $\min_s$ winning ratio. We have not derived a closed form for $\bar{r}_t$. Two further care points apply when moving from the clean theorem to the sample-path behaviour of Fig. 6:

- *Expected level vs. typical trajectory.* For an iid multiplicative process $v_T = v_0 \prod_t r_t$, the expectation $\mathbb{E}[v_T] = v_0\,(\mathbb{E}[r])^T$ grows multiplicatively whenever $\mathbb{E}[r] > 1$, but the *median* (and any typical sample path) is governed by $\mathbb{E}[\log r]$, with median growth controlled by $\exp(T\,\mathbb{E}[\log r])$. By Jensen's inequality $\mathbb{E}[\log r] \leq \log \mathbb{E}[r]$, so $\mathbb{E}[r] > 1$ does not on its own imply $\mathbb{E}[\log r] > 0$.
- *Boundary truncation.* Equation (10) contains a $\bar{r}_t = 1$ fallback for epochs with no winning bids ($\mathcal{Y}_t = \varnothing$); this puts a point mass at unity on the distribution of r_t that the clean theorem does not capture.

Empirical bridge. Table 1 reports the empirical distribution of $\bar{r}_t$ extracted directly from the EDEM-regime experiments of Section 5.2, computed as $\bar{r}_t = v_{t+T}(h)/v_t(h)$ at every epoch boundary across all seeds. Three rows are reported per scenario: mean, median, and the share of epochs with $\bar{r}_t > 1$; an additional row reports $\mathbb{E}[\log \bar{r}_t]$, which is the relevant condition for sample-path exponential growth. The condition $\mathbb{E}[\log \bar{r}_t] > 0$ holds in every regime examined, with the no-balancer Run 6 the strongest at 0.041 per epoch.

Table 1: Empirical distribution of the per-epoch multiplier $\bar{r}_t$ in the EDEM regimes (Section 5.2), across 1500 epochs (10 seeds, 150 epochs/seed). The empty-$\mathcal{Y}_t$ point mass at 1 shows up as a low $\Pr(\bar{r}_t > 1)$ for the balanced regimes; in those cases the relevant growth quantity is $\mathbb{E}[\log \bar{r}_t]$, whose positivity certifies sample-path exponential growth.

Regime	$\mathbb{E}[\bar{r}_t]$	$\Pr(\bar{r}_t > 1)$	$\mathbb{E}[\log \bar{r}_t]$	median $\bar{r}_t$
Run 6 ($C_b = 0,\ \bar{\sigma} = 0.15$)	1.0429	83.5%	+0.0411	1.044
Run 7 ($C_b = +1,\ \bar{\sigma} = 0.15$)	1.0104	19.7%	+0.0098	1.000
Run 7 ($C_b = -1,\ \bar{\sigma} = 0.15$)	1.0153	25.4%	+0.0144	1.000
Run 8 ($C_b = -1,\ \bar{\sigma} \nearrow$)	1.0287	24.9%	+0.0236	1.000

Equation (8) should therefore be read as the *mechanism*: the same order-statistic logic that yields the closed form for the seller-side maximum also produces empirical $\mathbb{E}[\log \bar{r}_t] > 0$ for the buyer-min update we actually simulate. We do not claim that Eq. (8) proves Fig. 6; we claim that the closed-form theorem isolates the structural ingredient (maximum-order-statistic selection) that the empirical $\bar{r}_t$ inherits, and that Table 1 confirms the inheritance.

3.8 DE as a special case

The Dynamic Equilibrium model of Arbuzov [2018a] corresponds to the following restriction of EDEM:

- Estimation: $\varepsilon \sim \mathcal{U}[-\sigma_i, +\sigma_i]$, with $\sigma_i \sim \mathcal{U}[0, \bar{\sigma}]$ across agents (decimal convention; Section 3.2).
- Schedules: Q_s, Q_d are restored toward linear-in-price targets $\hat{Q}_s(p) = a_s + b_s p$, $\hat{Q}_d(p) = a_d + b_d p$ once per balance period T_B.

- Clearing: A is the rolling average of the last W completed sales. Let $k(t)$ count the completed sales up to tick t and P_j^{sale} denote the price of the j-th sale; then

$$p_t = \frac{1}{m_t} \sum_{j=k(t)-m_t+1}^{k(t)} P_j^{\text{sale}}, \qquad m_t = \min(W, k(t)). \tag{9}$$

This is the formulation that Section 5.1 simulates. The sale-event indexing of Eq. (9) matters: ticks without completed sales contribute nothing, and the average is taken over realised events rather than time. The implementation maintains the window as a fixed-length deque of the last W sale prices.

The speculative-market EDEM variant of Section 5.2 replaces Eq. (9) with a multiplicative per-epoch update on each home's value. Define:

$$\mathcal{W}_b(t) = \{s \in \mathcal{S}_t : b \text{ holds the winning bid on } s \text{ during the epoch ending at } t\},$$
$$\mathcal{Y}_t = \{b : \mathcal{W}_b(t) \neq \varnothing\}.$$

The per-epoch update is then

$$v_{t+T}(h) = v_t(h) \cdot \bar{r}_t \quad \forall h \in G, \qquad \bar{r}_t = \begin{cases} \dfrac{1}{|\mathcal{Y}_t|} \displaystyle\sum_{b \in \mathcal{Y}_t} \min_{s \in \mathcal{W}_b(t)} \dfrac{\beta_{b,s}^{(t)}}{v_t(h(s))} & \text{if } \mathcal{Y}_t \neq \varnothing, \\ 1 & \text{if } \mathcal{Y}_t = \varnothing. \end{cases} \tag{10}$$

The update is uniform across all homes in G. This is a strong *propagation assumption*: real-estate "comp" signals empirically diffuse non-uniformly, weighted by neighbourhood, amenities, and time-since-sale. Adopting a uniform multiplier keeps the model focused on the order-statistic mechanism of Section 3.7 rather than on diffusion geometry; we flag it as a candidate target for extension in Section 7.1.

3.9 Walrasian and Miller cases as limits

Setting $\sigma_i \equiv 0$ for all i collapses Eq. (1) to $E_{i,t}(h) = v_t(h)$: every agent agrees on the value of every home. The order-statistic drift of Eq. (8) vanishes ($\bar{r}_t \equiv 1$) and the rolling clearing in Eq. (9) delivers a constant price set by $\hat{Q}_s(p^*) = \hat{Q}_d(p^*)$. Under the additional assumption that the matching process and balancer introduce no frictions (instantaneous adjustment, costless encounters), EDEM approaches the Walrasian fixed-point equilibrium as a limiting case. Spatial frictions, patience, and the finite-population integerization of Eqs. (4) and (5) prevent exact equivalence at finite scales.

If we instead retain heterogeneous σ_i but freeze the dynamics ($T_B \to \infty$, single-period sale, no value update), we recover a *Miller-style static premium mechanism*: the N shares are assigned to the N most optimistic of the $K > N$ potential investors, so the clearing price reflects the optimism of the marginal optimist rather than the mean valuation. Exact equivalence to the model in Miller [1977] requires short-selling restrictions and the fixed-supply assumption; the mechanism here is the same upper-tail selection but in a different institutional framing.

The two regimes that empirical studies of divergence of opinion have struggled to reconcile (premium vs. discount; Section 2.2) appear in EDEM as two ends of a continuum indexed by the balancer parameters C_b, C_s and the ratio of estimation error to balancer strength. Section 5 demonstrates this empirically.

4 Implementation in Mesa

EDEM and DE are implemented in the Python agent-based modelling framework Mesa [Masad and Kazil, 2015], replacing the earlier NetLogo prototypes whose NetLogo source is preserved in the `drafts/netlogo/` subdirectory of the same repository. The choice of Mesa is pragmatic: its turtle–patch primitives map cleanly onto NetLogo's, but the surrounding Python ecosystem (NumPy, pandas, matplotlib, pytest) supports batched experimentation, deterministic seeding, and reproducible figure regeneration that the original NetLogo code cannot. The Mesa port, the NetLogo source, and the present paper are released together at `https://github.com/sibmike/dynamic-disequilibrium` — code under MIT, paper under CC-BY-4.0.

4.1 Spatial structure and agents

The market is staged on `mesa.space.MultiGrid(32, 32, torus=True)`, matching the wrapped 32×32 patch world of the original NetLogo. Per-cell state ($v_t(h)$, $v^*(h)$, last sale price, last sale tick) is stored as a lightweight `Home` dataclass attached to each grid cell, not as a Mesa Agent — treating the 1024 cells as Agents would multiply scheduler overhead without benefit, since cells never schedule behaviour.

Buyers and sellers are Mesa Agents:

- `Seller` is immobile (one per home), holds an ask price $a_t(s)$ and a patience timer, and accumulates incoming bids until the timer elapses.
- `Buyer` is mobile; each tick it executes a NetLogo-style *wiggle* (uniform heading change in $[-90°, +90°]$) followed by a unit forward step.

Bids are NetLogo undirected links in the original; we represent them as mirrored entries in Python dictionaries on each side (each agent's `bids` dict is keyed by the counterparty). The mirror is maintained in a single helper to avoid drift.

4.2 Market price and clearing

For DE, the rolling 25-sale market price of Eq. (9) is implemented as a `collections.deque(maxlen=25)` owned by the model. Sales are committed via a single `complete_sale` method that records the sale, advances the deque, removes both counterparties from the grid, and triggers the balancer to spawn a replacement pair (per the NetLogo `add_seller` / `add_buyer` idiom).

For the EDEM speculative-market variant, no sales are recorded. Instead the model maintains a per-tick `cycle_counter` initialised to `init_patience` $-\ 1$. When the counter reaches zero, the model fires Eq. (10) on every home in a single pass and resets the counter. The off-by-one initialisation is deliberate: it ensures the value update reads buyers' `is_yellow` flags one tick before the buyers' patience timers reset them, matching the order of operations in the NetLogo `to go` procedure.

4.3 Balancer

DE balances populations every T_B ticks by recomputing the linear target $\hat{Q}_s(p), \hat{Q}_d(p)$ at the current market price and adjusting by one agent in each direction. Spawning a replacement on every sale (the NetLogo behaviour) plus this slow drift toward the linear target keeps the population near $Q_s^* = Q_d^*$ without introducing artificial discrete jumps.

The EDEM balancer implements Eqs. (4) and (5) in finite-population form. The fractional balancer coefficient C_b is realised as $\lfloor |C_b| \rfloor$ deterministic swaps per epoch plus one Bernoulli swap with probability $|C_b| - \lfloor |C_b| \rfloor$; the sign of C_b controls direction. A population floor prevents the balancer from killing the last agent on either side; the 1-vs-many state that emerges with strong $|C_b|$ is kept as a feature of the model rather than papered over (Section 5.2.2).

4.4 Bid-acceptance rule (Cond. 2)

The original Arbuzov [2018a] prose stated the buyer's bid-acceptance rule with the strict inequality $\beta < \mathrm{mean}(\mathcal{B}_b)$, while the accompanying NetLogo source compared with the opposite sign. Empirically the prose-as-written rule causes a runaway price collapse, since it selects every sale to clear at the buyer's lowest bid; only the NetLogo-as-coded rule reproduces the stable equilibrium reported in the original Run 1. We adopt Eq. (6) (the NetLogo-as-coded rule, restated economically) as the canonical Cond. 2 and document this discrepancy explicitly in the source. The Mesa implementation exposes both rules as a parameter (`accept_rule="netlogo"` or `"prose"`) for sensitivity analysis.

4.5 Determinism, parallelism, and reproducibility

The Mesa `Model.rng` is seeded explicitly per run, and all agent-side randomness (epsilon draws, heading wiggles, balancer Bernoulli) draws from this single stream. Each experiment script in `python_simulation/experiments/` runs $N \geq 8$ independent seeds of the same parameter vector and stacks the per-tick model reporters into a single Parquet dataset; figures plot the median across seeds plus the 10th–90th percentile band. Pytest unit tests (`python_simulation/tests/`) cover the rolling market price, both Cond. 2 rules, the linear and fractional balancers, the EDEM epoch-timing trick, and the C_b-sign inversion.

4.6 Per-tick pseudocode

The complete tick is:

```
def step(self):  # Model.step
    self.agents.shuffle_do("step")      # buyers + sellers, randomised
    self.balancer.step(self)            # DE: every T_B; EDEM: per-epoch
    self.datacollector.collect(self)
```

where `Buyer.step` runs *buy* + *wiggle* + *move* in NetLogo order and `Seller.step` runs the patience-timer logic (timer decrement; on timeout, either drop the highest bid below ask and lower the ask, or offer to sign the highest bidder, applying Eq. (6)). The full source is under 600 lines including docstrings.

5 Experiments and Results

We report eight primary controlled experiments grouped into two regimes, plus a sensitivity-grid sweep (Run 9, Section 5.2.4) that anchors the conclusions across a 30-cell parameter grid. Section 5.1 (Runs 1–5) uses the Dynamic Equilibrium specialisation of EDEM with linear supply and demand schedules, discrete sales, and a rolling-mean clearing price; these isolate the effects of estimation noise, patience, agent density, and shocks under a textbook population balancer. Section 5.2 (Runs 6–8) uses the speculative-market EDEM variant with multiplicative value updates; these illustrate the bubble, balancer sign, and rising-error regimes that the static divergence-of-opinion literature cannot reach.

Unless stated otherwise, all runs use a 32×32 toroidal grid, balance period $T_B = 100$ ticks, sale window $W = 25$ sales, and the bid-acceptance rule Eq. (6). Each experiment is replicated across at least eight independent seeds; figures plot the median across seeds plus the 10th–90th percentile band, with individual seed trajectories overlaid faintly. Full parameter tables are in Section B.

5.1 Dynamic Equilibrium experiments

5.1.1 Run 1 – Stable equilibrium under favourable conditions

We begin from the textbook equilibrium $p^* = 100$, $q^* = 50$ defined by $\hat{Q}_s(p) = 0.5\,p$ and $\hat{Q}_d(p) = 100 - 0.5\,p$, with maximum estimation error $\bar{\sigma} = 5$ percent, patience drawn uniformly on $[0, 50]$, and balance period $T_B = 100$. Figure 1 shows the median market price hovering close to p^*: across ten seeds and 20,000 ticks, the maximum upward deviation is $+10.7\%$ and the maximum downward deviation -2.4%. The companion buyer/seller counts oscillate within ± 5 of the equilibrium quantity. The slight upward bias of the median trajectory is consistent with Arbuzov [2018a]'s reported $-4.9\%/+6.2\%$ envelope at 100,000 ticks; we attribute it to the bid-selection asymmetry discussed in Section 3.8, which the rolling average partially absorbs.

5.1.2 Run 2 – Imprecise valuations cause business cycles

Holding all other parameters at Run 1 values and increasing $\bar{\sigma}$ from 5 to 25 percent, the market loses its narrow-band stability. Figure 2 shows recurrent oscillations of the market price between roughly $+34\%$ and -17% of equilibrium, accompanied by counter-phased oscillations of buyer and seller populations: episodes of overbidding draw sellers into the market while buyers retreat, after which the seller surplus pushes prices back down and the cycle reverses. This is a pure consequence of estimation noise — there are no exogenous shocks — and it suggests that observed business cycles in real markets need not require any external driver beyond heterogeneous misvaluation, provided the divergence of opinion is large relative to the balancer's adjustment speed.

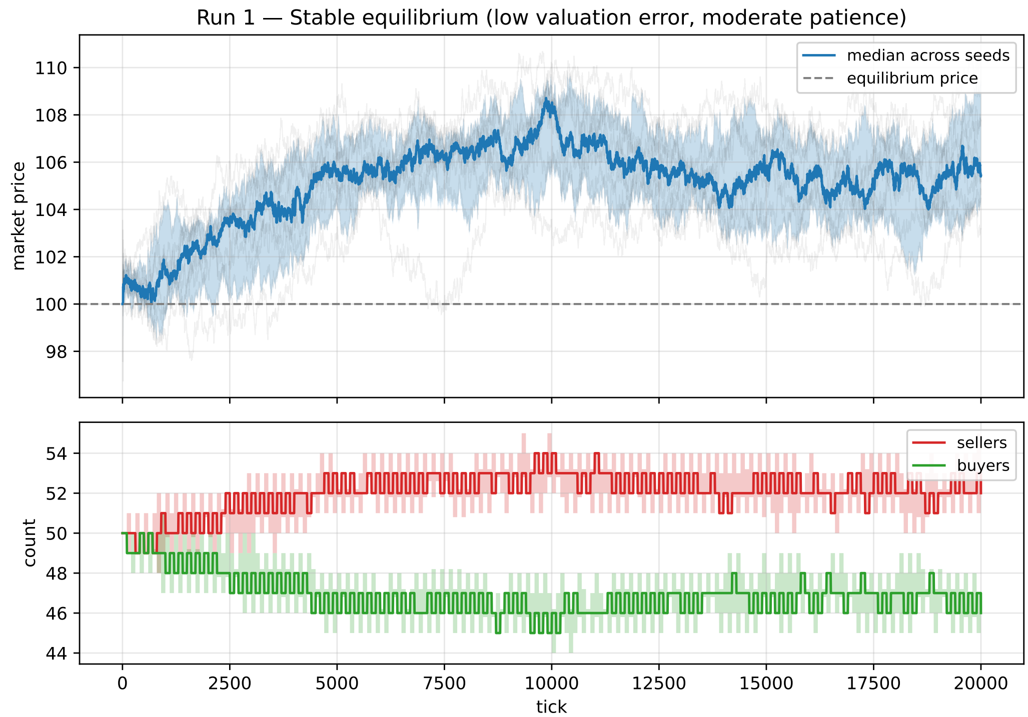

Figure 1: Run 1 – Stable equilibrium with low valuation error and moderate patience. Top: market price (median, with 10–90 percentile band; faint individual-seed traces). Bottom: agent counts. Dashed line: textbook equilibrium price.

5.1.3 Run 3 – Patience pushes price above equilibrium

Doubling the maximum seller patience from 50 to 100 ticks (holding $\bar{\sigma}$ at the Run 1 value of 5%) produces a durable upward shift in the market price. Figure 3 shows the median price climbing monotonically to roughly $+14\%$ above p^* over 20,000 ticks and showing no sign of reverting. The seller population settles at ~ 57 (about $+14\%$ of q^*) and the buyer population at ~ 41 (-18%). Patience increases sellers' time-on-market and thus their willingness to wait for richer bids; the realised sale prices tilt up; the rolling average follows. Arbuzov [2018a] reported a similar $+10\%/+10\%/-20\%$ shift, attributing it to seller forbearance on stagnating real-estate markets where neither side can quickly exit.

5.1.4 Run 4 – Low agent density pushes price below equilibrium

Shifting the demand intercept downward by 50 units yields a new textbook equilibrium $(p^*, q^*) = (50, 25)$. With only ~ 25 buyers and ~ 25 sellers on a 32×32 grid, encounters are sparse: a typical buyer's wiggle-and-step trajectory rarely crosses a seller's home before the seller's patience timer expires. Sellers therefore lower ask prices repeatedly without receiving adequate bids, and the realised sale prices — and hence the rolling market price — settle *below* the textbook equilibrium. Figure 4 shows a steady downward drift; the mean realised price across the second half of the simulation is $p \approx 41$, an 18% shortfall against the textbook equilibrium of 50, with episodes reaching as low as $p \approx 37$ (a 26% shortfall at the trough). The seller population shrinks to ~ 19 and the buyer population grows to ~ 30 as the balancer chases the depressed price. Patience and density are partial substitutes (Section 2.3): both proxy for time-on-market.

5.1.5 Run 5 – Shock handling and transitional markets

The previous four runs evolve under fixed parameters. Run 5 introduces shocks via the `set_demand/set_supply` hooks (Section 4.6) and demonstrates how the model absorbs them. Figure 5 reports two scenarios.

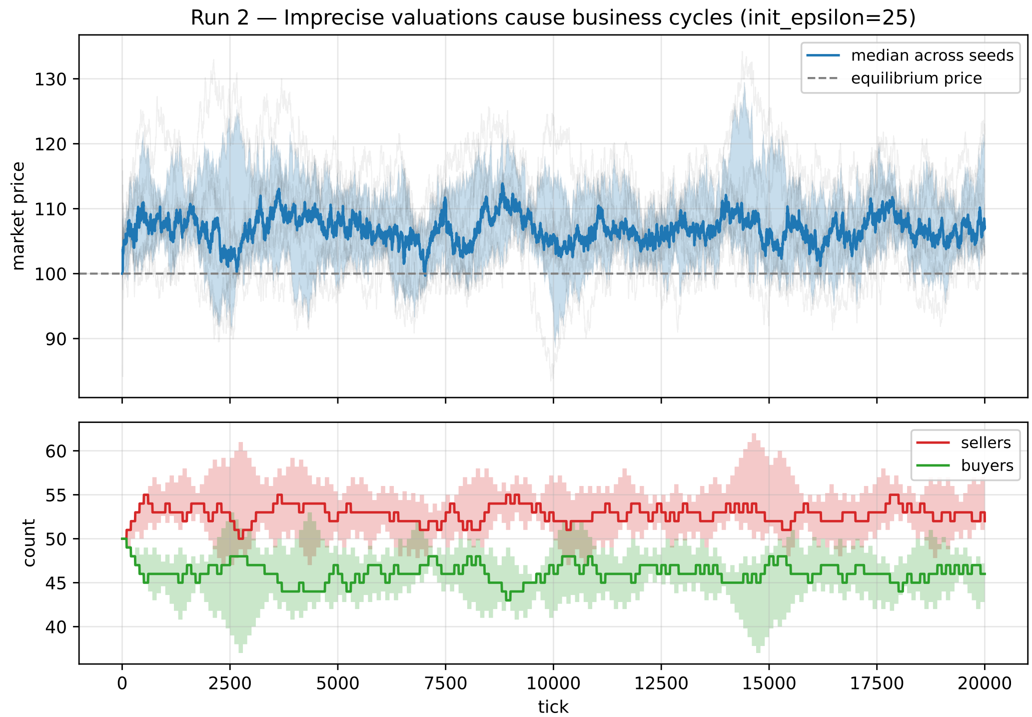

Figure 2: Run 2 – Imprecise valuations ($\bar{\sigma} = 25\%$) generate endogenous business cycles. Same conventions as Fig. 1.

Scenario A (single shock). The market begins in the Run 1 equilibrium. At $t = 3000$ the demand intercept drops to 50, collapsing the textbook equilibrium to $p^* = 50$. The realised price descends slowly — it has not finished adjusting by $t = 7000$ when seller patience is raised to 165. The combination of a lowered demand and elevated patience continues to drive the market downward (because the price is still adjusting from the previous $p^* = 100$), and even by $t = 12000$ the realised price has not settled. Doc 2's claim that high patience would lift the price *back to* $p^* = 50$ presumes a fully-adjusted starting state that is not reached within twelve thousand ticks; the figure makes visible the more general Arbuzov [2018a] observation that "adjustment to single shock takes a long time."

Scenario B (transitional market). The demand intercept toggles between 125 and 75 every 2000 ticks. The textbook equilibrium accordingly steps between $p^* = 125$ and $p^* = 75$, but the realised market price never settles at either: each shock is absorbed only partially before the next arrives. The figure confirms a market *in constant transition from one unknown state to another* — the empirical foothold for Section 3.5's claim that supply and demand are sample paths of a stochastic process, not deterministic schedules.

5.2 Speculative-market EDEM experiments

The remaining three runs use the EDEM variant of Section 3.8 with multiplicative value updates Eq. (10). We track the dimensionless ratio $v_t(h)/v^*(h)$ averaged over all homes; a value of 1 indicates fair valuation, values above 1 indicate the market has priced homes above their unobservable fair value.

5.2.1 Run 6 – Speculative bubble at $C_b = 0$

With balancer disabled ($C_b = 0$) and divergence of opinion $\bar{\sigma} = 15\%$, the market exhibits a clean log-linear bubble. Figure 6 shows the median value-to-true ratio multiplying by $\sim 421\times$ over 3000 ticks (150 epochs), with the band tightly enclosing the median: the bubble is deterministic in shape and stochastic only in slope. Buyer and seller populations are frozen at 20 each, since $C_b = 0$ admits

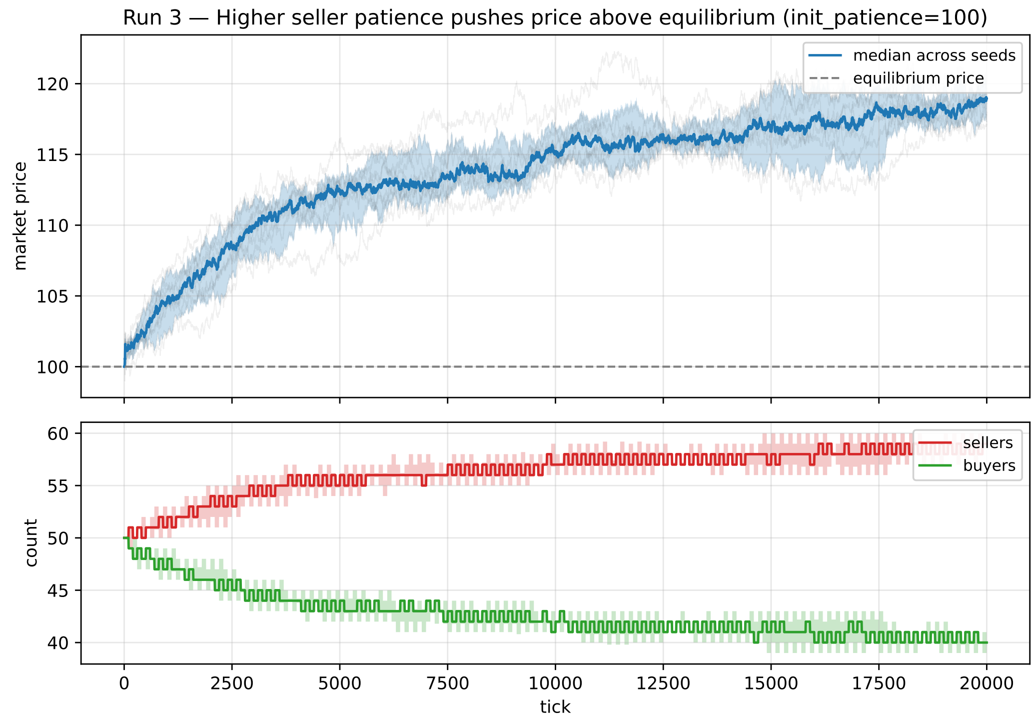

Figure 3: Run 3 – Doubling seller patience to 100 ticks pushes the market price about 14% above the textbook equilibrium and sustains it there. Same conventions as Fig. 1.

no balancer adjustment. The mechanism is the bid-selection asymmetry of Section 3.8: each epoch's average winning ratio $\bar{r}_t$ is > 1 in expectation because winning bids are the maxima of multiple noisy draws, and Eq. (10) compounds this multiplicatively over epochs.

5.2.2 Run 7 – Balancer sign shapes (but cannot eliminate) the bubble

Holding $\bar{\sigma}$ at 15% and varying $C_b \in \{+1, 0, -1\}$ yields three qualitatively distinct trajectories overlaid in Fig. 7. The mean-reverting balancer ($C_b = +1$) dampens the bubble most aggressively: the value ratio plateaus near $4\times$ as the seller population grows to roughly 40 and the buyer population shrinks to one. The trend-following balancer ($C_b = -1$) admits a larger bubble than $+1$, ending near $7.7\times$, with mirrored populations (one seller, ~ 40 buyers). The no-balancer case ($C_b = 0$) dominates both, ending at $\sim 421\times$.

This last observation refines the framing in Arbuzov [2018b], which described the trend-following case as the most explosive (the "Silicon Valley real estate" or "bitcoin" regime). In the Mesa simulation the most explosive case is the no-balancer one: removing one population pole entirely (as trend-following does) reduces the multiplicative pressure on each remaining seller, because there are simply fewer sellers whose values can be inflated. The trend-following case is more aggressive than mean-reverting, but neither is as aggressive as no balancer.

The companion finding is that no setting of C_b in this parameter grid restores $v_t/v^* = 1$. Bid-selection asymmetry introduces a positive bias that even a strong mean-reverting balancer cannot fully cancel within the timescales considered.

5.2.3 Run 8 – Trend-following balancer with rising error

Doc 2's stylised "bitcoin" or "Silicon Valley" regime combines trend-following ($C_b = -1$) with a divergence of opinion that itself grows over time, modelling the historical observation that late-stage bull markets see widening dispersion in analyst forecasts. We initialise $\bar{\sigma}_0 = 0.05$ (i.e. 5%) and let $\bar{\sigma}_{t+T} = \bar{\sigma}_t + 0.005$ each epoch (equivalently, $+0.5$ percentage points), so $\bar{\sigma}$ reaches 0.80 (80%) by

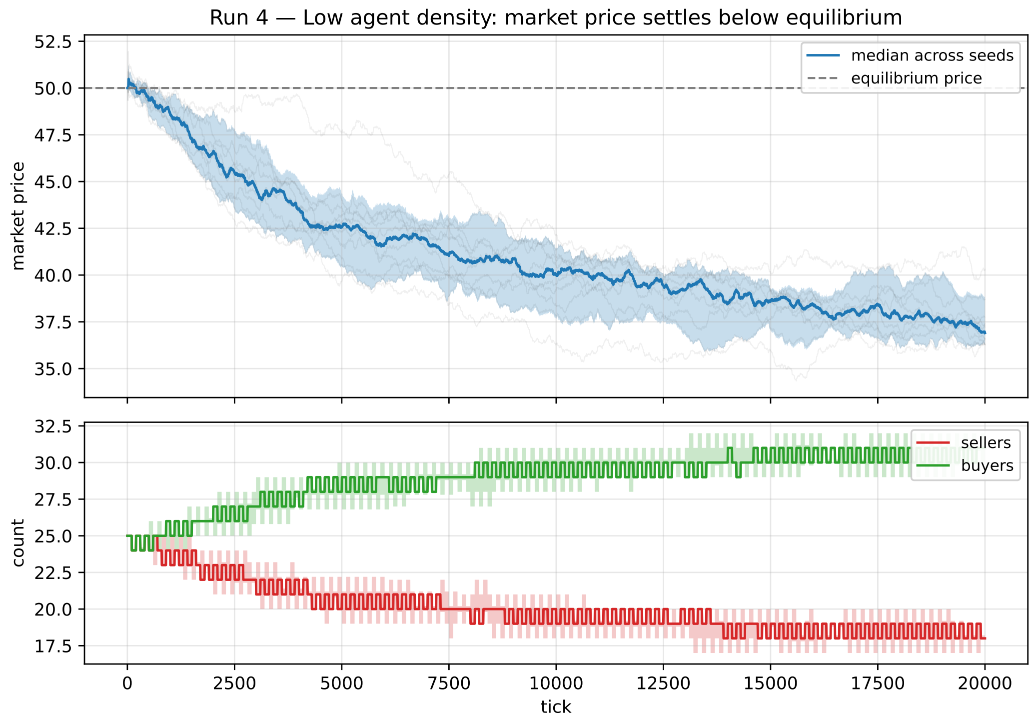

Figure 4: Run 4 – Halving the demand intercept reduces equilibrium quantity to 25. Sparse encounters cause sellers to lower ask prices repeatedly; the realised market price settles about 18% below the textbook equilibrium. Same conventions as Fig. 1.

the end of the run. Figure 8 shows the resulting trajectory: the value ratio climbs to $\sim 30\times$ over 3000 ticks. The trajectory is super-linear in linear time but visibly concave-down in log-time, indicating sustained exponential rather than double-exponential growth at this parameter setting; we conjecture that genuine double-exponential acceleration would emerge for larger $\bar{\sigma}$ growth rates than we explore here. The population panel mirrors Run 7's $C_b = -1$ behaviour: sellers drained to one, buyers saturated at ~ 44.

5.2.4 Run 9 – Sensitivity grid

The eight primary runs reported above each fix a single parameter vector. To rule out the possibility that the order-statistic drift of Section 3.7 is an artefact of one well-chosen calibration, we sweep C_b and $\bar{\sigma}$ over a 5×6 grid ($C_b \in \{-1, -0.5, 0, +0.5, +1\}$, $\bar{\sigma} \in \{5, 10, 15, 20, 25, 30\}\%$), running five seeds per cell for 1500 ticks. Figure 9 reports the median terminal v_t/v^* across seeds, on a $\log_{10}$ colour scale.

Three observations stand out:

1. **The drift is universal in this range.** Every cell of the grid has $v_t/v^* > 1$, and 100% of cells exceed $1.5\times$. The order-statistic mechanism is not a corner-case.
2. **Bubble territory is large.** 43% of cells exceed $10\times$ within 1500 ticks; growth at $\bar{\sigma} = 30\%$ reaches $235\times$ in the no-balancer column.
3. $C_b = 0$ **dominates within each** $\bar{\sigma}$ **column.** This confirms the Run 7 finding (Section 5.2.2) at five distinct $\bar{\sigma}$ values: removing the balancer is more explosive than either mean-reverting or trend-following balancers, across the whole grid.

The grid does not constitute a closed-form proof that the drift survives *all* parameter regimes; in particular, very large $|C_b|$ values that drain one side of the population to a single agent within a few epochs may behave differently. Section 6.3 discusses the limits of the linear-balancer intervention more carefully.

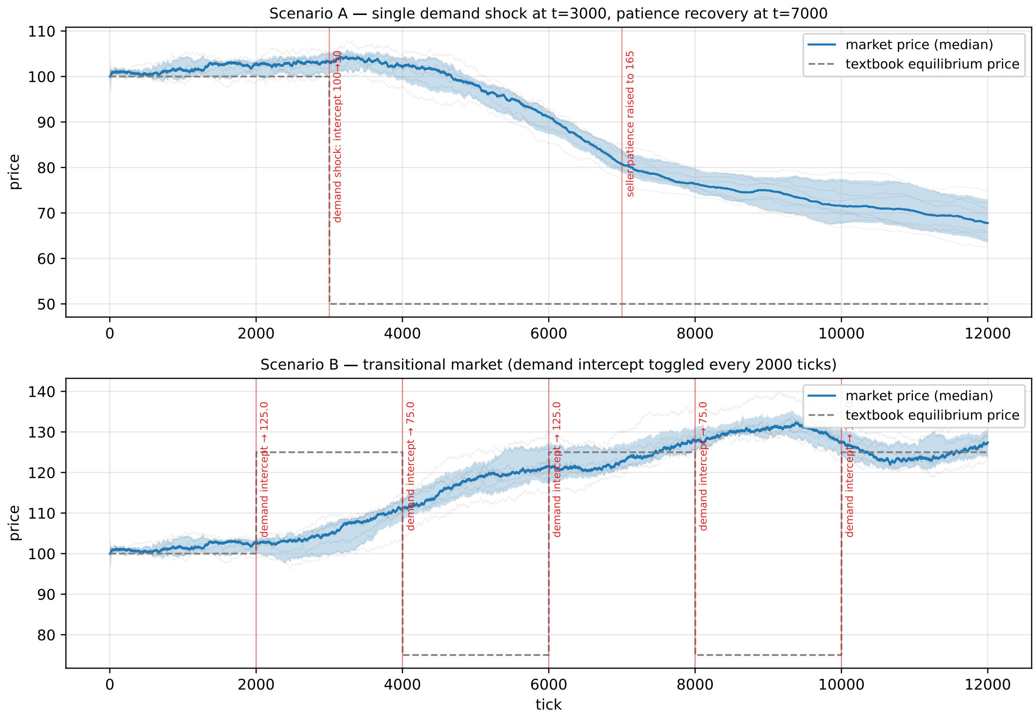

Figure 5: Run 5 – Two shock scenarios. Top: a single demand shock at $t = 3000$ followed by a patience boost at $t = 7000$; the market absorbs the shock slowly and has not fully adjusted by the end of the simulation. Bottom: a transitional market in which the demand intercept toggles every 2000 ticks. The market price never converges to either textbook equilibrium. Red dashed verticals mark shock events.

Table 2: Eight EDEM runs: parameters and observed outcomes at the end of each simulation. Median across ≥ 8 seeds. DE runs report median deviation of the rolling market price from the textbook equilibrium; EDEM runs report the dimensionless v_t/v^* ratio.

Run	**Manipulated parameter**	**Outcome metric**	**Regime**
1	$\bar{\sigma} = 5\%$, patience ≤ 50	-2% to $+11\%$ band	band-stable
2	$\bar{\sigma} = 25\%$	-17% to $+34\%$, periodic	business-cycle
3	patience ≤ 100	$+14\%$ persistent	overshoot
4	demand intercept -50	-18% persistent	undershoot
5	shock at $t = 3000$, then re-shocks	price chases moving target	transitional
6	EDEM, $C_b = 0$, $\bar{\sigma} = 15\%$	$v_t/v^* = 421\times$	runaway bubble
7	EDEM, $C_b \in \{-1, 0, +1\}$	7.7, 421, 4.2$\times$ (in C_b order)	balancer-shaped
8	EDEM, $C_b = -1, \bar{\sigma} : 5 \to 80$	$v_t/v^* = 30\times$	sustained super-linear

5.3 Summary across runs

Table 2 consolidates the end-of-run statistics. The eight primary runs sweep three parameter axes: estimation noise ($\bar{\sigma}$), seller patience, and agent density (Runs 1–4); a shock schedule (Run 5); and the EDEM balancer parameter C_b with a fixed or growing $\bar{\sigma}$ (Runs 6–8). Together they delineate six qualitatively distinct steady-state regimes – band-stable, business-cycle, persistent overshoot, persistent undershoot, runaway exponential, and constant transition – all reachable from the same underlying agent ruleset.

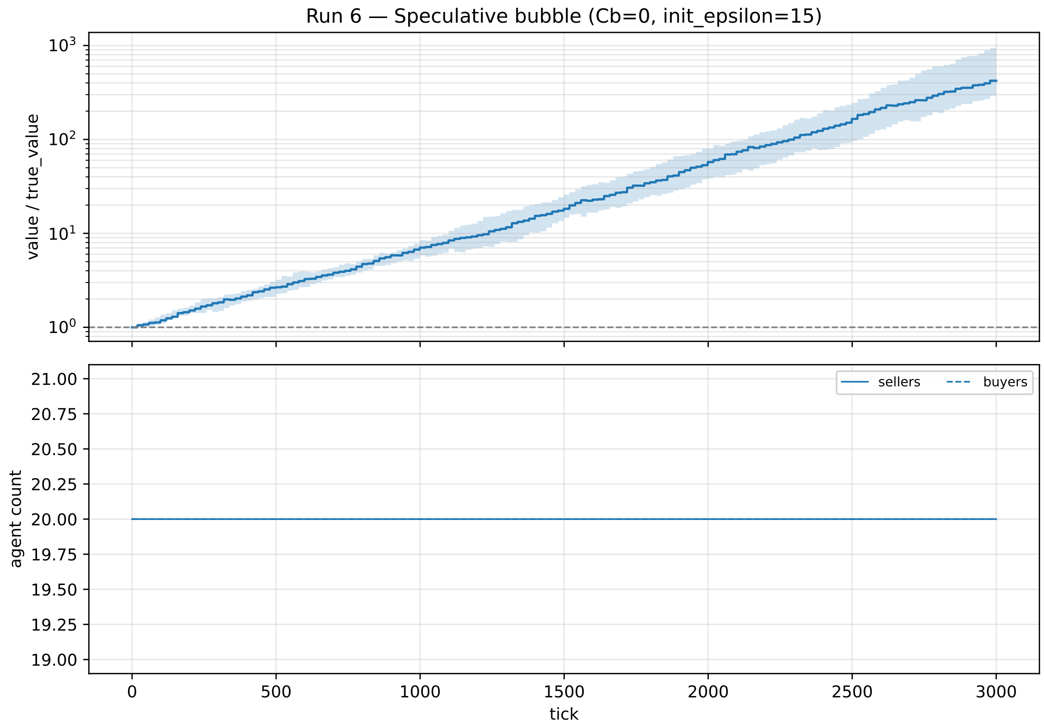

Figure 6: Run 6 – Speculative bubble at $C_b = 0$, $\bar{\sigma} = 15\%$. Top: dimensionless market value v_t/v^* on a logarithmic axis; the trajectory is approximately linear in $\log v$, indicating exponential growth. Bottom: agent populations remain at $(20, 20)$.

6 Discussion

The eight primary experiments plus the Run-9 sensitivity grid together support three substantive claims about how real markets generate the price patterns that empirical research has found contradictory.

6.1 Time-on-market is the missing third dimension

Runs 3 and 4 show that two distinct manipulations — raising seller patience and lowering agent density — both produce stable deviations from the textbook equilibrium price. Their effect signs are opposite (Run 3 overshoots, Run 4 undershoots), but their mechanism is the same: both alter the typical *time on market* of an unsold home. Patience increases the upper bound on time on market directly; density increases it indirectly by decreasing encounter rates per unit time. In each case the mismatch between the balancer's price-to-quantity correction and the underlying microstructural adjustment shows up as a persistent equilibrium offset.

This suggests that the textbook supply-and-demand pair $\hat{Q}_s(p)$ and $\hat{Q}_d(p)$ omits a state variable. A more accurate characterisation would be $\hat{Q}_s(p, \tau)$ and $\hat{Q}_d(p, \tau)$ where τ is expected time on market. Empirically, real-estate datasets that include days-on-market information vastly outperform those that do not for short-run forecasting; our model offers a generative explanation for that finding.

6.2 Bid-selection asymmetry, not estimation bias, drives bubbles

The estimation function Eq. (1) has zero mean error: an agent's expected valuation is exactly the home's current value. Yet the EDEM speculative regime (Runs 6–8) produces unbounded upward drift even with this unbiased estimator. The mechanism is not behavioural bias but a statistical asymmetry: each seller selects the *maximum* bid received, and the maximum of N noisy estimates of a fair value is, in expectation, larger than the fair value itself. Eq. (10) compounds this asymmetry multiplicatively across epochs.

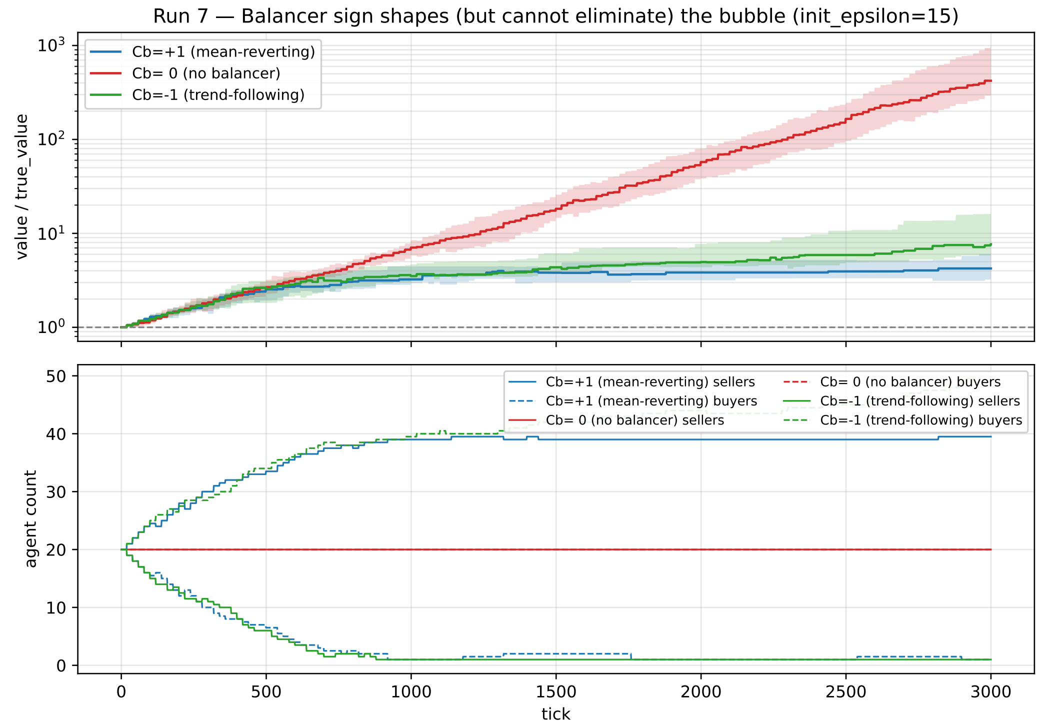

Figure 7: Run 7 – Balancer-coefficient sweep at $\bar{\sigma} = 15\%$. Three regimes are overlaid: mean-reverting ($C_b = +1$), no balancer ($C_b = 0$), and trend-following ($C_b = -1$). The no-balancer case dominates; the trend-following case beats the mean-reverting case but is itself bounded above by the no-balancer case. Bottom panel shows the population swings induced by each balancer.

This refines the standard "divergence of opinion" story. Miller [1977]'s static premium emerges from constrained supply: the N shares are sold to the N most optimistic of $K > N$ investors. EDEM's dynamic premium emerges from *repeated* constrained supply: each epoch's winners feed back into the next epoch's anchor value. The static premium is a corner of the dynamic premium, recovered when the feedback loop is severed.

A practical implication is that empirical premium-vs-discount findings depend strongly on the sample's position along this feedback trajectory, which short-window studies cannot identify. Two studies that find opposite signs may both be correct characterisations of distinct epochs in the same underlying process.

6.3 Mean-reverting balancers cannot fully cancel the asymmetry

Run 7's three-way comparison shows that no setting of the EDEM balancer coefficient C_b within the parameter grid we examined restores $v_t/v^* = 1$, and the 30-cell sweep of Run 9 (Section 5.2.4) reproduces this finding at every combination of C_b and $\bar{\sigma}$ tested. The mean-reverting case ($C_b = +1$) plateaus near $4\times$ at $\bar{\sigma} = 0.15$; the trend-following case ($C_b = -1$) continues drifting upward; the no-balancer case dominates both. A corollary, qualified by the parameter range surveyed, is that *linear policy interventions cast as balancer adjustments will not, on the evidence assembled here, eliminate persistent mispricing.* We do not claim this as an impossibility result — the parameter grid is finite, and we have not explored extreme $|C_b|$ values where the population dynamics qualitatively change. The asymmetry is structural to the bidding mechanism; to neutralise it, an intervention would have to change either the selection rule (e.g., uniform random matching instead of max-bid) or the per-epoch update rule (e.g., a winsorised mean instead of the arithmetic mean). Section 7 sketches both alternatives.

The companion observation — that the no-balancer case dominates the trend-following case in our simulation — refines the framing in Arbuzov [2018b], which described the trend-following "Silicon

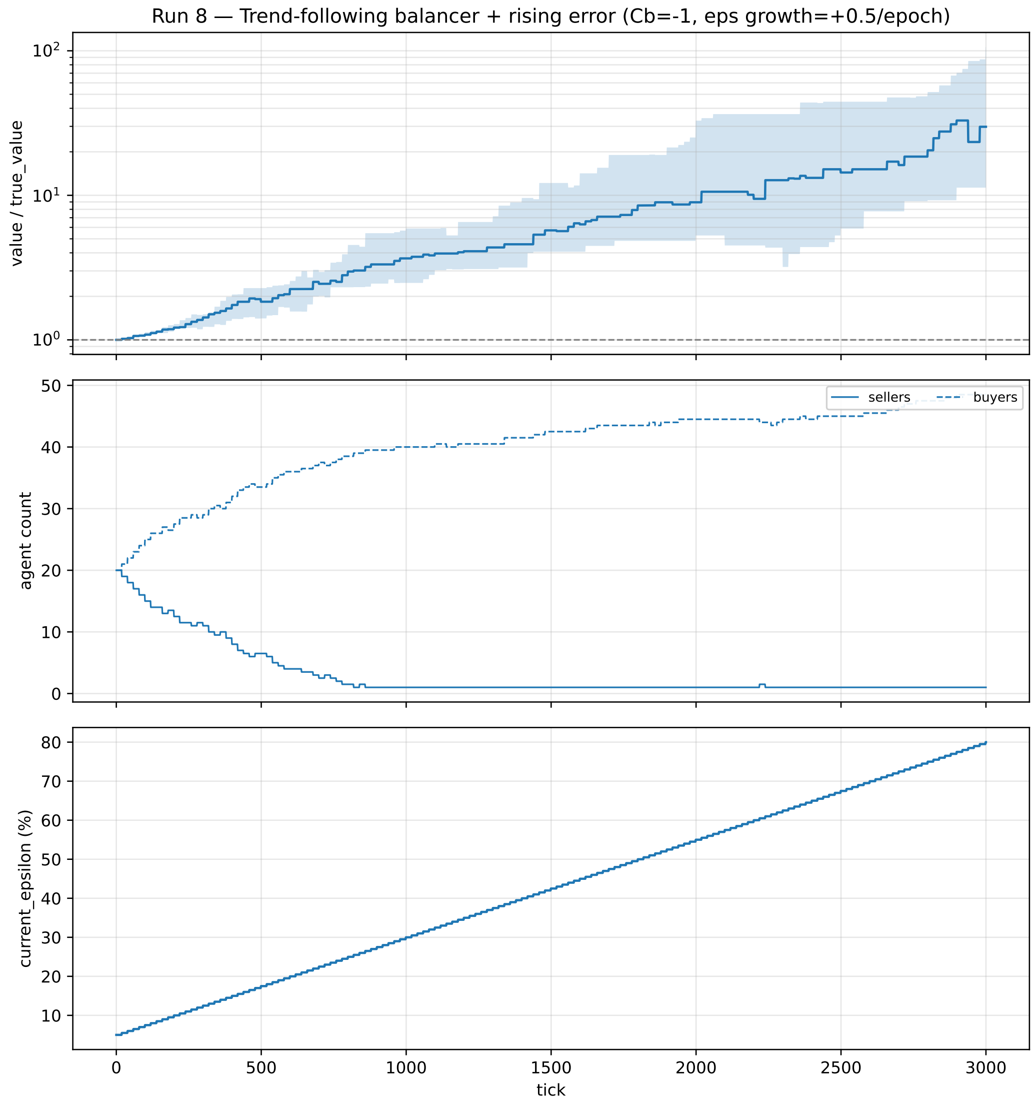

Figure 8: Run 8 – Trend-following balancer ($C_b = -1$) combined with linearly-growing divergence of opinion ($\bar{\sigma} : 5 \to 80$). Top panel: market value on a log scale. Middle: agent populations. Bottom: $\bar{\sigma}$ over time. The trajectory is sustained exponential rather than visibly double-exponential.

Valley real estate" regime as the most explosive. That framing turns out to overstate the role of the balancer's sign relative to its presence. We attribute the discrepancy to the agent-pool depletion that strong $|C_b|$ induces: trend-following removes sellers, but the remaining sellers are all that the multiplicative update rule can act on.

6.4 Implications for valuation algorithms

The same mechanism that drives EDEM's exponential bubble in Run 6 operates in any pricing system whose inputs are anchored to recent market-clearing prices. A machine-learning real-estate valuation algorithm trained on historical sale prices estimates not the home's fair value but the typical winning bid for a similar home — which is, by construction, the maximum of the relevant noisy estimates, and therefore systematically above the fair value. Deploying such an algorithm at scale risks providing the market with a coordinating signal that pulls realised winning bids upward, which then feeds back into the next training cycle.

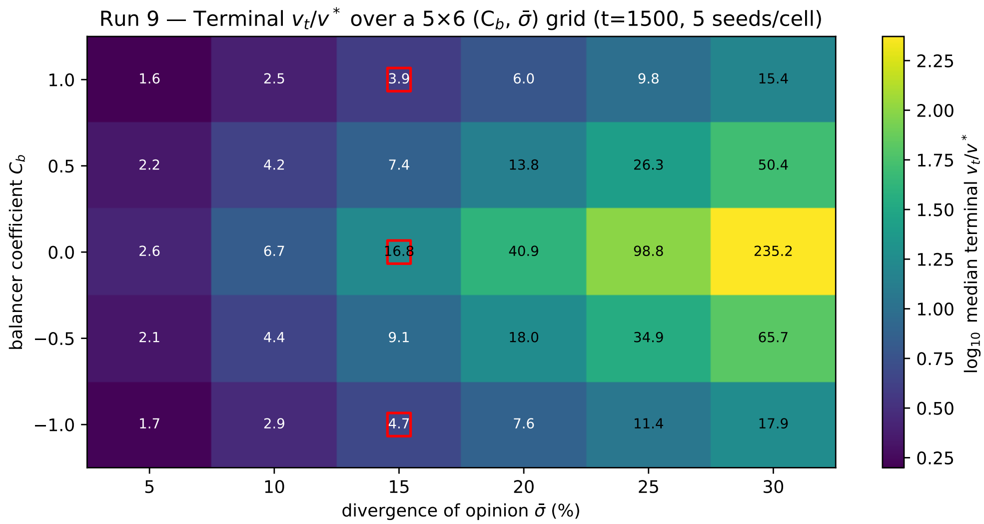

Figure 9: Run 9 – Sensitivity grid: median terminal v_t/v^* over a 5×6 sweep of balancer coefficient C_b versus divergence of opinion $\bar{\sigma}$, with 5 seeds per cell at $t = 1500$. Colour: $\log_{10}$ of the cell median; cell text: linear value. Red squares mark the cells corresponding to Runs 6 and 7 ($C_b \in \{-1, 0, +1\}$ at $\bar{\sigma} = 15\%$). The order-statistic drift produces $v_t/v^* > 1$ in every cell.

We do not advance EDEM as an explanation for any specific real-world mispricing event. The 2021 wind-down of Zillow's algorithmic-iBuying programme [Zillow Group, Inc., 2021] is one episode in which a historical-sale-price-anchored valuation system at scale incurred substantial losses; the public record is consistent with EDEM's order-statistic prediction but is also consistent with several other explanations (timing of housing-market regime changes, operational mispricing, capital-allocation constraints), and we have not attempted to discriminate between them. The robust theoretical finding is more modest: an unbiased ML valuation trained on historical clearing prices is, by construction, fitting the upper-order statistic rather than the population mean of plausible sale prices, and at deployment time it functions as a positively-biased estimator. A defensible deployment would correct for this — for example by predicting the median rather than the mean of plausible sale prices, or by using a censored-data likelihood that accounts for the offer-acceptance threshold — rather than treating the historical record as a clean training target.

6.5 Limitations

Several modelling choices restrict the scope of the findings.

- **Single asset class.** EDEM models a single class of indistinguishable homes; portfolio-level effects (cross-asset hedging, sector rotation) are absent.
- **No leverage.** Buyers carry no debt and sellers no mortgages; financialised real-estate markets that exhibited the most dramatic anomalies of the past two decades require leverage to model accurately [Geanakoplos et al., 2012, Baptista et al., 2016].
- **No macroeconomic environment.** The supply and demand schedules in Eqs. (4) and (5) respond only to price; they do not respond to interest-rate shocks, employment, or migration.
- **Adaptive vs. static estimators.** Agents do not learn from sale outcomes. A learning extension — Bayesian belief updating, or reinforcement-learning bidding — would either converge toward bias-corrected estimation (closing the bubble channel) or amplify it (opening a stronger one), depending on how counterfactual observations are weighted.
- **Run 5 timing.** The 12,000-tick budget for the single- shock scenario is too short for the market to reach its post-shock floor before the patience boost fires. Doc 2's specific claim that patience returns the price to the new textbook equilibrium would require a longer simulation. The figure as published makes the more general slow-adjustment claim that Section 2.3 actually requires.

7 Extensions and Future Work

The framework of Section 3 is deliberately minimal so that extensions can be plugged in without rewriting the agent skeleton. Five directions appear especially fruitful and we sketch each below.

7.1 Bias-corrected clearing

Section 6.3 argued that the bid-selection asymmetry embedded in the max-bid clearing rule cannot be cancelled by any linear balancer. Two replacements for the clearing functional A in Eq. (3) would address it directly:

- **Winsorised mean**: replace $\bar{r}_t$ in Eq. (10) with a 10% winsorised mean of the yellow buyers' winning ratios. This caps the influence of extreme single-epoch outliers without changing the agent rule set.
- **Selection-aware estimator**: substitute the median or a lower-percentile statistic of all bids (not only winning bids) for $\bar{r}_t$. This estimates the underlying value distribution rather than the maximum-order statistic.

Both modifications are one-line edits to `EDEMModel._end_epoch` in our reference implementation. We expect either to break the exponential bubble of Run 6 entirely while leaving the disequilibrium regimes of Runs 2–5 qualitatively intact, since those derive from schedule dynamics rather than from clearing asymmetry.

7.2 Realtors and commission structures

The DE seller is a strict utility-maximiser whose only friction is patience. Real-world residential transactions are intermediated by realtors whose payoff is a fraction of the sale price; this ties realtor incentives to seller incentives but injects an additional clock (the listing agreement) and a different bargaining geometry (realtors negotiate on behalf of multiple sellers in parallel). A realtor extension would replace the seller agent with an aggregated listing agent whose patience and ask-price strategies optimise expected commission across a portfolio. This is a natural setting for empirical calibration: realtor commission rates and listing durations are publicly observable, unlike the patience parameter of the bare DE seller.

7.3 Open vs. blind auctions

EDEM as posed is a blind-bid model: a buyer's bid does not depend on other buyers' bids on the same home. Open-bid auctions ("best-and-final" rounds, escalation clauses) substantially change the bid distribution by exposing buyers to one another's valuations. The likely effect is to amplify the bid-selection asymmetry in Section 3.8, because each buyer can revise upward in response to seeing competitors. Conversely, a sealed-bid auction with full information disclosure post-clearing would let buyers calibrate over time and could reduce the asymmetry. Both regimes are minor modifications to the buyer agent's bid procedure.

7.4 Construction and the supply side

The construction-market variant invertes the EDEM agent roles: *contractors* bid downward to win projects, with the lowest-bid contractor winning. The selection asymmetry then operates in reverse, biasing realised contract prices below the population fair value, and the dynamics of contractor entry / exit are governed by an asymmetric balancer (a contractor exits after sustained losses, a homeowner does not exit a partially-built project lightly). This is a one-class extension that captures the historically high bankruptcy rates of small construction companies as a stable emergent feature rather than as an exogenous shock.

7.5 Adaptive agents

EDEM agents are non-learning. Two adaptive extensions are obvious:

- **Bayesian estimation.** Each agent maintains a posterior over the home's fair value and updates it after every observed sale. This converges, in the limit, toward the bias-corrected regime sketched in Section 7.1, since the agent learns the selection structure of observed prices.
- **Reinforcement-learning bidders.** Each agent is a policy mapping local state (price level, time on market, recent sales) to bids. With a buy-low/sell-high reward structure, RL agents would in principle discover the bid-selection asymmetry and exploit it; this would be a computational analogue to the Zillow-Offers iBuying problem [Zillow Group, Inc., 2021] and would offer a stress test of any proposed fix in Section 7.1.

We have implemented none of these in the present codebase but each is a self-contained extension on top of the existing `Buyer`/`Seller` classes; pull requests are welcomed at the project repository.

8 Conclusion

We have presented EDEM, an agent-based framework that models supply and demand as the realisation of a coupled stochastic process driven by heterogeneous, error-prone agent valuations. The framework recovers the classical Walrasian fixed point and a Miller-style static premium mechanism as limiting cases, and admits a family of out-of-equilibrium regimes – band-stable, business-cycle, persistent overshoot, persistent undershoot, runaway bubble, and constant transition – that are reachable from the same agent ruleset by varying three parameters: the divergence-of-opinion scale $\bar{\sigma}$, the seller patience timer, and the population balancer coefficient C_b.

The Mesa replication of Section 5 supports three substantive empirical claims. First, *textbook supply-and-demand schedules omit a state variable*: seller patience and agent density both shift the realised market price away from the equilibrium predicted by linear schedules, and in opposite directions, while preserving the schedules themselves. Second, *bubbles do not require behavioural bias*; an unbiased estimator combined with a max-bid clearing rule produces the same multiplicative drift that observers attribute to investor optimism. Third, *linear balancer interventions cannot eliminate the drift*: in the parameter range we explored, no setting of C_b restores the textbook valuation, because the asymmetry is structural to the bidding mechanism.

An immediate corollary is for valuation algorithms. A machine-learning estimator trained on historical sale prices is fitting an upper-order statistic of plausible sale prices, not their mean; deployed at scale, it functions as a positively-biased estimator that may provide the market with a coordinating signal amplifying the very bubble it was meant to price. We do not advance this as a complete account of any specific historical episode (see Section 6.4). The robust theoretical recommendation is that clearing-price-anchored valuations be replaced with selection-aware procedures (median rather than mean of plausible sale prices, or censored-data likelihoods that account for the offer-acceptance threshold) before being used as deployment-time point estimators.

EDEM's framework is intentionally general: the eight primary experiments herein use a real-estate parameterisation, but the divergence of opinion premium they explain has been documented across asset classes, including initial public offerings [Miller, 1977, Doukas et al., 2006] and cryptocurrency markets. The Mesa implementation is open-source and the experiments reproduce to bit-equivalent figures from a single `pytest`-and-script invocation; we hope this lowers the barrier to extending the framework toward leveraged markets, adaptive agents, and empirically-calibrated parameter sweeps that the present exposition leaves for future work.

References

M. Arbuzov. Disequilibria in markets with stochastic processes: An agent-based model approach. Working paper, Department of Economics, San José State University, 2018a.

M. Arbuzov. Estimated dynamic equilibrium model: Supply and demand as a sample path of a stochastic process. Best Paper, EDEM Spring 2018, San José State University, 2018b.

R. Baptista, J. D. Farmer, M. Hinterschweiger, K. Low, D. Tang, and A. Uluc. Macroprudential policy in an agent-based model of the UK housing market. *Bank of England Staff Working Paper*, (619), 2016.

J. A. Doukas, C. F. Kim, and C. Pantzalis. Divergence of opinion and equity returns. *Journal of Financial and Quantitative Analysis*, 41(3):573–606, 2006. doi: 10.1017/S0022109000002544.

L. G. Epstein and T. Wang. Intertemporal asset pricing under Knightian uncertainty. *Econometrica*, 62(2):283–322, 1994.

J. Geanakoplos, R. Axtell, J. D. Farmer, P. Howitt, B. Conlee, J. Goldstein, M. Hendrey, N. M. Palmer, and C.-Y. Yang. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review*, 102(3):53–58, 2012.

V. Grimm, S. F. Railsback, C. E. Vincenot, U. Berger, C. Gallagher, D. L. DeAngelis, B. Edmonds, J. Ge, J. Giske, J. Groeneveld, A. S. A. Johnston, A. Milles, J. Nabe-Nielsen, J. G. Polhill, V. Radchuk, M.-S. Rohwäder, R. A. Stillman, J. C. Thélin, and S. Berger. The ODD protocol for describing agent-based and other simulation models: A second update to improve clarity, replication, and structural realism. *Journal of Artificial Societies and Social Simulation*, 23(2):7, 2020. doi: 10.18564/jasss.4259.

D. Masad and J. Kazil. Mesa: An agent-based modeling framework. In *Proceedings of the 14th Python in Science Conference (SciPy 2015)*, pages 51–58, 2015.

J. Mayshar. On divergence of opinion and imperfections in capital markets. *The American Economic Review*, 73(1):114–128, 1983.

R. C. Merton. A simple model of capital market equilibrium with incomplete information. *The Journal of Finance*, 42(3):483–509, 1987.

E. M. Miller. Risk, uncertainty, and divergence of opinion. *The Journal of Finance*, 32(4):1151–1168, 1977. doi: 10.1111/j.1540-6261.1977.tb03317.x.

S. D. Moffitt. *The Strategic Analysis of Financial Markets, Volume I: Framework*. World Scientific, 2017. Earlier circulated as "Why Markets are Inefficient: A Gambling Theory of Financial Markets For Practitioners and Theorists," Feb. 22, 2017.

R. J. Shiller. From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*, 17(1):83–104, 2003.

H. R. Varian. Divergence of opinion in complete markets: A note. *The Journal of Finance*, 40(1): 309–317, 1985.

U. Wilensky. NetLogo. http://ccl.northwestern.edu/netlogo/, 1999. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.

J. T. Williams. Capital asset prices with heterogeneous beliefs. *Journal of Financial Economics*, 5(2): 219–239, 1977.

Zillow Group, Inc. Zillow group reports third quarter 2021 financial results, announces wind down of Zillow Offers. Press release, Nov. 2021. URL https://investors.zillowgroup.com/investors/news-and-events/news/news-details/2021/Zillow-Group-Reports-Third-Quarter-2021-Financial-Results-Announces-Wind-Down-of-Zillow-Offers default.aspx.

A ODD Protocol

This appendix documents EDEM in the standard ODD protocol for agent-based models [Grimm et al., 2020].

A.1 Purpose and patterns

EDEM's purpose is to characterise the conditions under which an agent-based real-estate market reaches a stable equilibrium price, and to identify the mechanisms that prevent equilibrium when they do not. Patterns the model is designed to reproduce: stable band-bounded equilibria; endogenous business cycles; persistent shifts in the realised price relative to the textbook equilibrium under altered patience or density; multiplicative price drift in the absence of a balancer.

A.2 Entities, state variables, and scales

Entities. **Buyers** and **Sellers** (`mesa.Agent` subclasses) and **Homes** (per-cell dataclass instances on a 32×32 toroidal grid).

State variables.

- Buyer: epsilon (estimation-error bound), delay, current bid, heading. EDEM variant additionally: lowest-bid-to-value ratio, yellow flag.
- Seller: epsilon, patience timer, current ask price, dictionary of received bids. EDEM additionally: number of bids, total bids, best-bid-to-value, best-bid-to-true-value.
- Home: market price, last-sale tick, last-sale price, fair value v^*, current value v.
- Model: rolling window of last 25 sale prices (DE), epoch counter (EDEM), current cross-population epsilon (EDEM).

Spatial and temporal scales. One spatial unit = one home; one tick is the smallest temporal unit. Each tick triggers one round of buyer movement and one round of seller bid-processing. DE balance period is 100 ticks; EDEM epoch is 20 ticks.

A.3 Process overview and scheduling

Per tick, in order:

1. Each agent steps once (in randomised order over both classes). A Buyer's step posts bids, then wiggles, then steps forward. A Seller's step processes patience and best-bid logic.
2. The model invokes the balancer: DE recomputes linear targets every 100 ticks; EDEM fires the per-epoch update every 20 ticks (using the cycle-counter trick of Section 4.2 to preserve buyer `is_yellow` flags).
3. The data collector samples model-level reporters.

A.4 Design concepts

Basic principles. EDEM extends the dynamic divergence-of-opinion premise of Miller [1977] to a multi-period setting; agents act on noisy estimates of an unobservable fair value.

Emergence. The price level, agent population dynamics, and macroscopic regimes (band-stable, business-cycle, persistent shift, bubble, transitional) emerge from the local interaction of agent estimation, bidding, and patience timers; none of these phenomena are coded in directly.

Adaptation. Sellers lower ask prices when their patience timer elapses without high-enough bids; this is the only adaptive behaviour. Buyers do not adapt: they place bids drawn from a fixed distribution and accept by Cond. 2 (Eq. (6)).

Objectives. Sellers maximise realised sale price subject to patience. Buyers commit only to offers whose realised bid is at or above their own running benchmark over their outstanding bids (Eq. (6)) — a selection rule, not a price-minimisation objective. Neither agent class has an explicit utility function; objectives are implicit in their action rules.

Learning. None in the baseline model (Section 7.5 sketches an adaptive extension).

Prediction. None.

Sensing. Buyers sense whether a seller is on their patch; sellers sense the bids in their own bid table. Neither senses the global market price directly.

Interaction. Direct: a buyer that lands on a seller's patch posts a bid via a mirrored-dictionary update on both sides. Indirect: the average ask price (DE) and the rolling sale-price average (DE) influence newly-spawned agents' starting prices.

Stochasticity. Used for: agent placement at setup; buyer heading initialisation and per-tick wiggle; epsilon draws per agent; per-bid epsilon-error draws; victim selection in the balancer; Bernoulli draws for fractional C_b swaps in EDEM. All draws come from a single seeded `numpy.random.Generator` owned by the model.

Collectives. None.

Observation. The data collector records, per tick: market price (DE: rolling-25 mean; EDEM: average home value), agent counts, the equilibrium price implied by current schedules, the rolling-window fill count, and (EDEM) the current epsilon. Output is stacked into Parquet datasets, one per seed, for downstream analysis.

A.5 Initialisation

DE: spawn `equi_qnty` sellers at random unique cells with patience $\sim \mathcal{U}[0, \texttt{init\_patience})$ and ask price drawn from $\mathcal{U}[-\bar{\sigma}, +\bar{\sigma}]$ around p^*; spawn `equi_qnty` buyers at random cells. EDEM: 20 sellers and 20 buyers, all homes initialised with $v_0(h) = v^*(h) = 100$.

A.6 Input data

None. EDEM is a closed model; all dynamics are endogenous. The shock hooks in Run 5 are scheduled in the experiment script, not loaded from external data.

A.7 Submodels

Estimation function. Eq. (1) with $\varepsilon \sim \mathcal{U}[-\sigma_i, +\sigma_i]$ and $\sigma_i \sim \mathcal{U}[0, \bar{\sigma}]$.

Bid acceptance. Eq. (6).

Market price. DE: Eq. (9), the rolling mean of the last $W = 25$ sale prices. EDEM: Eq. (10), the per-epoch multiplicative update.

Balancer. DE: linear-target restoration every T_B ticks. EDEM: Eqs. (4) and (5) in finite-population form, with fractional C_b realised as integer + Bernoulli swaps and a population floor of one agent per side.

Agent spawn / exit. On a sale, both counterparties leave the market and a fresh pair is spawned by the balancer with patience drawn from $\mathcal{U}[50, \texttt{init\_patience})$ for sellers and immediate readiness for buyers.

B Parameter Tables for All Runs

The Mesa source files in `python_simulation/experiments/` are the authoritative parameter manifest; the tables below are reproduced verbatim from the model-construction kwargs.

B.1 Common parameters across all runs

Parameter	Value
World size	32×32 toroidal
Seeds per run	≥ 8 (typically 10)
Bid-acceptance rule	`netlogo` (Eq. (6))

B.2 Dynamic Equilibrium runs (1–5)

Run 5 introduces shocks via the model's `set_demand` and patience-rebinding hooks. Scenario A schedules a single demand shock at $t = 3000$ followed by a patience boost at $t = 7000$; Scenario B toggles the demand intercept between 125 and 75 every 2000 ticks.

Parameter	Run 1	Run 2	Run 3	Run 4	Run 5
Supply intercept a_s	0	0	0	0	0
Supply slope b_s	0.5	0.5	0.5	0.5	0.5
Demand intercept a_d	100	100	100	50	$100 \to 50$ (A) $\{125, 75\}$ (B)
Demand slope b_d	-0.5	-0.5	-0.5	-0.5	-0.5
Max valuation error $\bar{\sigma}$ (%)	5	25	5	5	5
Max patience	50	50	100	50	$50 \to 165$ (A)
Balance period T_B	100	100	100	100	100
Sale window W	25	25	25	25	25
Ticks per seed	20,000	20,000	20,000	20,000	12,000
Equilibrium price p^*	100	100	100	50	100 (initial)
Equilibrium quantity q^*	50	50	50	25	50 (initial)

B.3 Speculative-market EDEM runs (6–8)

Parameter	Run 6	Run 7	Run 8
Initial $\bar{\sigma}$ (%)	15	15	5
$\bar{\sigma}$ growth / epoch	0	0	$+0.5$ pp
Balancer coefficient C_b	0	$\{+1, 0, -1\}$	-1
Initial buyers / sellers	20 / 20	20 / 20	20 / 20
Epoch length (`init_patience`)	20	20	20
Ticks per seed	3,000	3,000	3,000
True value $v^*(h)$ (uniform across h)	100	100	100

B.4 Reproduction

From a clean checkout:

```
cd python_simulation && pip install -e ".[dev]"
pytest                              # 42 unit tests
for f in experiments/run*.py; do
    python "$f"
done                                # writes paper/figures/
cd ../paper && bash build.sh        # builds main.pdf
```

End-to-end reproduction takes approximately twenty minutes on a 2024-era laptop.