



Minimal Non-Equilibrium Models of Socio-Economic Systems: From the French Housing Market to the US Labor Market

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique

École doctorale n°626 École doctorale de l'Institut Polytechnique de Paris (EDIPP)
Spécialité de doctorat: Mécanique des fluides et des solides, acoustique

Thèse présentée et soutenue à Palaiseau, le 23 septembre 2025, par

ANTOINE-CYRUS BECHARAT

Composition du Jury :

Marc Barthelemy CEA & EHESS	Président
Alberto Rosso CNRS & Université Paris-Saclay	Rapporteur
Elsa Arcaute University College London	Rapporteur
Pierpaolo Vivo King's College London	Examineur
Michael Benzaquen CNRS & École Polytechnique	Directeur de thèse
Jean-Philippe Bouchaud Académie des Sciences & CFM	Co-directeur de thèse

A Eve.

Foreword and acknowledgements

This journey began long ago, though I never imagined I would one day be writing the acknowledgments for a PhD in statistical physics.

For as long as I can remember, both sides of my family have spun stories of École Polytechnique. Drawn to mathematics and physics—and following the path of my father and great-grandfather—I entered the school in 2018 after two formative years of preparatory classes at Lycée Louis-le-Grand.

It still took time at Polytechnique to discover where my interests truly lay. I gravitated toward courses in mathematics and theoretical physics, which eventually led me to Jean-Philippe Bouchaud's M2 course at École Normale Supérieure, *From Statistical Physics to Social Sciences*. That experience set the stage for a PhD in statistical physics at École Polytechnique under the guidance of Michael Benzaquen and Jean-Philippe Bouchaud.

The same academic year concluded with a research internship within the Complexity Economics group at the Institute for New Economic Thinking in Oxford, working with Doyne Farmer and José Moran on an extension of R. Maria del Rio-Chanona's model—collaborations that continued throughout my three-year doctoral journey.

First and foremost, I express my deepest gratitude to Michael Benzaquen and Jean-Philippe Bouchaud. However tentative my ideas, they always found time to listen, critique, and guide me toward a finished project. Collaborating with Jean-Philippe—and glimpsing his remarkable intellect at work—has been a privilege I will never deserve.

I am likewise profoundly indebted to J.Doyne Farmer, whose penetrating insight and intellectual generosity illuminated every discussion we shared on labor-market dynamics. Those conversations were at once inspiring and instructive; his great mind taught me more than words can convey.

I am grateful to my defense jury—Alberto Rosso, Elsa Arcaute, Marc Barthélémy, and Pierpaolo Vivo—for generously agreeing to be part of it, and especially to Alberto Rosso and Elsa Arcaute for acting as thesis referees. I also thank Marc Barthélémy and Jean-Pierre Nadal, whose participation in my *comité de suivi* over these three years provided invaluable counsel and guidance.

My heartfelt thanks also go to my co-authors José Moran, Ruben Zakine, Jérôme Garnier-Brun, and R.Maria del Rio-Chanona for their insight, patience, and support. Working alongside such gifted researchers was a genuine pleasure.

A special word for Guillaume Maitrier, who accompanied me from our section at La

Courtine during military training to a desk barely a meter away for three years. His energy and humor lightened many a difficult day. I am also grateful to Swann Chelly, whose companionship carried me from the “section football” of École Polytechnique all the way to the doctoral chair.

I have been fortunate, too, to share daily life with the entire Econophysics Lab—Jérôme, Ruben, Samy, Cecilia, Nirbhay, Salma, Max, Pierre, Elia, Fabian, Victor, Thomas, Jutta, and Natascha—whose lunchtime debates, group meetings, and retreats were both interesting and joyful, the best mix one could hope for in a research group.

Finally, I am deeply grateful to the Chair of Econophysics & Complex Systems, the Laboratoire d’Hydrodynamique, and École Polytechnique for their steadfast support and for granting me the freedom to follow my curiosity wherever it led. I extend the same gratitude to Capital Fund Management for hosting me and offering an equally vibrant and stimulating intellectual environment. I am likewise indebted to the Institute for New Economic Thinking at the University of Oxford for its warm hospitality during my master’s research internship and throughout my three doctoral years.

None of this would have been possible without the unwavering support of my family. To my parents, who have always supported and encouraged me in every endeavor; to my brother and sister, whose literary tastes I hoped to honor with the epigraphs scattered through these pages; and to my wider family, whose love stretches from the splendor of Isfahan to the wilds of La Jasse de Bernard beneath the Cévennes—your encouragement and love have been greatly appreciated throughout this journey.

My deepest gratitude goes to Eve Gadou, who unfailingly stood by my side—supporting me even through my grumpier moments. Her unflagging optimism and joy are a constant inspiration.

To my friends—too numerous to name—your companionship has revived my spirits whenever they faltered.

I am profoundly grateful to you all.

Publications and collaborations

- ▶ Zakine et al. (2024) [1] R. Zakine, J. Garnier-Brun, **A.-C. Becharat** and M. Benzaquen. *Socioeconomic agents as active matter in nonequilibrium Sakoda-Schelling models*. Physical Review E 109.4 (2024), p. 044310.
- ▶ Becharat, Benzaquen, and Bouchaud (2025) [2] **A.-C. Becharat**, M. Benzaquen, J.-P. Bouchaud. *Diffusive Nature of Housing Prices*. Physical Review Letters 135 (2025), p. 107401
- ▶ **A.-C. Becharat**, J. Moran, R.M. del Rio-Chanona and J.D. Farmer. *An out-of-equilibrium agent-based model of the labor market*.
- ▶ **A.-C. Becharat**, J. Moran, R.M. del Rio-Chanona and J.D. Farmer. *A simple yet effective disequilibrium model of the Beveridge curve*. In prep.

Résumé substantiel en Français

Cette thèse explore la dynamique, souvent non-intuitive, des systèmes socio-économiques en combinant la rigueur formelle de la physique statistique avec la souplesse des modèles multi-agents. L'idée directrice est simple : des interactions locales — échanges d'information de voisinage, ajustements de prix ou règles décisionnelles limitées — suffisent à produire des régularités globales et des transitions qualitatives que les approches économiques classiques peinent à anticiper. En articulant analyses mathématiques, simulations à grande échelle et confrontation systématique aux données, nous montrons comment des mécanismes micro-fondés, élémentaires mais réalistes, génèrent des phénomènes agrégés persistants, structurés et parfois non efficients.

La première partie revisite le modèle de ségrégation de Sakoda-Schelling sur un réseau. Nous considérons des agents qui évaluent localement une densité perçue et ajustent leurs choix d'emplacement selon une fonction d'utilité non linéaire. À partir de ce comportement microscopique, nous menons une analyse de stabilité linéaire, réalisons des simulations numériques et introduisons des équations hydrodynamiques stochastiques. Ce passage du microscopique au macroscopique éclaire la nature hors d'équilibre du système et l'inscrit dans le cadre plus général de la matière active, où des entités auto-propulsées engendrent des motifs collectifs durables. Nous montrons qu'une transition de phase apparaît pour une large gamme de paramètres, et qu'elle appartient à la classe d'universalité du modèle d'Ising en deux dimensions. Autrement dit, malgré des règles d'interaction propres au système étudié, les exposants critiques et les lois d'échelle observés coïncident avec ceux d'un modèle canonique de la physique, signe que la structure des interactions l'emporte sur les détails idiosyncratiques. L'identification d'un tel comportement critique robuste fournit un langage commun pour comparer, étudier et anticiper l'évolution de ce système complexe.

La seconde étude de cette première partie porte sur le marché immobilier français. Nous proposons une équation de diffusion stochastique, micro-fondée de manière heuristique à travers un mécanisme d'ajustement local des prix, afin de décrire un tel système : nous postulons que les changements se propagent entre communes voisines à la manière d'une « chaleur » sociale, plutôt que de s'imposer instantanément et uniformément. Calibrée sur environ un demi-siècle de données (1970–2022), cette équation reproduit des régularités empiriques clés : variogrammes spatiaux à croissance logarithmique, autocorrélations longue portée et persistance pluriannuelle des chocs économiques. Une analyse par réponse impulsionnelle fournit une valeur plausible pour le coefficient de diffusion et met en évidence la forte hétérogénéité

spatiale de l'amplitude des chocs. Ces résultats confirment des conjectures anciennes sur la nature diffusive des prix immobiliers, tout en soulignant le caractère non efficient de ce marché : l'information locale ne se propage pas instantanément, et des décalages spatiaux et temporels substantiels subsistent, ouvrant la voie à des politiques d'aménagement ou de régulation différenciées selon les territoires.

La deuxième partie de cette thèse se concentre sur le marché du travail via un modèle d'agents zéro-intelligence. Les demandeurs d'emploi postulent aléatoirement à des offres compatibles avec leurs compétences, tandis que les entreprises adaptent heuristiquement le nombre et le type d'annonces. Malgré cette simplicité, les simulations reproduisent des traits stylisés robustes : délais de coordination entre postes vacants et candidats, asymétrie entre embauches et licenciements, et mise en relation non instantanée des demandeurs d'emploi avec les offres. À partir de ces dynamiques discrètes, nous introduisons rigoureusement des équations différentielles continues décrivant l'évolution conjointe des stocks (demandeurs d'emploi, postes vacants) et des flux (embauches, séparations). Ce passage du microscopique au macroscopique explicite les déphasages mécaniques du processus d'appariement et montre qu'ils suffisent à générer l'hystérésis observée sur la courbe de Beveridge, sans recourir à des hypothèses simplificatrices du type anticipation rationnelle ou optimisation globale par les agents. Autrement dit, le caractère cyclique de la courbe de Beveridge provient d'engorgements logistiques endogènes plutôt que de chocs exogènes permanents.

Enfin, nous extrayons de ce cadre un modèle dynamique minimal hors d'équilibre, calibré sur les données américaines de 1951 à 2023. En prenant comme entrée une trajectoire exogène du taux de postes vacants, le modèle reconstitue fidèlement la trajectoire du taux de chômage, en rendant compte du déphasage observé et de l'amplitude des réponses aux chocs économiques. Sa forme analytique simple en fait un outil parcimonieux permettant d'explorer des scénarios contrefactuels (variations de coûts d'embauche, dispositifs d'incitation, chocs de productivité) et d'identifier les horizons temporels sur lesquels les interventions sont les plus efficaces compte tenu des frictions d'appariement.

L'ensemble de ces résultats permet de mettre en relief une conclusion transversale : dans des systèmes socio-économiques complexes, la localité des interactions et la non-instantanéité des ajustements suffisent à produire des comportements collectifs non-triviaux, des transitions de phase et une mémoire à long terme, que l'on peut caractériser et prévoir grâce aux outils de la physique statistique. Ainsi, cette thèse offre un pont opérationnel entre modèles micro-fondés simples et phénomènes macro-émergents, et propose des cadres analytiques parcimonieux utilisables pour éclairer des décisions de régulation, tout en représentant fidèlement la richesse empirique des données et en limitant le nombre d'hypothèses simplificatrices de modélisation.

Foreword and acknowledgements	iii
Publications and collaborations	v
Résumé substantiel en Français	vi

MOTIVATION AND THEORETICAL BACKGROUND 1

1. Introduction	2
1.1. From statistical physics to agent-based models of socio-economic systems	5
Utility theory and decision rules	5
The curse of dimensionality	7
What is a good model?	10
The Sakoda-Schelling model	11
1.2. Classical macroeconomics frameworks and the need for agent-based models	13
On bounded rationality	14
A very famous macroeconomics feature: The Beveridge curve	15
The Diamond-Mortensen-Pissarides framework	16
Other perspectives in agent-based models	20
1.3. This thesis	20
2. Theoretical Foundations	22
2.1. The Langevin equation	22
Physical origin	22
Free particle solution and the Ornstein-Uhlenbeck process	23
The overdamped (Smoluchowski) limit	23
Diffusion and mobility	24
Generalization to higher dimensions	24
2.2. From the Master equation to the diffusion equation	25
Symmetric notation and conservation	25
Continuum limit	25
Kramers-Moyal expansion	26
2.3. Variograms and correlograms: tools for characterizing fluctuations	27
The variogram	27
From variogram to correlogram	28
2.4. From a stochastic partial differential equation to the correlogram	28

PART I: HYDRODYNAMIC MODELING OF SOCIO-ECONOMIC SYSTEMS: SEGREGATION DYNAMICS AND HOUSING-PRICE DIFFUSION 32

3. An out-of-equilibrium Sakoda-Schelling model	33
3.1. Introduction	33
3.2. A Sakoda-Schelling occupation model	35
Setup	35
In or out of equilibrium?	36
Microscopic simulations	37
3.3. Field theory and the local-move approximation	41
General description	41
Linear stability analysis	42
Comparison to microscopic simulations	43
Local-move approximation	44

3.4.	Generalized-thermodynamic construction	45
	The generalized-thermodynamic expansion	45
3.5.	Further socioeconomic considerations	47
	Two populations	47
	Housing market	49
3.6.	Discussion	50
4.	The diffusive nature of housing prices	54
4.1.	Introduction	54
4.2.	Model description	56
	Introducing a stochastic differential equation	56
	Computing the spatial variogram	57
4.3.	Comparison with empirical data	60
	Static analysis	60
	Temporal analysis	61
4.4.	Extension: coupling population and prices	63
	Model setting	63
	Stability analysis	64
4.5.	Conclusion	65
PART II: FROM AGENT-BASED MODELS TO MINIMAL DISEQUILIBRIUM MODELS		68
5.	An out-of-equilibrium agent-based model of the labor market	69
5.1.	Introduction	69
5.2.	Literature review	69
	The Diamond-Mortensen-Pissarides framework	70
	Agent-based models of labor markets	73
5.3.	The dynamics of the Beveridge curve	79
	Convergence to the steady state	79
	Out-of-equilibrium dynamics	80
5.4.	A generic dynamical model	82
5.5.	Empirical results	85
5.6.	Conclusion	92
6.	A simple yet effective disequilibrium model of the Beveridge curve	95
6.1.	Introduction	95
6.2.	Our conceptual framework	97
	Behavior of profit-maximizing firms	97
	A quantitative model for vacancies and unemployment	99
	Qualitative explanation of the Beveridge curve cycle	100
6.3.	Empirical tests of the predictions of the model	101
	Qualitative tests of predicted lag and tail behavior	101
	Time series model	104
	Comparison to the DMP model	104
	Comparison of matching functions	106
	Evidence for bounded rationality	107
	Measuring matching rates and adjusting the model to make better predictions	107
6.4.	Conclusion	109

CONCLUSIONS	112
7. Perspectives & Conclusions	113
Bibliography	116
 APPENDIX	 129
A. Part I: An out-of-equilibrium Sakoda-Schelling model	130
A.1. Coarsening exponent	130
A.2. Monte Carlo simulations for 2nd-order phase transitions	131
A.3. Local mean field description and LSA	131
A.4. Local versus non-local PDEs	132
A.5. Linear stability for two coupled populations	133
A.6. LSA for two populations with local moves	134
A.7. Introducing a non-linear price field	135
 B. Part I: The diffusive nature of housing prices	 136
B.1. Analytical derivation of the diffusive term	136
B.2. Computation of the temporal variogram	137
B.3. Spatial variogram in the ballistic case	140
B.4. Additional Plots	142
B.5. Estimating the diffusion constant D	145
B.6. Deriving the coupled linear stability criterion	145
 C. Part II: A zero intelligence agent-based model of the labor market	 147
C.1. Additional Plots	147
C.2. Stationary point and equilibrium in the ABM	149
C.3. Additional information on time scales in the ABM	150
C.4. Asymmetric behavior	150
 D. Part II: A simple yet effective disequilibrium model of the Beveridge curve	 152
D.1. Data analysis	152
D.2. Fitting procedure and error evaluation for unemployment	152
D.3. Additional plots	155

MOTIVATION AND THEORETICAL BACKGROUND

Introduction

1.

*When the Light of Lights appeared, they understood:
the Simorgh was themselves, and they the Simorgh.
The Conference of the Birds- Farid al-Din Attar, 1177.*

Applying concepts from statistical physics to socio-economic and financial systems is a long-standing idea. One of the earliest and most famous examples is Brownian motion. In 1827, the Scottish botanist Robert Brown observed that pollen grains suspended in water undergo erratic movement. Although each grain's motion is governed by deterministic Newtonian laws, its trajectory appears random—an observation that extends to many physical and socio-economic systems.

Inspired by this, Louis Bachelier—in his 1900 thesis *Théorie de la spéculation* [3], written under the supervision of Henri Poincaré—modeled stock-market prices using a stochastic differential equation, paving the way to the wide use of such equations to describe the subtle interplay between deterministic laws and randomness that rule many real-world systems:¹

«Les influences qui déterminent les mouvements de la Bourse sont innombrables, des événements passés, actuels ou même es-comptables, ne présentant souvent aucun rapport apparent avec ses variations, se répercutent sur son cours. À côté des causes en quelque sorte naturelles des variations, interviennent aussi des causes factices : la Bourse agit sur elle-même et le mouvement actuel est fonction, non seulement des mouvements antérieurs, mais aussi de la position de place. La détermination de ces mouvements se subordonne à un nombre infini de facteurs : il est dès lors impossible d'en espérer la prévision mathématique. [...] Mais il est possible d'étudier mathématiquement l'état statique du marché à un instant donné—c'est-à-dire d'établir la loi de probabilité des variations de cours qu'admet à cet instant le marché.»²

Bachelier's work introduced Brownian motion in finance five years before Einstein formalized it in his landmark paper [4], making it one of the first applications of statistical-physics methods to an economic system.³

Using physical tools to describe socio-economic systems with many interacting parts—an economy or urban population, for example—is natural. Just as fluids can be modeled by continuous fields in hydrodynamics [5], one can describe human mobility via a density field and hydrodynamic equations. This macroscopic approach discards microscopic details, making it suitable for systems with too many agents to track individually.

Moreover, as Philip Anderson emphasized in “More Is Different” [6], physics is fundamentally concerned with symmetry—“the existence of

1: Equations of this sort have long been applied to classic physical processes—ranging from turbulent fluid flows to the growth of active-matter interfaces—and, as we will demonstrate in this manuscript, they also provide a powerful framework for modeling socio-economic dynamics, including both housing and labor markets.

2: “Countless factors influence stock-market movements: past, present, or even anticipated events—often with no obvious connection to price changes—feed into market behavior. Alongside these “natural” causes of fluctuation, there are also artificial ones: the market reacts to itself, so current movements depend not only on earlier trends but also on prevailing market positioning. Because an infinite array of factors shapes these movements, mathematical prediction is impossible. [...] However, it is still possible to study the market's static state at a given moment—that is, to determine the probability distribution of price changes that the market admits at that instant.”

3: In his 1905 “*annus mirabilis*”, Einstein's paper on Brownian motion explained the seemingly erratic motion of pollen grains as the result of innumerable collisions with water molecules, hence formalizing the existence of atoms and molecules at a time where this was still debated. That same year, he published his theory of special relativity, explained the photoelectric effect—earning him the 1921 Nobel Prize in Physics—and formulated the famous mass-energy equivalence. Together, these breakthroughs transformed physics and laid the foundation for modern science.

different viewpoints from which the system appears the same." Yet, when large numbers of particles interact, symmetry is often broken: the macroscopic state exhibits lower symmetry than the underlying microscopic laws, so knowledge of the latter alone cannot predict key large-scale behaviors. As in Anderson's words, "*the whole becomes not only more than but very different from the sum of its parts,*" a reality exemplified by phenomena such as superconductivity or antiferromagnetism.

In the same vein, complex systems composed of simple interacting units can give rise to emergent properties that defy deduction from individual-level rules. A striking illustration is the human brain: billions of neurons each transmit basic electrical signals, yet their collective interactions produce consciousness, creativity, and abstract thought. This principle of emergence is equally pertinent to socio-economic systems, where aggregate dynamics cannot be inferred solely from the behavior of individual agents.

This stands in stark contrast to the *representative and rational agent* paradigm in economics, which reduces many interacting agents to a single perfectly rational individual—endowed with infinite computational power and always choosing the globally optimal action. However, removing heterogeneities and interactions in such models denies the possibility of destabilizing feedback loops, making these models unrealistic to explain the numerous crises observed in economic history, as eloquently presented in [7]. This led many researchers to think about a new way of describing large collections of interacting agents, as it is relevant to understand socio-economic systems.

Indeed, early work by Blumer [8] showed that social phenomena arise from collective behavior rather than isolated individuals. Sumpter [9] demonstrated that structures in animal groups, from termite mounds to human crowds, can emerge from simple local rules rather than individual complexity. These insights motivated the development of agent-based models (ABM) in socio-economics, where one studies the aggregate behavior of numerous heterogeneous and interacting agents. Epstein and Axtell's seminal book [10] presented ABMs as a tool for studying how individual behaviors produce system-level patterns, especially in out-of-equilibrium settings. This approach abandons homogeneous and perfectly rational agents for heterogeneous agents with *bounded rationality*⁴ and limited information, challenging classical equilibrium models and introducing a realistic framework for capturing the complex, emergent dynamics of socio-economic systems, as advocated in Axtell and Farmer's impressive review [12].

Finally, economists have long debated whether markets—composed of self-interested agents—naturally converge to an optimal state. Adam Smith's "invisible hand"⁵ posits that decentralized decision-making guides the economy toward collective welfare [13]. Yet the agent-based segregation models introduced by Sakoda and Schelling [14, 15]—in which households freely decide whether to stay in or leave their current location to maximize a simple "satisfaction" score—illustrate how elementary local preferences can generate unexpected large-scale outcomes. In certain parameter ranges the entire system becomes locked in a configuration where no agent is satisfied,

4: The concept of bounded rationality, introduced by Herbert Simon in the 1950s [11], will be examined in greater detail later in this introducing chapter.

5: The invisible hand metaphor encapsulates the notion that unseen forces guide the economy toward a globally optimal state, achieved through the actions of individuals who, paradoxically, seek only to maximize their personal satisfaction.

exposing the limits of Adam Smith’s “invisible hand” even in this rudimentary coordination problem. Similar failures of agents coordination are discussed for other settings in [7].

Hence, the aim of this thesis is twofold. First, we show how stochastic differential equations, formulated directly at an aggregate level, can be employed to capture and reproduce key macro-economic regularities. Second, we demonstrate how agent-based models, which start from explicit micro-level rules, can likewise be coarse-grained into effective equations that govern the evolution of aggregate variables, thus providing an alternative—and complementary—route to understanding socio-economic dynamics.

Central to this work—and indeed to the broader ambition of applying statistical physics to socio-economic systems—is the challenge of deriving macroscopic phenomena from microscopic rules, as eloquently framed in [16]. This mirrors Thomas Schelling’s insight that individual *micromotives* collectively generate *macrobehaviors* in socio-economic contexts [17].

This introductory chapter is structured as follows. We begin by motivating the use and issues of stochastic differential equations (SDEs) and agent-based models (ABMs) to represent socio-economic systems. We then introduce the well-known Sakoda-Schelling model, which will be examined in greater depth later in the thesis. Finally, we present the influential Diamond-Mortensen-Pissarides [18, 19, 20] framework for modeling the labor market. This model forms the foundation of the second part of the thesis, where we propose aggregate differential equations, built upon an agent-based model of the labor market, as extensions that go beyond the classical formulation.

1.1. From statistical physics to agent-based models of socio-economic systems

As noted earlier, an agent-based model comprises a population of heterogeneous, interacting entities—“agents”—that may represent households, firms, or other decision-makers. Each agent follows simple micro-level rules—typically aimed at improving a proxy of individual satisfaction, the choice of which we discuss later in this chapter—while operating within an explicit socio-economic environment. Crucially, aggregate behavior is *not* specified a priori: it *emerges* from the micro-level interactions. As emphasized in [12], ABMs are constructed *bottom-up*: any coherent macro-level outcome obtained from a simulation is produced endogenously through agent interactions, thereby lending this modeling approach a greater degree of realism. Depending on parameter values, the system may converge to a Nash equilibrium⁶ or remain in perpetual motion as agents continuously adjust to one another.

Moreover, the increasing availability of computational power allows ABMs to dispense with many ad-hoc simplifications historically imposed for the sake of analytical tractability. In particular, they can relax assumptions such as instantaneous market clearing⁷ and perfect efficiency—cornerstones of conventional economic models—thereby offering a more flexible framework for studying economies that are frequently out of equilibrium.

Utility theory and decision rules

We can now draw a direct parallel between statistical physics and socio-economic systems by treating economic or social agents as akin to interacting particles. Just as microscopic forces between atoms give rise to observable macroscopic phenomena in physics, the local decisions and interactions of individual agents can collectively generate large-scale patterns in markets and societies. Many of the effects familiar in physical systems—such as phase-transitions—have clear analogues in economics⁸. This perspective thus equips us with a rich toolbox for analyzing endogenous crises and other emergent dynamics. Before turning to these complex, system-wide questions, however, we must first establish a firm understanding of how individual agents behave and how to formulate an appropriate modeling strategy.

This naturally leads us to *utility* theory. How, in practice, can one represent an agent’s satisfaction, that will be the basis of his decision making and hence the evolution of the system? The standard approach is to introduce a *utility function* $u(X)$, where X denotes a state of the system accessible to the agent, and $u(X)$ serves as a numerical proxy for their level of satisfaction when being in the said state [22]. In analogy with physics—where systems tend to minimize a global energy function—agents are then assumed to seek the state X that *maximizes* $u(X)$.

It is important to emphasize, however, that in this framework each agent’s utility depends not only on their own choices but also on

6: This cornerstone of game theory characterizes an equilibrium as a state in which each agent’s outcome is optimized in light of the choices made by the others; it was first proposed by Antoine-Augustin Cournot in his 1838 *Recherches sur les Principes Mathématiques de la Théorie des Richesses* [21] and later extended by the famous economist John Nash.

7: Market clearing occurs when supply equals demand for any good or service—for example, when the number of job vacancies matches the number of unemployed workers in a labor market.

8: One may interpret densely populated neighborhoods in the Sakoda-Schelling model as analogous to *liquid* phases, while more sparsely populated areas resemble *gaseous* phases. This analogy allows us to investigate how variations in model parameters may trigger phase transitions between these states. Similarly, endogenous economic crises can be likened to natural shocks such as earthquakes, offering a metaphorical framework to explore abrupt, systemic disruptions.

those of other agents. This stands in stark contrast to the representative-agent paradigm, where one assumes a single, isolated decision-maker. Once we posit that agents aim to maximize $u(X)$, the pressing question becomes: *how* do they make their decisions? Under the assumption of perfect rationality, agents would simply pick the strict maximizer of u . Yet this assumption has been the subject of extensive debate, and has led to the concept of *bounded rationality* [11], wherein agents do not possess complete information and infinite computing power, but rather decide based on “*the body of dynamic, idiosyncratic information that constitutes their knowledge*”, in the words of Axtell and Farmer [12].

A convenient way to capture boundedly rational decision-making is via stochastic choice rules. The most widely used is the so-called *logit rule* [23, 24], which specifies that an agent selects alternative α with probability

$$P_\alpha = \frac{\exp(\beta u_\alpha)}{\sum_{\gamma \in \Lambda} \exp(\beta u_\gamma)}. \quad (1.1)$$

Here u_α is the utility associated with choice α , Λ is the full set of available options, and β is the *intensity of choice* parameter that quantifies the degree of rationality. In fact, β serves a role analogous to the inverse temperature $1/T$ in classical thermodynamics, while u corresponds to the negative of the system’s energy. This correspondence renders the logit choice rule mathematically similar to the well-known Gibbs-Boltzmann distribution in statistical mechanics.

In the limit $\beta \rightarrow 0$, we have $\exp(\beta u_\alpha) \approx 1$ for all α , so

$$P_\alpha = \frac{1}{|\Lambda|}, \quad (1.2)$$

where $|\Lambda|$ is the cardinal of Λ , i.e. each choice is equally likely, reflecting purely random behavior. Conversely, as $\beta \rightarrow +\infty$, the probability concentrates entirely on the utility maximizer:

$$P_\alpha = \begin{cases} 1, & u_\alpha = \max_{\gamma \in \Lambda} u_\gamma, \\ 0, & \text{otherwise.} \end{cases} \quad (1.3)$$

In this regime, agents again behave as strict utility maximizers.⁹ Justifications for the logit rules are discussed in depth in [7], and we will come back to this in chapter 3. A comprehensive review of the key concepts underlying utility functions and the logit rule is also provided in [25].

Lastly, it is important to introduce a particularly extreme form of bounded rationality known as *zero-intelligence* agents [26]. As explained by Axtell and Farmer [12], these agents do not engage in forward-looking optimization or strategic planning; instead, they simply respond to the most recent signals in their environment. What makes zero-intelligence models so compelling is that, despite the agents’ lack of sophisticated reasoning, they often reproduce key aggregate patterns with remarkable fidelity. In chapters 5 and 6, we will demonstrate exactly this in the context of the labor market. These re-

9: In the thermodynamic analogy, $\beta \rightarrow 0$ corresponds to $T \rightarrow +\infty$, i.e. maximal thermal agitation, whereas $\beta \rightarrow +\infty$ corresponds to $T \rightarrow 0$, i.e. a completely “frozen” system.

sults further bolster the case for agent-based modeling and call into question the continued reliance on the fully rational representative-agent framework.

The curse of dimensionality

Now that we have moved beyond the representative-agent paradigm and embraced truly agent-based models—where hundreds, thousands, or even millions of heterogeneous agents interact according to simple micro-rules—a natural concern is *tractability*. If every agent's state evolves according to its own equation, must we really solve an enormous, tightly coupled system to recover the aggregate behavior? And, even if we could, how can we explore a model that depends on dozens of parameters without running head-long into the *curse of dimensionality*¹⁰?

A powerful way to tame the second problem is to identify and eliminate *irrelevant directions* in parameter space, as eloquently described in [28]. Suppose our model's fit to data is measured by the squared-error loss

$$L(\boldsymbol{\theta}) = \sum_t [y_{\text{model}}(t; \boldsymbol{\theta}) - y_{\text{data}}(t)]^2, \quad \boldsymbol{\theta} = (\theta_1, \dots, \theta_p). \quad (1.4)$$

Around the best-fit point $\hat{\boldsymbol{\theta}}$, a second-order Taylor expansion gives

$$L(\boldsymbol{\theta}) \approx L(\hat{\boldsymbol{\theta}}) + \frac{1}{2} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^T H (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}), \quad (1.5)$$

where the Hessian matrix $H \in \mathbb{R}^{p \times p}$ has entries

$$H_{ij} = \left. \frac{\partial^2 L}{\partial \theta_i \partial \theta_j} \right|_{\hat{\boldsymbol{\theta}}}, \quad i, j = 1, \dots, p. \quad (1.6)$$

Because H is real and symmetric, it admits an orthonormal eigen-decomposition

$$H \mathbf{v}_k = \lambda_k \mathbf{v}_k, \quad \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p, \quad k = 1, \dots, p. \quad (1.7)$$

This eigen-decomposition then allows to identify:

- **Stiff directions:** eigenvectors \mathbf{v}_k with λ_k large. Small moves along these combinations cause L to increase sharply.
- **Sloppy directions:** eigenvectors with λ_k near zero. Large moves here barely change L , so these parameter combinations are effectively irrelevant for aggregate behavior.

By projecting out or fixing the sloppy directions—in other words, by restricting $\boldsymbol{\theta}$ to the subspace spanned by the first $m \ll p$ eigenvectors—we dramatically reduce the effective dimension of the calibration problem and gain both computational speed and interpretability. In practice, we do not pursue formal sloppiness analyses in this thesis, since our models are deliberately kept parsimonious, in the spirit of *Ockham's razor*¹¹ and parameter-selection is therefore straightforward. For instance, we computed the Hessian for the housing-market model in chapter 4, but found no parameters that could be discarded; conse-

10: Bellman (1957) first introduced this term in his work on dynamic programming [27]. In our case, this highlights that the computational effort required to explore a parameter space grows exponentially with the number of parameters.

11: Often referred to as the principle of parsimony, Ockham's razor is encapsulated in the Latin maxim "*Entia non sunt multiplicanda praeter necessitatem*," which translates as "entities should not be multiplied beyond necessity." Philosophically, it urges us to prefer explanations that rely on the fewest assumptions, a sentiment commonly paraphrased as "when faced with two competing theories, choose the simpler one." This principle is attributed to 14th-century English philosopher and theologian William of Ockham.

quently, that calculation is omitted from the final text.

Addressing the first issue is more challenging. A physicist might instinctively seek to derive a single macroscopic evolution equation—often formulated as a global stochastic differential equation—that governs the aggregate dynamics of the system. For instance, consider a Sakoda-Schelling-type occupation model in which agents move around a city according to simple local rules. One would like to “coarse-grain” those millions of microscopic trajectories into a continuum equation describing how the *density* of agents, $\rho(x, t)$, evolves over space and time. This is exactly the *hydrodynamic* strategy sketched in the introduction, and later developed in chapter 3—but it must be justified.

Drawing on the work of Gualdi et al. [29], in the case of the synchronization transition,¹² one finds that many collective phenomena observed in fully detailed agent-based simulations persist even after those microscopic rules are replaced by much simpler, “bare-bone” equations. In other words, when you *zoom out* and describe the system in terms of smoothly varying density and velocity fields rather than tracking each individual actor, the same large-scale cooperative patterns—waves of consensus, clustering, or shock-like transitions—continue to appear. This remarkable robustness lends strong support to the idea that socio-economic interactions can be fruitfully modeled using hydrodynamic or fluid-dynamical analogies: by setting aside the details of each person’s decision process, one still captures the essence of how groups organize and evolve.

Likewise, Silverberg et al.[30] showed that the chaotic movements of attendees at a heavy-metal concert can be modeled in much the same way as molecules in a gas. By translating measures of crowd density and local “pressure” into the language of kinetic theory, they successfully anticipated large-scale phenomena—like the emergence of empty corridors, dense flow channels, or sudden collective surges¹³—and found these predictions to match video observations. Altogether, such findings reinforce a unifying principle: complex group behaviors often reduce to the dynamics of a continuous field, offering clear qualitative understanding and reliable quantitative forecasts without the need to track every individual interaction.

Indeed, in a pioneering study, Hughes [31] treated pedestrian crowds explicitly as a *thinking fluid*, writing down Navier–Stokes-type equations¹⁴ enhanced by functions aimed at capturing human responses to density and obstacles. When the crowd density is high enough—so that the typical distance between two people is much smaller than the domain of interest—the group behaves like a compressible, elastic medium in which pressure waves can propagate, and mechanical effects such as asphyxiation can dominate. In two dimensions, this equation takes the form

$$-\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x} \left(\rho g(\rho) f^2(\rho) \frac{\partial \phi}{\partial x} \right) + \frac{\partial}{\partial y} \left(\rho g(\rho) f^2(\rho) \frac{\partial \phi}{\partial y} \right) = 0 \quad (1.8)$$

where

- ϕ is the *remaining travel time* for an agent, akin to a physical potential,

12: Synchronization—for example the coordinated firing of neurons in the brain or clapping in opera halls—is both familiar and captivating. The authors detected a similar collective rhythm in their agent-based macroeconomic model, manifesting as regular spikes of unemployment that correspond to waves of mass bankruptcies. This emergent synchronization prompted them to develop an even more minimal model, which likewise exhibited the same synchronized bursts, underscoring the phenomenon’s resilience across different levels of abstraction.

13: For instance, the so-called *wall of death*, where concertgoers split into two opposing lines and then charge into each other, produces distinct dynamical states that their continuum equations capture remarkably well.

14: First formulated by Claude-Louis Navier and George Gabriel Stokes in the early 1800s to describe the behavior of viscous fluids, these equations remain fundamental today for simulating everything from airflow around aircraft wings to ocean currents and weather systems.

- $f(\rho)$ gives the agents' *speed* as a function of local density ρ ,
- $g(\rho)$ is a *discomfort factor* that slows movement as the crowd becomes denser.

Two striking historical examples of dense human flows are the Battle of Azincourt¹⁵ and the annual Muslim *Hajj* pilgrimage to Mecca in Saudi Arabia. In each scenario, vast numbers of individuals converge within confined geometry, streaming along multiple intersecting routes under strict spatial constraints.

At Azincourt, French knights advanced in dense ranks only to be forced into a narrowing corridor between English battalions, battered by a hail of arrows from tightly arrayed English archers on their front and flanks, while fresh French lines charged behind them. The resulting compression waves overwhelmed the leading troops, causing mass casualties far beyond direct combat losses. Similarly, during the 1990 “Stoning of the Devil” ritual¹⁶ of the *Hajj*, streams of worshipers funnel onto a multi-level bridge. When a localized clog formed, the built-up rearward force propagated through the crowd, resulting in tragic loss of life.

Both cases vividly illustrate how, whether in medieval battlefields or modern religious gatherings, collective human motion under confinement can generate shock-like density waves and catastrophic outcomes whenever local interactions produce global compression. These examples illustrate the relevance of a large-scale aggregate partial differential equation—here, the Navier–Stokes equation—for understanding and analyzing complex system behavior, which is the main focus of this thesis.

Recent studies have also leveraged the powerful *mean-field-game* framework [32]¹⁷ to describe how pedestrians compete for space in dense crowds and to predict the collective response to perturbations such as the passage of a cylindrical obstacle—an abstraction of a vehicle moving through the crowd [33].

Moreover, as elaborated in [16], one can recognize striking similarities between the movement of human crowds and other self-organizing systems found in the natural world. Consider, for instance, a murmuration of starlings sweeping across the evening sky (see Figure (1.1) for an illustration), a congregation of fireflies synchronizing their bioluminescent flashes, or tens of thousands of spectators collectively rising and sitting to form a stadium “Ola.” In all of these examples, each individual—whether bird, insect, or person—responds only to the behavior of its immediate neighbors. Yet, through these simple, local interactions, coherent waves and large-scale patterns emerge that often obey the same kinds of continuum equations one uses to describe fluids. A landmark investigation into starling flocks [34] provided the first quantitative evidence that such coordination does not require awareness of every other member of the group. Instead, each bird tracks and aligns itself with only six or seven of its closest peers. Despite this remarkably limited horizon, the flock as a whole remains tightly knit and can dynamically reshape itself—say, to evade a predator—without ever isolating an individual. This discovery underscores how a network of simple, neighbor-to-neighbor rules can give rise to the robust, flexible collective behaviors we observe in both

15: Fought on 25 October 1415 during the Hundred Years’ War between France and England, this engagement ended in a disastrous defeat for the French army. In the aftermath, England maintained dominance for nearly fourteen years, until the lifting of the siege of Orléans under Joan of Arc reversed the course of the conflict.

16: During this rite, pilgrims symbolically stone three pillars—the *jamarāt*—located on the outskirts of Mecca, thereby reenacting Abraham (Ibrahim)’s rejection of Satan. On a single day, millions cross the Jamarat Bridge; in 1990 a sudden blockage triggered a deadly pressure surge, leading to hundreds of fatalities.

17: Mean-field game (MFG) theory analyzes the strategic behavior of vast populations of interacting agents. Lying at the intersection of game theory, stochastic analysis and optimal control, it extends the *mean-field* idea from statistical physics, where the influence of any single particle vanishes as the particle count tends to infinity. Similarly, in a mean-field game each agent solves an individual optimization problem that depends on the collective state of the population. The framework was pioneered by the French mathematicians Jean-Michel Lasry and Pierre-Louis Lions.



Figure 1.1.: Photograph of a flock of starling birds murmuring.

nature and human gatherings.

What is a good model?

Having seen how one can tackle the curse of dimensionality—condensing a detailed, micro-level description into a tractable, field-like continuum formulation (stochastic fluctuations included)—we must now ask: what makes for a truly useful model? Some recent efforts have aimed to recreate the Hungarian macro economy in exhaustive, “one-to-one” detail [35], simulating four million individual households, a central bank, firms, credit networks, and more, all finely calibrated against microeconomic data. While such high-resolution agent-based constructions are undeniably impressive, they carry enormous computational and energy costs. One may also reasonably question the practical value of constructing models whose complexity rivals that of the phenomena they aim to represent—a trap that Borges vividly illustrates in his cautionary tale [36].¹⁸ Even simpler frameworks like the “Mark 0” model [37] which captures many of the aggregate macro economic phenomena, contains a dozen parameters to tune and calibrate, echoing our earlier discussion about parsimony versus complexity.

This raises a deeper question: should our ambition be to build ever more faithful digital replicas of economic reality? Certainly, there is great merit in such endeavors—they can reveal policy impacts, systemic vulnerabilities, and fine-grained distributional effects that coarse descriptions miss. Yet, from the viewpoint of statistical physics, our aim is subtly different. As Jean-Philippe Bouchaud has argued [16], models fall into several broad categories, each serving a distinct purpose. For example:

Phenomenological models: These models strive to reproduce observed patterns by choosing flexible mathematical forms that closely fit the data. A paradigmatic example is the Black–Scholes–Merton framework, which treats asset prices as a geometric Brownian motion—hence assuming normally distributed returns. This Gaussian hypothesis inherently downplays the chance of extreme events, even though systems ranging from financial markets to seismic activity and tsunami propagation exhibit fat-tailed behavior. In such contexts, the

18: In “On Exactitude in Science” (1946), Borges satirizes an empire in which cartography has become so obsessively precise that only a map drawn at the scale of the empire itself is considered adequate.

probability of a rare, high-impact occurrence decays far more slowly than an exponential law, giving rise to *black swans*¹⁹ whose likelihood and impact are vastly understated by a purely Gaussian model. Put differently, these approaches might predict a catastrophic earthquake to be orders of magnitude less probable than what the empirical record actually shows.

Metaphorical models: Rather than mirroring every detail, these models distill a handful of stylized rules or heuristics—sometimes seemingly simplistic—into compact equations whose nontrivial outcomes illuminate underlying mechanisms. As Bouchaud writes:

“These models do not seek to explain an observation in detail, but rather to highlight mechanisms that would be difficult to identify without a precise mathematical framework. The hypotheses are crudely stylized, sometimes even seemingly absurd, yet their consequences are non-trivial, hard to guess without the model, and appear to correspond to real phenomena.”

The research carried out in this thesis clearly belongs to the realm of the *metaphorical* framework, even when we employ agent-based models. Although they are often prized for their ability to recreate rich, detailed portraits of real-world behavior, in practice we frequently distill those high-resolution dynamics into a handful of macroscopic, field-like equations. As previously discussed, coarse-graining serves two important purposes. First, it dramatically lowers the computational burden, enabling us to explore large parameter spaces and long-time behavior that would be infeasible if every individual agent’s state had to be tracked. Second, by focusing on aggregated variables—mean densities, average flows, or global order parameters—we gain clearer insight into the fundamental mechanisms driving phenomena across diverse socio-economic contexts. In this way, the full complexity of the micro-level rules informs, but does not obscure, the emergent laws that govern collective dynamics. By embracing simplicity and abstraction, we uncover the fundamental drivers of collective behavior in economic systems without getting lost in the thicket of micro-level detail.

The Sakoda-Schelling model

One of the most celebrated socio-economic *metaphorical* frameworks is the so called Schelling segregation model. Its origins trace back to James Sakoda’s 1971 publication in the *Journal of Mathematical Sociology* [15], which pioneered the use of agent-based modeling within a socio-economic context. Independently, in that same year, Thomas Schelling introduced essentially the same mechanism [14], later elaborating in his influential book on how individual agents’ *micromotives* can give rise to unexpectedly complex *macrobehaviors* [17].

These models are typically implemented on a two-dimensional square lattice, where each lattice site represents a “house” and the lattice as a whole represents a city of size $N \times N$. A fixed total of n agents occupies some of these sites, so that the global agent density remains $\rho_0 = \frac{n}{N^2}$. Each agent’s decision to move is driven by a utility function $u(\rho)$ that depends solely on the *perceived density* ρ of its local neighborhood—

19: The metaphor of the *black swan* dates back to 2nd century Roman poet Juvenal: “*rara avis in terris nigroque simillima cygno*” (Satire VI), meaning “a bird as rare upon the earth as a black swan.” At this time, black swans were considered not to exist. Real black swans were first observed in Australia in the late 1600s. Nassim Taleb later revived the term [38] to denote extreme, unforeseen events whose very possibility we tend to deny—an attitude exemplified by the centuries-long European belief that black swans could not exist simply because none had ever been seen.

i.e., the fraction of neighboring sites that are occupied. In most formulations, $u(\rho)$ attains its maximum at some intermediate target density ρ^* , reflecting the idea that individuals prefer to live in neighborhoods neither too empty nor too crowded. An agent will relocate from its current site to an empty site following a given decision rule which is a function of the utility of both sites $u(\rho_{\text{new}}), u(\rho_{\text{old}})$, as discussed previously.

Formally, consider an $N \times N$ lattice containing N^2 possible residences. Exactly n of these sites are occupied by indistinguishable agents, so that

$$\rho_0 = \frac{n}{N^2}$$

is conserved over time. Define each agent's utility as $u(\rho)$, which is assumed to peak at $\rho = \rho^*$. At each time step, vacant site j (with local density ρ_j) is proposed to an agent at site i (with local density ρ_i). If the agent accept this move, he will occupy site j and site i becomes empty. Thus each transition modifies the global state or the system.

Surprisingly, even when every agent's utility function is maximized at the same intermediate density ρ^* , the long-term steady-state configuration often exhibits large, nearly homogeneous clusters. In certain regions of parameter space—particularly when agents tolerate only slight deviations from ρ^* —one finds that the lattice partitions into dense clusters of agents and zones that are either nearly empty. In other words, rather than settling into the “ideal” neighborhood density ρ^* everywhere, the system collapses into a highly condensed state in which local densities far exceed the preferred value ρ^* .

A natural question is: how should one define an agent's *perceived density* on this lattice? Grauwin et al. [39, 40] propose subdividing the entire $N \times N$ grid into Q disjoint “neighborhood blocks.” Each block q contains H total houses, and if there are n_q agents in block q , then its block density is

$$\rho_q = \frac{n_q}{H}.$$

Agents only evaluate the utility of any proposed site j by considering the block density $\rho_{q(j)}$ of the block in which j lies.

Introducing this block structure allows the definition of a *block-level free energy*

$$f(\rho) = -\frac{1}{\beta}[\rho \ln(\rho) + (1 - \rho) \ln(1 - \rho)] + \int_0^\rho u(\rho') d\rho', \quad (1.9)$$

where β is the intensity of choice (inverse temperature) parameter. When each agent's moving cost corresponds to a change in the global system-wide quantity ΔV , and if that ΔV can be derived from a free energy, the dynamics satisfy detailed balance and behave like an equilibrium statistical-mechanics model. In this setting, one can analytically find the equilibrium system configuration by minimizing this free energy. Grauwin et al. show that under these conditions, the system can get stuck, at equilibrium, in a suboptimal configuration, leading to the observed clustering in Figure 1.2: starting from a random initial distribution (left panel), the system converges to a state in which some blocks (neighborhoods) are completely vacant while others ap-

proach density $\rho_q \approx 0.7$, even though $\rho^* = 0.5$.




Figure 1.2.: Example of the Schelling model on a 100×100 lattice subdivided into 10×10 blocks. Each site is either occupied (red) or empty (blue). The utility function peaks at $\rho^* = 0.5$. The left panel shows a random initial configuration; the right panel displays the steady state, in which entire blocks are either empty or over-populated with local density $\rho_q \approx 0.7 > \rho^*$.

In chapter 3, we demonstrate that these segregation phenomena persist even when one replaces the equilibrium dynamics with more general non-equilibrium decision rules. In other words, the tendency for agents to condense into suboptimal, high-density clusters remains robust under a wide variety of updating protocols.

Nonetheless, classical Schelling-type occupation models omit a crucial real-world ingredient: housing prices. In standard economic theory, markets allocate resources and adjust prices so that supply and demand reach equilibrium. One could thus argue that introducing a functioning housing market into the Schelling framework would allow agents to compare purchase or rental prices across neighborhoods, potentially enabling them to achieve configurations that maximize their utilities—rather than becoming trapped in overly concentrated clusters. A well-functioning market should, in principle, mitigate the unanticipated segregation effects observed in the pure occupation model (though some empirical anomalies remain, e.g. [41]).

In chapter 4, after extensive analytical and empirical analysis of the French housing market, we demonstrate that coupling the Schelling model with a similar housing prices evolution equation does not invariably prevent condensation.

1.2. Classical macroeconomics frameworks and the need for agent-based models

Following our discussion of how statistical physics provides powerful tools for analyzing socio-economic systems—and having introduced the Sakoda-Schelling model as a paradigmatic, *metaphorical* framework for exploring patterns of social segregation—we now shift our focus toward macroeconomics. This transition is natural, since the first part of this thesis applies physics-inspired techniques to the French housing market, and the second part undertakes a comprehensive study of the US labor market. A macroeconomic perspective is hence essential for capturing the collective phenomena that arise when vast numbers of agents interact.

By definition, macroeconomics is the branch of economics devoted to understanding the behavior of the economy *as a whole*, rather than analyzing individual sectors or firms in isolation. It concentrates on aggregate quantities such as gross domestic product (GDP), national unemployment rates, inflation, and total output. The canonical starting

point for modern macroeconomic thought is John Maynard Keynes's seminal 1936 work, *The General Theory of Employment, Interest and Money*²⁰ [42]. Prior to Keynes, the prevailing view held that markets would clear instantaneously, ensuring full employment. Keynes challenged this assumption by demonstrating that aggregate demand might fall short of the level required to employ all willing workers, thereby creating persistent unemployment.

In the second part of this thesis, we study the US labor market, where millions of job seekers and employers interact through a decentralized matching process. Here, macroeconomics becomes especially relevant: to explain how aggregate unemployment persists, we must move beyond individual decision rules and instead study the overall flow of workers into and out of employment and the role of external shocks. By adopting the macroeconomic framework pioneered by Keynes—one that emphasizes aggregate demand and the possibility of sustained disequilibrium—we link microscopic agent-based dynamics to the broad, system-level patterns observed empirically.

On bounded rationality

To understand how local interactions among economic agents give rise to macro-level outcomes—such as aggregate output, employment dynamics, and systemic stability—we must reconceptualize macroeconomic analysis. At the heart of this transformation is a critical reappraisal of the classical assumption that agents possess perfect foresight and unlimited computational capacity—the *rational expectations* paradigm introduced by Robert Lucas [43]. Under this framework, individuals are presumed to form expectations by fully integrating all available information and the correct economic model, thereby making decisions that optimally anticipate future states.

Herbert A. Simon challenged this idealization with his concept of *bounded rationality* [11], asking how real-world decision-makers behave when the strict conditions of *rational choice theory* are unattainable. In these settings, agents employ *heuristics*²¹ rather than solving fully specified optimization problems.

As J. Dooyne Farmer eloquently puts it, socio-economic systems reside in *the wilderness of bounded rationality* [44], accessible via two contrasting gateways. The first, the *rationality* gate, embodies traditional economic models presuming agents with infinite computational capacity and foresight. The second, the *zero-intelligence* gate, dispenses with such assumptions entirely, specifying only that agents follow simple heuristic rules.

In Part II of this thesis, we show that zero-intelligence models not only reproduce fundamental empirical regularities but, in many instances, surpass classical rational-expectations frameworks—underscoring the importance of this alternative route through the wilderness of bounded rationality.

In the following section, we examine the well-known macroeconomic relationship called the Beveridge curve, which we analyze in detail in chapters 5 and 6.

20: Keynes's book inaugurated what is now known as the *Keynesian Revolution*. In it, he argued that economies do not automatically restore full employment following an external shock, because wages and prices may adjust only slowly or remain rigid. As a result, prolonged periods of under-employment and economic downturn can occur. This insight anticipates our second-part analysis of the US labor market, where we explore its inherently nonequilibrium dynamics.

21: From the Ancient Greek for “to discover,” heuristics are mental shortcuts or “rules of thumb” that guide judgment when facing complexity or uncertainty. A classic relatable example is the “shortest-line” heuristic at the grocery store. Faced with several checkout lanes, you don't calculate the exact number of items each customer has or predict each cashier's speed—instead, you simply choose the lane that looks shortest or moves fastest in that moment. This rule of thumb usually works well enough under time pressure and limited information, even though it isn't a perfectly optimal calculation of total wait time.

A very famous macroeconomics feature: The Beveridge curve

The Beveridge curve ²², also known as the U-V curve, is a very well known stylized fact of labor markets. It shows that vacancies (V) and unemployment (U) tend to move in opposite directions: increasing one of them tends to decrease the other and vice-versa. Although this antagonistic behavior between the two variables is broadly understood, the mechanisms that explain its precise shape, and in particular how it shifts during business cycles, remain a hotly debated topic.

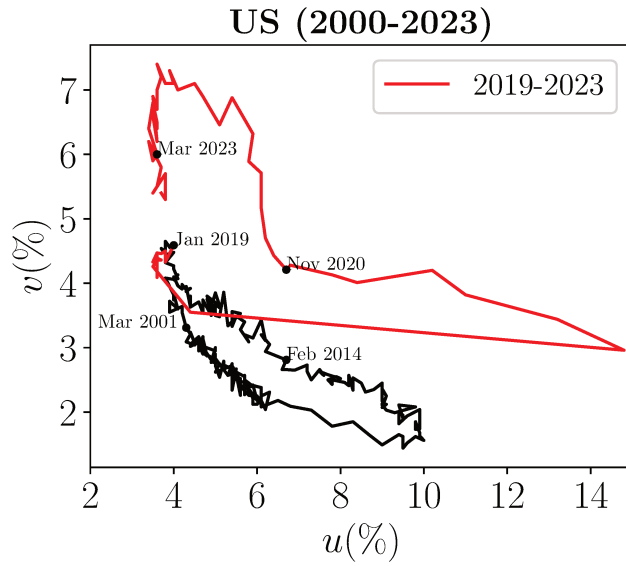


Figure 1.3.: Empirical U.S. Beveridge curve, 2000–2023. The 2019–2023 (COVID and post-COVID) observations are shown in red to highlight the structural shift, while the 2000–2019 segment is plotted in black. The curve clearly follows an anticlockwise trajectory, and displays two distinct cycles corresponding to two different economic modes.

During a boom-bust (succession of recoveries and recessions) business cycle the Beveridge curve is similar to a semi-closed ellipse, with a trajectory that is always counter-clockwise. This means that for a fixed vacancy rate the corresponding level of unemployment is higher during the economy's recovery than it was during its recession. As a result getting back to pre-recession unemployment levels is difficult, which is called a *jobless recovery* [46]. This means that even if policy makers manage to stimulate more job vacancies, the economy may still have persistent high unemployment after crises (see [47]). Understanding the shape and dynamics of the Beveridge curve is also important because it might help us understand the *unemployment-inflation trade-off* ²³ [50, 51].

In chapter 5, we propose a novel mechanism to explain the Beveridge curve using an (almost) zero intelligence disequilibrium model. This is an alternative to the standard economic equilibrium and rational agents framework. Our work here adds more evidence to models in complexity economics, which argues that disequilibrium and bounded rationality are often needed to explain dynamic economic phenomena (see [52, 53] and [16]). Further empirical evidence suggests that out-of-steady-state dynamics play a significant role in shaping the Beveridge curve, as shown by [54]. This motivates us to follow the footsteps of past work that has used agent-based models to study labor markets using features such as bounded rationality [55] or network effects [56, 57]. Our contribution is to use such a model to explain

23: This idea traces back to New Zealand economist A.W. Phillips, who in 1958 demonstrated a negative relationship between wage-rate changes and unemployment [48]. Building on that insight, Paul A. Samuelson and Robert M. Solow (1960) extended the analysis to show a similar inverse relationship between inflation and unemployment, thereby formalizing the well-known unemployment-inflation trade-off [49]. In practical terms, their result implies that periods of higher inflation tend to coincide with lower unemployment rates, and conversely, reductions in inflation often coincide with rising unemployment. However, the persistence of this effect—particularly whether it vanishes in the long run—remains the subject of ongoing debate.

the behavior of the labor market during business cycles.

The framework we will discuss is a zero-intelligence (in the sense of [26]) agent-based model which is a simplified, aggregated version of an earlier model [57]. Agents seeking jobs follow simple rules of thumb: unemployed workers apply at random to jobs compatible with their skills, and firms post vacancies when they need to hire more people to meet a production target or remove them when demand is reduced. Although we base our analysis on the work of [57]—which was the first agent-based model to explicitly study the Beveridge cycle—the key mechanisms we argue drive the Beveridge cycle are commonplace in agent-based modeling.

Crucially, our explanation does not require firms to be rational agents, capable of computing the optimal number of vacancies they should post, although it is not incompatible with rational behavior from firms. The key mechanism in explaining the cycle in the Beveridge curve lies in reconsidering the comparative statics approximation that is most commonly used in modeling labor markets. We argue instead that the time scales required for the different agents in the economy – firms and workers in this case – to coordinate can be much longer than what is typically assumed. This means that *both* the unemployment and vacancy rate do not immediately adjust to their long-term equilibrium value. If the economy is subject to exogenous fluctuations that happen so quickly that they do not allow it to reach its equilibrium, then it is effectively out of its steady state. We model explicitly how the dynamics are different during a boom and a bust, an asymmetry related to the fact that hiring and firing workers is not done at the same speed. In short, the assumption that agents do not coordinate instantaneously means that the economy is not in its point of equilibrium; when supplemented with asymmetric dynamics when employment is above or below the equilibrium target, this is sufficient to explain the cycle in the Beveridge curve. Lastly, we show that the same qualitative behavior can be obtained even in an aggregate model with a standard Cobb-Douglas²⁴ matching function. The key concept of matching function will be elaborated on in more details in chapter 5.

We will then follow up on this work in chapter 6 by developing a natural extension of that framework—driven by the aggregated equations inferred from our initial agent-based model—into a parsimonious dynamic disequilibrium model that shows striking agreement with US empirical data.

The Diamond-Mortensen-Pissarides framework

The main classical framework used to understand the Beveridge curve is the well-known Diamond-Mortensen Pissarides (DMP) framework [20, 19, 18], which differs from previous work, such as implicit contract models [58, 59], bargaining models [60] or efficiency wage models [61], by explicitly modeling vacancy evolution without assuming market clearing.

In particular, Pissarides proposed [18] a partial equilibrium model in which workers, employed or unemployed, go from one category to the other by being separated from their current employer or matched with

24: Formulated and empirically tested in the early twentieth century by mathematician Charles Cobb and economist Paul Douglas, the function captures the relationship between the two inputs L and K through the functional form: $L^\beta K^{1-\beta}$.

a vacancy. The matching process is hindered by frictions that prevent unemployed workers from being instantaneously matched with available vacancies. This component of the framework is now present in nearly all labor market models, where the job search process and the matching function capture most of the attention.

However, the framework implies a series of additional assumptions that are rarely questioned in the literature. These assumptions revolve around the way in which firms set vacancies: it is assumed that firms follow a standard optimization procedure, where they weigh the cost of posting a vacancy and hiring a worker versus the profits and losses associated with having an unused machine that could be sold at a given price. This view implicitly assumes that this arbitrage can be done in very liquid markets, where the value of an idle machine can be readily established and where they can be easily exchanged. In reality, evaluating the price of these factors is a costly and time-consuming task, given that they cannot be easily traded. Firms are also assumed to have the capability to carry out all of these computations within an arbitrarily small time-frame, as implied by the continuous-time setting of the model. These hypotheses imply a level of coordination between firms that means that the *free-entry condition* applies²⁵, and so the economy is in an equilibrium where firms can post and remove vacancies without any cost or friction.

Therefore, in this framework only unemployment follows explicit dynamics, which will be outlined in chapter 5. The solution to the firms' optimization problem means that they set their vacancies proportionally to the unemployment rate at all times. The labor market converges to a stationary state with an unemployment rate that depends on exogenous factors, like the workers' bargaining power, their productivity distribution or the efficiency of the job search and match process.

Within the DMP framework, it has been argued that the labor market reaches its stationary state within a timescale of the order of one month [62], after either an exogenous shock or given a change in the economic conditions described above. This has important consequences in the modeling choices for the Beveridge curve. A common choice is to do comparative statics (as do [62]), where the Beveridge curve cycle is explained as a succession of stationary states of the labor market after changes in economic conditions. The nature of the shifts in the curve is then explained by several mechanisms, all of which picture the economy as being perpetually close to a steady state. For example, [46] suggest a crowding-out effect, with employed workers having more incentives to look for jobs after recessions than unemployed workers. [63] proposes demand externalities as an explanation for shifts in the absence of shocks, and [64] show that sectoral misallocation and heterogeneous speed in recovery for different sectors may also lead to outward movements in the Beveridge curve. Such interpretations through comparative statics amount to explaining the shifts in the Beveridge curve with structural changes in the labor market [65]. The outward shift of the curve is then interpreted through *hysteresis*: the unemployment rate does not increase in a recovery by just retracing its footsteps because the economic equilibria it follows are structurally different.

25: This premise is embedded within the *perfect competition* framework, where any prospective firm can enter the market and begin producing and selling (and open vacancies), thereby preventing monopolies from persisting.

In fact, a recent empirical study for the US Beveridge curve [66] suggests that each shift for each business cycle has its own specific explanation: An influx into the worker force between 1970-2009, a decrease in matching efficiency due to skill mismatch during the Great Recession [67], and a diversity of reasons during the COVID-19 pandemic, including the rise in quits during the Great Resignation [68]. Thus, a static framework with shifting points of equilibrium that change depending on specific economic changes would align well with this narrative.

This is further reinforced by more recent empirical work that shows lasting structural changes after periods of crisis. Bermejo et al. examine unemployment and vacancies levels across 24 European countries from 2010 to 2020 [69], finding evidence of enduring structural effects in the labor market following the COVID-19 pandemic. Another study conducted by Alana et al. [70] scrutinized the labor market of the G7 nations over the period 2002 to 2023, revealing a decreasing trend in the spread between unemployment and vacancies levels. The study shows that unemployment levels exhibit mean-reverting characteristics, potentially validating the theory that unemployment reverts back to a natural rate, which arguably strengthens the case for comparative statics. Other empirical work also argues in favor of a structuralist hypothesis after studying the labor markets of 23 OECD countries [71]. Nonetheless, while all of these papers use state of the art data and statistical techniques, there are substantial disagreements in their explanation of the Beveridge curve. The diversity of factors needed to explain each shift in the curve make it clear that these theories are not parsimonious – there is no common explanation with predictive power.

Some concerns have also been raised about the dynamical properties of the DMP framework. Shimer has, for example, objected that it is not able to account for the empirical fluctuations in unemployment and job vacancies during business cycles [72]. Mortensen proposed a dynamic version of the DMP model [73], where the equilibrium solutions of the dynamical system suggest that multiple equilibria are theoretically possible when parameter values are chosen to align reasonably with empirical estimates. This framework implies the potential for deterministic endogenous limit cycles in unemployment, but maintains the assumptions described in the DMP model.

On the other hand, Bocquet [74] proposes a model with a granular description of labor mobility frictions, where *occupational bottlenecks* prevent the system from reaching a stationary state quickly. In his model, the adjustment time of the labor market after a shock can be of the order of ten months, which is ten times larger than that argued for by Michaillat et al. [62]. It can also be much longer if the shock specifically targets bottlenecks in the network. Such a long adjustment time-scale makes the comparative static framework moot. If the time it takes the labor market to reach its stationary state is of the same order as the time during which economic conditions can change, then dynamics may play an important role, as also argued in [75].

This notion of bottlenecks was also picked up by a recent data analysis of the French labor market conducted by Knicker et al. [76], re-

vealing that up to 93% of occupations are bottlenecks in the occupational mobility network. Crucially, they studied the properties of the occupation-to-occupation transition matrix and the stationary distribution it implies. They compute the stationary number of workers within each occupation implied by the transition matrix, and find that nearly none of the occupations were close to this implied stationary value. In many cases, the discrepancy is quite large. This means that a picture of a stationary-state labor market is in conflict with empirical data, which seems instead to suggest a picture where the structure of the labor market changes at a speed that is commensurate with the time it would take it to reach its stationary state.

This is what is advocated by proponents of agent-based models, which allow for a study of dynamical economic systems through computer simulation. In the specific case of the labor market, ABMs have already been used to study similar questions. [64] uses such a framework with heterogeneous agents and firms with imperfect information to obtain productivity mismatches, a factor that could help in understanding job matching conditions leading to the Beveridge cycle. A similar idea is also explored in [56], using an agent-based model where agents are segregated into labor sectors, and face time frictions when transitioning between sectors. However, none of them explicitly studied the Beveridge cycle.

In ABMs, a traditional approach to modeling hiring and firing was initially outlined in [77]. In this approach, firms forecast their production needs based on factors such as the historical demand they have faced or expected inflation, making autonomous decisions without coordination with other economic actors. Firms adjust their workforce according to anticipated production changes – increasing hires when growth is required and reducing their staff when they anticipate a drop in demand. This picture is radically different from the picture of firms that are able to rationally find their optimal behavior taking into account all the possible decisions of workers or of other firms.

The modeling framework of ABMs has proven to be very robust and has been adopted by many models that demonstrate complex behavior, including endogenous crises and tipping points [37] and strong non-equilibrium dynamics [75]. It has also been used in models that analyze the economic recovery after the COVID pandemic [78] or the ensuing inflation dynamics [79]. More recently, this framework was also adopted in models that have been successfully calibrated to actual macroeconomic time series, thereby showing their relevance in real-world economic analysis [80, 81].

Finally, more recent macroeconomic ABMs make quantitative time series predictions that are comparable to those of standard macro-models. [80, 81]

The idea of modeling the hiring/firing in the labor market as a myopic process was combined with a detailed microscopic model of the matching process in an occupational network by del Rio-Chanona et al. [82, 57], and used to study the Beveridge cycle. Our contribution to the literature in this thesis is to make the links between the ABM and DMP framework explicit, to propose an explanation that shows how a very simple ABM can reproduce the Beveridge curve cycle

and to show how our model makes several predictions about the Beveridge curve. We do this by proposing a detailed study of the model described in [82], but we highlight that the mechanisms that are key to our findings are common to all the models in the ABM framework described above. Our model is also compatible with the findings of [74] and [76]. In agreement with [74], we find that the labor market can be very slow to react to certain shocks, but our mechanism and framework are different from theirs. Our model also explains recent finding that labor markets can be far from their stationary state [76].

Other perspectives in agent-based models

Having discussed agent-based models of the macro-economy, and the *metaphorical* Sakoda-Schelling model, one can conclude this chapter by discussing other interesting use and findings of ABMs.

Recent work have designed agent-based models of the housing market in order to explain its link with social segregation. For example [83] concluded that segregation patterns can be observed even with the simplest parameter setting in an agent-based model of the housing market. Moreover, [84] showed how understanding this link could be very helpful to test and apply effective policies to prevent such segregation, in the same vein as [85] where the effectiveness of macroprudential policies is tested on an agent-based model of the UK housing market. Interestingly, [86] showed that social segregation is strongly linked with social influence. Earlier studies, like that of [87], have also explored the inefficiency of the housing market concerning its responsiveness to external news and the potential diffusion of such effects. This is aligned with the work of Borghesi and Bouchaud [88, 89], who proposed to explain the spatial trends observed in turning out rates in French elections, from 1992 to 2009, with a noisy diffusive cultural (or influence) field. Such diffusion equation is a remarkable feature of social modeling, as it has been introduced multiple times in the past to describe socio-economic systems [90, 91], and will be the corner stone of our work presented in chapter 4.

Finally, Seara et al. [92] recently demonstrated the effectiveness of employing a hydrodynamic framework to model socio-economic systems, as previously proposed. This approach, which abstracts away from individual-level details and instead focuses on collective, large-scale phenomena, remains analytically manageable. By training neural networks on US census data, they were able to show—specifically in the context of urban mobility—that the sociohydrodynamic model successfully captured how, in the words, again, of Thomas Schelling, individual *micromotives* can give rise to emergent *macrobehaviors*.

1.3. This thesis

Our endeavor here aims at contributing to the field of econophysics and complexity economics, advocating for non-equilibrium models grounded in bounded rationality and heterogeneous agents, as well as the tools of stochastic differential equation to elucidate dynamic socio-economic phenomena (as outlined for example in [52, 53, 16]).

Within this framework, my research makes two principal contributions. First, it deepens our quantitative understanding of French socio-economic data, with a particular focus on the housing market. By integrating advanced statistical techniques and models inspired by statistical physics, we are able to characterize price evolution in a rigorous, data-driven manner. Second, it broadens insight into labor-market dynamics by employing both zero-intelligence agent-based models and more parsimonious, dynamical-disequilibrium *metaphorical* models in the physicist's tradition. Together, these approaches illuminate how macroeconomic aggregates—such as unemployment rates or housing prices—emerge from the interaction of agents with limited information and adaptive behavior.

Chapter 3 presents a novel reinterpretation of the Sakoda–Schelling model by importing techniques from hydrodynamics and Active Matter. We demonstrate that the well-known condensation phenomena persist under non-equilibrium dynamics and across a broad class of decision rules. This work was carried out in collaboration with R. Zakiné, J. Garnier-Brun, and M. Benzaquen [1].

In Chapter 4, we introduce a model based on a stochastic differential equation to describe and analyze the French housing market in detail. We show that our framework naturally reproduces the empirically observed long-range correlations. Moreover, we find that coupling this housing-market model to the Sakoda–Schelling dynamics does not necessarily inhibit condensation. This research was conducted under the supervision of M. Benzaquen and J.-P. Bouchaud [2].

Chapter 5 addresses the classic macroeconomic stylized fact known as the Beveridge curve, in particular its counter-clockwise, elliptic motion over business-cycle phases. Building on an agent-based labor-market simulation, we derive a minimal aggregate description comprising two coupled partial differential equations and show that it successfully reproduces the observed cycle-dependent shape. This chapter is based on work with J. Moran, R.M. del Rio-Chanona, and J.D. Farmer.

Finally, in chapter 6, we refine our aggregate approach to propose an even more parsimonious model that not only captures the Beveridge-curve dynamics but also delivers accurate out-of-sample forecasts of unemployment conditional on vacancies and GDP. This further validates the utility of *metaphorical* models founded on simple *heuristics*. This work was developed jointly with J. Moran, R.M. del Rio-Chanona, and J.D. Farmer.

Key results and conclusions are summarized at the end of each chapter, with some technical derivations and supplementary calculations deferred to dedicated appendices.

Theoretical Foundations

2.

*sic itur ad astra.*¹
Aeneid-Virgil, circa 29-19 BC.

This chapter aims to introduce the primary theoretical tools that underpin the subsequent chapters of this thesis. These discussions are largely introductory and are not intended to replace a dedicated course in statistical physics.

2.1. The Langevin equation

As already underlined in the introduction, stochastic differential equations are indispensable tools for analyzing complex socio-economic phenomena. Their strength lies in combining deterministic forces with random fluctuations, thereby capturing both systematic trends and intrinsic uncertainties. To establish the mathematical foundations that will be used throughout Part I of this thesis, we now return to several key ideas from statistical physics and develop them in greater detail.

Physical origin

A Brownian particle immersed in a fluid simultaneously experiences two broad classes of forces: (i) long-range or slowly varying forces of external origin, such as gravity or electromagnetic fields, and (ii) rapidly fluctuating impacts caused by incessant collisions with the surrounding fluid molecules. Paul Langevin's seminal insight¹ was to split the second category into two parts: an averaged viscous drag, corresponding to friction, and a stochastic contribution. Because the correlation time of molecular hits, τ , is typically several orders of magnitude smaller than any macroscopic time scale, the noise can be idealized as $\tau \rightarrow 0$. Hence, the stochastic component (called the *Langevin noise*) is traditionally modeled as *Gaussian white noise*.

For simplicity, we restrict attention to motion in a single spatial dimension X . The Langevin equation for the particle's velocity $V(t) = \dot{X}(t)$, which corresponds to Newton's equation of motion enhanced with the stochastic contribution described previously, reads

$$m \frac{dV(t)}{dt} = -\gamma V(t) + F(X) + \eta(t), \quad V(t) = \frac{dX(t)}{dt}, \quad (2.1)$$

where m denotes the particle mass, γ is the viscous-drag coefficient, $F(X)$ is an externally prescribed deterministic force, and $\eta(t)$ represents the stochastic kicks delivered by the fluid. The latter is charac-

1: P. Langevin, *Sur la théorie du mouvement brownien*, (1908).

¹"Thus one goes to the stars."

terized by its two-point correlator

$$\langle \eta(t)\eta(t') \rangle = \sigma^2 \delta(t - t'), \quad (2.2)$$

with σ^2 fixing the noise strength and $\delta(\cdot)$ the Dirac delta function. All higher-order cumulants of $\eta(t)$ vanish, reflecting the assumed Gaussian statistics.

Free particle solution and the Ornstein-Uhlenbeck process

If the external force is switched off [$F(X) = 0$], solving the linear first-order ordinary differential equation yields

$$V(t) = V(0) e^{-\frac{\gamma}{m}t} + \frac{1}{m} \int_0^t dt' e^{-\frac{\gamma}{m}(t-t')} \eta(t'). \quad (2.3)$$

The factor $\tau_v = m/\gamma$ sets the characteristic time over which the particle “forgets” its initial velocity. For times $t \gg \tau_v$ the first term decays exponentially, and $V(t)$ becomes a stationary Gaussian random variable with zero mean and autocorrelation

$$\langle V(t)V(t') \rangle = \frac{\sigma^2 \tau_v}{2m^2} \exp[-|t - t'|/\tau_v], \quad (t, t' \gg \tau_v). \quad (2.4)$$

The resulting stationary process is known as the *Ornstein–Uhlenbeck* process and constitutes the archetype of colored noise with finite correlation time that we will extensively use in Part-I of this thesis.

At thermal equilibrium, kinetic-energy equipartition demands $m\langle V^2 \rangle = k_B T$, where T is the absolute temperature and k_B the Boltzmann constant. Matching this requirement to the long-time limit of the above correlator leads to the celebrated *fluctuation–dissipation theorem*:

$$\frac{\sigma^2 \tau_v}{2m^2} = \frac{k_B T}{m} \implies \sigma^2 = 2\gamma k_B T. \quad (2.5)$$

Notice that σ is independent of the mass of the particle m —the amplitude of the thermal kicks is determined solely by the solvent temperature and viscosity.

The overdamped (Smoluchowski) limit

In numerous experimental and economic contexts the inertia of the moving entity is negligible compared with viscous damping, i.e. $\tau_v = m/\gamma$ is the shortest time in the problem. Letting $m \rightarrow 0$ while keeping γ fixed converts the velocity dynamics into an instantaneous constraint and leaves a first-order stochastic differential equation for the position,

$$\gamma \dot{X}(t) = F(X) + \eta(t). \quad (2.6)$$

Choosing units such that $\gamma = 1$ simplifies notation:

$$\dot{X}(t) = F(X) + \eta(t). \quad (2.7)$$

Hence the time derivative of $X(t)$ consists of a deterministic drift, $F(X)$, plus a stochastic contribution, $\eta(t)$, whose variance is fixed by

the fluctuation-dissipation theorem, $\sigma^2 = 2k_B T$. Standard functional-integration arguments show that an entire trajectory $\{X(t)\}_{0 \leq t \leq t_f}$ is weighted by

$$P(\{X(t)\}; 0 \leq t \leq t_f) \propto \exp\left[-\int_0^{t_f} dt \frac{[\dot{X}(t) - F(X)]^2}{4k_B T}\right], \quad (2.8)$$

up to a normalization constant.

Diffusion and mobility

If no systematic force acts on the particle [$F(X) = 0$], the above stochastic differential equation reduces to pure diffusion. Integrating twice shows that

$$\langle X^2(t) \rangle = 2Dt, \quad D = k_B T, \quad (2.9)$$

demonstrating that the motion is characterized by a diffusion coefficient D . Restoring physical units, one recovers the more familiar expression $D = k_B T / \gamma$. By contrast, when a constant force F_c is applied, the mean displacement grows linearly,

$$\langle X(t) \rangle = \frac{F_c}{\gamma} t, \quad (2.10)$$

which defines the particle's *mobility*, equal to $1/\gamma$. The equality $D = k_B T / \gamma$ is the celebrated *Einstein relation* of 1905.²

Generalization to higher dimensions

For motion in two or three dimensions ($d = 2$ or 3), each Cartesian component evolves according to

$$\gamma \frac{dX_\mu(t)}{dt} = F_\mu(\vec{X}) + \eta_\mu(t), \quad \mu = 1, \dots, d, \quad (2.11)$$

where $\vec{X} = (X_1, \dots, X_d)$ and $\eta_\mu(t)$ is the μ -th component of the noise. Isotropy of the fluid implies

$$\langle \eta_\mu(t) \eta_\nu(t') \rangle = 2\gamma k_B T \delta_{\mu\nu} \delta(t - t'), \quad (2.12)$$

so fluctuations are mutually uncorrelated across directions.

In summary, the Langevin framework provides a clear and flexible description of Brownian motion by combining deterministic forces, viscous dissipation, and stochastic fluctuations into a single differential equation. In its overdamped limit, this formulation seamlessly reduces to classical diffusion theory. These concepts will form the theoretical foundation for the socio-economic models examined in the first part of this thesis.

2: By relating the diffusion constant to directly observable macroscopic parameters, Einstein succeeded in inferring the approximate size of individual atoms and estimate Avogadro's number—the count of atoms in a mole, or equivalently the gram-molecular mass of a gas.

2.2. From the Master equation to the diffusion equation

The *Master equation* is a bookkeeping identity that tracks how probability (or, more generally, any conserved scalar field) flows between discrete states. To fix ideas let us imagine a one-dimensional lattice with equally spaced nodes labeled by integers $n \in \mathbb{Z}$ and lattice spacing a . A quantity³ $\psi_n(t)$ resides at each site and can hop to the two neighboring nodes with transition rates

$$W(n \rightarrow n+1), \quad W(n \rightarrow n-1).$$

Because the only mechanism that changes the content of a site is these nearest-neighbor exchanges, the net balance reads

$$\begin{aligned} \frac{d\psi_n(t)}{dt} = & - \left[W(n \rightarrow n+1) + W(n \rightarrow n-1) \right] \psi_n(t) \\ & + W(n+1 \rightarrow n) \psi_{n+1}(t) + W(n-1 \rightarrow n) \psi_{n-1}(t). \end{aligned} \quad (2.13)$$

The Master equation (2.13) is nothing more than a discrete “mass conservation” equation: the first bracket measures what leaves node n , while the last two terms inject whatever enters from its neighbors.

3: In applications ψ might represent a genuine probability, a particle density, a concentration, or any scalar field obeying local conservation.

Symmetric notation and conservation

It is convenient to introduce the shorthand

$$\Gamma_{n,\ell} \equiv W(\ell \rightarrow n),$$

so that $\Gamma_{n,\ell}$ quantifies the influence⁴ of site ℓ on site n . Writing (2.13) with this notation gives

$$\frac{d\psi_n(t)}{dt} = \sum_{\ell} \Gamma_{n,\ell} \psi_{\ell}(t) - \sum_{\ell} \Gamma_{\ell,n} \psi_n(t), \quad (2.14)$$

where the sums collapse to $\ell = n \pm 1$ for a nearest-neighbor lattice but remain fully general otherwise. Equation (2.14) emphasizes a key structural point: the change of ψ_n at time t is the difference between an *influx* from all other sites and an *outflux* that ψ_n itself sends out.

4: Note that we use the term *influence* here rather than “probability rate of transition” to prepare for the housing market model in chapter 4.

Continuum limit

To recover a partial differential equation one now lets the spacing a go to zero. Define the continuous coordinate $x = na$ and write the field as

$$\psi(x, t) \equiv \psi_n(t),$$

keeping the same symbol for simplicity. Similarly the discrete kernel $\Gamma_{n,\ell}$ becomes a continuous influence kernel

$$\Gamma(x | x').$$

With these replacements the lattice sum turns into an integral, and (2.14) transforms into

$$\partial_t \psi(x, t) = \int dx' \Gamma(x | x') \psi(x', t) - \int dx' \Gamma(x' | x) \psi(x, t). \quad (2.15)$$

For notational convenience introduce the jump variable $y = x - x'$ and define

$$t(y | x) \equiv \Gamma(x | x - y).$$

Substituting $x' = x - y$ in (2.15) yields the compact *continuum Master equation*

$$\partial_t \psi(x, t) = \int dy t(y | x) \psi(x - y, t) - \left[\int dy t(y | x) \right] \psi(x, t). \quad (2.16)$$

Kramers–Moyal expansion

Assume that typical jumps $|y|$ are small compared with the macroscopic scales of interest, so that the field remains smooth in x . Expanding $\psi(x - y, t)$ in a Taylor series

$$\psi(x - y, t) = \psi(x, t) - y \partial_x \psi(x, t) + \frac{y^2}{2} \partial_x^2 \psi(x, t) + \dots$$

and inserting into (2.16) produces the Kramers–Moyal hierarchy. Truncating after the second order—appropriate when the kernel is sufficiently compact—gives

$$\partial_t \psi(x, t) = -\partial_x [R_1(x) \psi(x, t)] + \frac{1}{2} \partial_x^2 [R_2(x) \psi(x, t)], \quad (2.17)$$

where the first two jump moments are

$$R_1(x) = \int dy y t(y | x), \quad (2.18)$$

$$R_2(x) = \int dy y^2 t(y | x). \quad (2.19)$$

Symmetric kernel. If the influence kernel is *symmetric*, $t(y | x) = t(-y | x)$, every odd moment vanishes and in particular $R_1(x) = 0$. Equation (2.17) then reduces to

$$\partial_t \psi(x, t) = \partial_x^2 [D(x) \psi(x, t)], \quad (2.20)$$

with the spatially dependent diffusion coefficient

$$D(x) = \frac{1}{2} R_2(x) = \frac{1}{2} \int dy y^2 t(y | x).$$

Thus the familiar diffusion equation emerges as the natural continuous surrogate of a symmetric jump process when only the first two jump moments are retained.

Uniform- D simplification

In many practical situations the fine-scale variability of $t(y \mid x)$ is either unknown or negligible. Assuming the second moment is approximately constant, $D(x) \approx D$, (2.20) simplifies even further to the classic homogeneous diffusion equation

$$\partial_t \psi(x, t) = D \partial_x^2 \psi(x, t). \quad (2.21)$$

The procedure we have followed—from discrete hops, through the Master equation, to the Kramers–Moyal expansion—demonstrates precisely how a macroscopic diffusive law inherits its parameters from microscopic jump statistics. The same pathway will be reused throughout the thesis whenever we coarse-grain discrete, agent-level dynamics into smooth, continuum descriptions.

2.3. Variograms and correlograms: tools for characterizing fluctuations

In many empirical settings—ranging from geophysics to financial time-series analysis—the object of interest is a one-dimensional record $X(t)$, whether t is a temporal or a spatial coordinate. We introduce in this section *variograms* and *correlograms*, useful statistical tools that we extensively use in the first part of this thesis.

We assume that $X(t)$ possesses *stationary increments*: the statistics of the difference

$$\Delta_\tau X \equiv X(t + \tau) - X(t)$$

depend on the *lag* τ but not on the absolute origin t . Invariance under translations implies that possible long-term drifts have been removed, i.e.

$$\langle \Delta_\tau X \rangle_t = 0, \quad (2.22)$$

where $\langle \cdot \rangle_t$ denotes an average over t .

The variogram

A natural way to quantify the typical amplitude of fluctuations over a scale τ is to look at the mean square increment, known as the *variogram*:

$$V(\tau) \equiv \langle [X(t + \tau) - X(t)]^2 \rangle_t. \quad (2.23)$$

Because the mean increment vanishes by (2.22), $V(\tau)$ plays the role of a lag-dependent variance, whose square root $\sigma(\tau) = \sqrt{V(\tau)}$ is a natural measure of how much the signal $X(t)$ changes over time.

A familiar benchmark is standard Brownian motion, for which the variogram grows linearly,

$$V_{\text{BM}}(\tau) = 2D |\tau|, \quad (2.24)$$

with D the diffusion coefficient. Three important alternatives can be identified by comparing $V(\tau)$ to $|\tau|$:

- *Superdiffusion* (persistent walks): $V(\tau)$ increases *faster* than $|\tau|$; increment signs are positively correlated.
- *Normal diffusion*: growth is linear, cf.(2.24).
- *Subdiffusion* (anti-persistent walks): $V(\tau)$ grows *slower* than $|\tau|$; successive increments tend to compensate each other.

If the process is subject to a restoring force, $V(\tau)$ can saturate at large lags. The Ornstein–Uhlenbeck process,

$$\dot{X} = -\Omega X + \eta(t), \quad \langle \eta(t)\eta(t') \rangle = \sigma^2 \delta(t - t'), \quad (2.25)$$

provides a textbook illustration. When $\Omega t \gg 1$ the dynamics becomes stationary⁵; its variogram reads

$$V_{OU}(\tau) = \frac{\sigma^2}{\Omega} (1 - e^{-\Omega|\tau|}), \quad (2.26)$$

asymptotically converging towards a finite plateau $V_\infty = \sigma^2/\Omega$.

5: As opposed to “stationary increments” only.

From variogram to correlogram

Whenever (2.26) converges, the (auto)covariance of the *level* returns,

$$C(\tau) = \langle X(t + \tau) X(t) \rangle_t, \quad (2.27)$$

is well defined. A simple identity links $C(\tau)$ to the variogram:

$$C(\tau) = \frac{1}{2} [V(+\infty) - V(\tau)]. \quad (2.28)$$

Equivalently,

$$V(\tau) = 2[C(0) - C(\tau)]. \quad (2.29)$$

For unbounded processes such as Brownian motion, $V(+\infty)$ diverges and the right-hand side is ill defined—hence $C(\tau)$ does not exist even though $V(\tau)$ does. In such cases one may instead examine the covariance of the *increment or velocity*,

$$C_v(\tau) = \langle \dot{X}(t + \tau) \dot{X}(t) \rangle_t = C_v(-\tau), \quad (2.30)$$

which remains finite. The variogram and C_v are connected through the relation

$$C_v(\tau) = \frac{1}{2} \frac{d^2}{d\tau^2} V(\tau). \quad (2.31)$$

2.4. From a stochastic partial differential equation to the correlogram

To conclude this chapter we show, on a concrete example, how one *actually* computes a correlogram. The exercise is short yet central, because the same Fourier–space machinery will be reused throughout Part-I of the thesis. All of the following calculations are derived in details in chapter 4.

Step 1: the model. Consider a scalar field $\psi(\mathbf{r}, t)$ obeying

$$\partial_t \psi = D \Delta \psi - \kappa \psi + \eta(\mathbf{r}, t) + \xi(\mathbf{r}), \quad (2.32)$$

where

D diffusion constant,

κ mean-reversion rate,

η Ornstein-Uhlenbeck noise: $\langle \eta(\mathbf{r}, t) \eta(\mathbf{r}', t') \rangle = \frac{A^2}{T a^2} e^{-|t-t'|/T} g_a(|\mathbf{r} - \mathbf{r}'|)$,

ξ frozen (static) disorder: $\langle \xi(\mathbf{r}) \xi(\mathbf{r}') \rangle = \frac{\Sigma^2}{a^2} g_a(|\mathbf{r} - \mathbf{r}'|)$.

The bell-shaped kernel g_a has range a and normalisation $2\pi \int_0^\infty dr r g_a(r) = a^2$. At distances $\ell \gg a$ one may replace g_a by a Dirac delta function.

Step 2: Fourier solution. Fourier transforming (2.32) gives the linear ODE

$$\partial_t \psi_{\mathbf{k}} = -(Dk^2 + \kappa) \psi_{\mathbf{k}} + \eta_{\mathbf{k}} + \xi_{\mathbf{k}},$$

whose stationary solution reads

$$\psi_{\mathbf{k}}(t) = \int_{-\infty}^t e^{-(Dk^2 + \kappa)(t-\tau)} [\eta_{\mathbf{k}}(\tau) + \xi_{\mathbf{k}}] d\tau.$$

Step 3: two-point spectrum. Because η and ξ are independent, the correlation $\langle \psi_{\mathbf{k}}(t) \psi_{-\mathbf{k}}(t') \rangle$ splits into two additive terms. In the long-time limit one finds

$$\langle \psi_{\mathbf{k}}(t) \psi_{-\mathbf{k}}(t') \rangle = \underbrace{\frac{A^2}{2T} \frac{e^{-|t-t'|/T} - e^{-(Dk^2 + \kappa)|t-t'|} / [T(Dk^2 + \kappa)]}{(Dk^2 + \kappa)^2 - T^{-2}}}_{\text{fast noise } \eta} + \underbrace{\frac{\Sigma^2}{(Dk^2 + \kappa)^2}}_{\text{static noise } \xi \text{ (2.32*)}}.$$

Step 4: real-space correlogram. Set $t = t'$ to focus on spatial statistics and inverse-transform only over \mathbf{k} . Isotropy allows a polar change of coordinates so that the angular integral produces a Bessel function, $(2\pi)^{-1} \int_0^{2\pi} d\theta e^{ik\ell \cos \theta} = J_0(k\ell)$. Keeping the physically relevant window $1/\ell^* \ll k \ll 1/a$ and neglecting subdominant terms yields

$$\begin{aligned} \eta\text{-part: } C_\eta(\ell) &= \frac{A^2}{4\pi D} \int_{1/\ell^*}^{1/a} \frac{dk}{k} J_0(k\ell) \approx -\frac{A^2}{4\pi D} \log \frac{\ell}{\ell^*} + K(\ell), \\ \xi\text{-part: } C_\xi(\ell) &= \frac{\Sigma^2}{2\pi D^2} \int_{1/\ell^*}^{1/a} \frac{dk}{k^3} J_0(k\ell) \approx \frac{\Sigma^2}{8\pi D^2} \ell^2 \log \frac{\ell}{\ell^*} + K'(\ell). \end{aligned}$$

Here $K(\ell)$ and $K'(\ell)$ are bounded remainders. Finally, because the spatial variogram is $\mathbb{V}(\ell) = 2[C(0) - C(\ell)]$, we obtain the leading behavior

$$\mathbb{V}(\ell) \simeq \frac{A^2}{2\pi D} \log \frac{\ell}{\ell^*} - \frac{\Sigma^2}{4\pi D^2} \ell^2 \log \frac{\ell}{\ell^*} + \text{const.} \quad (a \ll \ell \ll \ell^*). \quad (2.33)$$

Hence, equation (2.33) demonstrates how logarithmic scalings emerge from a simple partial differential equation, and we will come back to

this model in more details in chapter 4.

In summary, the variogram quantifies the magnitude of signal variation across scales and serves as a foundation for the multiscale diagnostics developed later in this thesis.

Key Takeaways

1. **Microscopic–Macroscopic Bridge:** The Langevin formalism combines external forces and thermal noise in a single stochastic differential equation, providing a direct link between particle-level dynamics and continuum descriptions.
2. **From Hops to PDEs:** Starting with a discrete Master equation and performing a Kramers–Moyal expansion shows step-by-step how macroscopic partial differential equations are inherited from local interactions.
3. **Variogram:** The variogram measures the amplitude of fluctuations across scales, allowing to diagnose specific behavior such as diffusion, persistence, or confinement in real data.

**PART I: HYDRODYNAMIC MODELING OF
SOCIO-ECONOMIC SYSTEMS:
SEGREGATION DYNAMICS AND
HOUSING-PRICE DIFFUSION**

An out-of-equilibrium Sakoda-Schelling model

3.

Ses yeux s'attachèrent presque avidement entre la colonne de la place Vendôme et le dôme des Invalides, là où vivait ce beau monde dans lequel il avait voulu pénétrer. Il lança sur cette ruche bourdonnante un regard qui semblait par avance en pomper le miel, et dit ces mots grandioses : «À nous deux maintenant !» Il revint à pied rue d'Artois, et alla dîner chez madame de Nucingen.
Le Père Goriot-Honoré de Balzac, 1835.

This chapter is adapted from Zakine et al. (2024) [1], *Socioeconomic agents as active matter in nonequilibrium Sakoda-Schelling models*. The original publication was a collaborative effort with R. Zakine, J. Garnier-Brun and M. Benzaquen. My contributions include assisting in the initial framing of the problem, in the microscopic simulations, in the linear stability analysis and the model's extensions. Sections mainly due to the work of R. Zakine and J. Garnier-Brun have been removed.

3.1. Introduction

Will a collective system in which individuals share the same common goal ever reach an optimal state? This nontrivial question is at the very core of strong debates among economists, notably because the notion of “optimal state” is intrinsically political and most often ill-defined. Despite the common idea that a system made of agents individualistically improving their outcome will spontaneously converge by the action of the “invisible hand” to an optimal collective state, simple models have been shown to contradict this belief [93, 94, 95]. A well documented example of such system is the celebrated Schelling model [14]. The latter can be considered to be a variant of the model previously¹ introduced by Sakoda [15], and will thus be referred henceforth as the Sakoda-Schelling model. To understand some aspects of urban segregation in post-WWII American cities, and more widely of urban and social dynamics, both authors proposed simple lattice models of idealized cities. Each site, representing an accommodation, can be empty or occupied by an agent belonging to one of two sub-populations in the system. Interestingly, Schelling observed that when introducing a slight preference for agents to be surrounded by neighbors of their own group, the system evolves towards configurations with completely segregated regions. While in fact not very well suited to explain urban segregation, which is intimately related to past and present public policies rather than self-organization [97, 98], the model illustrates how the *micromotives* of the agents may lead to unanticipated *macrobehavior* [17].

1: To be perfectly precise, the first mention of Sakoda's model can be traced back to his unpublished PhD thesis completed in 1946, while Schelling's work can be found in a 1969 working paper [96]. In any case, there is no reason to think either author took inspiration from the other, the objective of the papers being clearly quite different.

Along the years, the Sakoda-Schelling model has attracted further attention of statistical physicists [99, 100, 101, 102], due to its simple microscopic rules, its paradoxical macroscopic consequences and its unconventional non-local particle moves. To the usual end of bridging the gap from *micro* to *macro*, mappings onto equilibrium systems were suggested [103], but with limited analytical results. To gain a more in-depth understanding of the mechanism through which individual choices may lead to sub-optimal collective outcomes, Grauwin *et al.* introduced a modified version of the Schelling model with a single type of agent occupying a lattice divided in pre-defined neighborhoods, or blocks [39]. In this occupation model, the agents now base their decisions on the neighborhood density, which is identical for all the agents in a given block. This fixed neighborhood structure then allows to describe analytically the steady state as the minimizer of a free energy, and to recover a nontrivial phase with suboptimal jam-packed neighborhoods. Subsequent works have then explored variations of these different models focusing on the effect of altruistic agents [104], dynamics close to criticality [105, 106, 107] or habit formation [108].

Even in the seemingly simpler occupation problem of Grauwin *et al.* [39], several questions persist, both from the socioeconomic and statistical physics perspectives. In particular, the role of the specific decision rule and the precise nature of neighborhoods on the phenomenology of the model remain unclear. Indeed, to allow for the standard techniques of statistical mechanics to be applicable, the choice of the neighborhoods and the dynamics is very constrained, see [7]. As will be discussed in detail, most non-trivial decision rules lead the system out of thermodynamic equilibrium, requiring calculations that are not always readily tractable. As it is extremely difficult to empirically determine how economic agents actually make decisions, the physics-inspired theoretical analysis of toy models has a significant part to play, in particular to determine the robustness of qualitative findings to specific modeling choices. Besides, as argued in [7] and by some of us in [109], the intrinsically individualistic nature of agent-specific moves in socioeconomic models means that the description of collective behaviors as the minimization of some global energy is often not possible. Understanding simple out-of-equilibrium dynamics as those that arise from the decision rules presented here is therefore also necessary from the methodological point of view.

The purpose of this chapter is to assess, within a general Sakoda-Schelling like occupation model, whether and how the sub-optimal concentration of agents in overly dense regions still occurs out of equilibrium. Most importantly, we relax the assumption of taking a specific decision rule, and no longer require pre-defined block neighborhoods as in [39]. The resulting heterogeneity of interactions in our model then requires the use of out-of-equilibrium statistical mechanics techniques, the progress of which in the last decade can be credited to active matter theory. Overall, we find that the phenomenology of the model is largely unaffected by its nonequilibrium nature, suggesting that the tendency of agents to aggregate sub-optimally is robust to large classes of decision rules. This being said, our analysis highlights interesting theoretical subtleties, notably related to the non-monotonicity of the utility functions considered, that may, in turn,

contribute to the understanding of other complex physical systems.

This chapter is organized as follows. In Sec. 3.2 we introduce a Schelling-like occupation model, in which we keep the utility function and decision rule as general as possible to allow for nonequilibrium dynamics. We then perform a numerical analysis of the model. In Sec. 3.3 we present a mean-field description of the dynamics, and determine the region in parameter space where condensation necessarily occurs. In Sec. 3.4 we show how the dynamics can be mapped on the Active Model B [110], which is considered to be the natural nonequilibrium extension of the Cahn-Hilliard field relaxation [111]. This mapping notably allows to compute the phase densities of the concentrated states. In Sec. 3.5 we propose some relevant generalizations of the model, namely with two different populations and a housing market. Finally, in Sec. 3.6 we discuss the implications of our study and conclude.

3.2. A Sakoda-Schelling occupation model

Setup

Consider a city structured as a two-dimensional rectangular lattice composed of $M = L_x \times L_y$ sites (or houses). Each site can be occupied by at most one of the $N (\leq M)$ agents living in this city. On each site of coordinate $\mathbf{r} = (i, j)$, the occupation field n takes the value $n(\mathbf{r}) = 1$ if the site is occupied, $n(\mathbf{r}) = 0$ if it is vacant. It is assumed that each agent k wants to maximize their own utility u_k , which depends on the local density of agents around them. Typically, it is natural to think that people like to gather in relatively dense areas to benefit from the city life, but not too dense as crowding might degrade the quality of life. Agents estimate the local density by averaging the occupation field with a probability-density-function kernel G_σ , where σ stands for the interaction range. The kernel is assumed to be isotropic and identical for all agents. The smoothed occupation field \tilde{n} at site \mathbf{r} is thus given by the discrete convolution

$$\tilde{n}(\mathbf{r}) = \sum_{\mathbf{r}'} G_\sigma(\mathbf{r} - \mathbf{r}') n(\mathbf{r}'). \quad (3.1)$$

At each time step, an agent k can decide to move out from their occupied site \mathbf{r}_k and to settle on a new, randomly chosen, empty site \mathbf{r}'_k where the utility $u[\tilde{n}(\mathbf{r}'_k)]$ – quantifying the agent’s satisfaction – might exceed their previous utility $u[\tilde{n}(\mathbf{r}_k)]$. We assume that the decision to move to the new site is a function of the utility difference $\Delta u_k \equiv u[\tilde{n}(\mathbf{r}'_k)] - u[\tilde{n}(\mathbf{r}_k)]$. While the very existence of the utility function is debatable from a behavioral standpoint [112], classical economics has traditionally taken agents to be strict utility maximizers, meaning the move will be accepted if $\Delta u_k > 0$ and rejected otherwise. In order to mitigate this assumption, a common approach is to introduce a stochastic decision rule of the form

$$\mathbb{P}(\mathbf{r}_k \rightarrow \mathbf{r}'_k) = f_\Gamma(\Delta u_k), \quad (3.2)$$

where the function f_Γ is larger than $\frac{1}{2}$ whenever $\Delta u_k > 0$. Typically,

f_Γ is a positive and monotonic function of the utility difference, with $\lim_{x \rightarrow -\infty} f_\Gamma(x) = 0$ and $\lim_{x \rightarrow +\infty} f_\Gamma(x) = 1$ [24]. The parameter $\Gamma \geq 0$, known as the *intensity of choice*², or simply the *rationality*, quantifies the propensity of agents to go for strict utility maximizing. In particular, $\Gamma \rightarrow 0$ corresponds to random decision making, while $\Gamma \rightarrow +\infty$ means perfectly rational agents.

2: In the introductory chapter of this thesis, we denoted this parameter as β , underscoring its connection with the inverse temperature in thermodynamics.

In reality, the specific shape of the function f_Γ is unknown. In the socio-economics literature, it is most of the time taken as the logistic function

$$f_\Gamma(x) = \frac{1}{1 + e^{-\Gamma x}}, \quad (3.3)$$

defining the so-called *logit rule* [23, 24]. The various reasons and justifications of this decision rule are discussed and summarized in [7]. In a nutshell, it can be motivated axiomatically [23], or by the fact that f_Γ is a maximum entropy distribution and therefore optimizes an exploration-exploitation tradeoff when the cost associated with information scales as $1/\Gamma$ [113, 114]. As empirical evidence supporting this choice remains extremely scarce, its popularity is in reality largely motivated by convenience [112]. Indeed, many calculations are made possible thanks to the fact that it preserves detailed balance with respect to the Gibbs-Boltzmann measure in the particular case where agents' utility change also coincides with a global utility difference [115]. In this context, $T \equiv 1/\Gamma$ can naturally be interpreted as the temperature, or "social temperature", of the system. In the following, the function f_Γ will be left unspecified, unless stated otherwise. In the Monte Carlo simulations we will notably use the logit rule for simplicity.

The last ingredient to specify is the utility function u of the agents. As stated above, we assume that the utility depends on the locally smoothed occupation \tilde{n} only, and that it is non monotonic. As in Ref. [39], we assume that the utility is maximal for some density $\rho^* \geq \frac{1}{2}$. We specifically choose for the simulations

$$u(x) = -|x - \rho^*|^\alpha, \quad (3.4)$$

with $\alpha > 0$, see Fig. 3.1(a), but theoretical computations below will keep u unspecified.

In or out of equilibrium?

As mentioned above, an often unspoken motivation for the use of the logit rule in the modeling of socioeconomic systems is that it may satisfy detailed balance. Indeed, as described by Grauwin *et al.* [39, 40] or in [7] in a more general setting, if one manages to find a system-wide energy-like function \mathcal{H} such that

$$\begin{aligned} \Delta u_k &= \mathcal{H}([n(\mathbf{r})], n(\mathbf{r}'_k) = 1, n(\mathbf{r}_k) = 0]) \\ &\quad - \mathcal{H}([n(\mathbf{r})], n(\mathbf{r}'_k) = 0, n(\mathbf{r}_k) = 1]), \end{aligned} \quad (3.5)$$

then the usual tools of equilibrium statistical mechanics can be used. The steady-state distribution of agents is notably identified as the minimum of the free energy, which is a Lyapunov function of the dynamics prescribed by the logit rule.




Figure 3.1: (a) Utility function $u(\rho) = -|\rho - \rho^*|^\alpha$ for $\rho^* = 0.5$, $\alpha = \{0.5, 1, 2\}$. Panels (b), (c) and (d) show snapshots of the stationary state for these different utility functions, starting from the same homogeneous profile at $\rho_0 = 0.5$. Here $\Gamma = 100$, $\sigma = 3$ and $L_x = L_y = 100$. The stationary density ρ_d in the dense phase is $\rho_d = 0.575(5)$ for $\alpha = 0.5$ in (b), $\rho_d = 0.575(5)$ for $\alpha = 1$ in (c), and $\rho_d = 0.585(5)$ for $\alpha = 2$ in (d). These bulk densities are all significantly higher than the density ρ^* for which agents maximize their utility. Note the accumulation of agents at the edge of the empty domain in (b) and (c), see Sec. 3.4.

At the agent level, the existence of such a global quantity is usually the symptom of either altruistic individuals (that voluntarily maximize some collective satisfaction) or of a central planner (that constructs individual rewards towards a collective objective). Outside of these two cases, the existence of a free energy when agents are individualistic is in fact restricted to a limited number of carefully chosen models (see [109] for a related discussion in the context of microeconomics). In the literature of Schelling-like models, taking a city divided in neighborhoods or blocks [39], where agents share the same utility, yields such a free energy description (which is importantly not a simple aggregation of individual utilities). In our model, however, this is no longer true.

An explicit proof that our model violates Kolmogorov's criterion [116]—largely following the arguments of J. Garnier-Brun—is provided in our original article [1].

Microscopic simulations

Having established the out-of-equilibrium nature of our model, we start by performing numerical simulations to assess whether the concentration of agents in overly dense regions is generic and robust to different shapes of the utility function. Here, all numerical simulations are performed on a two-dimensional grid with periodic boundary conditions. The utility is maximal for $\rho^* = 1/2$. For the sake of simplicity, here we use the logit decision rule and a truncated Gaussian kernel

$$G_\sigma(\mathbf{r}) = \begin{cases} \frac{1}{N_\sigma} e^{-\frac{1}{2\sigma^2}\|\mathbf{r}\|^2}, & \text{if } \|\mathbf{r}\| \leq 4\sigma, \\ 0, & \text{otherwise,} \end{cases} \quad (3.6)$$

where N_σ enforces the normalization of the kernel.




Figure 3.2.: (a) Typical dense domain size $L_d(t)$ during coarsening as a function of time t . A unit of time is defined as N Monte Carlo steps, where N is the number of agents. $L_d(t)$ is averaged over 5 independent simulations. (b), (c), (d) and (e) show snapshots at different times. Starting from a disordered configuration, we quench the system at low temperature, or high rationality, corresponding to $T \simeq T_c/6$. Parameters: $L_x = L_y = 600$, $\rho_0 = 0.3$, $\sigma = 1$, $\alpha = 3/2$, $T = 0.01$.

Phase separation

For large system size $L_x, L_y \gg \sigma$, we explore the behavior for different global densities $\rho_0 = N/(L_x L_y)$ and for various rationality parameters. Numerical results are qualitatively similar for all the values of α we tested, ranging from $\alpha = 0.5$ to $\alpha = 2$, see Fig. 3.1. The phenomenology can be summarized as follows. When rationality is low ($\Gamma \rightarrow 0$, $T \rightarrow \infty$), the stationary state remains homogeneous because agents settle at random. When rationality is high, agents may aggregate in dense clusters, which can surprisingly be more crowded than what agents' utilities prescribe. This was already discussed in [39] where the authors point out that the homogeneous state is actually an unstable Nash equilibrium, even though all agents maximize their utility. The destabilization occurs as one agent randomly moves to another region (with no regard to the effect it may have on the other agents utilities), which decreases the average density at their original site and increases the average density where they settle. Agents in the lower-density region will eventually move to gain the utility they lost when their neighbors moved out. This dynamics will eventually empty some regions, in which agent's return becomes statistically less and less probable. The final state, where a dense phase and an empty phase coexist, is a stable Nash equilibrium.

One can quantify the condensation dynamics when starting from the homogeneous state and taking high rationality. The system undergoes a spinodal decomposition where dense clusters grow and merge until there is one large dense cluster only, as shown in Fig. 3.2. The final cluster topology ultimately depends on noise realization and on the box dimensions. We measure the cluster size $L_d(t)$ as a function of time t using the radial structure factor (see App. A.1). We find $L_d(t) \sim t^{1/z}$, with the dynamical exponent $z \in [2, 3]$, reminiscent of the coarsening exponent observed in a 2D Ising system with long-range Kawasaki dynamics [117, 118, 119, 120]. Interestingly, and consistent with the findings of [120] in the low temperature region, our results suggest an exponent closer to the local Kawasaki dynamics result $z = 3$ (see Fig. 3.2(a)), despite long-range particle displacements.

Critical point and critical exponents

The complete phase separation that occurs when rationality is high indicates the use of the order parameter $m \equiv \rho_d - \rho_g$, where ρ_d, ρ_g are

the average densities of the dense and “gas” (dilute) phases, respectively. At the critical point (ρ_c, T_c) , we expect a second-order phase transition where m goes to 0 with power-law scaling

$$m \underset{\tau \rightarrow 0^+}{\sim} \tau^\beta, \quad (3.7)$$

where $\tau = (T_c - T)/T_c > 0$ defines the rescaled temperature difference, and β is the order-parameter critical exponent. Measuring the critical exponents allows one to determine to which universality class the system belongs to, providing precious information on the system behavior at large scales. Since simulations are carried out in finite systems, measuring the critical point with precision requires numerical ruse. We follow the approach that has been extensively used to measure critical exponents in systems undergoing a Motility-Induced Phase Separation (referred to as MIPS) [121, 122, 123], see App. A.2.

Simulations are performed in a rectangular domain of size $L_x \times L_y$, with $L_x = 3L_y$, with periodic boundary conditions to keep flat interfaces between a stripe of liquid (dense phase) and a stripe of gas (dilute phase). Starting with the dense phase in the center of the system, we track the center of mass such that we always compute the densities in the bulk of each phase. To compute the local density inside the bulk of each phase, we consider square boxes of size $\ell = L_y/2$, centered either in 0 in the gas bulk or centered in $L_x/2$ in the dense bulk (Fig. A.1). The local density in each box fluctuates and it is given by $\rho_b = N_b/\ell^2$ with N_b the number of agents in the box b in a given realization of the system. The distribution of the density in the system is thus bimodal for $T < T_c$ and unimodal when the system is homogeneous. Defining

$$\Delta\rho = \frac{N_b - \langle N_b \rangle}{\ell^2}, \quad (3.8)$$

where the $\langle \cdot \rangle$ stands for averaging on the four boxes and on independent realizations of the simulation, we compute the celebrated Binder cumulant [124, 125, 126]

$$Q_\ell(\Delta\rho, T) = \frac{\langle (\Delta\rho)^2 \rangle}{\langle (\Delta\rho)^4 \rangle}, \quad (3.9)$$

for a given box size ℓ and a given temperature T . For ℓ large enough, the curves $Q_\ell(T)$ all intersect in $T = T_c$ where the behavior of the system is universal. It is important to mention that the critical density is not known *a priori*. It has to be assessed beforehand to ensure that the system, as T changes, goes through the critical point, where the phase transition is of second-order type. To locate ρ_c , we compute the Binder cumulant at fixed temperature, close to the estimated critical point, for various densities ρ_0 . The critical density then corresponds to the maximal fluctuations of $\Delta\rho$, translated in a peak of the Binder cumulant, see Fig. 3.3(a). Once the critical point is precisely located, additional critical exponents can be measured. Notably, defining the susceptibility χ as

$$\chi \equiv \frac{\langle (N_b - \langle N_b \rangle)^2 \rangle}{\langle N_b \rangle} = \frac{\langle (\Delta\rho)^2 \rangle}{\langle N_b \rangle} \ell^4, \quad (3.10)$$

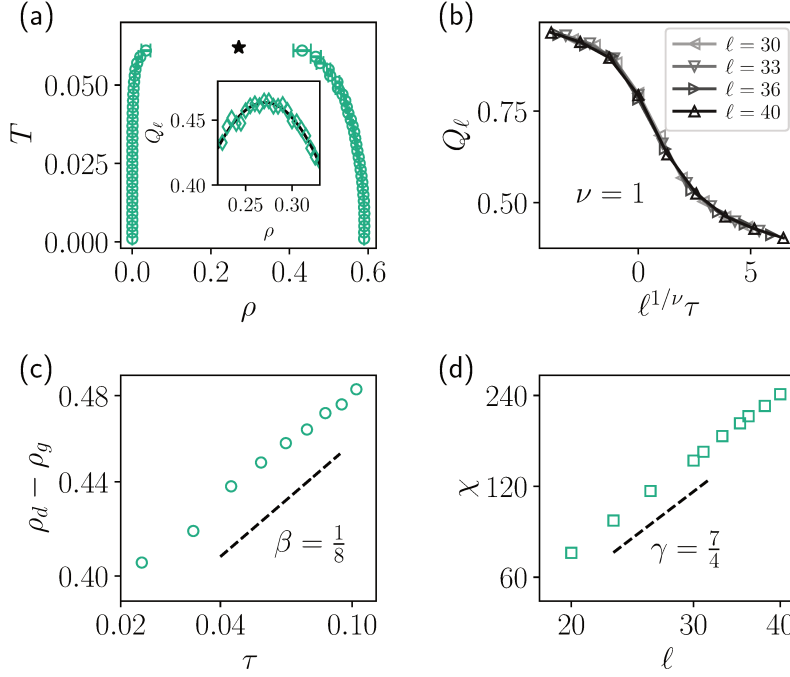


Figure 3.3: Numerical experiments for $\sigma = 1$, $\alpha = 3/2$. (a) Binodal densities measured for $L_x = 200$ and $L_y = 66$ ($\ell = 33$), inset showing the Binder cumulant as a function of the density and fitted (continuous line) to determine the critical density. (b), (c) and (d) show the numerical measurements of the critical exponents close to the critical point $(\rho_c, T_c) = (0.271, 0.0620)$ determined using various system sizes ranging from $\ell = 20$ to $\ell = 40$.

Model	ρ_c	T_c	β	γ
Ising 2D (exact)	0.5		0.125	1.75
$\alpha = 1/2$, $\sigma = 1$	0.309(5)	0.0983(5)	0.120(8)	1.71(5)
$\alpha = 3/2$, $\sigma = 1$	0.271(5)	0.0620(2)	0.119(5)	1.74(5)

Table 3.1: Critical density and exponents for nonequilibrium Sakoda-Schelling model for $\alpha = 1/2$ and $\alpha = 3/2$.

one obtains

$$\chi \sim \ell^{\gamma/\nu}, \quad \left. \frac{dQ_\ell}{d\tau} \right|_{\tau=0} \sim \ell^{1/\nu}, \quad (3.11)$$

at the critical point.

We report in Fig. 3.3 the various results on the critical point and on the critical exponent for $\sigma = 1$ and $\alpha = 3/2$. Using the Binder cumulant, one identifies the critical point at $\rho_c = 0.271(5)$ and $T_c = 0.0620(2)$, where the uncertainty on the last digit appears in the parentheses. The phase diagram in space (ρ, T) is shown in Fig. 3.3(a), the black star indicates the critical point and the circular markers show the densities of the coexisting phases: they define the *binodal* frontier. The exponent β is directly measured from the order parameter m as function of reduced temperature τ , at a fixed system size $L_x = 220$ [see Fig.A.2(c)]. From the Ising-2D ansatz, we check that $\nu = 1$ yields a neat collapse of the Binder cumulant, see Fig. A.2(b). The exponent γ is obtained by varying the system size at the critical temperature T_c and assuming $\nu = 1$ [see Fig.A.2(d)]. We report in Table 3.1 the values found for the critical exponents in the cases $\alpha = 3/2$ (Fig. 3.3) and $\alpha = 1/2$ (not shown here). They differ by less than 5% from the 2D-Ising static exponents. These results enjoin us to assert with a high degree of confidence that the model considered here belongs to the 2D-Ising universality class. Since the system is out of equilibrium and particle displacements can be of infinite range, recovering the Ising universality class is *a priori* nontrivial. However, finding other critical exponents would have been surprising since the ingredients at play are the ones

of the Ising model, namely, short-range and isotropic interactions, a homogeneous medium and a two state degree of freedom (sites are empty or occupied). This result must also be put into perspective with the recent debate on the universality class(es) of systems undergoing MIPS [121, 123, 122, 127, 128], and their associated active field theories [129, 130]. Notably here, our interaction kernel G_σ provides a so-called *quorum-sensing* interaction, like that found in assemblies of bacteria [131]. The particle dynamics is however quite different for bacteria and for our agents. The remaining of the chapter shall be devoted to establishing a quantitative relation between our Sakoda-Schelling occupation model and the field-theory descriptions of MIPS. The first step along this path is to formulate a mean-field approximation of our model.

3.3. Field theory and the local-move approximation

General description

The computation starts by writing the expectation of the occupation number $n_{r,s+1} \equiv n(r, s+1)$ of site r at time $s+1$, conditioned on the previous configuration $\{n_{r,s}\}$. Averaging over multiple realizations of noise and using a mean-field approximation in which all correlation functions factorize, one obtains

$$\begin{aligned} \langle n_{r,s+1} \rangle - \langle n_{r,s} \rangle &= (1 - \langle n_{r,s} \rangle) \sum_{r' \neq r} \langle n_{r',s} \rangle f_\Gamma(\Delta u_{r' \rightarrow r}^s) \\ &\quad - \langle n_{r,s} \rangle \sum_{r' \neq r} (1 - \langle n_{r',s} \rangle) f_\Gamma(\Delta u_{r \rightarrow r'}^s), \end{aligned} \quad (3.12)$$

where $\Delta u_{r \rightarrow r'}^s \equiv u(\langle \tilde{n}_{r',s} \rangle) - u(\langle \tilde{n}_{r,s} \rangle)$. For convenience, we take the continuous time and continuous space limit, following the common procedure to obtain a mean-field description of exclusion processes on lattices (see e.g. [132]). The average occupation number $\langle n \rangle$ is now described by the density ρ , while the spatially smoothed average occupation number $\langle \tilde{n} \rangle$ is described by the field $\phi \equiv G_\sigma * \rho$. The master equation for the occupation probability then takes the form of a noiseless hydrodynamic equation, in our case:

$$\begin{aligned} \partial_t \rho(x, t) &= [1 - \rho(x, t)] \int dy \rho(y, t) w_\Gamma([\phi], y, x, t) \\ &\quad - \rho(x, t) \int dy [1 - \rho(y, t)] w_\Gamma([\phi], x, y, t), \end{aligned} \quad (3.13)$$

with the transition rate from y to x explicitly given by

$$w_\Gamma([\phi], y, x, t) = \omega f_\Gamma(u(\phi(x, t)) - u(\phi(y, t))), \quad (3.14)$$

where ω is homogeneous to an inverse time scale and f_Γ is left unspecified. Equation (3.13) is valid in any dimension, but, for simplicity, we will work out the mean-field computations in dimension 1 in space. This can be justified *a posteriori* when we compare the mean-field (MF) to the Monte Carlo (MC) simulations. Let us also mention that the dimension does not play a role in determining the phase densities in the

steady state of coarse-grained field theories (Allen-Cahn [133], Cahn-Hilliard [134], etc.).

Integrating Eq. (3.13) over space, one immediately sees that the total density $\int \rho$ is conserved. One can also check that in the very specific case where $u(\phi)$ is linear in ϕ , one can build a free-energy functional that is a Lyapunov function of the non-local MF dynamics, ensuring a convergence towards local minima and preventing limit cycles and oscillatory dynamics. This is a natural consequence of the fact that detailed balance is satisfied at microscopic level. For more information about how we construct this free energy and show that the dynamics is relaxational, (mainly due to R. Zakine) see [1].

Linear stability analysis

In the general case, we would like to understand how the homogeneous state becomes unstable. To do so, we consider a small perturbation around the homogeneous state: $\rho(x, t) = \rho_0 + \rho_1(x, t)$, with ρ_1 the perturbation. By linearity of the convolution, one has $\phi(x, t) = \rho_0 + \phi_1(x, t)$, with $\phi_1 \equiv G_\sigma * \rho_1$. A Taylor expansion of Eq. (3.13) combined with mass conservation (i.e $\int_D \rho_1 = \int_D \phi_1 = 0$, where D is the full domain), finally yields:

$$\begin{aligned} \partial_t \rho_1(x, t) = & 2\Omega \rho_0(1 - \rho_0) f'(0) u'(\rho_0) \phi_1(x, t) \\ & - \Omega f(0) \rho_1(x, t), \end{aligned} \quad (3.15)$$

with Ω the full domain size. Defining the Fourier transform for any field h as $\hat{h}(k) = \int dx e^{-ikx} h(x)$, one obtains

$$\partial_t \hat{\rho}_1(k, t) = \Lambda(k) \hat{\rho}_1(k, t), \quad (3.16)$$

$$\Lambda(k) = \Omega f_\Gamma(0) \left(2\rho_0(1 - \rho_0) \frac{f'_\Gamma(0)}{f_\Gamma(0)} u'(\rho_0) \hat{G}_\sigma(k) - 1 \right). \quad (3.17)$$

This last equation shows that the homogeneous state is unstable if there exists a mode k^* such that

$$2\rho_0(1 - \rho_0) \frac{f'_\Gamma(0)}{f_\Gamma(0)} u'(\rho_0) \hat{G}_\sigma(k^*) > 1. \quad (3.18)$$

The manifold for which the inequality becomes an equality defines the spinodal in the phase diagram (ρ_0, \cdot) . In particular, for any monotonically decreasing kernel $G_\sigma(|x|) \in L_2(\mathbb{R})$, one has $\hat{G}_\sigma(0) > |\hat{G}_\sigma(k)|$, such that for large system size, the stability of the homogeneous state is given by the stability of modes $k \rightarrow 0$, and the spinodal is thus defined by the equation

$$2\rho_0(1 - \rho_0) \frac{f'_\Gamma(0)}{f_\Gamma(0)} u'(\rho_0) = 1. \quad (3.19)$$

Note that this criterion is generic as it only depends on the decision rule through $f_\Gamma(0)$ and $f'_\Gamma(0)$. The simulations also reveal the existence of a bistable region in the vicinity of this spinodal. This is the binodal region, where hysteresis and bistability can notably be observed, and which can be fully characterized in the case of an equilibrium system [101]. Here however, there is a priori no free energy one can rely

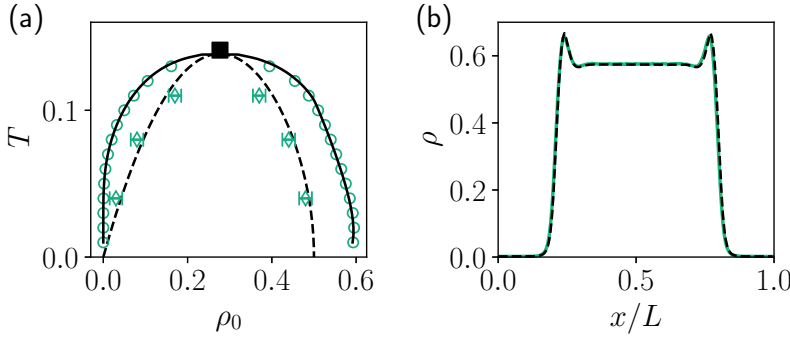


Figure 3.4.: Comparison between Monte Carlo simulations and mean-field results for $\alpha = 3/2$, $\sigma = 7$ and $L_x = 200$, $L_y = 66$ ($\ell = 33$). (a) Phase diagram in the (ρ_0, T) plane. The mean-field binodal (continuous black line) is given by measuring the densities of the bulk of each plateau in a phase separated state. The circles are the bulk averaged densities in Monte Carlo (MC) simulations. The dashed black line represent the mean-field spinodal, which is obtained analytically from the linear stability analysis (see Eq. (3.18)), with $\hat{G}_\sigma(k_1) = e^{-\sigma^2 k_1^2/2}$ and $k_1 = 2\pi/L_x$. The diamonds indicate the loss of stability of the homogeneous state in the MC simulations. The black square is the critical point for $\sigma/L \rightarrow 0$. (b) Averaged density profile $\rho(x)$ from MC simulations (continuous green line) for $\rho_0 = 0.35$, $T = 0.05$. The dashed black line is the stationary solution of the mean-field equation Eq. (3.13) for the same parameters, solved on a grid of step size 1 with a Euler explicit scheme.

on to describe the nucleation scenario and to obtain the densities of the phase-separated state.

Comparison to microscopic simulations

The MF prediction is expected to be accurate for systems with high connectivity, which here corresponds to large σ . In the following, we shall take the limit $L \rightarrow +\infty$, $\sigma \rightarrow +\infty$ with $\sigma/L \rightarrow 0$ to obtain mean-field predictions that are independent of both σ and L , and perform numerical simulations as close as possible to this scaling regime.

The first analytical prediction of the MF description is the spinodal, that determines the onset of instability of the homogeneous state, see Eq. (3.19). The spinodal is the dashed line in the (ρ, T) phase diagram in Fig. 3.4(a). To check the prediction, we start in the MC simulations from a uniformly distributed configuration of agents for three different values of temperature, $T = 0.04, 0.08, 0.11$, and we detect the frontier across which the homogeneous profile either coarsens, or needs a nucleation event to converge to the separated state. This frontier is marked with the diamonds, which agrees with the MF prediction.

Second, the MF dynamical Eq. (3.13) can be solved numerically with an Euler explicit scheme. From the numerical solution, one obtains the densities of the bulk of each phase when a phase separation occurs: these densities define the binodal, the continuous line in Fig. 3.4(a). These MF phase densities are perfectly recovered by the MC simulations (circles). In addition, one can compare the steady-state average density profile from MC simulations to the mean-field stationary density, which superimpose almost exactly, see Fig. 3.4(b).

As previously stated, the MF predictions fail for small values of σ . The phase diagram in Fig. 3.3(a) is for instance obtained for $\sigma = 1$, and indeed strongly differs from the MF solution. For $\sigma = 1$, we notably iden-

tify the critical point at $(\rho_c, T_c) = (0.271, 0.0620)$, whereas the MF predicts $(\rho_c, T_c)_{MF} = (0.2763, 0.1418)$, where, as expected, $T_c^{\sigma=1} < T_c^{MF}$.

Local-move approximation

To make progress into the identification of a possible effective free energy functional, it may be convenient to consider slightly modified dynamics where jumps are now only authorized in the direct neighborhood of the agents. Indeed, considering an evolution enforcing a *local* mass conservation will allow for more familiar partial differential equations (PDEs) and field theoretic approaches on conserved scalar fields. Here, the absence of macroscopic density currents in the steady state, both in MC simulations and in the MF solution suggests that the system generically converges to a stationary stable fixed point, where the details of the dynamics become inconsequential. In addition, when the majority of agents have aggregated in a single dense cluster in the steady state, it is unlikely that they would perform moves outside of the bulk, in low-density regions, since the utility there is minimal. The local-move approximation, as it strongly simplifies the description, thus appears natural.³

Following the Taylor expansion outlined in App. A.3, the local mean-field dynamics is given by

$$\partial_t \rho = f_\Gamma(0) \partial_x^2 \rho - 2f'_\Gamma(0) \partial_x [\rho(1 - \rho) \partial_x u], \quad (3.20)$$

which can be rewritten as the canonical equation for the mass-conserving dynamics

$$\partial_t \rho = \partial_x [M[\rho] \partial_x \mu([\rho], x)], \quad (3.21)$$

with the mobility operator $M[\rho] = \rho(1 - \rho)$, stemming from the non-overlapping nature of the agents, and with the chemical potential $\mu = \mu_{ent.} + \mu_{util.}$ where

$$\mu_{ent.} = f_\Gamma(0) \log \left(\frac{\rho}{1 - \rho} \right) \quad (3.22)$$

$$\mu_{util.} = -2f'_\Gamma(0) u[\phi(x)]. \quad (3.23)$$

The first contribution to the chemical potential $\mu_{ent.}$ is purely local and accounts for entropy in the system where agents cannot overlap. The second contribution $\mu_{util.}$ encodes the drive from agents' utility. This term exhibits non-locality with respect to the field ρ , and as a consequence, cannot be expressed as a functional derivative of any free energy, in general [135, 136, 137]. However, in the particular case of a linear utility in ϕ , one again recovers that $\mu_{util.} + \mu_{ent.}$ can be written as the functional derivative of the free energy \mathcal{F} given [1] and, as a consequence, the dynamics (3.21) becomes a gradient descent [138]. Let us emphasize that, here again, the decision rule is kept general, and that the entire local dynamics only depend on it through $f_\Gamma(0)$ and $f'_\Gamma(0)$.

Performing the linear stability analysis on the dynamics with local moves (see App. A.3), we find that the criterion for the homogeneous solution to be unstable is identical to that given in Eq. (3.19), when moves are global. Also, the stationary density profiles computed ei-

3: Dynamically, the coarsening exponent $z \simeq 3$ displayed in Fig. 3.2(a), and which is also observed in a Cahn-Hilliard relaxation dynamics can also be invoked to support the idea of local moves.

ther with the local, or with the non-local MF PDEs for the same parameters are identical, as shown in App. A.4. Both these observations therefore allow us to confirm the relevance of the local-move approximation to characterize the system in the long-time limit.

Finally, note that the local hydrodynamic equations can also be obtained using the path integral approach on a lattice [139], which, in passing, provides the fluctuating hydrodynamics:

$$\partial_t \rho = \partial_x \left[\rho(1 - \rho) \partial_x \mu([\rho], x) + \sqrt{\rho(1 - \rho)} \xi \right], \quad (3.24)$$

where $\xi(x, t)$ is a Gaussian white noise with zero mean and with $\langle \xi(x, t) \xi(x', t') \rangle = 2f_\Gamma(0) \delta(t - t') \delta(x - x')$. We then remark that when the utility is linear, the stochastic field evolution describes a complete equilibrium dynamics, irrespective of the choice of the decision rule: A rule that breaks detailed balance at the microscopic level can still lead to an equilibrium field theory after coarse graining. Similar findings had been pinpointed in active matter models [140, 136]. While not central to the present work, the fluctuations can be studied in more detail, providing information on the nucleation scenarii and on transition paths between macroscopic states for instance [141, 142, 143, 144]. The study of the associated functional Fokker-Planck equation using the tools described in [145, 146] may also be an interesting perspective for future works. In the case of non-local moves, the formalism from [139] cannot be straightforwardly adapted, since the local gradient expansion of the jump rates in the action breaks down. Establishing an appropriate fluctuating hydrodynamic description in the case of non-local dynamics is therefore an open problem.

3.4. Generalized-thermodynamic construction

Since the previous section has shown that the phase separation is well described by the local-move approximation, we can now use the machinery of field theory for scalar active matter (e.g. Active Model B), as developed in [110, 147, 148]. This mapping will notably allow us to obtain the binodal densities for some class of utility functions detailed below.

The generalized-thermodynamic expansion

Even though μ in Eq. (3.21) cannot be written as the functional derivative [136, 135, 137], the dynamics can be analyzed by resorting to a gradient expansion. Indeed, expanding the chemical potential up to $O(\nabla^4, \rho^2)$ terms yields

$$\mu[\rho] = g_0(\rho) + \lambda(\rho)(\nabla \rho)^2 - \kappa(\rho) \nabla^2 \rho + O(\nabla^4, \rho^2), \quad (3.25)$$

with g_0, λ, κ local function of the field ρ , and a generalized thermodynamic mapping [147, 149] can yield the prediction of the binodal densities.

For simplicity, we will now assume that the smoothing kernel is a

Gaussian distribution of zero mean and variance σ^2 . In Fourier space, the smoothed field is given by $\hat{\phi}(k) = \hat{\rho}_k \exp(-\sigma^2 k^2/2)$, which can be truncated to leading order:

$$\hat{\phi}_k \simeq \hat{\rho}_k \left(1 - \frac{k^2 \sigma^2}{2} + O(\sigma^4 |k|^4) \right). \quad (3.26)$$

In real space, this translates into $\phi = \rho + \frac{\sigma^2}{2} \nabla^2 \rho + O(\nabla^4, \rho)$. This allows us to further expand the $\mu_{util.}$ given in Eq. (3.23). To leading order in the $O(\nabla, \rho)$ expansion, one has

$$\mu_{util.} = -2f'_\Gamma(0) \left[u(\rho) + \frac{\sigma^2}{2} u'(\rho) \partial_x^2 \rho + O(\partial_x^4, \rho) \right]. \quad (3.27)$$

Combining this expansion of $\mu_{util.}$ with the entropic contribution $\mu_{ent.}$, it is now possible to identify the different terms in Eq. (3.25), namely:

$$g_0(\rho) = -2f'_\Gamma(0)u(\rho) + f_\Gamma(0) \log \left(\frac{\rho}{1-\rho} \right); \quad (3.28)$$

$$\lambda(\rho) = 0; \quad \kappa(\rho) = f'_\Gamma(0)\sigma^2 u'(\rho). \quad (3.29)$$

This identification enables us to follow up to the next step, which is finding the proper function $R(\rho)$ and the generalized functional $\mathcal{G}[R]$ by means of which the dynamics will be given by

$$\partial_t \rho(x, t) = \partial_x \cdot \left[M[\rho] \partial_x \frac{\delta \mathcal{G}}{\delta R(x, t)} \Big|_{R(\rho)} \right]. \quad (3.30)$$

A double-tangent construction on $\mathcal{G}[R]$ then provides the binodal densities [147]. Since $\lambda(\rho) = 0$, the differential equation that the function R must satisfy (see [147, 149]) is

$$\kappa(\rho) R''(\rho) = -\kappa'(\rho) R'(\rho), \quad (3.31)$$

which simplifies into $(\kappa R')' = 0$, where the $'$ denotes the derivative with respect to ρ .

Then, because the function R is bijective, it can thus be inverted to get $\rho(R)$, and this allows us to define a new chemical potential $g[R] \equiv \mu[\rho(R)]$. The functional $\mathcal{G}[R]$ is then obtained by integrating $g[R]$ on each domain and by gluing together the two integrated parts at . Explicitly using the notations introduced in [149], we have

$$g = \frac{\delta \mathcal{G}}{\delta R}, \quad \mathcal{G} = \int dx \left[\Phi(R) + \frac{\kappa}{2R'} (\partial_x R)^2 \right], \quad (3.32)$$

where $\Phi(R)$ defines a generalized free energy density verifying

$$\frac{d\Phi}{dR} = g_0[\rho(R)], \quad (3.33)$$

with g_0 defined in Eq. (3.28). The double-tangent construction or the Maxwell construction on $\Phi(R)$ then yields the binodal densities.

For a thorough account of this method—including the challenges posed by a sign change in κ and the details of the double-tangent construction (both mainly due to R. Zakine)—see [1].

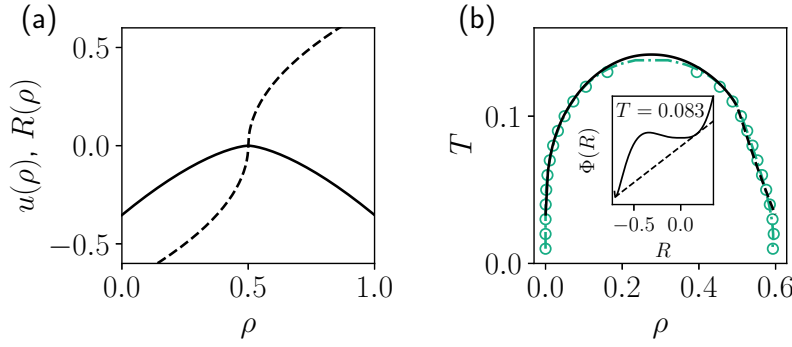


Figure 3.5: (a) Utility $u(\rho)$ (solid line), and change of variable $R(\rho)$ (dashed line) for $\alpha = 3/2$ and $\rho^* = 1/2$. (b) Comparison between the semi-analytical prediction (dark line) and the binodal densities both obtained via the Monte Carlo simulations (green circles) and solving numerically the mean-field Eq. (3.13) (green dot-dashed line). The decision rule here is the logit function, whose values in 0 are $f(0) = 1/2$ and $f'(0) = 1/4$. Inset: double-tangent construction on $\Phi(R)$ for $T = 0.083$. The dense phase is given by $R_d = 0.170$, yielding $\rho_d = 0.53$, and the gaseous phase is given by $R_g = -0.69$, yielding $\rho_g = 0.022$.

Figure (3.5)(a) shows the change of variables to $R(\rho)$, while the inset of Fig.(3.5)(b) presents a specific instance of the effective free-energy density Φ . By performing the Maxwell construction numerically over a range of temperatures, we obtain a semi-analytical form for the binodal curve, which agrees superbly with our Monte Carlo results in Fig.(3.5)(b). Crucially, although the interaction range σ enters into κ , it cancels out in the expression for Φ , meaning the coexistence densities in the thermodynamic limit are unchanged—confirming that clustering inefficiencies are not an artifact of finite system size.

3.5. Further socioeconomic considerations

Two populations

A natural extension of the problem is to restore some diversity among agents, as initially considered by both Sakoda and Schelling. Here we consider two types of interacting agents (say A and B), with possibly different utility functions, which could for example represent higher and lower revenue individuals, or city dwellers and business professionals, etc.⁴ A central question in this case is whether the system reaches fixed points, or if more complicated dynamics can persist in the long time limit, especially if the two populations have competing interests. Recent work has been devoted to studying nonreciprocal interactions between different kinds of particles, exhibiting the wealth of possible dynamical behavior when particle displacements are local [153, 154]. An interesting question in our setup is for instance: do propagating waves (or frustrated states) survive when nonlocal moves are allowed? Indeed, one may expect that enforcing local displacement constitutes a dynamical constraint that drives the system in a particular way. Allowing for nonlocal moves may change the dynamics of how the frustrated states are resolved. One may think of three major types of interactions:

- First, a cooperative interaction where agents A and agents B may maximize their utility when agents of opposite type are found in their neighborhood. This kind of interaction will typically lead to homogeneous well-mixed systems, or to some condensation into a dense phase where agents are well-mixed, but since frustration is not implemented in the microscopic rules, we reasonably expect stationary states.

4: The recent work [92] on urban segregation in the United States has brought to our attention the existence of surveys [150, 151, 152] confirming the idea that different sub-populations may require markedly different utility functions.




Figure 3.6.: Snapshots of the system for two frustrated interaction parameter choices. (a) Stationary demixing in a region where LSA presents complex eigenvalues. The agent A phase still contains some B agents. Parameters: $c_1 = -2$, $c_2 = 1$, $\sigma = 3$, $\bar{\rho}_A = 0.2$, $\bar{\rho}_B = 0.5$, $\Gamma = 10$. (b) Chaotic propagation of polarized blobs in a region where LSA presents pure real eigenvalues (null imaginary part). Parameters: $c_1 = -2$, $c_2 = 0.5$, $\sigma = 7$, $\bar{\rho}_A = 0.6$, $\bar{\rho}_B = 0.2$, $\Gamma = 100$. For both (a) and (b), $L_x = L_y = 300$. Movies are available online [155].

- Second, each agent type may decide to settle among peers and/or avoid agents of the other type in their surroundings. One should then expect a complete phase separation into two domains, one displaying a majority of As and, the other, a majority of Bs. Whether the $A - B$ phase separation additionally displays some condensation depends on the self-affinity of each agent type.
- Third, frustrated situations in which A settles with A but wants to avoid B agents, while B agents would like to gather and settle close to A. In this situation, we may expect non stationary patterns, stemming from the fact that all agents cannot be satisfied at the same time.

With this last situation in mind, we have considered the following utility functions (u_A for A agents and u_B for B agents):

$$u_A(x, [\phi_{A,B}]) = -|\phi_A(x) - \rho^\star|^2 + c_1 \phi_B(x) \quad (3.34)$$

$$u_B(x, [\phi_{A,B}]) = -|\phi_A(x) - \rho^\star|^2 + c_2 \phi_B(x), \quad (3.35)$$

where $c_1 < 0$ translates the fact that As are fleeing from B, and $c_2 > 0$ translates the fact that Bs have a tendency to gather with Bs. The term $-|\phi_A - \rho^\star|^2$ enjoins both populations to settle among A populated areas. Of course, the specific shape of utilities taken here may be restrictive and can be easily generalized.

The extension of the mean-field dynamics to this two population problem is rather straightforward. Writing $\rho_A(x, t)$ (resp. $\rho_B(x, t)$) the density of agents A (resp. B) at location x and time t , and denoting the total density by $\rho(x, t) \equiv \rho_A(x, t) + \rho_B(x, t)$, we now have an evolution equation of the form

$$\begin{aligned} \partial_t \rho_A(x, t) = & [1 - \rho(x, t)] \int \rho_A(y, t) w_{\Gamma_A}([\phi_{A,B}], y, x, t) dy \\ & - \rho_A(x, t) \int [1 - \rho(y, t)] w_{\Gamma_A}([\phi_{A,B}], x, y, t) dy, \end{aligned} \quad (3.36)$$

and, by symmetry, a similar equation for B. The transition rates depend on the utility function of each agent type and are *a priori* agent specific. Denoting $u_Z(x) \equiv u_Z(x, [\phi_{A,B}])$ (with $Z = A$ or B), we set

$$w_Z([\phi_{A,B}], y, x, t) = \omega_Z f_{\Gamma_Z}[u_Z(x) - u_Z(y)], \quad (3.37)$$

where ω_Z and Γ_Z can be agent dependent.

In App. A.5, we perform the linear stability analysis of the homogeneous state. As expected, in the frustrated two-population system, unstable modes can display temporal oscillations. However, these oscillations may stop when nonlinear terms become relevant, and the system may end up in a stationary phase separation (similar to classical demixing in equilibrium systems), as displayed in Fig. 3.6(a). Reciprocally, non-oscillating growing modes at the linear level may give rise to propagating structures and waves when nonlinearities become important, as shown in Fig. 3.6(b) (see Supplementary Material in [155]). In our system, and at odds with recent work [153, 154], the oscillatory nature of the non-homogeneous steady state cannot be predicted from a simple linear stability analysis about the homogeneous solution.

A thorough analysis of the emerging behaviors in the multi-population system would require more work, beyond the scope of the present work. Still, it is remarkable that, here as well, the linear stability analysis in the case of local jumps yields exactly the same instability conditions as the nonlocal dynamics ones (see results of Appendices A.5 and A.6). As a consequence, we expect that some results of the recent works [153, 154] should be relevant, to some extent, to describe our multi-population system.

Housing market

A common and reasonable criticism of the kind of model developed here is that, while the perceived density may be a significant factor in the decision making process of agents, the price of a house should also necessarily be taken into account. Indeed, in classical economics, the market is usually considered to be the mechanism through which the optimal allocation of goods and assets occurs (despite some contradicting empirical evidence e.g. [41]). As a result, one could rightfully argue that a housing market is necessary to ensure that agents eventually reach a steady state where their utility is maximal, at odds with what we have observed in the condensed phase.

Incorporating pricing in the model is not trivial, however, and there are a number of ways in which this could be done. A common approach in the modeling of socioeconomic systems is to introduce an agent-dependent budget and to constrain agents' moves based on such budget, as done in [86] for example. Realistically, this budget should then be heterogeneous among the agents (e.g. Pareto distributed for instance). While relevant and interesting, this agent-specific dependence as well as its formulation as a hard constraint would require a different modeling approach, and is unlikely to be tractable analytically. The alternative that we take here is to consider that if a given move leads to excess costs for an agent, its utility would decrease. We may then conveniently stay within the modeling framework of our model and assume that such a price field ψ is an increasing function of the smoothed density field ϕ , such that houses are more expensive in dense neighborhoods— or in other words: prices are driven by demand only. In its most general form, we propose the price-adjusted utility

$$\bar{u}(\phi) = u(\phi) - u_p[\psi(\phi)], \quad (3.38)$$

where u_p is the price penalty, assumed to be an increasing function

of the price, and therefore, of the density ϕ . This penalty is then, by construction, expected to drive the system away from condensation.

For concreteness, one can consider the example of a linear penalty term on the utility function, and that the price grows linearly with the local (smoothed) density, such that:

$$\bar{u}(\phi) = u(\phi) - \gamma\phi, \quad (3.39)$$

with $\gamma > 0$. Interestingly, introducing such a coupling boils down to introducing a pair-wise repulsive interaction between agents. The condition to observe condensation is then shifted by γ , and reads explicitly

$$2\rho_0(1 - \rho_0) \frac{f'_\Gamma(0)}{f_\Gamma(0)} (u'(\rho_0) - \gamma) > 1. \quad (3.40)$$

Clearly, condensation cannot occur anymore if the price penalty on the utility is too important.

This example therefore illustrates that it is indeed possible for an appropriate housing market to destroy the condensation observed in our continuous model. However, it is important to note that the outcome (without the price penalty) for agents is not necessarily improved by this brutal homogenization through the price field. Besides, it can be argued that the effect of price should not be as trivial as a utility penalty. In fact, other models have used price as a proxy for social status, in which case agents are actually more attracted by the most expensive site they can afford [86]. In App. A.7 we also consider the case of a more complex price field dependence. This actually led to performing a comprehensive study of housing price correlations in empirical data and the careful construction of an adequate pricing dynamics, which is presented in chapter 4.

3.6. Discussion

Let us summarize what we have achieved in this chapter. We have introduced a neighborhood-less extension of the Sakoda-Schelling model for the occupation of a lattice representing a city. In this version of the model, the agents attempt to maximize a utility that is a function of their *perceived* local density, and are most satisfied when they are located in a site where such density is of an intermediate value, i.e. neither empty nor too crowded. Having that agents only consider their own site dependent utility, and that their utility is non-linear, drives the system out of equilibrium. As a result, the system can no longer be studied by constructing a free energy directly from an aggregate system-wide utility function, as was done by Grauwin *et al.* [39] for agents inhabiting predefined neighborhoods or blocks in which the utility is identical for all. Using numerical simulations as well as a mean-field description of the nonequilibrium dynamics, we have established that the apparent disparity between *micromotives* and *macrobehaviors* initially observed by Schelling [14] is robust to the absence of neighborhoods and to the out-of-equilibrium nature of our extension. Interestingly, we find that the transition between the fully homogeneous state and the phase-separated one likely belongs to the

2D Ising universality class, a debated topic in the active matter literature regarding the very similar Motility Induced Phase Separation (MIPS) phenomenon. Taking advantage of the similarity between our problem and the Active Model B (describing MIPS), we predict the local density in the bulk of the concentrated phase, confirming that the majority of agents find themselves in a sub-optimal situations with a perceived density exceeding the ideal value.

While seemingly technical, the fact that the original observations of Schelling is robust to out-of-equilibrium dynamics actually carries far reaching consequences, in our opinion. Indeed, as discussed in Sec. 2, equilibrium descriptions of socioeconomic problems require the decision rule $f_{\Gamma}(x)$ to be the “logit” function. This very specific choice is a common source of criticism, as any results are then *a priori* uncertain to hold for other decision rules. Here, on the other hand, our out-of-equilibrium description presents no such restriction, as all calculations have been written as generally as possible and, interestingly, turn out to only depend on $f_{\Gamma}(0)$ and $f'_{\Gamma}(0)$. While the numerical results presented here are those using the classical choice of the logit function for lack of a more plausible decision rule, one could readily adapt the outcomes following behavioral evidence that another function is more appropriate, and we expect the entire phenomenology of our model to hold for any other sigmoid function. More generally, we believe that this robustness to other decision rules holds for a large number of socioeconomic models that have been described using the methods of statistical physics [109, 156, 157, 158, 159]. Of course, subtleties around the dynamics, such as the relaxation time towards the steady state or the coarsening dynamics discussed here, will inherently be affected by the specific choice that governs transition rates. This being said, we have observed a remarkable similarity in the local and non-local versions of our model for which the dynamics are yet qualitatively very different. It is therefore possible that there also exists some degree of universality in the dynamical behavior of different decision rules, at least at the mean-field level.

Going back to the Sakoda-Schelling model, we have also introduced some generalizations that we believe are natural and relevant. First, the introduction of different sub-populations is interesting, as the exhibited dynamical patterns are impossible to observe in an equilibrium version of the model. Second, we have seen that introducing a linear dependence of the price on the density has the effect of delaying the transition, eventually killing it off completely if the price penalty in the utility function is strong enough. As previously stated, however, this mechanism remains very simple. In order to determine more plausible effects of a housing market, a thorough analysis of real estate transactions appears to be a promising avenue, in particular given the increasing availability of open datasets in this area in major European cities. The extensive study of French data presented in the next chapter, will allow us to couple this continuous Sakoda-Schelling model with a plausible housing market model in the last chapter of this part of the thesis. Note finally that the recent preprint by Seara et al. [92] revisits the problem of urban segregation in the United States and proposes a two-population model very similar to ours. Studying the validity of hydrodynamic descriptions of the two-population problem

using the census data brought forth in this chapter appears to be an important perspective. Such a comparison could notably be necessary to highlight the role of ingredients not taken into account in existing models – such as public policy and economic inequalities – in the emergence of strong geographic disparities.

Key Takeaways

1. **Out-of-equilibrium Robustness:** Removing fixed neighborhoods and permitting nonlinear, site-specific utilities drives the system out of equilibrium, yet the classic Schelling result—micromotives generating unexpected macro-segregation—remains intact.
2. **Ising-Class Transition:** The shift from a homogeneous to a phase-separated state in our neighborhood-less model falls into the 2D Ising universality class, linking social segregation to broader active-matter phenomena like MIPS.
3. **Minimal Decision-Rule Dependence:** Unlike equilibrium analyses that require a logit transition function, our mean-field theory depends only on $f_{\Gamma}(0)$ and $f'_{\Gamma}(0)$, implying full phenomenological robustness across a vast class of decision rules.
4. **Extensions and Heterogeneity:** Introducing several subpopulations enriches the system's dynamics, and incorporating a density-dependent price term recovers a similar condensation criterion as in the simpler model—further contradicting the “invisible-hand” paradigm.

The diffusive nature of housing prices

4.

Ils se turent. La route plate, bordée de peupliers, étirait devant eux un fragment du libre univers. L'aventurier de la puissance et l'aventurier du savoir marchaient côte à côte.
L'OEuvre au noir-Marguerite Yourcenar, 1968.

This chapter is taken from Becharat, Benzaquen, and Bouchaud (2025) [2] *Diffusive Nature of Housing Prices*. The figures and accompanying text have been lightly adapted to suit the context of this thesis, and a new section now couples the price equation to the preceding Sakoda–Schelling analysis. The original publication was written under the supervision of M. Benzaquen and J.-P. Bouchaud.

4.1. Introduction

Complex spatial patterns often result from a subtle interplay between random forcing and diffusion, like for example surface growth [160] or fluid turbulence [161]. One can also expect such competition between heterogeneities and diffusion to take place in socio-economic contexts. For example, word of mouth leads to spreading of information or of opinions. Provided the spreading mechanism is local enough (i.e. before the advent of social media), the large scale description of such phenomena is provided by the diffusion equation that leads to specific predictions for the long-range nature of spatial correlations of voting patterns, which seems to be validated by the analysis of empirical data [88, 162, 163].

One may argue that housing prices should display similar patterns. Indeed, it is intuitively clear that the price of real estate in a given district is affected, among many other factors, by the price of the surrounding districts, through a sheer proximity effect. This is enough to generate a diffusion term in any coarse-grained description of the spatio-temporal evolution of prices – see below and Appendix (B.1) for more precise statements. The aim of this work is to present such a phenomenological description of the dynamics of the *price field* in a given region of space, and to compare analytical prediction to empirical data using spatially resolved transaction prices in France for the period 1970 to 2022 – see Fig. 4.1 for a visual representation of the price field that motivates our analysis. We will find what we consider to be rather remarkable agreement with theory, in view of the minimal amount of modeling ingredients. In particular, the logarithmic dependence of spatial correlations, characteristic of two-dimensional diffusion, is clearly visible in the data at all scales (see Fig. 4.3 below).




Figure 4.1.: Spatial transaction log-prices p distribution in France in 1970 (left) and in 2022 (right). We use a sigmoid transformation of the log prices rescaled by their mean and divided by their standard deviation in order to highlight price differences. As seen in this plot, high prices are concentrated around France's principal cities and on the coasts and mountains, but the price pattern clearly displays spatial diffusion. Data from [164].

Due to its potent macroeconomic and systemic risk implications, housing prices have long been studied by economists, see [165]. One of the most famous description of the housing market is through the Hedonic prices hypothesis (see e.g. [166]), which states that *goods are valued for their utility-bearing attributes. Hedonic prices are defined as the implicit prices of attributes and are revealed from observed prices of differentiated products and the specific amounts of characteristics associated with them.* In essence, we shall argue that real-estate prices in the vicinity of a given location is one of these characteristics.

There is also a great body of empirical literature highlighting the links between the housing market prices and, for example, violence [167] or school grades [168]. This has naturally led to models of the housing market using reasonable assumptions. In particular, recent agent-based models of the housing market have been designed to explain price dynamics [165], or its link with social segregation. Ref. [83] observed that segregation patterns can be observed even with the simplest parameter setting in an agent-based model of the housing market. Ref. [84] showed how such models could be very helpful to test and apply effective policies to prevent social / racial segregation, in the same vein as Ref. [85] where the effectiveness of macro-prudential policies is tested on an agent-based model of the UK housing market. Interestingly, [86] showed that social segregation is also strongly linked with social influence.

Concerning spatial patterns, *local* studies from the mid-1990's have suggested the potential importance of spatial diffusion effects. For example, Clapp & Tirtiroglu [169] find evidence of local price diffusion from their empirical study of the metropolitan of Hartford, Connecticut. Pollakowski & Ray [87] confirms these results at the local level, and conclude that housing prices are inefficient: *If housing markets were efficient, [...] shocks would either be confined to one area, in which case information transfer is irrelevant, or affect a number of areas, in which case the price changes should occur nearly simultaneously, not one after another.* These authors also note that price changes are auto-correlated in time (a feature that we will explicitly include in our theoretical model), which is a further sign of price inefficiency. Indeed, properly anticipated prices should not be predictable [170].

As we argue below, such *local* diffusion of prices is expected to create *long-range correlations* in the price field both in space and in time, which we observe in the data. Although the presence of spatial corre-

lations was noticed in [171], their analysis did not provide any theoretical framework to elucidate the origin of these correlations. In particular, no mention was made of their long-range character, let alone of their specific logarithmic dependence that we establish empirically and justify theoretically in the following. Other socio-economic variables are known to be long-range correlated [162], with far-reaching consequences on the statistical significance of many results in spatial economics, as forcefully argued in [172].

4.2. Model description

Introducing a stochastic differential equation

Our theoretical framework aims at modeling the dynamics of the housing price field in a similar spirit as for the dynamics of opinions or intentions [173, 174, 88, 175], or of criminal activities [176]. We introduce a two-dimensional field $\psi(\mathbf{r}, t)$ which represents the deviation from the (possibly time dependent) mean of the log-price of housing around point \mathbf{r} at time t . We then posit that such a field evolves in time according to the following stochastic partial differential equation

$$\frac{\partial \psi(\mathbf{r}, t)}{\partial t} = D \Delta \psi(\mathbf{r}, t) - \kappa \psi(\mathbf{r}, t) + \eta(\mathbf{r}, t) + \xi(\mathbf{r}), \quad (4.1)$$

where Δ is the Laplacian operator, D a diffusion coefficient, κ a mean-reversion coefficient, $\eta(\mathbf{r}, t)$ a Langevin noise with zero mean and short range time and space correlations, and $\xi(\mathbf{r})$ a static random field with zero mean and short range correlations. The correlators of these terms are assumed to be of the following type:

$$\begin{aligned} \langle \eta(\mathbf{r}, t) \eta(\mathbf{r}', t') \rangle &= \frac{A^2}{T a^2} e^{-|t-t'|/T} g_a(|\mathbf{r} - \mathbf{r}'|); \\ \langle \xi(\mathbf{r}) \xi(\mathbf{r}') \rangle &= \frac{\Sigma^2}{a^2} g_a(|\mathbf{r} - \mathbf{r}'|), \end{aligned} \quad (4.2)$$

where $g_a(r)$ is a bell-shaped function that decays over length scale a , such that $2\pi \int_{r>0} g_a(r) r dr = a^2$. Note that in terms of dimensions, $[A^2] = [D] = [L^2 T^{-1}]$, $[\kappa] = [T^{-1}]$ and $[\Sigma] = [L T^{-1}]$.

The four different terms of Eq. (4.1) capture the following features: (i) the diffusion term describes the proximity effect alluded to in the introduction and documented in Refs. [169, 87]: pricey districts tend to progressively gentrify; conversely, rundown districts lower the market value of their surroundings. (A more technical version of this argument is given in Appendix (B.1)). (ii) The mean-reversion term can be seen as a coupling between local log-prices and the mean log-price, here set to zero, and can be thought of as the result of long-range economic forces that keep prices within a country more or less in sync through the effect of e.g. migrations, policies or wealth inequalities. (iii) The time-dependent noise term η models all idiosyncratic shocks affecting the “hedonic” variables determining the price of properties – for example the creation of a local metro or train station, of a pedestrian zone, or adverse shocks like increase in local crime, floods, etc. The impact of such shocks is often drawn out in time, so we assume η to be auto-correlated with a decay time T , in line with the observa-

tions reported in [87]. (iv) The time-independent stochastic term ξ is meant to represent persistent biases in the local quality of life in different regions, due to e.g. geographical features (close to the sea-shore, or to river banks, etc.). For simplicity, We have assumed that the spatial correlation lengths of both η and ξ are equal to the same value a .

Now, Eq. (4.1) makes detailed predictions for the spatial and temporal correlations of the field $\psi(\mathbf{r}, t)$. To wit, the spatial variogram $\mathbb{V}(\ell, 0) := \langle (\psi(\mathbf{r}, t) - \psi(\mathbf{r}', t))^2 \rangle_{|\mathbf{r}-\mathbf{r}'|=\ell}$ can be explicitly computed in the range $\max(a, \sqrt{DT}) \ll \ell \ll \ell^*$ (where $\ell^* := \sqrt{D/\kappa}$), as we will show in the next section.

Computing the spatial variogram

Let us consider again Eq. (4.1) and Eq. (4.2). For the rest of the calculations we assume

$$|\mathbf{r} - \mathbf{r}'| = \ell \gg a, \quad \implies \quad \frac{1}{a^2} g_a(|\mathbf{r} - \mathbf{r}'|) \approx \delta(|\mathbf{r} - \mathbf{r}'|).$$

Moreover, the space-time correlation function can be written as

$$\begin{aligned} \mathbb{C}(|\mathbf{r} - \mathbf{r}'|, |t - t'|) &= \langle \psi(\mathbf{r}, t) \psi(\mathbf{r}', t') \rangle \\ &= \iint e^{-i\mathbf{k}\cdot\mathbf{r} - i\mathbf{k}'\cdot\mathbf{r}'} \langle \psi_{\mathbf{k}}(t) \psi_{\mathbf{k}'}(t') \rangle \frac{d\mathbf{k}}{(2\pi)^2} \frac{d\mathbf{k}'}{(2\pi)^2}. \end{aligned} \quad (4.3)$$

Here $\psi_{\mathbf{k}}(t)$ satisfies, in Fourier space,

$$\frac{\partial \psi_{\mathbf{k}}(t)}{\partial t} = -(D|\mathbf{k}|^2 + \kappa) \psi_{\mathbf{k}}(t) + \eta_{\mathbf{k}}(t) + \xi_{\mathbf{k}}, \quad (4.4)$$

whose formal solution is

$$\psi_{\mathbf{k}}(t) = \psi_{\mathbf{k}}(0) e^{-(D|\mathbf{k}|^2 + \kappa)t} + \int_0^t e^{-(D|\mathbf{k}|^2 + \kappa)(t-\tau)} [\eta_{\mathbf{k}}(\tau) + \xi_{\mathbf{k}}] d\tau. \quad (4.5)$$

Because of the two fields η and ξ - assumed to be independent - we will separate the calculation for the correlation function into two contributions. In the long time limit, the first contribution in Fourier space, coming from field η , is:

$$\int_0^t \int_0^{t'} dt_1 dt_2 e^{-(D\mathbf{k}^2 + \kappa)(t-t_1) - (D\mathbf{k}'^2 + \kappa)(t'-t_2)} \langle \eta_{\mathbf{k}}(t_1) \eta_{\mathbf{k}'}(t_2) \rangle, \quad (4.6)$$

leading to:

$$\frac{A^2(2\pi)^2}{T} \int_0^t \int_0^{t'} dt_1 dt_2 e^{-(D\mathbf{k}^2 + \kappa)(t-t_1) - (D\mathbf{k}'^2 + \kappa)(t'-t_2)} e^{-\frac{|t_1-t_2|}{T}} \delta(\mathbf{k} + \mathbf{k}'). \quad (4.7)$$

We find, in the long time limit, that the integral yields in Fourier space:

$$\frac{A^2(2\pi)^2}{2T} \left[\frac{e^{-(D\mathbf{k}^2 + \kappa)|t'-t|}}{2(D\mathbf{k}^2 + \kappa)(D\mathbf{k}^2 + \kappa + \frac{1}{T})} + \frac{e^{-\frac{|t-t'|}{T}}}{(D\mathbf{k}^2 + \kappa)^2 - \frac{1}{T^2}} - \frac{e^{-(D\mathbf{k}^2 + \kappa)|t'-t|}}{2(D\mathbf{k}^2 + \kappa)(D\mathbf{k}^2 + \kappa - \frac{1}{T})} \right]. \quad (4.8)$$

This can be condensed as:

$$\frac{A^2(2\pi)^2}{2T((D\mathbf{k}^2 + \kappa)^2 - \frac{1}{T^2})} \left[e^{-\frac{|t-t'|}{T}} - \frac{e^{-(D\mathbf{k}^2 + \kappa)|t'-t|}}{T(D\mathbf{k}^2 + \kappa)} \right]. \quad (4.9)$$

Similarly, we can compute the contribution for the correlation function coming from field $\xi(\mathbf{r})$:

$$(2\pi)^2 \Sigma^2 \int_0^t \int_0^{t'} dt_1 dt_2 e^{-(D\mathbf{k}^2 + \kappa)(t-t_1) - (D\mathbf{k}'^2 + \kappa)(t'-t_2)} \delta(\mathbf{k} + \mathbf{k}'). \quad (4.10)$$

This yields, in the long time limit:

$$\frac{(2\pi)^2 \Sigma^2}{(D\mathbf{k}^2 + \kappa)^2}. \quad (4.11)$$

We come back to the first contribution (coming from field η) in Fourier space for the space time correlation function:

$$\frac{A^2(2\pi)^2}{2T((D\mathbf{k}^2 + \kappa)^2 - \frac{1}{T^2})} \left[e^{-\frac{|t-t'|}{T}} - \frac{e^{-(D\mathbf{k}^2 + \kappa)|t'-t|}}{T(D\mathbf{k}^2 + \kappa)} \right]. \quad (4.12)$$

We now focus on the static behavior of this term, hence imposing $t = t'$. This yields:

$$\frac{A^2(2\pi)^2}{2T((D\mathbf{k}^2 + \kappa)^2 - \frac{1}{T^2})} \left[1 - \frac{1}{T(D\mathbf{k}^2 + \kappa)} \right]. \quad (4.13)$$

Using notations $|\mathbf{k}| = k$, $\mathbf{k} \cdot (\mathbf{r} - \mathbf{r}') = k\ell \cos(\theta)$ and notation \mathbb{C}_η to describe the contribution from η to the correlation function, it comes in polar coordinates:

$$\mathbb{C}_\eta(\ell, 0) = \frac{A^2}{2T(2\pi)^2} \int dk \int d\theta e^{-ik\ell \cos(\theta)} \frac{k}{((Dk^2 + \kappa)^2 - \frac{1}{T^2})} \left[1 - \frac{1}{T(Dk^2 + \kappa)} \right]. \quad (4.14)$$

The integral is defined for $1/\ell^* \ll k \ll 1/a$, which ensures that $Dk^2 \gg \frac{D}{\ell^2} = \kappa$. We can hence neglect the mean-reversion term in the computation. Moreover, we can neglect $D^2 k^4$ in favor of $\frac{1}{T^2}$ if $Dk^2 < \frac{1}{T}$, hence if $\ell > \sqrt{DT}$. This is typically the regime that we consider for this study, since we estimate (see in the main text) $\sqrt{DT} \approx 13$ km, so we assume here that this term is negligible. Finally, we can identify the Bessel function

$$\frac{1}{2\pi} \int_0^{2\pi} d\theta e^{ik\ell \cos(\theta)} = J_0(k\ell) = J_0(-k\ell), \quad (4.15)$$

so:

$$\mathbb{C}_\eta(\ell, 0) \approx \frac{A^2}{4\pi} \int_{1/\ell^*}^{1/a} dk \frac{J_0(k\ell)}{Dk}. \quad (4.16)$$

The Bessel function can be expanded for $k\ell \rightarrow 0$, and yields $J_0(k\ell) \approx 1 - \ell^2 k^2/4 + o(k^4 \ell^4)$. Moreover, the Bessel function decays to zero when $k\ell \gg 1$, concentrating the integral towards its lower bound. This gives,

up to constant contributions:

$$\mathbb{C}_\eta(\ell, 0) \approx -\frac{A^2}{4\pi D} \log \frac{\ell}{\ell^*} + K(\ell) \quad (4.17)$$

with correction term $K(\ell)$. Similarly, we can compute the contribution from field ξ :

$$\mathbb{C}_\xi(\ell, 0) = \frac{\Sigma^2}{2\pi D^2} \int_{1/\ell^*}^{1/a} dk \frac{J_0(k\ell)}{k^3} = \frac{\Sigma^2}{2\pi D^2} \ell^2 \int_{\ell/\ell^*}^{\ell/a} du \frac{J_0(u)}{u^3}. \quad (4.18)$$

In order to have a non-constant contribution here, we must go to the second order in the expansion of the Bessel function towards the lower bound of the integral. This yields:

$$\mathbb{C}_\xi(\ell, 0) \approx \frac{\Sigma^2}{2\pi D^2} \ell^2 \int_{\ell/\ell^*}^{\ell/a} du \frac{1 - \frac{u^2}{4}}{u^3}, \quad (4.19)$$

which finally yields, up to constant terms:

$$\mathbb{C}_\xi(\ell, 0) \approx \frac{\Sigma^2}{8\pi D^2} \ell^2 \log \frac{\ell}{\ell^*} + K'(\ell) \quad (4.20)$$

with correction $K'(\ell)$. Furthermore, the variogram is defined as $\mathbb{V}(\ell, 0) = 2\langle \psi(\mathbf{r}, 0)^2 \rangle - 2\mathbb{C}(\ell, 0)$. Hence, summing both contributions yields:

$$\mathbb{V}(\ell, 0) \approx \frac{A^2}{2\pi D} \log \frac{\ell}{\ell^*} - \frac{\Sigma^2}{4\pi D^2} \ell^2 \log \frac{\ell}{\ell^*} + C, \quad (4.21)$$

where C is a constant.

Note that the first term is the familiar logarithmic correlation of the Gaussian free-field in two dimensions, see e.g. [177]. For $\ell \gtrsim \ell^*$, the variogram reaches a plateau value.

Similarly, the temporal variogram $\mathbb{V}(0, \tau) := \langle (\psi(\mathbf{r}, t) - \psi(\mathbf{r}, t + \tau))^2 \rangle$ can be computed, but the final expression is cumbersome and depends on the relative position of three time scales: κ^{-1} , the correlation time T and the typical diffusion time $S = a^2/D$ over length scale a , see (B.2). There are typically four regimes, a short time regime where $\mathbb{V}(0, \tau) \propto \tau^2$ that reads

$$\mathbb{V}(0, \tau) = \frac{A^2}{16\pi D} \log \left(\frac{1 + \frac{T}{S}}{1 + \kappa T} \right) \frac{\tau^2}{T^2}, \quad \tau \ll T, S \quad (4.22)$$

followed by two intermediate regimes where $\mathbb{V}(0, \tau) \propto \tau$ and $\log \tau$, and finally a saturated regime for $\kappa\tau \gg 1$.

In the next sections, we will compare these predictions to empirical data, with good overall agreement. We will find that the spatial variogram is well described by a pure logarithm, i.e. the first term of Eq. (4.21) – this allows us to determine the ratio A^2/D . With the same value of A^2/D , we then fit the temporal variogram with reasonable values of T and S .

4.3. Comparison with empirical data

We conducted extensive empirical analyses based on two data sources. The first one is accessible online via the DVF (Demande de Valeur Foncière) website, and displays every housing market transaction in France between 2018 and 2022. This data include the price of the property, its surface and its spatial coordinates. This allows us to study both transaction prices and prices per square meter, up to the granularity of a given point in space. The second data source comes from [164], where the authors compiled a wealth of socio-economic indicators, spanning from 1970 to 2022 ¹, including housing market prices, but the dataset only contains average transaction prices per *communes* in France up to 2022 and average prices per square meter per *communes* from 2014 to 2022. ² Even though the latter data source is much less granular than the DVF dataset, its time span of 52 years allows us to investigate the temporal variogram of prices, and so this is the data we focus on in the main text. (The DVF data only span 5 years, which will turn out to be of the same order of magnitude as the correlation time T of the noise). For empirical findings on prices per square meters from DVF, which are fully consistent with those established for transaction price, see Appendix (B.4).

We first show a color map of transaction log-prices $p := \log P$ across France in Figure (4.1), sourced from [164], to compare the spatial distribution of prices in France over the past five decades, a key aspect of our investigation. Indeed, one can see that the price distribution in France is far from uniform, and reveals spatial diffusion around big cities, coastal regions or ski resorts.

Then, it is interesting to study the distribution of individual transaction log-prices p , unconditionally over the whole of France. Using the DVF data base, we find that the distribution of prices has a double hump shape, probably reflecting the superposition of two different price distributions for cities and for the countryside, see Fig. B.1. We show in Fig. 2 of the Appendix (B.4), a comparison between the distribution of prices in the *département* of la Creuse (chosen to represent a typical countryside district) and in Paris, highlighting the mixture of two distributions seen in the global price distribution for the whole of France. The tail of the distribution of the transaction prices decays as $P^{-1-\mu}$ with $\mu \approx 1.5$, implying that the variance of the transaction prices is mathematically infinite. This should be compared to the Pareto tail of the wealth distribution in France, which decays with a similar exponent [178]. The distribution of prices per square meters does not have the same shape, but has again a similar power-law tail, as shown in (B.4).

Static analysis

We now shift our focus to the spatial correlations of the logarithm of prices, which we characterize by the equal-time variogram $\mathbb{V}(\ell, 0)$ defined above. The square-root of this quantity measures how different the (log-)prices are when considering two properties a distance ℓ away.

In order to determine the order of magnitude of the diffusion constant

1: For the specific case of the housing market. Other socio-economic indicators cover an even longer time span. We in fact found similar logarithmic correlations for, e.g., the alphabetization rate in France.

2: The housing market data compiled by [164] for the years 2014-2022 comes from the DVF database, and is averaged per *communes*.

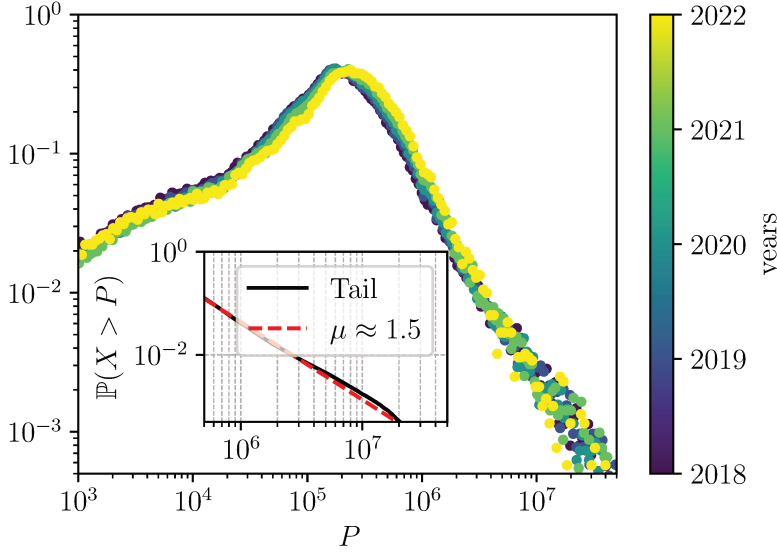


Figure 4.2.: Distribution of all transaction log-prices $p := \log P$, for the 5 years of DVF data. Note the double hump shape, reflecting a mixture of two distributions, corresponding to prices in cities and prices in the countryside. The right tail, for property prices above 500,000 Euros, corresponds to a power-law tail for prices as $P^{-1-\mu}$ with $\mu \approx 1.5$.

D we analyze the propagation of price “shocks” induced by the opening of a TGV (Train à Grande Vitesse) train station in various cities, as shown in Fig. (4.4). We find that D is of order $50 \text{ km}^2/\text{year}$, corresponding to prices adapting to a local shock on a scale of 7 km after a year. This leads to a value of $A^2 \approx 2\pi \times 0.19D \sim 60 \text{ km}^2/\text{year}$. We will comment on this value below, after having discovered that the noise amplitude A^2 is in fact space dependent.

The reader must have noticed that although the slopes of the variograms are the same at all scales, they are shifted up and down in the y-direction. This is expected if one accounts for measurement noise. Indeed, the “true” price field $p(\mathbf{r}, t)$ is approximated here by an empirical average over the chosen cells of transaction prices. The larger the cell size and the smaller the dispersion of prices within each cell, the smaller such idiosyncratic contributions to the difference of prices for two neighboring cells.

Finally, note that the spatial variograms do not seem to reveal any departure from the $\log \ell$ behavior predicted by the first term of Eq. (4.21), except at large distances where finite size and boundary effects start playing a role. Comparing the two terms of Eq. (4.21), one concludes that the second term remains negligible provided $\ell \lesssim D/\Sigma$. Choosing $D = 50 \text{ km}^2/\text{year}$ and $\ell = 500 \text{ km}$, this holds provided $\Sigma \lesssim 0.1$, i.e. whenever idiosyncratic effects lead to persistent differential of price variations of at most 10% after a year and over 1 km. We believe that this is indeed an upper bound to such idiosyncratic effects.

Temporal analysis

Turning to the temporal variogram of prices, there are two different empirical definitions for such an object, which should lead to similar results if the system is (statistically) spatially homogeneous. One ($\mathbb{V}_1(\tau)$) is to compute the temporal variance of local price changes $p(\mathbf{r}, t) - p(\mathbf{r}, t + \tau)$ over the full time period, which is then averaged over \mathbf{r} . The second ($\mathbb{V}_2(\tau)$) is to remove from $p(\mathbf{r}, t)$ the spatial average of the log-price at time t , i.e. $\bar{p}(t) = \langle p(\mathbf{r}, t) \rangle_{\mathbf{r}}$, and then compute




Figure 4.3.: Spatial variogram for the log price field $p(\mathbf{r})$ averaged over the period 2018-2022 for France as a whole, its *régions*, *départements* and cities, with their respective cross-sectional variability highlighted in shaded colors and the averages for each scale as filled circles. The black dashed lines have a slope equal to 0.19 for all scales, corresponding to $A^2/D = 1.2$. The different offsets in the y direction corresponds to the measurement noise contribution to the empirical field $p(\mathbf{r})$. The inset shows the comparison between $\mathbb{V}_1(\tau)$ and $\mathbb{V}_2(\tau)$ for the empirical data, in a log-log representation. We also show (in red) the fit found for $\mathbb{V}_1(\tau)$ with our theoretical equation. Note that the short time behavior of $\mathbb{V}_1(\tau)$ is in-between τ and τ^2 , indicating a non-zero correlation time T . We find $T = 3.5$ years, $S = 1$ year, $D = 50\text{km}^2$ per year and $A^2/D \approx 1.2$. The observed shift between $\mathbb{V}_1(\tau)$ and $\mathbb{V}_2(\tau)$ is a consequence of strong spatial heterogeneities, see the appendix B.4. Note that with 50 years of data, only the first 10 years of lags are reliable.

the average of $[p(\mathbf{r}, t) - \bar{p}(t) - (p(\mathbf{r}, t + \tau) - \bar{p}(t + \tau))]^2$ over both t and \mathbf{r} . For a statistically homogeneous system, these two procedures lead to comparable results. However, as shown in Fig. 4.3, our data reveals strong differences between $\mathbb{V}_1(\tau)$ and $\mathbb{V}_2(\tau)$, which can be accounted for by assuming that the variance A^2 of the driving noise η is space dependent: $A^2 \rightarrow A^2(\mathbf{r})$. In this case, spatial correlations lose their translation invariance but if one insists on computing them as a function of $\ell = |\mathbf{r} - \mathbf{r}'|$, one recovers Eq. (4.21) with A^2 replaced by its spatial average $\langle A^2 \rangle_{\mathbf{r}}$, see Appendix (B.4).

Now, it turns out that in the presence of spatial heterogeneities, the temporal variogram $\mathbb{V}_1(\tau)$ is also given by Eq. (4.22) with $A^2 \rightarrow \langle A^2 \rangle_{\mathbf{r}}$, see Fig. 5 of Appendix (B.4). Hence we focus our attention to $\mathbb{V}_1(\tau)$ and attempt to fit it with our theoretical formula (see Appendix (B.2)) with T, S as adjustable parameters, with $\langle A^2 \rangle_{\mathbf{r}}/D$ fixed and set to 1.2, close to the value inferred from spatial variograms. (D itself has negligible influence on the goodness-of-fit, only the ratio A^2/D matters). The optimal values are then found to be $S = 1$ year, corresponding to a correlation length for shocks $a = \sqrt{DS} = 7$ km, and a correlation time of $T = 3.5$ years, such that $\sqrt{DT} = 13$ km. The order of magnitude of A^2 is expected to be $a^2/T \sim 30 \text{ km}^2/\text{year}$, a factor two times smaller than expected if $D = 50 \text{ km}^2/\text{year}$, but not unreasonable in view of the crudeness of our model and the possibility to change the value of parameters without substantially affecting the joint goodness-of-fit of spatial and temporal variograms. For example, choosing $\langle A^2 \rangle_{\mathbf{r}}/D = 1.3$ leads to $T = S = 2.5$ years and in this case $a^2/T \sim 50 \text{ km}^2/\text{year}$. Note that the short-time regime of $\mathbb{V}_1(\tau)$ is a sign that price changes are persistent, which is inconsistent with the hypothesis that the housing market is “efficient” [87]. In view of the large transaction costs incurred when buying a house, this is hardly surprising.

Finally, in order to account for the empirical difference between the

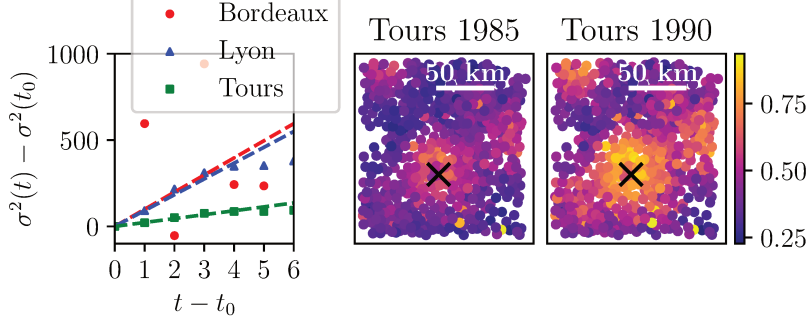


Figure 4.4. Left: Impulse response function. We measure how housing prices diffuse around a newly opened TGV station establishing a fast train connection to Paris at time t_0 from Lyon, Tours and Bordeaux (see (B.5) for a precise definition of σ^2). The x-axis is in years and the y-axis in km^2 . From the slope of the linear initial regime we infer the diffusion constant D : $D \approx 90 \text{ km}^2/\text{year}$ for Lyon, $D \approx 22 \text{ km}^2/\text{year}$ for Tours and $D \approx 100 \text{ km}^2/\text{year}$ for Bordeaux, roughly in agreement with our selection of $D \approx 50 \text{ km}^2/\text{year}$ for the whole of France. Middle and right: visual representation of the price field centered around the city of Tours, marked by a black cross. The isotropic diffusion of the price field after the opening of the TGV train station in 1985 is clearly observable.

two temporal variograms $\mathbb{V}_1(\tau)$ and $\mathbb{V}_2(\tau)$, one needs to introduce rather strong spatial heterogeneities in the noise amplitude A^2 , that must vary by a factor of 10 depending on the considered region, see Fig. 5 in Appendix (B.4). This is not very surprising in view of the very different structure of the housing market in international cities like Paris or Nice and the remote, low density regions like Lozère. A generalized version of our model, Eq. (4.1), that properly accounts for geographical heterogeneities that make both D and A^2 space dependent, would however require a different, much more granular calibration strategy.

4.4. Extension: coupling population and prices

Model setting

Building on our neighborhood-less Sakoda–Schelling model discussed in chapter 3, which admits a hydrodynamic description at the aggregate level, we now replace the raw density field by its logarithm, denoted $\rho(R, t)$, to parallel our earlier use of the log-price field $\psi(R, t)$. We then posit that prices and population evolve jointly as

$$\partial_t \psi(R, t) = D \Delta \psi(R, t) + \mathcal{F}[\rho(R, t)] - \kappa \psi(R, t) + \eta(R, t), \quad (4.23)$$

where \mathcal{F} encodes the feedback from population onto prices and η is the space-time dependent stochastic noise. For simplicity, we drop here the static stochastic contribution.

Next we couple this to a time-evolution equation for $\rho(R, t)$. Adopting the local-move approximation in the master equation—shown in chapter 3, which agrees well with numerical simulations and empirical data as shown in [92]—and using field-theoretic techniques inspired from [136, 179, 135] and detailed in Appendix A.3, one arrives at

$$\partial_t \rho(R, t) = f_\beta(0) \Delta \rho - 2 f'_\beta(0) \nabla[\rho(1 - \rho) \nabla u(\rho)] + \eta'(R, t), \quad (4.24)$$

where

$$f_\beta(x) = \frac{1}{1 + e^{-\beta x}}$$

is the logit transition rate, $u(\rho)$ is the agents' utility and $\eta'(R, t)$ is a stochastic space-time dependent noise. Here $f_\beta(0)$ and $f'_\beta(0)$ act like effective diffusion coefficients, evaluated at zero argument and its derivative with respect to the intensity of choice parameter β ³.

3: Here, we revert to the notation β for the intensity-of-choice parameter.

To close the system, we assume a linear coupling of the form

$$\mathcal{F}[\rho] = \alpha(\rho - \bar{\rho}(R)),$$

where $\bar{\rho}(R)$ is a local density threshold above which prices cease rising. Thus (4.23) becomes

$$\partial_t \psi = D \Delta \psi + \alpha[\rho(R, t) - \bar{\rho}(R)] - \kappa \psi + \eta(R, t). \quad (4.25)$$

In the noise-free, steady-state limit without diffusion we recover

$$\psi = \frac{\alpha}{\kappa} [\rho - \bar{\rho}(R)],$$

in agreement with the linear density-price coupling assumption made in chapter (3).

Putting everything together yields the coupled evolution system

$$\begin{cases} \partial_t \psi(R, t) = D \Delta \psi + \alpha[\rho(R, t) - \bar{\rho}(R)] - \kappa \psi(R, t) + \eta, \\ \partial_t \rho(R, t) = f_\beta(0) \Delta \rho - 2 f'_\beta(0) \nabla[\rho(1 - \rho) \nabla(u(\rho) - \lambda \psi)] + \eta'. \end{cases} \quad (4.26)$$

Here $u(\rho) = |\rho - \rho^*|^2$ and we have added a price penalty $-\lambda \psi$ to the utility, as agents will not want to live in neighborhoods where prices are too high, with λ measuring housing-market elasticity.

Stability analysis

To test whether introducing ψ can restore the “invisible hand,” we linearize (4.26) around the homogeneous state (ψ_0, ρ_0) , writing

$$\psi(R, t) = \psi_0 + \delta\psi, \quad \rho(R, t) = \rho_0 + \delta\rho,$$

and expand the nonlinear term as

$$\nabla[\rho(1 - \rho) \nabla u(\rho)] \approx \rho_0(1 - \rho_0) u'(\rho_0) \Delta \delta\rho.$$

Transforming to Fourier space and carrying out the calculation detailed in Appendix (B.6), we obtain, for a particular choice of parameters, the stability criterion (spinodal curve, in the (ρ_0, β) plane)

$$\beta < \frac{2 f_\beta(0)}{\rho_0(1 - \rho_0)(u'(\rho_0) - \frac{\alpha\lambda}{\kappa})}.$$

This analytic result was derived by fixing all parameters except β and ρ_0 . We then performed a suite of numerical simulations in which β and ρ_0 were varied freely to map out both the spinodal (stability boundary) and the binodal (coexistence curve). Notably, when the coupling

parameter between the price and density fields, α , is of the same order as the mean-reversion rate κ , the stability condition matches the one obtained in the simpler setting of chapter 3. As shown in Fig. (4.5), the simulation-derived curves coincide precisely with the analytic predictions, confirming that—even in the presence of a dynamic price field—the system can undergo condensation into a sub-optimal cluster.




Figure 4.5.: Phase diagram for the coupled system with the set of parameters : $\alpha = 1, D = 10, \kappa = 1, \lambda = 1$, with order parameter $q = \rho_l - \rho_g$ where ρ_l and ρ_g are the average density values in the liquid and gas phases. The black curve is the spinodal $\beta_{sp}(\rho_0)$ which characterizes the stability frontier, obtained analytically, and the black empty dots are representing the binodal of the system, meaning the region where both phases coexist, obtained via numerical simulations.

4.5. Conclusion

In conclusion, we have proposed a simple, general dynamical model for the spatial evolution of housing prices inspired from the robust statistical regularities found in the French data, in particular the logarithmic dependence on distance of the spatial variogram of prices. Indeed this is a signature of two-dimensional diffusing fields driven by random noise, captured by our stylized model, Eq. (4.1), which was already used in the past to model spatial regularities in voting patterns [88, 162]. Note that a model where prices propagate in a ballistic way ($r \sim t$) instead of diffusing ($r \sim \sqrt{t}$) would lead to completely different ($\sim e^{-\ell}$ instead of $\log \ell$) spatial correlations, see Appendix (B.3). The temporal fluctuations of prices can be accounted for within the same framework, provided the shocks are persistent over a time scale that we find to be around 3 years. The data also suggests, not surprisingly, that the amplitude of the price shocks is spatially heterogeneous, with a large variation span. The order of magnitude of the diffusion constant was estimated through an impulse response analysis. Other dimensional parameters obtained from fitting the spatial and temporal correlations appear to be of reasonable order of magnitude.

Our study thus quantifies the diffusive nature of housing prices that was anticipated long ago [169, 87], albeit on more restricted, local data sets. The possibility of describing the spatio-temporal dynamics of housing prices is clearly interesting from many standpoints, in particular for land use planning and territorial development policies or for real estate investment strategies as a response to “shocks”, like the opening of a fast train or a metro station. Extending our analysis to other spatial socio-economic variables would also shed light on the mechanisms underlying diffusion of socio-cultural traits, as sug-

gested in [175]. Lastly, we demonstrate that integrating the Sakoda–Schelling segregation dynamics with the housing-market diffusion equation still produces condensation, once again contradicting Adam Smith’s notion of the “invisible hand.”

Key Takeaways

1. **Logarithmic Spatial Variograms:** Empirical variograms of French housing prices grow as $\log \ell$, a hallmark of two-dimensional diffusion under white noise, rather than the exponential decay one would obtain from a ballistic propagation model.
2. **Unified Spatio-Temporal Model:** A single diffusion equation (Eq. (4.1)) captures both the logarithmic spatial correlations and the temporal price fluctuations, provided shocks persist on a characteristic timescale of about three years.
3. **Heterogeneous Shock Amplitudes:** The magnitude of price-shock noise varies substantially across space, reflecting locally diverse market conditions and reinforcing the need for spatially resolved modelling.
4. **Diffusion Constant Estimation:** An impulse-response analysis yields a diffusion coefficient of the correct order of magnitude, validating our stylized framework against real-world data.
5. **Policy and Investment Implications:** By quantifying how prices diffuse—e.g. after a new transport link—our model provides actionable insights for land-use planning, territorial development, and real estate strategies.
6. **Persistent Condensation under Coupling:** Even when the Sakoda–Schelling model is coupled to a housing market, the system continues to condense into suboptimal clusters, refuting the idea that decentralized interactions necessarily yield efficient, self-regulating outcomes.

PART II: FROM AGENT-BASED MODELS TO MINIMAL DISEQUILIBRIUM MODELS

An out-of-equilibrium agent-based model of the labor market

5.

*Quoi qu'il arrive n'oublie jamais ceci : notre existence
s'écoule en quelques jours. Elle passe comme le vent du
désert. Aussi, tant qu'il te restera un souffle de vie, il y a
deux jours dont il ne faudra jamais t'inquiéter : le jour
qui n'est pas venu, et celui qui est passé.
Avicenne ou La Route d'Ispahan—Gilbert Sinoué,
1989.*

This chapter is a collaborative effort with *J. Moran, R.M. del Rio-Chanona* and *J.D. Farmer*. The drafting of the manuscript was carried out in collaboration with the co-authors.

5.1. Introduction

This chapter advances a near-zero-intelligence, out-of-equilibrium model that reproduces the Beveridge curve, offering an alternative to the canonical equilibrium-with-rational-agents paradigm outlined in the introduction. By situating the analysis within complexity-economics perspectives—where disequilibrium and bounded rationality are essential for dynamic phenomena [52, 53, 16]—we build on evidence that off-steady-state adjustments shape the curve [54]. Extending earlier agent-based approaches that incorporate bounded rationality [55] and network effects [56, 57], we show how such a framework can account for labor-market behavior over the business cycle.

This chapter is organized as follows. Section (5.2) surveys the labor-market literature, focusing on the Diamond–Mortensen–Pissarides paradigm and the agent-based model of [57]. Section (5.3) derives the aggregate dynamics of that agent-based framework. Section (5.4) then proposes a parsimonious reduced-form specification—with a Cobb–Douglas matching function—to identify the minimal mechanisms that reproduce the Beveridge cycle. Finally, Section (5.5) tests the model's predictions against data from several OECD economies, including the United States.

5.2. Literature review

We begin by recapping the canonical Diamond–Mortensen–Pissarides (DMP) model introduced earlier, then describe the agent-based labor-market framework of [57] that underpins our subsequent analysis.

The Diamond-Mortensen-Pissarides framework

A cornerstone of the search and matching models of labor markets is the Diamond-Mortensen-Pissarides (DMP) framework (see [18, 19, 20]), which models the dynamics of unemployment and vacancies by describing how firms post vacancies, how employed workers become separated from their job and how unemployed workers are matched with a job.

In that framework, the differential equation for unemployment is generically written as

$$\frac{du}{dt} = s(1 - u(t)) - m(u(t), v(t)), \quad (5.1)$$

where $u(t)$ and $v(t)$ are the unemployment and vacancy rates, s is the state-independent separation rate, and m is a *matching function* that determines the rate at which many unemployed workers are matched with vacancies.

The matching function is usually assumed to be increasing and concave with respect to u and v (indicating intuitively that the number of matches per unit time increases with the number of unemployed people and the number of vacancies) and is also assumed to have constant returns to scale with respect to u and v [180] (meaning that it satisfies $m(\lambda u, \lambda v) = \lambda m(u, v)$ for any λ). Therefore, the probability that a worker is matched with a vacancy¹ is $m(u, v)/u = m(1, \theta)$, where $\theta = v/u$ is the *labor market tightness*.

The equilibrium level of labor market tightness, denoted θ^* , depends on the profit maximization strategies of both firms and workers. A pivotal assumption in this context is the free entry condition, indicating that any vacancy created will be filled and that firms continue to post vacancies as long as the expected payoff for hiring workers remains positive.

Implicitly, this assumes that firms are not only rational enough to compute the expected payoff, but that they have enough information to do so accurately. Other standard assumptions include a well-determined interest rate, known to all firms, that determines the arbitrage between current costs and anticipated future profits. Firms are assumed to have perfect knowledge of their expected future profits, as outlined by [18]. This results in dynamics for vacancies in which the only non-divergent solution implies that vacancies are consistently adjusted to maintain a constant market tightness, since firms have perfect foresight and have all access to the same information. This lets them act in perfect coordination in a way that achieves the optimal tightness ratio. The consequence of all of this is that firms observe the unemployment rate and post vacancies that satisfy

$$v = \bar{\theta}u, \quad (5.2)$$

establishing the unique perfect-foresight adjustment path.

The equilibrium reached, corresponding to the stationary state where $\frac{du}{dt}$ is equal to zero, has characteristics that may shift due to economic shocks, but classical models presume that the system moves from one equilibrium point to another without considering out-of-equilibrium effects.

1: Finding an unfilled vacancy is however not a guarantee of being hired: firms will only consider candidates with a productivity above a certain threshold (applications from workers with low productivity are rejected). Similarly, workers may have a reservation wage which would lead them to reject certain offers. Specifying the details of these considerations then allows the modeler to determine the matching function m .

Separation of timescales

The DMP framework implicitly models the labor market with different timescales. The firms decide on their vacancies based on the unemployment level and do the necessary arbitrage to reach equilibrium at a *fast* timescale, which allows vacancies to be adjusted quickly enough to ensure that $v(t) = \bar{\theta}u(t)$ at all times. The constant value of $\bar{\theta}$ is maintained through optimization by firms who are constrained by exogenous parameters, such as workers' bargaining power, their reservation wages, the productivity distribution and interest rates, in the absence of structural changes. The exogenous parameters that affect θ change on a timescale that is assumed to be slow compared to the time it takes for vacancies and unemployment to reach equilibrium.

This separation of timescales is crucial for the explanations of the Beveridge curve cycle put forth by [18]. This curve represents points on the (u, v) plane, and it is called a cycle because during business cycles (that is, periods during which the economic environment changes cyclically over periods of the order of 3 – 12 years) the labor market has a trajectory that appears as an anti-clockwise loop on that plane. In this interpretation, each point on the Beveridge curve is at the intersection of the $v = \bar{\theta}u$ curve and the curve corresponding to $\frac{du}{dt} = 0$, as shown in Figure 5.1.

The emergence of the anti-clockwise cycle is explained by the following thought experiment. Starting from a given point (E1) with an associated value of $\bar{\theta} = \bar{\theta}_0$, imagine a positive output shock causing $\bar{\theta}_1 > \bar{\theta}_0$.² These changes shift the position of the two curves defining the stationary equilibrium, $v = \bar{\theta}u$ and $\frac{du}{dt} = 0$.³ Because of the increase in $\bar{\theta}$, the stable trajectory line $v = \bar{\theta}u$ shifts to the left and the locus $\frac{du}{dt} = 0$ shifts inwards, setting a new stationary point for the labor market in the (u, v) plane (E2), as shown in Figure 5.1. Because vacancies instantaneously adapt to output shocks, the economy first reaches the new $v = \bar{\theta}u$ curve, then relaxes along this curve towards E2. Now, assume that there is a negative output shock of the same magnitude as the previous one. Following the same process, the economy will reach the new stable trajectory line $v = \bar{\theta}u$, which is exactly the same as the original one, and then relax towards its intersection with the new locus $\frac{du}{dt} = 0$, which is also the same as the original one (E1). This is explained in detail by [18].

The argument made in the DMP framework is that, because shocks happen on a slow timescale relative to the speed for reaching the equilibrium between vacancies and unemployment, the labor market reacts very quickly and immediately adjusts to the new stationary point. The whole argument can be extended to obtain a Beveridge cycle curve during boom-bust cycles. Thus, the frictions of the job market that create the cycle are implicitly hidden in the dependence of the solutions of the equation $\frac{du}{dt} = 0$ on the parameters governing the job market, in particular the matching function. Parameters such as the matching function's elasticities, vacancy posting costs, and bargaining power are difficult to pin down empirically due to a combination of a lack of directly observable data and identification challenges, thereby limiting the predictive power of the DMP model.

2: Because $\bar{\theta}$ measures the number of vacancies per job seeker, an increase in $\bar{\theta}$ makes it easier to find and apply to a vacancy, and therefore increases the matching efficiency.

3: In particular, $\frac{du}{dt}$ is a function of m which is itself a function of $\bar{\theta}$. This ensures that the $\frac{du}{dt} = 0$ curve moves after an output shock.




Figure 5.1: $\frac{du}{dt} = 0$ curves and $v = \bar{\theta}u$ curves. The positive output shock, because of the increase in $\bar{\theta}$, shifts the stable trajectory line $v = \bar{\theta}u$ to the left and the locus $\frac{du}{dt} = 0$ inwards. The labor market first reaches the new $v = \bar{\theta}u$ line, then converges to the new intersection of these two curves. The intersection points E1 and E2 correspond to the locus of the labor market before and after the shock, highlighting an anti-clockwise trajectory in the (u, v) plane. We reproduce here Figure 1 from [18].

Recent empirical predictions using the DMP framework

[62, 181] have used the DMP framework to make empirical predictions, postulating that the matching function is a Cobb-Douglas function, with $m(u, v) \propto v^\eta u^{1-\eta}$. The use of the Cobb-Douglas function is justified by its scale invariance and its ability to capture diminishing marginal returns to vacancies and unemployment, which ensures that additional job openings or job seekers generate progressively fewer new matches. Another important reason is analytical tractability.

Using the hypothesis of a fixed tightness of the job market, $v/u = \theta$, Eq. (5.1) is simplified to become

$$\frac{du}{dt} = s(1 - u(t)) - fu(t), \quad (5.3)$$

where $f = \frac{m(u(t), v(t))}{u(t)}$ is a constant representing the number of matches per unemployed worker per unit time. The combination of the Cobb-Douglas assumption and the fixed value for θ result in a constant matching rate of $f \propto \theta^\eta$. This equation can be solved to obtain an exponential convergence to a stationary state, namely $u(t) = u_b + Ce^{-(s+f)t}$ with C a constant depending on the initial conditions and the stationary value $u_b = \frac{s}{s+f}$. The economy therefore converges to the stationary state at an exponential rate with a timescale of $(s + f)^{-1}$. If this timescale is small, it means effectively that an exogenous shock knocks the labor market away from its stationary state for a very short time, thereby justifying an analysis through comparative statics.

[181] use estimated empirical values of $s = 3.2\%$ per month and $f = 55.8\%$ per month to obtain a convergence speed of $s + f = 59\%$ per month. This leads to exponential convergence with a half-life of a little more than a month. They conclude that the job market converges extremely quickly, with a time scale that is much smaller than the average duration of business cycles. Because the parameter f , which determines the rate at which workers are matched to vacancies, is of the order of 50% per month, this means that on average an unemployed worker has the order of a 50% chance of getting a job every month. This prediction is used as a validation for the assumption that vacancies adjust instantaneously made in the DMP framework. In particular, such a rapid convergence of the economy to its stationary state implies that

there should be no phase lag between unemployment and vacancies; they should always remain maximally anti-correlated.

Shortcomings of the DMP framework

Recent studies, such as [74], challenge the arguments of [181]. After calibrating his model on French labor market data, [74] determines a half-life of 10 months after a symmetric shock⁴ which is of the same order of magnitude as our own result. Our work here supports the conclusion that the relative timescale is much longer than a month. [74] finds that the half-life for non-symmetric shocks, i.e. those that target some occupations more than others, can be up to 150 months, an order of magnitude bigger than that of symmetric shocks, which further challenges the results of [181].

Moreover, the DMP framework does not take into account the labor network. A worker who is unemployed and has the skills required for job i needs time to acquire the skills needed to convert to a new job j ⁵. This creates frictions in the labor market and leads to a slower theoretical convergence time scale which depends on the ease with which workers can transition from one occupation to the next. Intuitively, this means that the presence of occupational bottlenecks prevents the economy from settling into a new stationary state quickly after a shock. Formally, this can be related to the spectral properties of the transition network of the occupational mobility network. As discussed in more detail in the next sections, the conclusions are in good agreement with the work of [74].

A detailed empirical study of the properties of the occupational transition matrix of the French labor market was also carried out in [76]. This matrix A_{ij} is defined by the fraction of workers in occupation i that transitioned to j . If the job market was close to a stationary state, one would expect the number e_i of workers in occupation j to satisfy $\sum_j e_i A_{ij} = \sum_j e_j A_{ji}$, so that the number of workers leaving occupation j matches the number of workers entering that occupation.

Solving for the values of e_i satisfying this equation leads to the implied stationary number of workers e_i^* . If the labor market was close to its stationary state, as is often assumed, then the data should show $e_i \approx e_i^*$. [76] find however that this is not at all the case, and that nearly all occupations have values that are significantly different from their implied stationary values.

Agent-based models of labor markets

In contrast with the DMP framework defined previously, agent-based models, particularly those in line with the framework started by [77], have labor dynamics that are tied to firms' adaptive expectations of market demand, which determine their production targets.

In models that simulate economies out of equilibrium, the absence of market coordination means that agents must make decisions without full knowledge of the actions and intentions of other agents. These decisions, whether rational or heuristic, are determined by backward looking expectations rather than by equilibrium conditions imposed

4: Meaning a shock that has the same impact on all occupations, which is also the case for the results presented later in this work when studying the agent-based model at the aggregate level.

5: In general, workers cannot immediately transition to any occupation, which is already in contradiction with the arguments presented in the DMP framework.

on the economy, such as market clearing or competition/no arbitrage conditions that would align the actions of all the agents. After firms have determined their production targets, they assess their labor levels to determine if they are understaffed or overstaffed. This allows them to adjust their workforce accordingly, firing excess employees if overstaffed or hiring additional labor to meet production needs. This process is not instantaneous and is subject to frictions that may arise from delays inherent to the hiring/firing process, such as administrative hurdles or contractual obligations, or from limited availability of workers in the hiring pool. Firms may also be reluctant to react strongly to perceived temporary fluctuations: for example, a firm that determines it has twice the amount of workers it needs may be reluctant to fire half of its workforce if it perceives that a drop in demand, and therefore a drop in its target production, will be only temporary.

This approach to labor dynamics aligns with the perspectives broadly adopted in the agent-based modeling literature. In most of these models (for example, [29, 75, 80]) the hiring process is quite straightforward: firms select their hires directly from a pool of unemployed workers, typically without distinguishing between occupations. Other models, such as [77], explicitly model the matching process to fill vacancies. As in the DMP framework, hiring becomes a multi-step process where firm post a vacancies and unemployed workers apply for jobs.

Although the features we describe above are prevalent in many ABMs, we focus in this section specifically on the model proposed by del Rio-Chanona et al. [57], since it is the first to highlight the presence of a Beveridge cycle within such a framework. (Its original purpose was to study how changes in occupational demand could affect unemployment, but the Beveridge curve emerged as a side-effect). This model integrates vacancy posting with a categorization of firms and workers into different industries and occupations, offering a more granular view of labor dynamics based on parsimonious assumptions. We provide a description of the model below and explain how it generates Beveridge cycles.

Description of the model

In this model, workers have an occupation and may be employed or unemployed at any given time. A worker goes from a state of being employed to being unemployed when they become separated from their job, something that happens at random. Unemployed workers seek new jobs, but may only apply to jobs that correspond to their skill set. The jobs available to a given worker are determined by an *occupational mobility network*, which dictates the frequency with which a worker may move from one occupation to another.⁶ This network imposes transition rates, determining the probability A_{ij} with which a worker initially in occupation i may apply to a job in occupation j . This is the only behavioral assumption in the model: for the sake of simplicity, it neglects both geography and wage dynamics.

Each occupation i is described by its employment rate $e_{i,t}$, its unemployment rate $u_{i,t}$ and its vacancy rate $v_{i,t}$; the latter is fully determined by firms. These rates are linked by the conservation equation

6: We note however that the same phenomenology can be reproduced with the Mark-0 model introduced by [37], a fact that was pointed to us by the authors of [79].

$$u_{i,t} + e_{i,t} = 1.$$

The model distinguishes two types of demand for occupational labor. $d_{i,t} = e_{i,t} + v_{i,t}$ is the *posted demand* for workers in occupation i , whereas $d_{i,t}^+$ is the *target demand* for workers, which is determined by expectations of future economic conditions.

This is meant to represent the mechanism through which firms adjust their hiring targets in agent based models. Indeed, in such models a firm f is able to attain a given level of production $Y_{f,t}$ that is a function of its current capital, intermediate inputs, and, crucially, its current workforce $L_{f,t}$. For example, if there is an increase in the demand for its good or an increase in its production costs, firm i changes its production target $Y_{i,t}^+$, which requires a different target workforce, $L_{i,t}^+$. The procedure through which the firm computes its target production and target workforce levels can be fully rational or can be obtained using rules of thumb, but the common point in ABMs is that this procedure is done with only partial information from the economy and without full coordination between economic agents. As a consequence, there is necessarily a lag between a firm's target workforce $d_{i,t}^+$ and its actual workforce $d_{i,t}$.⁷ The realized/target demands can thus be seen as the aggregate demand from all firms for a given occupation.

Although most ABMs do not explicitly model vacancy dynamics, this model explicitly models both vacancies and employment, while treating changes in labor demand as a phenomenological modeling of firms' decisions. The decisions they make impact hiring/firing rates through the discrepancy between the realized demand $d_{i,t}$ and the target demand $d_{i,t}^+$.

With this in mind, the basic dynamics of the model can be expressed in the style of the DMP framework [18]. Here, the temporal evolution of the unemployment rate, for instance, is determined solely by spontaneous separations and the formation of job matches, resulting in the following equations:

$$\begin{aligned} e_{i,t+1} - e_{i,t} &= -\omega_{i,t+1} + \sum_j f_{ji,t+1} \\ u_{i,t+1} - u_{i,t} &= \omega_{i,t+1} - \sum_j f_{ij,t+1} \\ v_{i,t+1} - v_{i,t} &= v_{i,t+1} - \sum_j f_{ji,t+1}, \end{aligned} \tag{5.4}$$

where $\omega_{i,t+1}$ represents the number of workers separated between t and $t+1$, $v_{i,t+1}$ denotes the number of vacancies created, and $f_{ij,t+1}$ is the number of unemployed workers who previously worked in occupation i and have been matched with a job in occupation j . This corresponds to the matching function term $m(u, v)$ in the DMP framework. These equations also ensure the preservation of the total workforce, $L(t) = e_{i,t} + u_{i,t} = \text{cst}$.

To capture the mechanism described above—where firms fire workers when employment exceeds their target and hire when below it—we specify the firing rate in our model as:

7: In some cases, it may not be possible for firms to reach their objectives simultaneously, leading to an economy with endogenous fluctuations [75].

$$\omega_{i,t+1} = \delta_u e_{i,t} + (1 - \delta_u) \gamma_u \left[d_{i,t} - d_{i,t}^+ \right]^+. \quad (5.5)$$

The first term, $\delta_u e_{i,t}$, represents spontaneous separations: any worker can be fired with probability δ_u . On average, this corresponds to $\delta_u(1 - u_{i,t})$ workers being fired per iteration, mirroring the $s(1 - u)$ term in [18]. The second term arises when the current number of positions $d_{i,t}$ exceeds the target $d_{i,t}^+$. In that case, firms additionally lay off workers at a rate $(1 - \delta_u) \gamma_u \left[d_{i,t} - d_{i,t}^+ \right]^+$, where $[\cdot]^+ = \max(0, \cdot)$ ensures this term only applies if the firm is above its employment target, and γ_u sets the speed of adjustment. This can be viewed as a linear approximation, valid when $d_{i,t} - d_{i,t}^+$ is sufficiently small.

Similarly, the number of vacancies is adjusted depending on how far firms are below their hiring targets. A vacancy is added either through a random spontaneous term, corresponding to a rate δ_v , or because of the discrepancy between realized and target demands. This leads to

$$v_{i,t+1} = \delta_v e_{i,t} + (1 - \delta_v) \gamma_v \left[d_{i,t}^+ - d_{i,t} \right]^+. \quad (5.6)$$

Unlike the DMP framework, vacancies here do not instantly adjust to new macroeconomic conditions; instead, they evolve in time, just as the employment does. Crucially, these equations imply that if firms greatly overshoot their employment targets, their firing rate spikes – but if the target later shifts quickly, it will take time to hire enough workers to reach the new target.

Finally, in this model, $d_{i,t}^+$ (the target demand) is treated as exogenous. It may be driven by macroeconomic shocks (such as booms and busts or crises like the COVID-19 pandemic) or by technological shocks (e.g., automation reducing demand in certain occupations). We therefore think of this model as the study of the labor market in isolation, responding to the rest of the economy and neglecting the feedback from labor to the rest of the economy.

We impose a sinusoidal exogenous target demand to develop an intuitive understanding of the model's core mechanism. The sinusoidal pattern is chosen to reflect typical business cycle frequencies observed in the US.

Finally, the term $f_{ij,t+1}$ captures the *matching process* that describes how unemployed workers get hired. In the model, unemployed workers from occupation i submit applications to occupation j and are hired with some probability. The total number of applications sent from i to j , denoted $s_{ij,t+1}$, is

$$s_{ij,t+1} = u_{i,t} \frac{v_{j,t} A_{ij}}{\sum_k v_{k,t} A_{ik}}, \quad (5.7)$$

where A_{ij} is the *adjacency matrix* of the occupation network and reflects how strongly workers in i prefer to apply to jobs in j .

Intuitively, the probability that an unemployed worker in i applies to a vacancy in j is proportional to $v_{j,t} A_{ij}$. Note that, in contrast with

other models, workers do not take wages explicitly into account here, although among other things, A_{ij} could depend on the relative wages of occupations i and j . This allows us to focus on frictions arising from mismatches between workers' skills and vacancy requirements. Moreover, our model assumes *wage rigidity*, a standard premise in classical economics as outlined in [182], meaning that wages remain essentially fixed in the short term. In the DMP framework, wage dynamics are also only partially addressed, as they are set via flexible Nash bargaining.

The number of workers moving from occupation i to j equals the total number of applications, times the probability that these applications are accepted. Each occupation receives $s_{j,t+1} = \sum_i s_{ij,t}$ applications for $v_{j,t}$ vacancies. In the limit of a large number of agents, it is possible to show that the expected fraction of vacancies filled is $1 - \exp(-s_{j,t+1}/v_{j,t})$, so we finally obtain

$$f_{ij,t+1} = \frac{s_{ij,t+1}v_{j,t}}{s_{j,t+1}} \left(1 - e^{-s_{j,t+1}/v_{j,t}}\right). \quad (5.8)$$

This expression is the number of applications from industry i to industry j times the fraction of applications that are matched. As $s_{ij,t+1} \propto u_{i,t}$, the flow of workers is an increasing function of the labor tightness ratio $\theta = v/u$. The fraction of matched workers can therefore be written as $p(\theta) = \left(1 - e^{-\frac{1}{\theta}}\right)\theta$, an increasing function of θ . This is precisely the same urn matching process introduced by [180], but supplemented with an occupational network.

At the aggregate level, we define the matching function

$$m(u_{t+1}, v_{t+1}) = \sum_i \sum_j f_{ji,t+1}. \quad (5.9)$$

In the continuous-time limit, this becomes

$$m(u(t), v(t)) = \sum_i \sum_j f_{ji}(t) \quad (5.10)$$

With these components in place, the model is fully specified except for the exogenously determined target demand, $d_{i,t}^+$.

We now rewrite the model in continuous time⁸ and sum over all occupations to obtain differential equations on aggregate quantities (e.g. the aggregate unemployment is $u(t) = \sum u_i(t)$), and its rate of change is

$$\frac{du}{dt} = \delta_u(1 - u) + \gamma(1 - \delta_u) \max(0, D(t) - D^+(t)) - m(u, v). \quad (5.11)$$

The first term $\delta_u(1 - u)$ represents the total number of workers who are separated at random, independently of economic conditions. The second term, $F_u(D(t) - D^+(t)) = \gamma(1 - \delta_u) \max(0, D(t) - D^+(t))$, represents state-dependent hiring and firing, conditional on the gap between realized demand $D(t)$ and target demand $D^+(t)$. The function F_u describes the rule used by firms to increase or decrease the firing

8: Numerical checks confirm that reducing the time-step leads to a well-defined continuous limit with the same phenomenology as the discrete-time model.

rate depending on how much they deviate from their workforce objectives.

Our primary aim, in this chapter, is to develop an approximation for these equations that retains the main mechanisms of the model but is tractable enough to understand the model's aggregate behavior. We note that substituting a Cobb-Douglas matching function $Ku^\alpha v^{1-\alpha}$ for the urn-ball matching function previously discussed does not change the phenomenology of the model.

In chapter 6, we introduce a simpler, more intuitive out-of-equilibrium framework for analyzing and modeling the US labor market.

In contrast to the DMP framework - where forward-looking, rational firms calculate their ideal number of vacancies and adjust them instantaneously - our firms follow a myopic rule of thumb. Firms increase vacancies if they are below their employment target, and vice-versa, according to

$$\frac{dv}{dt} = \delta_v(1 - u) + \gamma(1 - \delta_v) \max(0, D^+(t) - D(t)) - m(u, v), \quad (5.12)$$

where again we introduce $F_v(D^+(t) - D(t)) = \gamma(1 - \delta_v) \max(0, D^+(t) - D(t))$.

The terms F_u and F_v in $\frac{du}{dt}$ and $\frac{dv}{dt}$ are not symmetric. Only one of them is non-zero at a given time t . This asymmetry is responsible for the cycle observed in the Beveridge curve and will be discussed further. The state-independent rates δ_u and δ_v drive the direction of the Beveridge curve cycle. An anti-clockwise cycle occurs when δ_u exceeds δ_v , which causes unemployment to lag vacancies when demand is changing (see Figure (C.3) for an example of the model's output when $\delta_u < \delta_v$). The inherent asymmetry between unemployment and vacancies is well known in the literature and has been verified empirically.

For completeness, other technical details (such as the derivation of the model's steady-state convergence) are provided in Appendix (C.2).

In a full simulation, the agent-based model typically includes around 500 different occupations with inter-occupational transitions. Although the two differential equations above might appear to be a substantial simplification relative to the full agent-based model, they cannot be solved analytically. This is because Eq. (5.10) expresses the matching function $m(u, v)$ as a sum over occupations, and there is no straightforward way to express this sum in closed form.

Replicating the Beveridge curve cycle

Although [62] emphasize that the empirical Beveridge curve roughly traces out straight lines in log-log scale, with a slope close to -1 in the vacancy-unemployment plane, deviations from this are apparent in the data. These deviations largely correspond to what we call the *cycle* in the Beveridge curve, i.e. to understand it, we need to understand why these deviations occur as they do. Following [57], we calibrate the parameters of the model to the US labor market and assume a sinusoidal cyclical change in demand with a frequency that roughly corresponds to the 2008 recession. As shown in Figure 5.2, the model

qualitatively replicates the observed behavior, with slopes of roughly -1 during the recession and -0.7 during the recovery.

We attribute this match with empirical data – together with the calibrated parameters in which $\delta_u > \delta_v$, ensuring an anti-clockwise cycle – to the relatively slow adjustment time of the economy. The essential ingredients to obtain a cycling Beveridge curve are that the workforce objectives targeted by firms change at a pace that is much faster than the time it takes the firms to reach them, and that $\delta_u > \delta_v$. When the economy contracts, firms can fire workers quickly; but when the economy recovers, firms are only able to increase their vacancies and hire at a slow pace. This asymmetry prevents the economy from simply retracing the steps in the recovery that it took in the recession.

Mimicking comparative statics with the ABM

We now use the agent-based model to highlight the difference between a traditional DMP framework using comparative statics with our dynamical disequilibrium approach. We also hope this discussion helps clarify how to design simpler, solvable models starting from a complicated ABM.

Suppose we wanted to explain the Beveridge curve with the ABM in terms of comparative statics. We can illustrate the idea by setting the parameters so that the system converges to its steady state quickly, then doing stepwise changes in the target demand. As seen in Figure (C.4), the system relaxes exponentially to each new steady state. This means that one could try to explain labor market trajectories via shifting demand targets.

But this approach would not produce a Beveridge *cycle*: simulating a crisis followed by a recovery in a quick-convergence regime, the system would simply retrace the crisis path in reverse during the recovery. The only way to produce a cycle would be to adjust the parameters on the recovery path so that they differed from those on the recession path. This could be done, but then one would need to explain why the parameters varied as they did. In contrast, as seen above, when we use realistic parameters that yield longer relaxation times, the ABM automatically generates a Beveridge cycle that can match data without any need for auxiliary assumptions or parameter adjustments.

5.3. The dynamics of the Beveridge curve

Convergence to the steady state

The fact that the convergence to the steady state is exponential suggests an approach for approximating (5.11) and (5.12) with a solvable model. In the comparative statics experiment described above, the convergence to the steady state is well-described by an exponential decay, of the form

$$D(t) = \alpha D_1^* + \alpha \left(D_0^* - D_1^* \right) \exp \left[-\frac{t}{\tau} \right], \quad (5.13)$$

where τ is the characteristic convergence time and D_0^\dagger, D_1^\dagger represent the target demands for the initial steady state and the final steady state. For realistic values in our simulations, $\alpha \approx 1$.

This leads us to model the behavior of the demand relaxation via a simple first order differential equation,

$$\tau \frac{D(t)}{t} + D(t) = \alpha D^\dagger(t), \quad (5.14)$$

where, in this case, $D^\dagger(t) = \mathbf{1}_{t \geq 0} D_1^\dagger + \mathbf{1}_{t < 0} D_0^\dagger$.⁹ By adjusting τ appropriately, we can match the convergence of the full network model with this simple first order differential equation very accurately. The value of τ reflects the friction of the labor market.

In the full ABM, the convergence time τ is determined by the relaxation rate γ and the efficiency of the matching process. The relaxation rate is an external parameter, while the matching efficiency depends on coordination within the search process and the alignment of skills. Coordination issues are modeled through an urn-ball matching function, whereas skill mismatches are determined by the network's structure. Less interconnected occupational networks lead to slower convergence times. This relationship is depicted in Figure C.2, illustrating that networks with smaller spectral gaps, indicating higher connectivity, are associated with reduced convergence times.

In an ideal scenario with a fully connected network – where any worker can transition to any occupation – and where frictions arise solely in the search-and-match process, the convergence time is approximately 5 months. If we introduce realistic network frictions, based on occupational mobility data that capture which occupations a worker can realistically transition to, this period increases to 9 months (refer to Figure (C.4)). These findings roughly align with those of [74], who estimated a 10-month convergence period for the French labor market under a homogeneous shock. This is in sharp contrast with [62]'s assumption of a 1-month convergence time. This result underscores the importance of both the relaxation rate and the network connectivity in determining labor market dynamics.

9: This equation may be reinterpreted as $D(t + t) = (1 - \frac{t}{\tau}) D(t) + \frac{t}{\tau} D^\dagger(t)$, i.e. the realized demand is an exponential moving average of the target demand. This aligns with the notion of sluggish adjustments in ABMs. Interestingly, the DMP framework assumes a similar sluggishness, albeit only for unemployment.

Out-of-equilibrium dynamics

Mathematical description

Assuming a sinusoidal change in demand, as before and substituting into Eq. (5.14), the general solution is

$$D(t) = D(0)e^{-\frac{t}{\tau}} + \int_0^t s \alpha D^\dagger(s) e^{-\frac{t-s}{\tau}}. \quad (5.15)$$

To get an understanding of the behavior of the system, suppose that during a business cycle the target demand changes as

$$D^\dagger(t) = A_1 + A_2 \cos(\omega t), \quad (5.16)$$

where $\omega := \frac{2\pi}{T}$ is the frequency of the oscillation and $A_1 > A_2$.

For large t , this yields




Figure 5.2: a) Output of the model when simulating the Beveridge curve with parameters fitted on US data, using the occupational mobility network, in log-log scale. The behavior of the target demand $D^\dagger(t)$ is shown in the inset, with an amplitude of order 2.5×10^5 and a period $T \approx 10$ years. The time steps correspond to 6.75 weeks and these parameters were chosen to reproduce typical US business cycles. The recession periods are represented with grey dots and the recovery periods are represented with magenta dots. The direction of time is highlighted with black arrows, showing that the curve cycle progresses counter-clockwise over time. During recession, the slope of the curve is approximately -1 whereas during recovery the slope is around -0.7. The empirical US Beveridge curve for years 2000-2019 is plotted in black. b) Area of the Beveridge curve cycle when we hypothetically vary the frequency of the demand shock, for both the occupational mobility network based on U.S. data (in red) and a fully connected network (in black), shown for comparison. High frequency shocks induce stronger hysteresis, causing faster Beveridge curve cycles. The curve for the occupational mobility network is above the one for the fully connected network, showing that the network plays a part in increasing the market frictions.

$$D(t) \approx \alpha A_1 + A_2(\omega) \cos(\omega t + \varphi(\omega)), \quad (5.17)$$

where $A_2(\omega)$ is the amplitude of the realized demand and $\varphi(\omega)$ is a phase shift, with¹⁰

$$\begin{cases} A_2(\omega) &:= \alpha \frac{A_2}{\sqrt{1+\omega^2\tau^2}} \\ \varphi(\omega) &:= -\arctan(\omega\tau). \end{cases} \quad (5.18)$$

This result has an intuitive interpretation. If the shock is too sudden, i.e. $T \ll \tau$ or equivalently $\omega\tau \gg 1$, then $A_2(\omega) \approx 0$ and the labor market does not oscillate at all because it cannot catch up to the cyclical driving due to changes in demand. For values of T of order τ , $\omega\tau = \mathcal{O}(1)$ and so the reaction of the labor market is to oscillate with a damped amplitude, and with a phase shift given by φ . In other words, if the lowest point of target demand is reached at a time $T/2$, the lowest realized demand (and therefore the highest level of unemployment) will be reached at $\frac{T}{2}(1 + \frac{\arctan(\omega\tau)}{2\pi})$. If $\tau \ll T$, then this reduces to $(T + \tau)/2$, indicating that the labor market takes roughly time τ to catch up. Since business cycles have timescales on the order of 3-12 years, whereas τ is the order of 9 months, the economy is normally in this limit. The functional form introduced here agrees well with the output of the model in Figure (C.5), which paves the way to building a simplified model.

10: The general solution to this equation can be obtained by performing the integral $\int_0^t ds \cos(\omega t) e^{-\frac{t-s}{\tau}} = \Re \left(\int_0^t ds e^{-\frac{t-s}{\tau} + i\omega t} \right)$, leading to the results in Eq. (5.18).

The simplified model above explains the asymmetric behavior of the system in the vacancy-unemployment plane, as explained in the Appendix. Periods of crisis are marked by a sharp decline in vacancy creation relative to firings, primarily due to the absence of the demand-related term, while firings simultaneously increase. In contrast, recoveries involve continued separations but are characterized by a sharp rise in hirings as the demand-related term becomes active. In both cases, spontaneous separation and hiring processes persist, but they are asymmetric because $\delta_u > \delta_v$. This asymmetry explains the differing behaviors of the dynamics during crises and recoveries, and why the economy does not simply oscillate in symmetric boom-and-bust cycles.

A key prediction from the model

As shown in (C.4), solving our aggregate model reveals that the difference between target demand and realized demand, $D^+(t) - D(t)$ is constant when the shock frequency is low and becomes an oscillating function as the shock frequency increases. Consequently, this results in no cycle in the Beveridge curve for low-frequency shocks, when the economy has time to adjust to the shock. On the other hand, higher frequency shocks induce a cycle, with the amplitude increasing with the frequency of the shocks.

As a result, higher business cycle frequency results in a greater mismatch between vacancies and unemployment, i.e. Beveridge cycles enclosing a greater area. This prediction is confirmed by numerical simulations, as illustrated in Figure (5.2) b) which indicates that the area of the Beveridge curve cycle is an increasing function of the shock frequency ω . This implies a direct relationship between the speed of the demand shock and the mismatch between vacancies and unemployment. The more rapid the shock, the more difficult it is for the economy to adjust.

In summary, the agent-based model accurately reproduces a cyclic Beveridge curve without relying on instantaneous adjustment assumptions, in contrast to classical models.

5.4. A generic dynamical model

In this section we take advantage of the simplified description of the demand dynamics to build a phenomenological model that approximates the aggregate behavior observed in the full agent-based model. We show that the counter-clockwise cycle of the Beveridge curve can be reproduced using an aggregate model with a standard Cobb-Douglas matching function, without having to rely on all the complexity of the agent-based model and without requiring an occupational network or any specific matching function (although we have checked that everything below holds also for a Pólya-urn type matching function as used in the agent-based model). This shows that the mechanism we propose is robust and generic, and only requires out-of-steady-state dynamics where vacancies and unemployment do not automatically adjust to their long-term value.

The key ingredients of our model can be defined as follows:

- A time dependent rate $s(t)$: there is a baseline separation rate s_0 that captures the regular turnover of the labor market due to random firings or quits. During periods of economic downturn, this rate increases to reflect higher job losses. This increase is modeled as an exogenous function driven by external economic factors.
- A time dependent rate $c(t)$: similarly, there is a baseline vacancy creation rate c_0 representing the constant need for firms to fill positions. During economic booms, this rate increases as firms expand and seek additional labor to boost production. This increase is also exogenously driven by economic factors that are not explicitly modeled here.
- The absence of instantaneous adjustments: unlike traditional models, where firms instantaneously adjust vacancies based on optimal foresight, our model assumes that both separation and vacancy creation rates adjust over time. This introduces frictions into the labor market dynamics. As a consequence, our model consists of a dynamical system of two coupled differential equations.

We define our model as

$$\begin{aligned}\frac{du}{dt} &= s(t) [1 - u(t)] - m(u(t), v(t)) \\ \frac{dv}{dt} &= c(t) [1 - v(t)] - m(u(t), v(t)),\end{aligned}\tag{5.19}$$

where the number of separations is assumed to be proportional to $1 - u(t)$, or to the total number of employed workers. Similarly, the number of created vacancies is assumed to be proportional to $1 - v(t)$, or the number of workers available if all the vacancies existing at time t were to be filled, i.e. a full employment economy.¹¹ For convenience we choose a matching function of the form $m(u, v) = Ku^\theta v^{1-\theta}$, i.e. a constant returns to scale Cobb-Douglas function with a matching efficiency K , or alternatively we use the Pólya urn function as in the ABM, $m(u, v) = K(1 - e^{-u/v})v$. We use a Cobb-Douglas matching function with $\theta = 1/2$, but other choices lead to the same qualitative results. The separation and vacancy-creation rates are modeled as functions of an exogenous cyclical factor, capturing economic booms and busts:

$$\begin{aligned}s(t) &= s_0 + A_u [\cos(\omega t)]^+ \\ c(t) &= c_0 + A_v [-\cos(\omega t)]^+\end{aligned}\tag{5.20}$$

where s_0 and c_0 are the baseline creation rates, A_u and A_v represent the sensitivity of the separation/vacancy creation rates to cyclical fluctuations, and ω is the frequency of the economic cycle (with period $T = \frac{2\pi}{\omega}$). As before, $[\cdot]^+ = \max(0, \cdot)$ indicates the positive part, which ensures that the separation rate increases during economic downturns, reflecting higher layoffs, while the vacancy creation rate increases during upturns to capture increased hiring effort. Thus, at the height of an economic boom the separation rate is s_0 , while the va-

11: This is a natural prescription for unemployment dynamics. For vacancies, we have also checked that other prescriptions, such as a term $c(t)v(t)$ or $c(t)(1 - u(t))$ leads to the same qualitative dynamics.

cancy creation rate is $c_0 + A_v$. In contrast, at the worst part of a crisis, the vacancy creation rate stays at its baseline c_0 while the separation rate increases to $s_0 + A_u$.

Emergence of the Beveridge curve cycle By simulating the system in Eq. (5.19) with the rates defined in Eq. (5.20), taking $A_u = A_v$ for simplicity and using a Cobb-Douglas matching function, we observe the emergence of the characteristic counter-clockwise cycle in the (u, v) plane, showing that this model is capable of reproducing the qualitative tear-drop shape of the curve (which is not present in the sinusoidal cycle driving the recession and recovery).

Although the specific form of the matching function affects the exact shape of the Beveridge curve and the employment dynamics, the emergence of the cycle is robust to different choices of θ (although the cycle can have a cross-over point for certain choices of parameters). An example set of simulations are shown in Figure (5.3) a). As for the ABM, we find that the model reacts with a lag with respect to the economic cycle: taking, e.g., the unemployment time-series $u(t)$ we observe that it oscillates with the same frequency ω as the cycle. Plotting the variables $(\cos(\omega t), u(t))$ in a plane, however, leads to an ellipsoid, which is expected if $u(t) \sim \cos(\omega t + \varphi)$ where φ is a phase-shift that changes depending on the choice of parameters.

In addition to this, we perform a simple experiment to show that the system takes some time to reach a steady-state value. Taking, e.g., $A_u = A_v = 0$, we initialize the system at values (u_0, v_0) and observe the time it takes to reach its stationary state, (u_∞, v_∞) . We find generically that this time becomes shorter for larger values of K , demonstrating that better matching efficiency allows the labor market to reach its equilibrium faster. If we set K to reach a convergence time of around a month we match the prediction of [62]. An example set of simulations are displayed in Figure (5.3) b). We perform a similar analysis that consists in replacing the economic cycle with a step function, using the following prescription for the separation and creation rates:

$$\begin{aligned} s(t) &= \begin{cases} s_0 & \text{for } t \leq T_0 \\ s_0 + A & \text{for } t > T_0 \end{cases} \\ c(t) &= \begin{cases} c_0 + A & \text{for } t \leq T_0 \\ c_0 & \text{for } t > T_0 \end{cases} \end{aligned} \quad (5.21)$$

representing a quickly changing economic environment, that suddenly switches from an economic boom with a high vacancy creation rate and a low separation rate, to a crisis environment where the opposite happens. We observe again a switching from one equilibrium to the other with a non-zero convergence time. As before, this time is shorter for a higher matching efficiency K .

Interpretation and implications Our simple model highlights that the Beveridge curve cycle can be reproduced even without having to introduce detailed agent-based mechanisms. The key drivers are time-varying separation and vacancy creation rates with some asymmetry,

with separations increasing during crises and not during booms, and vice versa for vacancies. This, coupled with the fact that the labor market does not converge instantaneously to its equilibrium, is enough to generate the cycle. In addition to this, the cycle's existence does not hinge on a particular matching function. The qualitative dynamics are therefore driven first and foremost by their out-of-steady-state nature. In particular, this description of the cycle is not incompatible with models that would see firms and workers as behaving rationally, albeit reacting to economic conditions that change much quicker than the time it takes for them to coordinate.

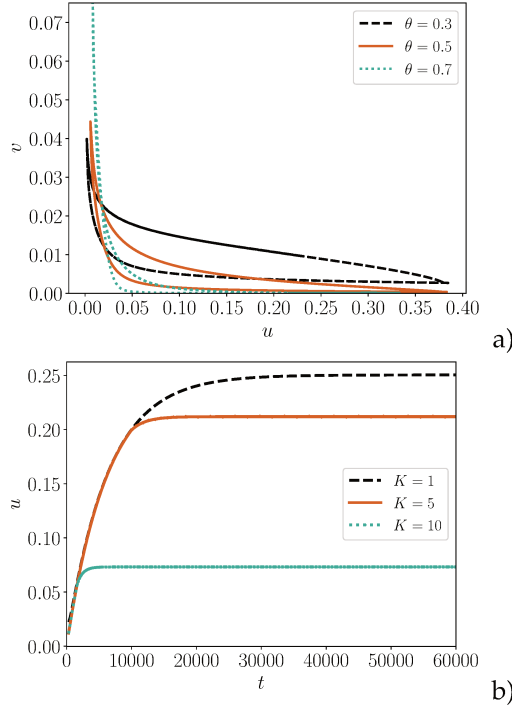


Figure 5.3: Outputs of the phenomenological model simulations with parameters s_0 , c_0 and ω taken as the same as for the ABM simulations, with a) $A_u = A_v = 6$, $K = 1$, for different values of θ . The tear-shaped cycle output is fairly robust when varying the Cobb-Douglas matching function parameter θ . b) $A_u = A_v = 0$, $\theta = 0.5$, for different values of the matching efficiency K . The time-dependent variable studied here is the unemployment rate $u(t)$, and the results are of course the same for any other time dependent variables of the model such as the vacancy rate $v(t)$ or the employment rate $e(t)$. As expected, the convergence time to the steady state becomes smaller when K increases.

5.5. Empirical results

When calibrated to the US labor market, our model predicts that vacancies lag unemployment by approximately 9 months. We determine this lag by computing the cross-correlation between the time series $v(t)$ and $u(t + \lambda)$, finding that the anti-correlation reaches its minimum at $\lambda = 9$ months (see Figure (5.4)a). Moreover, the predicted lag depends on the parameters δ_u and δ_v , as illustrated in Figure (5.4)b). Ensuring $\delta_u > \delta_v$ (to maintain the anti-clockwise behavior of the Beveridge curve) shows that the lag increases with the ratio δ_u/δ_v . This result suggests that countries with only minor differences between these two parameters – starting near a ratio of one and thus initially with zero lag – will exhibit shorter lags. Consequently, a more symmetric labor market implies a smaller area enclosed by the Beveridge curve cycle, indicating a more efficient economy. Finally, our simulations confirm that this lag is independent of the driving frequency of the target demand.

To assess the robustness of our predictions and the substantial lag between unemployment and vacancy rates, we examined Federal Re-

serve Economic Data (FRED) for the U.S. (2000–2023), data from Statistics Canada (2015–2023), and OECD data (2007–2023) for several countries. We display the time series for unemployment and vacancy rates for these countries and their corresponding Beveridge curves in Figures (5.5) and (5.6).

As expected, the unemployment rate and the vacancy rate are anti-correlated. We compute the cross-correlation between the two time series $v(t)$ and $u(t + \lambda)$, and find that the anti-correlation reaches its most negative value for roughly $\lambda = 5$ months for the U.S, $\lambda = 7$ months for Canada, $\lambda = 9$ months for Austria, and $\lambda = 14$ months for Poland, all of which align well with our prediction (see Figure (5.7)). These observations are consistent with recent empirical findings noted in the introduction of this thesis (see for example [70]).

By contrast, we obtain $\lambda = 55$ months for Luxembourg, which is significantly larger than our prediction, and $\lambda \approx 3$ months for the United Kingdom, Germany, and Switzerland, which, while several times greater than the predictions of earlier models, also deviates from our own prediction. This could be an artifact of the statistical noise in the data, but also because our model is calibrated on US data, which could result in deviations when predicting outcomes for economies that differ significantly from the US.

The results shown in Figure (5.7) are not smoothed. We have tested several smoothing methods to reduce the level of noise in the data, which all yield peaks in the negative correlation at lags between 6 and 11 months (except for Luxembourg, which persistently remains much larger). The lag is always many times larger than the 1 month adaptation timescale predicted by the DMP framework [181], and is always closer to the 10 month timescale predicted by [74] for the French labor market under symmetric shocks affecting all occupations. We also find that the behavior of the correlation function for the output of the model calibrated on the US labor market is remarkably close to the empirical correlation functions displayed for the US (see Figure (5.4)).

Turning back to the U.S. Beveridge curves from 2000–2019 (including the 2008 recession) and 2019–2023 (including COVID), we note that the COVID cycle enclosed a much larger area—consistent with our model’s prediction that faster business cycles produce larger loops. Although a thorough demonstration would require in-depth analysis across multiple cycles, our preliminary results support the model’s key conclusions.

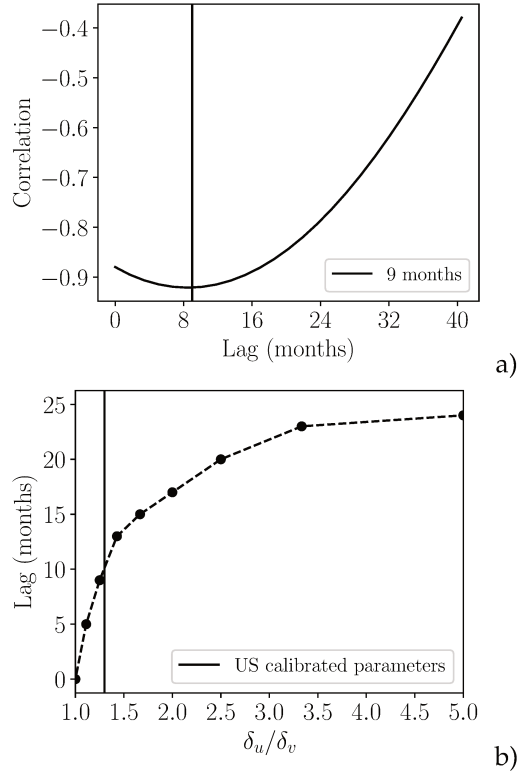


Figure 5.4: a) Simulated lagged-unemployment and vacancy rate correlation function from the model calibrated on the US labor market data. The correlation function peaks around approximately 9 months, as expected. b) Comparison between the measured lags at which the correlation function peaks and the ratio of parameters δ_u and δ_v in the model simulations. The ratio of parameters calibrated to the US labor market, corresponding to a measured lag of approximately 9 months, is highlighted in the figure.

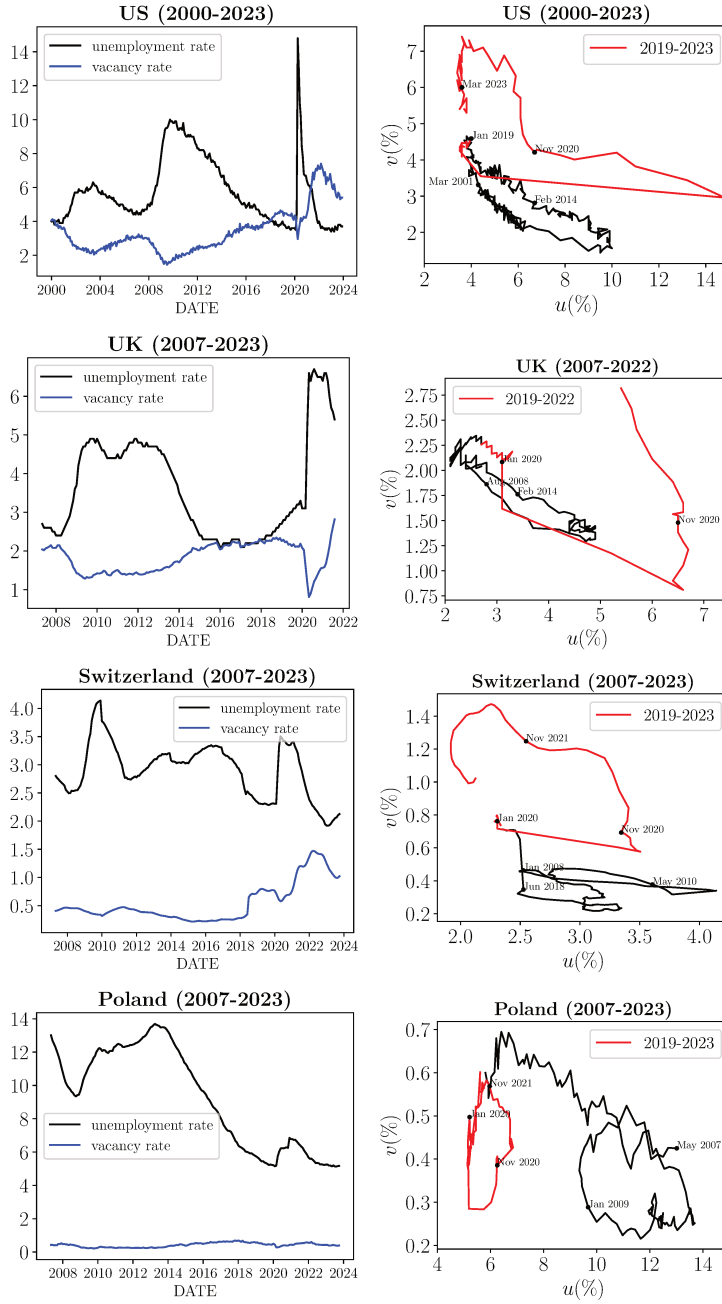


Figure 5.5: US, UK, Switzerland, and Poland unemployment and vacancies rate time series (left) and Beveridge curves (right). The time series span the years 2000-2023 and 2007-2023, respectively. The COVID period is highlighted in red on the Beveridge curves, with dates included to illustrate the anti-clockwise cycle direction, which is consistently observed whenever cycles are present.

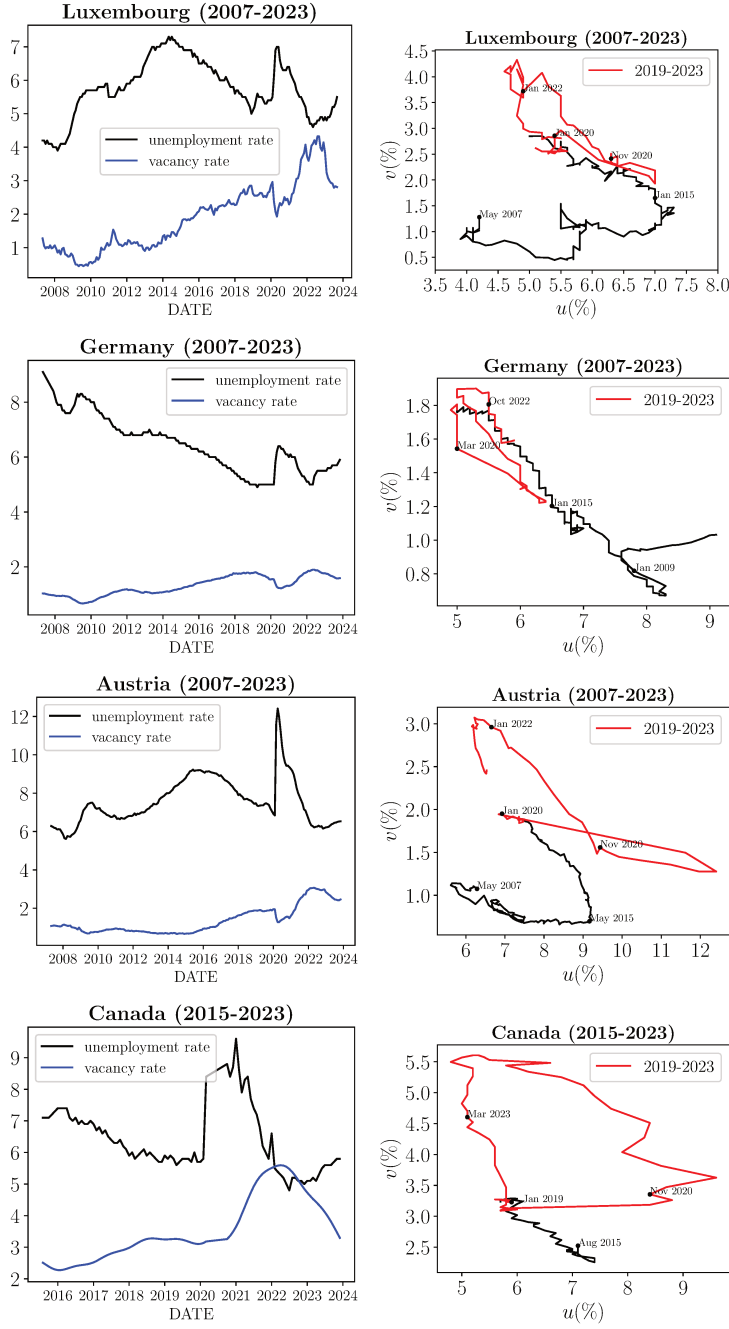


Figure 5.6.: Luxembourg, Germany, Austria, and Canada unemployment and vacancies rate time series (left) and Beveridge curves (right). The time series span the years 2007-2023 and 2015-2023, respectively. The COVID period is highlighted in red on the Beveridge curves, with dates included to illustrate the anticlockwise cycle direction, which is consistently observed whenever cycles are present.

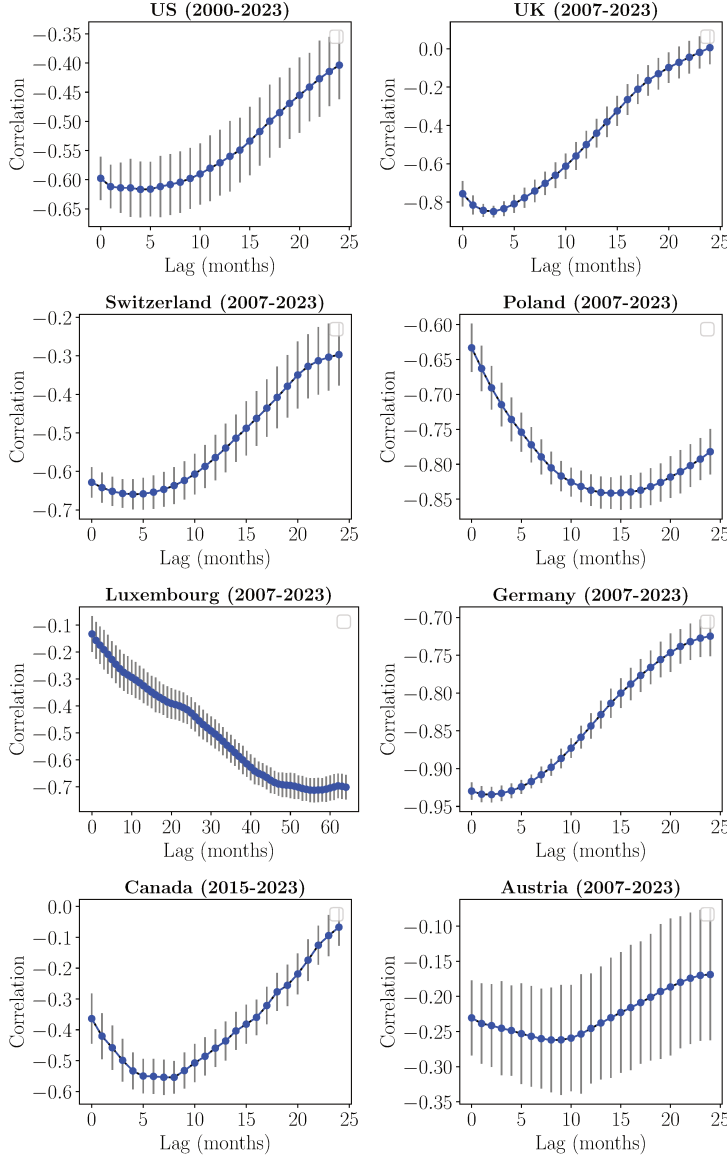


Figure 5.7.: US, Canada, and other OECD countries lagged-unemployment and vacancies rate correlation function, respectively for years 2000-2023, 2015-2023 and 2007-2023. The data used for this computation is not smoothed, and the plots include error bars corresponding to the calculation.

We also evaluate how well our model, even without any calibration, reproduces the US Beveridge curve during COVID. To do this, an alternative proxy for target demand is required, as the oscillatory function calibrated on pre-COVID US business cycles fails to capture structural changes in the economy. That oscillatory function performed well for pre-COVID business cycles—typically around 10 years long—because they were relatively regular in duration and pattern. For completeness, we will use our new proxy to reproduce both the pre- and post-COVID Beveridge curves.

Our chosen proxy is the seasonally adjusted U.S. GDP series from the FRED database (2000–2023). While it does not directly represent the target demand for firms’ hiring/firing decisions, and GDP does not correlate closely with vacancies, we assume it behaves similarly to labor demand. Equivalently, we assume GDP fluctuations serve as a reasonable indicator for fluctuations in labor demand, so the abrupt drop in GDP growth observed during COVID mirrors the sharp reduction in firms’ labor needs.

We denote this time series as y_t and apply a Hodrick-Prescott filter to decompose it into a trend component and a cyclical component as follows:

$$y_t = \tau_t + \hat{y}_t. \quad (5.22)$$

The cyclical part has zero mean, so we normalize it by its standard deviation:

$$x_t = \hat{y}_t / \sigma(\hat{y}_t). \quad (5.23)$$

Figure 5.8a shows the rescaled cyclical component of GDP. This replaces the original sine-based target demand in the model, which previously took the form

$$D_t^+ = D_0 \times (1 + a \sin(\omega t)), \quad (5.24)$$

with D_0 the initial demand and a is the amplitude of the cycle, calibrated to approximately 6.5% for the US labor market. We now write:

$$D_t^+ = D_0 \times (1 + ax_t), \quad (5.25)$$

keeping the parameter a fixed as a first guess.

We notice a lag in simulation outcomes relative to observed data, which is reduced by approximating target demand with a lagged GDP series. An 18-month lag provides the best fit to the data. (Replacing GDP with lagged vacancies did not improve the results.)

This is, admittedly, a rather crude first attempt. Nonetheless, as shown in Figure (5.8), the model captures the general behavior of the US labor market both before and after the COVID shock rather well. Indeed, the pre-COVID behavior of the Beveridge curve is correctly reproduced by the model, which also captures the abrupt increase in unemployment and the economy's cyclical return towards its pre-COVID position in the u - v plane after the COVID shock. The non-equilibrium, agent-based model thus appears relevant for understanding and studying various types of structural changes in the economy. The alignment of the model's results with empirical observations is therefore not merely a result of calibration.

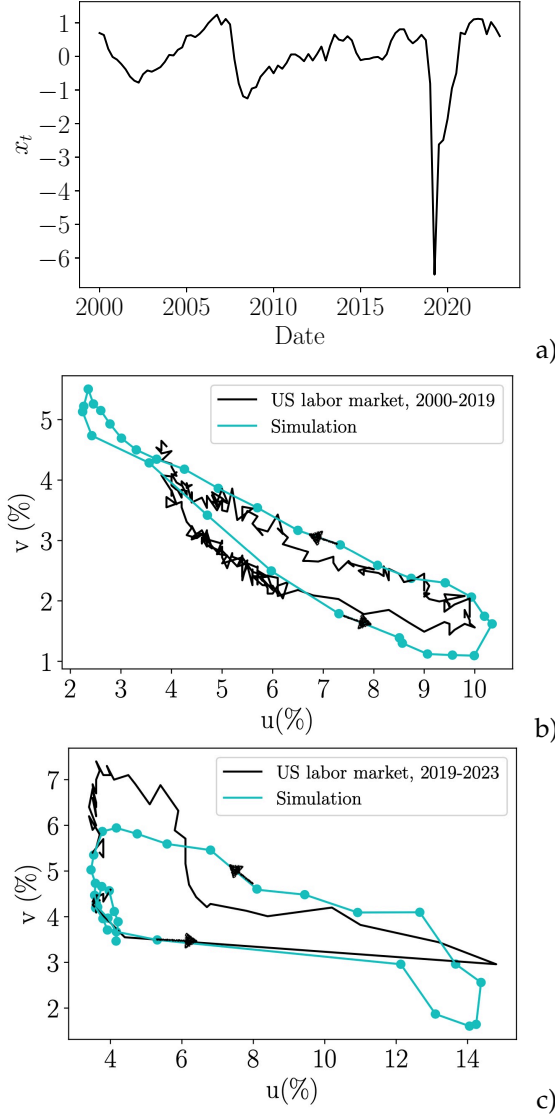


Figure 5.8: a) rescaled and demeaned GDP (x_t) for the US, for years 2000 ($t = 0$) to 2023, using a Hodrick-Prescott filter with a smoothing parameter equal to 1600. b) Outcome of the simulations of the model compared to empirical data, using US lagged (i.e., 18 months backward) GDP as a proxy for target demand, for the years 2000-2019. The US data is plotted in black, and the simulation in cyan. The anti-clockwise direction of time is highlighted on the simulation with black arrows. c) Outcome of the simulations of the model compared to empirical data, using US lagged (i.e., 18 months backward) GDP as a proxy for target demand, for the years 2019-2023. The US data is plotted in black, and the simulation in cyan. The anti-clockwise direction of time is highlighted on the simulation with black arrows.

5.6. Conclusion

Many theories have been proposed to explain the Beveridge curve and its characteristics. These range from classical rational expectations models addressing structural changes and hysteresis loops to agent-based models that integrate job-to-job transitions or sectoral dynamics. Additionally, there is a substantial body of empirical research.

Here we studied the Beveridge curve from an out-of-equilibrium, zero-intelligence standpoint, employing an agent-based model calibrated to accurately replicate the US Beveridge curve. In order to better understand why this model yields realistic-looking Beveridge cycles, we used linear response theory to construct a simple, approximate model that captures most of the features of the Beveridge curve. The reduced model makes it clear that the counter-clockwise Beveridge cycle observed in the ABM is caused by the fact that unemployment lags the demand for labor more than job vacancies do. For the U.S. this lag is the order of 9 months, which is large enough to account for the empirically observed behavior. The lags observed for other countries vary, but the smallest observed value is 3 months (for the UK) and most

countries show much larger lags.

Our approach aligns with the classical economics assertion that frictions stem from the matching process, introducing non-linearities in the model through the matching term $m(u, v)$. However, our approach highlights the role of demand-related terms in the evolution equations for vacancies and unemployment, contributing to the asymmetric behavior that leads to the Beveridge curve cycle. Notably, these terms capture the imperfect forecasting power of agents, in contradiction of rational expectations models.

We show that the network structure is one of the main causes of the slow convergence time of approximately 9 months. This is roughly in agreement with empirical data, and is roughly ten times longer than lags expected in some other models. Furthermore, we demonstrate the relevance of our model even when subjected to an uncalibrated shock. Specifically, by using lagged GDP as a proxy for target demand and without any additional calibration, our model approximately reproduces the shift observed in the Beveridge curve during COVID. The current results are hence suggestive that the Beveridge curve is a disequilibrium, boundedly rational phenomenon.

Key Takeaways

1. **Out-of-Equilibrium ABM:** A zero-intelligence agent-based model calibrated to U.S. data reproduces the counter-clockwise Beveridge curve without classical rational expectations or instantaneous clearing assumptions.
2. **Lag-Driven Cycle Mechanism:** Linear response computations show the Beveridge loop arises because unemployment lags vacancies by about nine months.
3. **Cross-Country Variation:** Empirical unemployment–vacancy lags range from roughly three months to over a year elsewhere, underscoring a universal disequilibrium mechanism with varying magnitudes.
4. **Matching Frictions & Demand Asymmetry:** Nonlinearities in the matching function $m(u, v)$ introduce frictions, and explicit demand terms create asymmetry, reflecting agents' imperfect forecasting.
5. **Network-Induced Slowness:** The structure of skill-compatible job networks drives a slow, nine-month convergence—an order of magnitude longer than models without explicit network frictions.
6. **Disequilibrium Interpretation:** Our findings support viewing the Beveridge curve as a boundedly rational, out-of-equilibrium phenomenon, challenging equilibrium-based policy prescriptions.

A simple yet effective disequilibrium model of the Beveridge curve

6.

*Cette obscure clarté qui tombe des étoiles.
Le Cid¹-Pierre Corneille, 1637.*

This chapter is taken from Becharat, Moran, del Rio-Chanona and Farmer (in prep): *A simple yet effective disequilibrium model of the Beveridge curve*. The drafting of the manuscript was carried out in collaboration with the co-authors.

6.1. Introduction

In the previous chapter, we simplified the ABM of the labor market proposed in [57] to construct an aggregate model and used it to study the Beveridge curve cycle in more detail. This allowed us to show that the Beveridge curve cycle occurs because the unemployment rate lags the vacancy rate due to slow convergence to equilibrium (at least in the model). Here we further develop this line of work, simplifying the model even more and explicitly comparing our predictions to time series data. We hence propose a novel mechanism to explain the Beveridge curve cycle using a simple dynamic disequilibrium model. We use some of the results from the DMP framework, but we also add to it and change it conceptually by abandoning comparative statics.

We alter the conceptual framework explaining the Beveridge curve cycle in several essential ways.

First, we take vacancies as given, and only attempt to understand how unemployment depends on the vacancy rate. Second, we attempt to explain as much as possible while keeping the matching function fixed in time. We show that a surprising amount of the variance can be explained this way. Third, we enhance the DMP formalism by taking state-dependent firings of workers into account. In our model, unemployment increases both due to state-independent separation of workers, which is almost constant in time, and state-dependent firings caused by drops in expected demand, which are highly variable.

Finally, we allow the economy to be truly out-of-equilibrium: In contrast to the DMP formalism, we do not assume a steady-state by requiring that the rate of change of unemployment is zero. As in the DMP formalism, unemployment decreases as workers are matched with jobs, but this happens slowly, on a timescale of several months to more than

¹ The title *Cid* comes from the Arabic honorific *sayyid*, meaning “lord” or “master.” Originally, *sayyid* was reserved for those regarded as descendants of the Prophet Muhammad.

a year. Firms can adjust vacancies and fire workers relatively quickly, but matching workers to jobs takes time. As a result, unemployment tends to lag vacancies. Changes in economic conditions affect unemployment and vacancies, but because unemployment adjusts slowly, *the economy is perpetually out of equilibrium*. We argue that this is the essential cause of the Beveridge curve cycle.

Our framework provides a very simple explanation of the unemployment gap between recessions and recoveries. When we go into recession workers are fired quickly and unemployment increases in pace with decreasing vacancies. As the economy recovers, vacancies increase but unemployment responds more slowly, so that at a given level of vacancies, unemployment is lower during the recovery than it was during the recession. Most of the lag occurs in the recovery phase; this is because workers can be fired quickly, but rehiring takes time. Thus, as the recession begins, increasing unemployment keeps pace with decreasing vacancies, but during recovery, it tends to lag.

Despite being very simple, our model makes several testable predictions. The first is that, even if we keep the matching function constant, we can still explain much of the Beveridge curve cycle. This allows us to test our model by making time series predictions of unemployment conditional on vacancies, and allows us to quantitatively explain much of the Beveridge curve cycle with minimal assumptions. The second prediction is that there is a lag between vacancies and unemployment, which causes counter-clockwise movements in the (u, v) plane. Our model predicts how this lag varies through time. Finally, all else equal, we predict that the size of the unemployment gap is an increasing function of the speed of the recession-recovery cycle. We test all of these predictions here and show that they are in good agreement with the data. To our surprise, we also find that results are relatively insensitive to the functional form of the matching rate.

Of course, our model doesn't explain unemployment perfectly. It is likely that the components it doesn't explain are due to changes in the matching rate. The errors of our model make it clear where further explanation is needed.

Our results are important from a methodological point of view. It is standard in economics to assume equilibrium from the outset. While there is a substantial older literature on disequilibrium models [183], and agent-based models have been successful in producing qualitative explanations of phenomena such as clustered volatility and housing bubbles, there are so far only a few disequilibrium models that have been shown to have quantitative predictive power [184, 185]. Our work provides an example where shifting the conceptual framework to allow the economy to be out-of-equilibrium both simplifies the explanation and results in substantially better predictions. We were inspired by a more complicated agent-based model for occupational labor diffusion [82] that explains the Beveridge curve cycle with similar ideas, but does so in a less parsimonious and predictive manner. This approach to building models based on bounded rationality has a long history in the agent-based modeling community [12].

Another virtue of our model is that it substantially lowers the level of rationality demanded of its agents. In the DMP framework the burden

of rationality for the agents in the model is high – employers and workers need to do complicated calculations to take each other into account. In our model, in contrast, firms just need to estimate the demand for their product and fire workers as needed to increase their profits by reducing costs. Workers are essentially zero intelligence agents who simply do their best to find appropriate jobs as quickly as they can. The predictive power of the model stems from its ability to estimate how far out of equilibrium the economy is, i.e. to predict the time-varying delay between vacancies and unemployment. Requiring less rationality of the agents makes our model more plausible from a behavioral point of view.

This work may also have broader macroeconomic implications. In control theory lags are strongly associated with oscillatory behavior, which can be chaotic [186]. Lags between demand and unemployment can affect production, and might play a role in generating endogenous business cycles.

This chapter is structured as follows. In section (6.2) we introduce our conceptual framework and propose our parsimonious disequilibrium model, and in section (6.3), we show that the predictions of our model are consistent with US labor market data, spanning 1951-2024.

6.2. Our conceptual framework

The DMP conceptual framework assumes sophisticated strategic reasoning by both firms and workers. We adopt a much simpler conceptual framework in which the only strategic reasoning is by firms, who form expectations about demand and post vacancies or fire workers as needed in order to maximize profits. A key element is that firing workers or posting vacancies can be done very quickly, while hiring requires matching vacancies with unemployed workers, which takes time. As a result, when expected demand changes, the economy is slow to reach equilibrium, and the speed with which this happens varies depending on the state of the economy. As we will see, it is essential to model unemployment and vacancies in dynamic terms. To illustrate our reasoning we give a few very simple examples.

Behavior of profit-maximizing firms

Permanent negative demand shock. Consider a large firm that produces a good or service with demand Y which suddenly experiences a permanent negative demand shock $-\Delta Y$. For convenience assume a simple production function $\Pi = KL$, where L is labor and K is everything else (K remains constant). The firm can adjust its labor force either by posting new vacancies or firing existing workers. In this case the firm can maximize profits by firing ΔL workers, so that

$$\frac{\Delta Y}{Y} = \frac{\Delta L}{L}. \quad (6.1)$$

If there are W workers in the whole economy, then as a result of the actions of this firm, unemployment in the whole economy will increase

by

$$\Delta u = \frac{\Delta L}{W} = \frac{L\Delta Y}{WY} = -\phi \frac{\Delta Y}{Y}, \quad (6.2)$$

where $\phi = L/W$ is the fraction of the total workforce employed by the firm prior to the shock.

Assuming the separation rate s of the firm remains constant, this will also cause a reduction in the rate at which it posts vacancies. If the number of vacancies is proportional to the number of workers, similar reasoning shows that this firm will cause vacancies to drop by $\Delta v = \phi \Delta Y/Y$. (In the long-run, if the shock were truly permanent, this would reduce the number of possible jobs in the economy, so that eventually vacancies would return to their former level, but we neglect this here).

Downward trend in expected demand. Suppose at time $t = 0$ the expected demand suddenly drops exponentially at rate α , so that the firm needs to adjust its workforce by $L(t) = L(0)e^{-\alpha t}$ in order to minimize costs. The firm can achieve this by firing workers and posting vacancies at an exponentially decreasing rate, so that its contribution to the rates of change of unemployment and vacancies in the whole economy will be

$$\frac{du}{dt} = \alpha(1 - \phi)e^{-\alpha t} \quad (6.3)$$

$$\frac{dv}{dt} = -\alpha(1 - \phi)e^{-\alpha t} \quad (6.4)$$

Positive demand shocks. In both of the examples above unemployment and vacancies respond to the shock immediately and move in opposite directions. However, this is only true for negative shocks. If a firm's expected demand suddenly experiences a positive shock, it cannot "un-fire" workers – its only option is to post vacancies. But hiring takes time, so it will necessarily experience an interval when it lacks the workers it would like to have in order to produce more and maximize its profits. This illustrates that *hiring and firing are inherently asymmetric*. While the dynamics of vacancies and unemployment will be similar during recessions, they will be different during recoveries, when unemployment will lag vacancies. As we will show, this is the essence of what causes the Beveridge curve.

The first two examples are deterministic and monotonic. More generally, we should consider stochastic and possibly non-monotonic behavior in future demand. Given the time delays for hiring and the time it takes to break in new workers, a firm may choose to retain workers during a downturn if it believes that it is temporary. It may be possible to work out a more general model under this reasoning, but the empirical work that we will present later suggests that firms use simple heuristics to adjust vacancies and to decide when to lay off workers, so it is not clear whether this would be empirically useful.

The reasoning above doesn't take into account heterogeneity and behavioral factors. In reality some sectors are much more likely to experience demand shocks than others. The construction industry is very sensitive to economic downturns, while government workers and nurses are not. One could potentially estimate industry-specific fac-

tors to take this into account. This is beyond the scope of this work, where we are taking an aggregate approach, but even at an aggregate level, heterogeneous industry-specific demand shocks could make the response non-linear. For example, if occupations such as nurses and government workers are always in demand, this imposes a floor on vacancies and a ceiling on unemployment.

It is also worth noting that there may be behavioral effects that are not profit maximizing. Large scale layoffs, for example, can be traumatic, and firms may decide to make downward adjustments to their workforce by posting fewer vacancies rather than firing workers, even if this causes delays in adjusting the workforce that are not myopically profit maximizing. This may be because management believes that this increases worker productivity in the long run (which is ultimately profit maximizing) or they may be hesitant to fire employees for humanitarian reasons.

A final caveat about our reasoning above is that the linearity assumption is likely only a crude assumption. There are many circumstances that might introduce nonlinearities, for example, if there is a minimum team size required to make something. In this case we would see all-or-nothing behavior, with no change in unemployment followed by a jump when the firm goes out of business. In general, whenever a firm goes out of business, for whatever reason, we should see non linear jumps in unemployment. The model above may nonetheless be a reasonable approximation since there are many firms in the economy.

A quantitative model for vacancies and unemployment

The reasoning above suggests that there is a key term missing from Eq. (5.1) corresponding to the fact that *firms fire workers in economic downturns*. Letting the expected demand be $D(t)$ and defining the function $[x]^- = 0$ when $x \geq 0$ and $[x]^- = x$ when $x < 0$, we propose the following model, often referred to as Model 1 in the rest of this work, for the dynamics of unemployment and vacancies,

$$\frac{du}{dt} = (1 - u) \left(s - C_u \left[\frac{dD}{dt} \right]^- \right) - Km(u, v) \quad (6.5)$$

$$\frac{dv}{dt} = (1 - v) \left(v_0 + C_v \frac{dD}{dt} \right) - Km(u, v), \quad (6.6)$$

Both equations have a similar form: On the right hand side there is a creation term and a destruction term. For unemployment the creation term can be broken down into a constant background separation rate s and an expectation-dependent firing rate $C_u [\frac{dD}{dt}]^-$, and similarly for vacancies there is a constant background rate v_0 and an expectation dependent rate $C_v dD/dt$. Both equations share the same matching function, corresponding to the fact that, if a worker is hired, unemployment and vacancies both drop by the same amount. The key difference between the two equations is that the dynamics of unemployment are highly nonlinear due to the $[\frac{dD}{dt}]^-$ term, which produces non-zero output only during downturns.

As a reminder, we write here the static and dynamic DMP equations,

referred to in the rest of this work as Model 2 and 3:

$$0 = s(1 - u(t)) - m(u(t), v(t)), \quad (6.7)$$

$$\frac{du}{dt} = s(1 - u(t)) - m(u(t), v(t)). \quad (6.8)$$

To make this operational we need to choose a matching function. Ignoring scale constants, we will consider three possibilities:

- ▶ Urn model. $m(u, v) = v(1 - e^{-u/v})$
- ▶ Cobb Douglas. $m(u, v) = u^\alpha v^{1-\alpha}$
- ▶ Simple. $m(u, v) = uv$.

The first two matching functions are standard. The “simple” matching function is useful as a point of comparison, both because it has an analytic solution for Eq. (6.8)¹ and because it provides the best empirical results. It can be motivated as follows: Suppose all workers and all jobs are equivalent, so that any worker can be matched to any job. If a worker and a job are chosen at random, the odds that an unemployed worker will match to a vacancy job are proportional to the product of u and v .

1: In this setting, $u = \frac{s}{s + Kv}$.

Qualitative explanation of the Beveridge curve cycle

To understand the properties of the model (corresponding to the coupled system constituted of Eq. (6.5) and Eq. (6.6)) and illustrate how it causes a Beveridge curve cycle, we consider an idealized recession and recovery where the expected demand varies sinusoidally, according to $D(t) = \sin \omega t$, as shown in Fig. (6.1).

This illustrates several qualitative predictions of the model. As the economy falls into a recession, employment and vacancies both adjust quickly, so the delay between them is small. In contrast, as the economy recovers, unemployment adjusts more slowly than vacancies because of the delays caused by the time required for matching unemployed workers to jobs. The result is that it takes longer for unemployment to reach a given vacancy target as the economy goes into the recession than it does when it recovers. *Note that the matching function is fixed*, illustrating that in contrast to the explanation in terms of comparative statics, changes in the matching function are not needed to generate the Beveridge curve.

Our model also predicts that the demand-sensitive term in the model should make the upper tail of unemployment heavier than the lower tail, whereas the tails of the distribution of vacancies should be less heavy-tailed and more symmetric. Similarly, it predicts that unemployment should be positively skewed, more so than vacancies or GDP. This is because the equation for du/dt in Eq. (6.5) has strongly nonlinear term boosting unemployment when the expected future demand drops, but does nothing when it increases. In contrast, for vacancies this term is the same for recessions and recoveries.

To test this we construct a synthetic GDP series g_t assuming that it is an asymmetric AR(2) random process with fat-tailed shocks calibrated to our 2000–2018 data. We then feed the generated GDP into

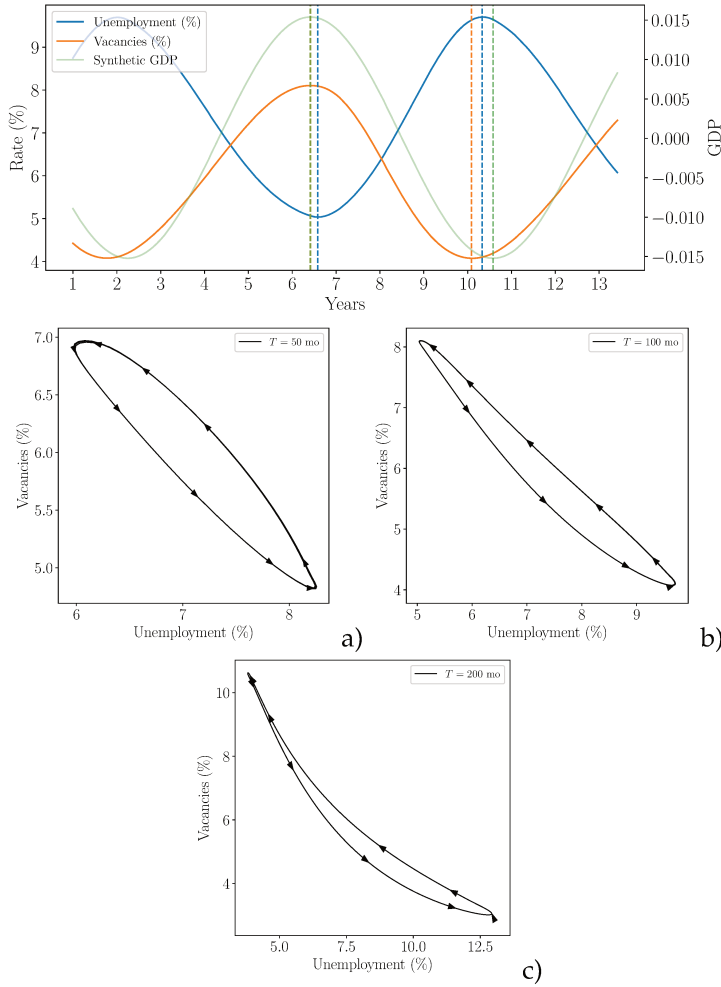


Figure 6.1: We show here realizations of our model assuming sinusoidal expected demand, corresponding to an idealized recession and recovery, with a simple matching function. The upper panel plots simulated unemployment (blue, left axis) and vacancy (orange, left axis) rates for 2000–2018, together with standardized GDP (green dashed, right axis), for a selected parameter set. Vertical dashed lines mark: (1) the unemployment trough and both vacancy and GDP peaks in 2004–06 (Peak 1, unemployment and vacancies are perfectly aligned, and lag GDP by one month), (2) the unemployment peak and both vacancy and GDP troughs in 2008–10 (Peak 2, unemployment lag vacancies by 3 months, and leads GDP by 4 months). Hence vacancies lead unemployment during a cycle. The three lower plots show v against u for three different sine wave periods ($T = 50$ months for a), $T = 100$ months for b) and $T = 200$ months for c)), illustrating the cycle’s dependence on the speed of recession and recovery. The model closely reproduces the characteristic tear drop trajectory of the Beveridge curve during the 2008 financial crisis. Moreover, as the frequency of GDP fluctuations increases (i.e. the cycle period T shortens), the simulated Beveridge cycle becomes progressively more rounded, as expected.

our matching-model for unemployment u_t and vacancies v_t and find that the sample skewness of unemployment consistently exceeds both the skewness of vacancies and that of GDP, in agreement with the empirical behavior that we will describe below.

6.3. Empirical tests of the predictions of the model

We now present several empirical tests of the predictions of our model.

Qualitative tests of predicted lag and tail behavior

As discussed above, our model predicts that the lag between unemployment and vacancies will be smaller as the economy goes into recession than it is when the economy recovers from recession. We test this prediction for the 2008 recession. As shown in Fig. (6.2), based on the peaks, at the beginning of the recession in 2007 unemployment leads vacancies by three months. At the trough of the recession vacancies begin to lead unemployment, but only by six months. By the time the economy fully recovers from the recession unemployment lags recession by about 33 months. Our model explains this asymmetry: As

the recession gets worse, workers are laid off as expected future demand falls. In contrast, recovery requires hiring new workers, which is a much slower process. For comparison we plot the characteristic matching time $1/m(u, v) = 1/(Kuv)$, which corresponds to the inverse exponential relaxation rate if the other terms in Eq. (6.5) and Eq. (6.6) are held constant. We measure this in units of months required for a one percent change in unemployment (based on matching alone). Before the crisis the characteristic time is about 40 months, but during the recovery it is about 30–35 months. Even though matching is faster during the recovery, unemployment moves slower than it did going into the recession because hiring is the primary mechanism of change.

We evaluate this prediction across three additional sample windows (1955–1965, 1970–1980, and 1985–2000; see Appendix (D.3)) and find that for the 1955–1965 and 1985–2000, the unemployment–vacancy lag is again shorter during downturns than during recoveries. For the 1970–1980 period, unemployment does not reach its former level so we cannot compute the corresponding lag.

Furthermore, across all four periods, the recessionary phases exhibit markedly steeper slopes in both unemployment and vacancy rates than the relatively gradual slopes observed during their subsequent recoveries. This is again something that our model predicts, and the results are summarized in Table (6.1).

Recovery slopes		
Period	Vacancies (%-pt/month)	Unemployment (%-pt/month)
1955–1965	0.019	–0.020
1970–1980	0.056	–0.065
1985–2000	0.020	–0.043
2000–2018	0.025	–0.066
Simulation	0.086	–0.101
Recession slopes		
Period	Vacancies (%-pt/month)	Unemployment (%-pt/month)
1955–1965	–0.068	0.193
1970–1980	–0.085	0.215
1985–2000	–0.047	0.074
2000–2018	–0.066	0.184
Simulation	–0.113	0.124

Table 6.1: Monthly slopes for vacancies and unemployment during recoveries (top) and recessions (bottom), by period (including simulation).




Figure 6.2.: Asymmetric lags in 2008 recession. Smoothed monthly unemployment (blue) and vacancies (orange) for the United States from 2000–2018, with the characteristic matching time $1/(Kuv)$ (black, right-hand scale) plotted in units of months required to achieve a one percent change based on matching alone. Vertical dashed lines represent key events. At the beginning of the recession in 2007 unemployment leads vacancies by three months, while at the trough of the recession in 2009–2010 vacancies lead unemployment by six months. Vacancies reach their pre-recession level in 2014 while unemployment doesn’t reach it until 33 months later, in 2017. We overlay dashed colored lines to distinguish the mean slopes of unemployment and vacancy rates during recessionary and recovery phases. Our analysis shows that, for both series, the average slope during recessions is roughly 2.5 times steeper than the subsequent recovery slopes.

We now test the model’s predictions about the tail behavior of the distribution of changes in unemployment and vacancies. In Table (6.2) we compute the upper and lower tail exponents for unemployment, vacancies and GDP (which is shown for comparison).

Series	α_{upper}	α_{lower}	Skew
Δu	1.4	1.8	10.5
Δv	2.3	2.6	−0.2
ΔGDP	2.2	1.9	−1.8

Table 6.2.: Tail exponents α and skewness (Skew) for US change in unemployment Δu , changes in vacancies Δv , and change in detrended GDP Δgdp . See Appendix for details of statistical estimation.

The tail exponent α corresponds to the highest moment of the distribution that exists (i.e. is less than infinity), so lower tail exponents indicate heavier tails. Both the upper and lower tail exponents of unemployment are less than 2, but as predicted, the upper tail exponent is larger than the lower tail exponent. In contrast, both tail exponents of vacancies are greater than two. Because unemployment tends to move in the opposite direction to vacancies and GDP, the upper tail exponent of unemployment (1.4) should be compared to the lower tail exponents of vacancies (2.6) and GDP (1.9). The lower tail exponent of unemployment is 1.8, compared to 2.3 for the upper tail of vacancies and 2.2 for GDP, indicating that movements of unemployment during recoveries are heavier tailed than the corresponding movement of vacancies, but as the model predicts, the difference is much less than during recessions. We also compute skewness, which is 10.5 for unemployment vs. −0.2 for vacancies and −1.8 for GDP, showing that unemployment is much more skewed than either of them.

By contrast, empirical estimates are much larger (GDP skew ≈ -1.8 , unemployment skew $\approx +10$), indicating that while amplified downside sensitivity can generate the qualitative ordering of skewness magnitudes, additional nonlinear mechanisms or structural asymmetries are required to match the observed extremes. More specifically, the intuitive logic holds: raising C_u amplifies unemployment’s response to

GDP downturns, shifting the series upward and increasing its positive skewness.

Time series model

In order to compare the theory to data we need a proxy for the expected demand $D(t)$. We have found that detrended GDP provides a useful approximation, which works best when we lag it by one month (we say more about this in Section (6.3)). Rewriting Eq. (6.5) and Eq. (6.6) as discrete time series models using the GDP proxy, assuming a monthly timescale and assigning the necessary free parameters C and K gives

$$u_{t+1} = u_t + (1 - u_t) (s - C_u [\Delta g_{t-1}]^-) - Km(u, v) + \xi_t, \quad (6.9)$$

$$v_{t+1} = v_t + (1 - v_t) (v_0 + C_v \Delta g_{t-1}) - Km(u, v) + v_t, \quad (6.10)$$

where $m(u, v)$ is a scale-free version of one of the three matching functions described above (which we will compare later). ξ_t and v_t are noise terms that take prediction errors into account, and Δg_{t-1} is the change in detrended GDP, i.e. $\Delta g_{t-1} = (g_t - g_{t-1})/g_{t-1}$.

For our empirical tests we use U.S. data from 1951 – 2023. We begin by fitting unemployment conditional on vacancies and GDP with the simple matching function using Eq. (6.5), as shown in Fig. (6.3). We have estimated $s = 0.00265$ independently based on a study of U.S. separation rates during 2000 – 2006 by Hobijn and Sahin [187] (see Appendix (D) for more details), so that the model only two free parameters other than the adjustment of the time lag for GDP.




Figure 6.3: A comparison of the model for unemployment from Eq. (6.5) (referred to as Model 1) fitted to U.S. data from 1951 - 2023, taking the vacancy and GDP time series as given and using the simple matching function $m(u, v) = Kuv$

The model does a remarkably good job of matching the dynamics of unemployment, matching the peaks well and the shape of the time series and the amplitude of its movements up and down, with a goodness of fit $R^2 = 0.37$, as seen of Fig. (6.3). There are, however, periods of many years to decades where the model is consistently above or below the empirical data. (We will revisit the implications of this later).

Comparison to the DMP model

As mentioned earlier, the DMP model is normally used to understand the factors that cause vacancies and the matching function to change, and to our knowledge, has not been used for time series predictions. This makes it difficult to make quantitative comparisons. However, we can at least compare to part of the DMP framework if we take Eq.

(6.7) and Eq. (6.8) literally and use them to predict unemployment conditional on vacancies. Eq. (6.7), which lets du/dt vary, is the same as Eq. (6.5) with the demand sensitive term omitted. Eq. (6.8) makes the equilibrium assumption $du/dt = 0$. To make a comparison, in absence of any knowledge of how the matching function changes, we keep it fixed. This is perhaps not in the spirit of the conceptual framework of the DMP, in which Eq. (6.7) is just used to justify convergence to equilibrium, and most of the dynamics is assumed to come from changes in the matching function, but it is a useful exercise for its own sake. We show fits of Eq. (6.8) and Eq. (6.7) (referred to as Models 2 and 3) to the empirical data in Fig. (6.4). We get $R^2 = 0.18$ for the dynamic model of Eq. (6.7) and $R^2 = 0.13$ for the static model of Eq. (6.8).




Figure 6.4.: A comparison of the “DMP” models of Eq. (6.8) (Model 2) and Eq. (6.7) (Model 3) to U.S. data from 1951 - 2023, taking the vacancy and GDP time series as given and using the simple matching function $m(u, v) = Kuv$.

We can similarly use Eq. (6.6) as a model for vacancies conditioned on unemployment and GDP, as shown in Fig. (6.5). As before, we use the *a priori* value of s from [Hobijn and Sahin \[187\]](#) so that there are only two free parameters. The fit is much worse than it was for unemployment, with $R^2 = 0.09$. This is still much better than doing the analogous procedure using Eq. (6.7), which has $R^2 \approx 0$ or Eq. (6.8), which has $R^2 < 0$. Interestingly, the errors made by the model are very different. In this case the model for vacancies does a much worse job of matching the peaks and the amplitude of the changes, but there are no long periods where the model is consistently high or low, as there were for the model of unemployment.

Finally, we estimate Eq. (6.5) and Eq. (6.6) simultaneously, yielding the result shown in Fig. (6.6). The fit of the joint model to vacancies is essentially the same as before, with $R^2 = 0.09$, while the fit to unemployment is somewhat worse, with $R^2 = 0.3$. Given that the fit to vacancies is relatively poor, and that the fit for unemployment now depends on vacancies, it is surprising that the loss in goodness of fit for vacancies is only $\Delta R^2 = 0.07$.

These findings prompt the question of why our unemployment specification outperforms the vacancy specification so markedly. One plausible explanation is that aggregate GDP growth is a poor proxy for firms’ forward-looking hiring plans: employers likely form expectations about future vacancies in a manner that differs from how they view macroeconomic output. Exploring those anticipatory mechanisms lies beyond the scope of our deliberately parsimonious disequilibrium model, which nonetheless succeeds in reproducing the observed unemployment trajectory given the vacancy and GDP paths.




Figure 6.5.: A comparison of the model for vacancies of Eq. (6.6) against U.S. data from 1951 - 2023, taking the unemployment and GDP time series as given and using the simple matching function $m(u, v) = Kuv$.

Figure 6.6.: A fit of the model given by Eq. (6.5) and Eq. (6.6) to the time series of vacancies and unemployment for the United States from 1951 - 2023. We use the simple matching function, which is kept fixed.

Comparison of matching functions

How important is the functional form of the matching function? In Fig. (6.7) we run a comparison fitting our model yielding unemployment conditional to vacancies and GDP corresponding to Eq. (6.5) to the data with each of the three matching functions mentioned above. The results are summarized in Table (6.3).

Matching function	C	K	R^2
$u \cdot v$	0.012	1.420	0.373
$\sqrt{u \cdot v}$	0.010	0.070	0.370
$(1 - e^{-u/v}) v$	0.009	0.120	0.134

Table 6.3.: Fitted parameters and R^2 for Model 1 under three matching functions.




Figure 6.7.: A comparison of the results from running our model with three different matching functions.

Interestingly, the "urn ball" matching function performs more poorly than the other two matching function, but the Cobb-Douglas and the "simple" matching function yield similar results, confirming us *a posteriori* in our choice of this simple class of matching functions for our simulations.

Evidence for bounded rationality

Our theoretical arguments in Section (6.2) were based on the assumption that firms maximize profits based on their expectations about future demand. But how do they form those expectations? The fact that we find the best results when GDP is lagged by one to two months (see Fig. (6.8)) suggests that firms fire employees after they discover they are not needed, rather than before. In other words, it suggests that firms use a simple backward-looking boundedly rational heuristic to formulate their estimates of future demand, in agreement with previous work [11, 188, 44]. Given that models cannot forecast changes in GDP very well, and that anyone who could forecast GDP well would quickly become very rich, the fact that firms use simple heuristics should not be surprising. Though there may be circumstances in which some firms have inside knowledge that allows them to forecast demand for their good, it seems that this is not typical.

A key advantage of our model is that the timing of its peaks aligns more closely with the data than do the alternative models. In Fig. (6.9), we mark several salient peaks and compare their timing in the model and in the data; on average, the model's peaks lead the observed peaks by about one month.




Figure 6.8.: Output of the R^2 score for Model 1 (simple matching function with lagged GDP) as a function of the lag. The R^2 attains its maximum at a lag of one month, although the value at a two-month lag is nearly identical.




Figure 6.9.: Comparison of Model 1's simulated peaks with those in the empirical series. The simulated peaks are highlighted in magenta, and the empirical peak in gray. On average, the simulated peaks occur approximately one month and a half earlier than the observed ones, although Model 1 still predicts peak timing very accurately.

Measuring matching rates and adjusting the model to make better predictions

Our model can also be used as a way of measuring labor tightness and using this to make better predictions. As seen in Fig. (6.3), the model does a good job of capturing the time dynamics of unemployment. For example, all of the predicted peaks line up well with the actual peaks. However, there are long periods where the model's predictions are consistently low or consistently high. From 1950 to 1970

the model's predictions are consistently high; from 1970 to 1992 they are consistently low, and from 2000 to 2008 they are consistently low. These shifts may occur because of non-stationarity, e.g. because the parameters of the model shift in these periods. While the shift could be a change in any of the three free parameters, DMP theory suggests that a more likely culprit is a change in the overall matching rate, as reflected by the parameter K . If this is indeed the cause, we can measure the overall matching rate by allowing K to shift as needed to better fit the time series. We let s be free in this exercise. Results are summarized in the following table:

Sample window	K estimate	Change vs. full (\times)
Full sample (1951–2023)	1.42	—
1950–1970	1.91	1.35
1970–1992	1.06	0.75
1992–2000	1.42	1.00
2000–2008	1.69	1.19
2008–2023	1.38	0.97

Table 6.4.: Estimated matching-efficiency parameter K by sample window




Figure 6.10.: Output of the simulations using Model 1 with the simple matching function, when vacancies and GDP are given and we let K be free for the five periods. As expected, the fits are much better for each windows, when we adjust the matching function parameter K .

Acknowledging that the overall matching parameter is non-stationary, we can improve our forecasts of unemployment by letting it adjust to respond to changes. We can do this by letting K vary with time in Eq. (6.5),

$$\hat{u}_{t+1} = u_t + (1 - u_t) (s - C_u[\Delta g_{t-1}]^-) - K_t m(u, v) \quad (6.11)$$

where we have written the left hand argument as \hat{u}_{t+1} and dropped the noise terms to indicate that these are the predicted values, in contrast to the observed values u_{t+1} and v_{t+1} . If we hypothesize that the errors are caused entirely by a mis-specification of K_t , the value which would have yielded the correct answer for the unemployment and vacancy predictions is

$$K_{t+1} = \lambda \left[K_t - \frac{u_{t+1} - \hat{u}_{t+1}}{m(u_t, v_t)} \right] + (1 - \lambda) K_t,$$

where \hat{u}_{t+1} is the one-step forecast produced with K_t . The value of λ can be adjusted to best capture the non-stationarity of the matching function.

We initialize this procedure by obtaining the triplet $\theta = (s, C, K)$, as in the previous section, fitting over the full sample, yielding $s = 0.0021$,

$C = 0.012$ and $K = 1.42$. We then keep s and C fixed to these values for the whole procedure.

The smoothing parameter $\lambda \in [0, 1]$ is selected by minimizing the in-sample one-step-ahead mean-square error; the optimum is $\lambda^* = 0.010$, corresponding to an adjustment time-scale for the matching function of about 100 months so north of 8 years.

We assess the adaptive specification in two ways. (i) *In-sample*: the full period 1951–2023 is used both for calibration and evaluation. (ii) *Out-of-sample*: the model is trained up to December 1987, then recursively re-estimated every five years while generating five-year-ahead unemployment forecasts. Table 6.5 reports the resulting performance metrics.

We also report the fitting and forecasting results in Fig. (6.11), and results are very good, once again confirming the accuracy of our model to fit and predict unemployment path given vacancies and GDP.

Sample	R^2	MAE	RMS	Corr.
In-sample adaptive fit (1951–2023)	0.71	0.006	0.009	0.891
Rolling forecasts (out-of-sample)	0.73	0.007	0.009	0.860

MAE and RMS are expressed in unemployment-rate fractions.

Table 6.5.: Goodness-of-fit and forecast accuracy

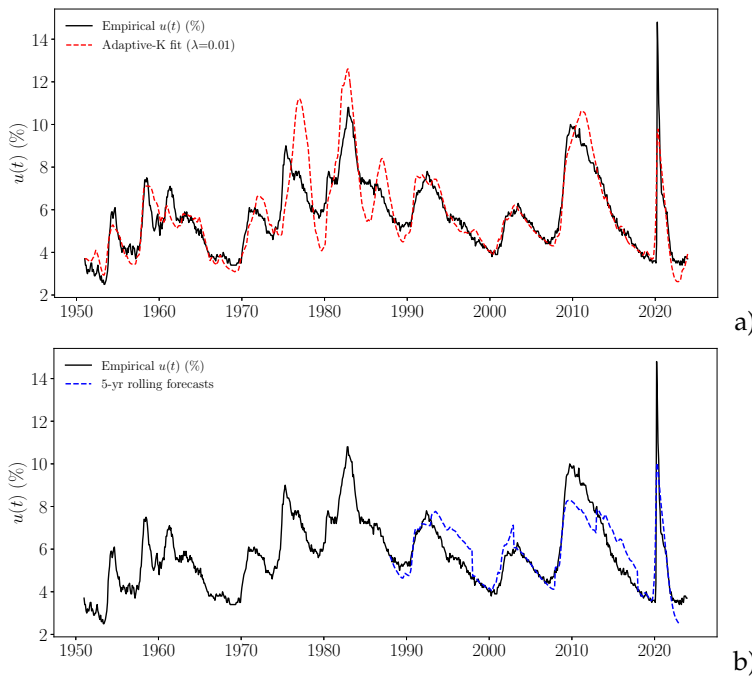


Figure 6.11.: Plot a) shows observed unemployment rate (black) against the adaptive- K_t in-sample fit (red dashed) for 1951–2023. Plot b) shows rolling five-year-ahead unemployment forecasts (blue dashed) compared with realized unemployment (black). Each forecast segment is produced with parameters re-estimated on the preceding window.

6.4. Conclusion

In this chapter, we have extended our earlier disequilibrium framework by dispensing with the classical comparative-statics assumption and embracing fully non-equilibrium, non-linear dynamics. The result is a remarkably parsimonious “bare-bone” model that, despite its simplicity, reproduces and forecasts the time-path of unemployment with high fidelity, conditional on observed vacancy rates. By holding

the matching function fixed throughout the sample, we demonstrate that structural re-estimation is unnecessary to capture the key empirical features of $u(t)$. Moreover, we introduce a novel demand-driven layoff term into the standard DMP unemployment equation, explicitly modeling how downward GDP shocks translate into spontaneous separations. Across a battery of tests, our zero-intelligence specification consistently outperforms both static and dynamic variants of the DMP model. Finally, this framework naturally generates the Beveridge curve dynamics, lending further support to the use of zero-intelligence models to advance our understanding of labor-market adjustment.

Key Takeaways

1. **Fully Non-Equilibrium Dynamics:** By abandoning, again, the comparative-statics assumption, we develop a genuinely non-linear, time-dependent framework that captures transitional paths rather than just steady states.
2. **Parsimonious “Bare-Bone” Model:** Despite its simplicity, the model accurately reproduces and forecasts the unemployment trajectory $u(t)$ conditional on observed vacancy rates, without any structural re-estimation.
3. **Demand-Driven Layoff Term:** Introducing a novel layoff mechanism driven directly by GDP shocks links downward demand impulses to spontaneous separations, enriching the standard DMP unemployment equation.
4. **Good Empirical Fit and Beveridge Dynamics:** Across a suite of tests, our zero-intelligence specification consistently outperforms both static and dynamic DMP variants, while naturally generating the observed Beveridge curve dynamics.

CONCLUSIONS

*A longdrawn carol, mournful, holy,
She chanted loudly, chanted lowly,
Till her eyes were darken'd wholly,
And her smooth face sharpen'd slowly,
Turn'd to tower'd Camelot:
For ere she reach'd upon the tide
The first house by the water-side,
Singing in her song she died,
The lady of Shalott.
The Lady of Shalott- Alfred, Lord Tennyson, 1832.*

The overarching aim of this thesis has been to study the complex, often counter-intuitive dynamics of socio-economic systems by combining the mathematical rigor of statistical physics with the adaptive flexibility of agent-based simulations. Across two complementary strands, we demonstrate that purely local interactions—whether governing continuous *fields* such as housing prices or guiding *individual agents* through simple, *boundedly rational* decision rules—can give rise to rich aggregate phenomena that classical equilibrium paradigms cannot explain.

Our approach is built on two mutually reinforcing pillars. First, we propose and analyze stochastic differential equations as a continuous, coarse-grained description of abstract socio-economic fields—in particular, regional housing markets. When micro-founded on simple, locally applied price-adjustment heuristics and microscopic noise terms, these equations reproduce hallmark spatial and temporal regularities—variograms growing logarithmically with distance, long-range autocorrelations and impulse-response dynamics—directly from their interaction kernels. Second, we study bottom-up agent-based models in which individuals follow transparent rules—relocating to maximize a local utility or applying randomly to skill-compatible vacancies—and aggregate their outcomes into macro-level patterns. By mapping these discrete simulations back onto our hydrodynamic equations, we recover well-known stylized facts—phase separation without predefined neighborhoods or unemployment–vacancy hysteresis—that perfectly coordinated, instantaneous-clearing models cannot produce.

By alternating between direct analysis of continuous fields and systematic derivation from many-agent dynamics, our frameworks capture the emergence of non-equilibrium steady states, path dependence and systemic inefficiencies—features that shatter the assumptions of instantaneous coordination and global market efficiency. These findings compel a reassessment of economic models and policy prescriptions predicated on equilibrium reasoning and advocate for the broader adoption of non-equilibrium, complexity-inspired methods in socio-economic research and policymaking.

In Part I we explore two applications of continuous stochastic equa-

tions to spatial markets:

Chapter 3 revisits the Sakoda–Schelling segregation model in a “neighborhood-less” lattice setting. Starting from an agent-based model in which each individual evaluates only the perceived density at their own site, we carry out a linear-stability analysis and large-scale simulations to map its out-of-equilibrium dynamics onto active-matter equations. We demonstrate that nonlinear, site-specific utilities alone suffice to drive phase-separation—segregation patterns—over a broad range of parameters, and we identify the transition as belonging to the 2D Ising universality class. This shows that the classical “invisible hand” can fail even when agents ignore any predefined neighborhood structure [39].

Chapter 4 turns to the French housing market. We propose a stochastic diffusion equation—micro-founded on simple price-adjustment heuristics—and calibrate its parameters against fifty years of regional price data. The resulting spatial variogram grows logarithmically with distance, in perfect agreement with the theory of two-dimensional diffusing fields driven by white noise, while temporal correlations reveal price-shock persistence over multi-year scales. An impulse-response analysis delivers a diffusion constant of plausible magnitude. These findings confirm long-standing conjectures about price diffusion [169, 87], reveal predictable long-range correlations, and highlight deviations from efficient-market behavior. Finally, by constructing a coupled hydrodynamic model that captures the interaction between the population-density field and the housing-market price field, we demonstrate that social segregation persists in this setting.

In Part II, we employ aggregate dynamic modeling of agent-based systems to investigate the Beveridge curve—the empirically observed relationship between unemployment and vacancies—which exhibits pronounced hysteresis over the business cycle.

Chapter 5 introduces a zero-intelligence agent-based model [57] (in the sense of [26]) in which unemployed workers apply randomly to skill-compatible vacancies and firms adjust their postings via straightforward, rule-of-thumb heuristics. From this discrete simulation we rigorously derive aggregate, continuous equations that isolate the effects of coordination lags, asymmetric hiring versus firing speeds, and the non-instantaneous matching process. These continuous disequilibrium laws bridge the micro–macro gap and demonstrate how purely mechanical delays can generate the observed Beveridge loop, without relying on rational-expectations assumptions.

Chapter 6 then distills these insights into a minimal dynamical disequilibrium model: a small system of coupled differential equations calibrated on US data spanning 1951–2023. Given an observed vacancy trajectory, this parsimonious model reproduces the unemployment path with high fidelity. Its simplicity makes it a powerful tool for policy analysis, capturing the essential out-of-steady-state mechanics of labor-market fluctuations while remaining analytically tractable. This framework naturally generates the Beveridge curve dynamics, lending further support to the view that, to advance our understanding of complex socio-economic systems, we must venture into the *wilderness of bounded rationality* from the zero intelligence gate [44].

By investigating two real-world systems—the spatial diffusion of French housing prices and the cyclical hysteresis of the US labor market—this thesis demonstrates that continuous stochastic equations grounded in micro-level heuristics uncover structural inefficiencies and predictable patterns that lie beyond the reach of classical equilibrium models. Our findings provide compelling evidence for the broader adoption of non-equilibrium, complexity-inspired methods in socio-economic research and policymaking.

Future research would aim at broadening the scope of this work by applying the same methodology to other socio-economic systems. In each case the goal would be to craft parsimonious, heuristics-driven *phenomenological* models that are rigorously calibrated to empirical data. By grounding every parameter in observation rather than abstraction, such models should yield clearer insight into the underlying mechanisms, provide testable quantitative forecasts, and offer decision-makers a transparent tool with which to evaluate policy options. In this way, the frameworks presented in this thesis aspire not only to deepen scientific understanding but also to inform practical interventions in complex socio-economic systems.

Bibliography

References in citation order:

- [1] R. Zakine et al. 'Socioeconomic agents as active matter in nonequilibrium Sakoda–Schelling models'. In: *Physical Review E* 109.4 (2024), p. 044310.
- [2] Antoine-Cyrus Becharat, Michael Benzaquen, and Jean-Philippe Bouchaud. 'Diffusive Nature of Housing Prices'. In: *Phys. Rev. Lett.* 135 (10 Sept. 2025), p. 107401. doi: [10.1103/f4vr-xdny](https://doi.org/10.1103/f4vr-xdny).
- [3] Louis Bachelier. 'Théorie de la spéculation'. In: *Annales scientifiques de l'École normale supérieure* 3.17 (1900), pp. 21–86.
- [4] Albert Einstein. 'Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen'. In: *Annalen der Physik* 17.8 (1905), pp. 549–560.
- [5] L.D. Landau and E.M. Lifshitz. *Fluid Mechanics*. Pergamon Press, 1987.
- [6] P. W. Anderson. 'More Is Different'. In: *Science* 177.4047 (1972), pp. 393–396.
- [7] J.-P. Bouchaud. 'Crises and Collective Socio-Economic Phenomena: Simple Models and Challenges'. In: *Journal of Statistical Physics* 151.3 (May 2013), pp. 567–606. doi: [10.1007/s10955-012-0687-3](https://doi.org/10.1007/s10955-012-0687-3).
- [8] Herbert Blumer. 'Social Problems as Collective Behavior'. In: *Social Problems* 18.3 (1971), pp. 298–306. (Visited on 02/22/2024).
- [9] D. J.T. Sumpter. *The principles of collective animal behaviour*. 2006. doi: [10.1098/rstb.2005.1733](https://doi.org/10.1098/rstb.2005.1733).
- [10] Joshua M Epstein. *Agent-Based Computational Models And Generative Social Science*. 1999.
- [11] Herbert A. Simon. 'A Behavioral Model of Rational Choice'. In: *The Quarterly Journal of Economics* 69.1 (1955), pp. 99–118. doi: [10.2307/1884852](https://doi.org/10.2307/1884852).
- [12] Robert L. Axtell and J. Doyne Farmer. 'Agent-Based Modeling in Economics and Finance: Past, Present, and Future'. In: *Journal of Economic Literature* 63.1 (Mar. 2025), pp. 197–287. doi: [10.1257/jel.20221319](https://doi.org/10.1257/jel.20221319).
- [13] Adam Smith. *The Wealth of Nations*. AuthorHouse, 2009.
- [14] T. C. Schelling. 'Dynamic models of segregation'. In: *Journal of Mathematical Sociology* 1.2 (1971), pp. 143–186.
- [15] J. M. Sakoda. 'The checkerboard model of social interaction'. In: *The Journal of Mathematical Sociology* 1.1 (1971), pp. 119–132. doi: [10.1080/0022250X.1971.9989791](https://doi.org/10.1080/0022250X.1971.9989791).

- [16] Jean-Philippe Bouchaud. *From Statistical Physics to Social Sciences: The Pitfalls of Multi-disciplinarity*. 2023.
- [17] Thomas C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton & Company, 2006.
- [18] Christopher A Pissarides. ‘Short-run equilibrium dynamics of unemployment, vacancies, and real wages’. In: *The American Economic Review* 75.4 (1985), pp. 676–690.
- [19] Dale T Mortensen. ‘The matching process as a noncooperative bargaining game’. In: *The economics of information and uncertainty*. University of Chicago Press, 1982, pp. 233–258.
- [20] Peter A. Diamond. ‘Wage Determination and Efficiency in Search Equilibrium’. In: *The Review of Economic Studies* 49.2 (1982), pp. 217–227. (Visited on 01/16/2024).
- [21] Antoine Augustin Cournot. *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Paris: L. Hachette, 1838.
- [22] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Vol. 1. New York: Oxford University Press, 1995.
- [23] R. D. Luce. *Individual Choice Behavior*. John Wiley, 1959.
- [24] S. P. Anderson, A. De Palma, and J.-F. Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, 1992.
- [25] Jerome Garnier-Brun. ‘Navigating radical complexity : The influence of disorder, nonrelaxational dynamics and learning on aggregate coordination’. Theses. Institut Polytechnique de Paris, Dec. 2023.
- [26] Dhananjay K. Gode and Shyam Sunder. ‘Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality’. In: *Journal of Political Economy* 101.1 (1993), pp. 119–137. (Visited on 01/29/2024).
- [27] Richard Ernest Bellman and Rand Corporation. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957, p. ix.
- [28] Karl Naumann-Woleske et al. ‘Exploration of the Parameter Space in Macroeconomic Models’. In: *Handbook of Complexity Economics*. In press. hal-03797418. 2025.
- [29] Stanisla0 Gualdi et al. ‘Endogenous Crisis Waves: Stochastic Model with Synchronized Collective Behavior’. In: *Physical Review Letters* 114.8 (Feb. 2015). doi: [10.1103/physrevlett.114.088701](https://doi.org/10.1103/physrevlett.114.088701).
- [30] Jesse L. Silverberg et al. ‘Collective Motion of Humans in Mosh and Circle Pits at Heavy Metal Concerts’. In: *Physical Review Letters* 110.22 (May 2013). doi: [10.1103/physrevlett.110.228701](https://doi.org/10.1103/physrevlett.110.228701).
- [31] Roger L. Hughes. ‘THE FLOW OF HUMAN CROWDS’. In: *Annual Review of Fluid Mechanics* 35.1 (2003), pp. 169–182. doi: [10.1146/annurev.fluid.35.101101.161136](https://doi.org/10.1146/annurev.fluid.35.101101.161136).
- [32] Jean-Michel Lasry and Pierre-Louis Lions. ‘Jeux à champ moyen. II – Horizon fini et contrôle optimal’. French. In: *Comptes Rendus. Mathématique* 343.10 (2006), pp. 679–684. doi: [10.1016/j.crma.2006.09.018](https://doi.org/10.1016/j.crma.2006.09.018).

- [33] Matteo Butano, Cécile Appert-Rolland, and Denis Ullmo. *Mean-Field Games Modeling of Anticipation in Dense Crowds*. 2024. URL: <https://arxiv.org/abs/2403.01168>.
- [34] M. Ballerini et al. 'Interaction ruling animal collective behavior depends on topological rather than metric distance: evidence from a field study'. In: *Proceedings of the National Academy of Sciences of the United States of America* 105.4 (Jan. 2008). Epub 2008 Jan 28, pp. 1232–1237. doi: [10.1073/pnas.0711437105](https://doi.org/10.1073/pnas.0711437105).
- [35] B. Mériño et al. *A High Resolution Agent-Based Model of the Hungarian Housing Market*. Technical Report. MNB Working Papers, 2022.
- [36] Jorge Luis Borges. 'On Exactitude in Science'. In: *A Universal History of Infamy*. Trans. by Norman Thomas di Giovanni Steinberg. Originally published in Spanish, 1952. New York: E. P. Dutton, 1962, pp. 16–17.
- [37] Stanislao Gualdi et al. 'Tipping points in macroeconomic agent-based models'. In: *Journal of Economic Dynamics and Control* 50 (Jan. 2015), pp. 29–61. doi: [10.1016/j.jedc.2014.08.003](https://doi.org/10.1016/j.jedc.2014.08.003).
- [38] Nassim Nicholas Taleb. *The Black Swan: The Impact of the Highly Improbable*. 2nd. 1st ed. 2007; read online [archive]. London: Penguin, 2010, p. 366.
- [39] S. Grauwin et al. 'Competition between collective and individual dynamics'. In: *Proceedings of the National Academy of Sciences* 106.49 (2009), pp. 20622–20626. doi: [10.1073/pnas.0906263106](https://doi.org/10.1073/pnas.0906263106).
- [40] S. Grauwin, F. Goffette-Nagot, and P. Jensen. 'Dynamic models of residential segregation: An analytical solution'. In: *Journal of Public Economics* 96.1-2 (2012), pp. 124–141.
- [41] Xavier Gabaix and Ralph SJ Koijen. *In search of the origins of financial fluctuations: The inelastic markets hypothesis*. Tech. rep. National Bureau of Economic Research, 2021.
- [42] John Maynard Keynes. *The General Theory of Employment, Interest and Money*. United Kingdom: Palgrave Macmillan, 1936, p. 472.
- [43] Robert E. Lucas. 'Expectations and the Neutrality of Money'. In: *Journal of Economic Theory* 4.2 (1972), pp. 103–124. doi: [10.1016/0022-0531\(72\)90142-1](https://doi.org/10.1016/0022-0531(72)90142-1).
- [44] J. Doyne Farmer. *Making Sense of Chaos: A Better Economics for a Better World*. Includes bibliographical references (pp. 283–307) and index. London: Allen Lane, 2024, pp. xviii + 361.
- [45] J. C. R. Dow and L. A. Dicks-Mireaux. 'The Excess Demand for Labour: A Study of Conditions in Great Britain, 1946–56'. In: *Oxford Economic Papers* 10.1 (1958), pp. 1–33. doi: [10.1093/oxfordjournals.oep.a040791](https://doi.org/10.1093/oxfordjournals.oep.a040791).
- [46] Jan Eeckhout and Ilse Lindenlaub. 'Unemployment Cycles'. In: *American Economic Journal: Macroeconomics* 11.4 (Oct. 2019), pp. 175–234. doi: [10.1257/mac.20180105](https://doi.org/10.1257/mac.20180105).
- [47] Peter A Diamond and Ayşegül Şahin. 'Shifts in the Beveridge curve'. In: *Research in Economics* 69.1 (2015), pp. 18–25.

- [48] William Phillips. 'The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom'. In: *Economica* (1958).
- [49] Paul A. Samuelson and Robert M. Solow. 'Analytical Aspects of Anti-Inflation Policy'. In: *American Economic Review* 50.2 (1960), pp. 177–194.
- [50] Ben Bernanke. 'Recent Developments in the Labor Market'. In: <http://www.example.com> (2012).
- [51] Andrew Figura and Christopher J. Waller. *What does the Beveridge curve tell us about the likelihood of a soft landing?* FEDS Notes 2022-07-29. Board of Governors of the Federal Reserve System (U.S.), July 2022. doi: [10.17016/2380-7172.3190](https://doi.org/10.17016/2380-7172.3190).
- [52] J. Doynne Farmer and John Geanakoplos. *The virtues and vices of equilibrium and the future of financial economics*. 2008.
- [53] Jean-Philippe Bouchaud, J. Doynne Farmer, and Fabrizio Lillo. *How markets slowly digest changes in supply and demand*. 2008.
- [54] Hie Joo Ahn and Leland D. Crane. 'Dynamic Beveridge Curve Accounting'. In: *Finance and Economics Discussion Series* 2020 (027 Mar. 2020). doi: [10.17016/feds.2020.027](https://doi.org/10.17016/feds.2020.027).
- [55] Matteo Richiardi. 'The Future of Agent-Based Modeling'. In: *Eastern Economic Journal* 43.2 (2017), pp. 271–287.
- [56] Michael Neugart. 'Labor Market Policy Evaluation with ACE'. In: *Journal of Economic Behavior and Organization* 67.2 (June 2008), p. 418. doi: [10.1016/j.jebo.2006.12.006](https://doi.org/10.1016/j.jebo.2006.12.006).
- [57] R Maria del Rio-Chanona et al. 'Occupational mobility and automation: a data-driven network model'. In: *Journal of The Royal Society Interface* 18.174 (2021), p. 20200898.
- [58] Costas Azariadis and Joseph E. Stiglitz. 'Implicit Contracts and Fixed Price Equilibria'. In: *The Quarterly Journal of Economics* 98 (1983), pp. 2–22. (Visited on 02/05/2024).
- [59] Sanford J. Grossman and Oliver D. Hart. 'Implicit Contracts Under Asymmetric Information'. In: *The Quarterly Journal of Economics* 98 (1983), pp. 123–156. (Visited on 02/05/2024).
- [60] Ian M. McDonald and Robert M. Solow. 'Wage Bargaining and Employment'. In: *The American Economic Review* 71.5 (1981), pp. 896–908. (Visited on 01/15/2024).
- [61] Janet L. Yellen. 'Efficiency Wage Models of Unemployment'. In: *The American Economic Review* 74.2 (1984), pp. 200–205. (Visited on 01/15/2024).
- [62] Pascal Michailat and Emmanuel Saez. 'Beveridgean unemployment gap'. In: *Journal of Public Economics Plus* 2 (2021), p. 100009. doi: [10.1016/j.pubecp.2021.100009](https://doi.org/10.1016/j.pubecp.2021.100009).
- [63] Florian Sniekers. 'Persistence and volatility of Beveridge cycles'. In: *International Economic Review* 59.2 (2018), pp. 665–698.
- [64] Gabriele Cardullo and Eric Guerci. 'Interpreting the Beveridge curve. An agent-based approach'. In: *Journal of Economic Behavior & Organization* 157 (2019), pp. 84–100.

- [65] Olivier J. Blanchard and Lawrence H. Summers. 'Hysteresis in unemployment'. In: *European Economic Review* 31.1 (1987), pp. 288–295. doi: [https://doi.org/10.1016/0014-2921\(87\)90042-0](https://doi.org/10.1016/0014-2921(87)90042-0).
- [66] Gadi Barlevy et al. *NBER WORKING PAPER SERIES THE SHIFTING REASONS FOR BEVERIDGE-CURVE SHIFTS*. 2023. URL: <http://www.nber.org/papers/w31783>.
- [67] Ayşegül Şahin et al. 'Mismatch unemployment'. In: *American Economic Review* 104.11 (2014), pp. 3529–3564.
- [68] R Maria del Rio-Chanona et al. 'Mental health concerns precede quits: shifts in the work discourse during the Covid-19 pandemic and great resignation'. In: *EPJ Data Science* 12.1 (2023), p. 49.
- [69] Lorenzo Bermejo, Maria Malmierca-Ordoqui, and Luis Alberiko Gil-Alana. 'Unemployment and COVID-19: an analysis of change in persistence'. In: *Applied Economics* 55.39 (2023), pp. 4511–4521.
- [70] Luis Alberiko A. Gil-Alana, María Jesús González-Blanch, and Carlos Poza. 'Labour market mismatches in G7 countries: a fractional integration approach'. In: *Applied Economics* (2024), pp. 1–17.
- [71] Tolga Omay, Muhammad Shahbaz, and Chris Stewart. 'Is there really hysteresis in the OECD unemployment rates? New evidence using a Fourier panel unit root test'. In: *Empirica* 48.4 (2021), pp. 875–901.
- [72] Robert Shimer. 'The cyclical behavior of equilibrium unemployment and vacancies'. In: *American economic review* 95.1 (2005), pp. 25–49.
- [73] Dale T Mortensen. 'Equilibrium unemployment dynamics'. In: *International Economic Review* 40.4 (1999), pp. 889–914.
- [74] Leonard Bocquet. 'The Network Origin of Slow Labor Reallocation'. working paper or preprint. June 2022.
- [75] Théo Dessertaine et al. *Out-of-Equilibrium Dynamics and Excess Volatility in Firm Networks*. 2021.
- [76] Max Sina Knicker, Karl Naumann-Woleske, and Michael Benzaquen. *Bottlenecks in Occupational Transitions: A Data-driven Taxonomy*. 2024. URL: <https://arxiv.org/abs/2407.14179>.
- [77] Edoardo Gaffeo et al. 'Adaptive microfoundations for emergent macroeconomics'. In: *Eastern Economic Journal* 34.4 (2008). Type: Journal article, pp. 441–463.
- [78] Dhruv Sharma et al. 'V-, U-, L- or W-shaped economic recovery after Covid-19: Insights from an Agent Based Model'. In: *PLOS ONE* 16.3 (Mar. 2021). Ed. by Stefan Cristian Gherghina, e0247823. doi: [10.1371/journal.pone.0247823](https://doi.org/10.1371/journal.pone.0247823).
- [79] Max Sina Knicker et al. 'Post-COVID inflation and the monetary policy dilemma: an agent-based scenario analysis'. In: *Journal of Economic Interaction and Coordination* (June 2024). doi: [10.1007/s11403-024-00413-3](https://doi.org/10.1007/s11403-024-00413-3).

- [80] Sebastian Poledna et al. 'Economic forecasting with an agent-based model'. In: *European Economic Review* 151 (2023), p. 104306.
- [81] Cars Hommes et al. *CANVAS: A Canadian Behavioral Agent-Based Model*. Staff Working Papers 22-51. Bank of Canada, Dec. 2022.
- [82] R Maria del Rio-Chanona et al. 'Supply and demand shocks in the COVID-19 pandemic: an industry and occupation perspective'. In: *Oxford Review of Economic Policy* 36.Supplement₁ (2020), S94–S137. DOI: [10.1093/oxrep/graa033](https://doi.org/10.1093/oxrep/graa033).
- [83] F. Feitosa and W. Zesk. 'Spatial Patterns of Residential Segregation: A Generative Model'. Working Paper. 2008.
- [84] M. Pangallo, J.-P. Nadal, and A. Vignes. 'Residential income segregation: A behavioral model of the housing market'. In: *Journal of Economic Behavior & Organization* 159 (2019), pp. 15–35.
- [85] R. Baptista, J. D. Farmer, and A. Uluc. *Macroprudential Policy in an Agent-Based Model of the UK Housing Market*. Tech. rep. Staff Working Paper 619. Bank of England, 2016.
- [86] L. Gauvin, A. Vignes, and J.-P. Nadal. 'Modeling urban housing market dynamics: Can the socio-spatial segregation preserve some social diversity?' In: *Journal of Economic Dynamics and Control* 37.7 (2013), pp. 1300–1321.
- [87] H. O. Pollakowski and T. S. Ray. 'Housing Price Diffusion Patterns at Different Aggregation Levels'. In: *Journal of Housing Research* 8.1 (1997), pp. 107–124.
- [88] C. Borghesi and J.-P. Bouchaud. 'Spatial correlations in vote statistics: a diffusive field model for decision-making'. In: *European Physical Journal B* 75.3 (2010), pp. 395–404.
- [89] Christian Borghesi. *Une étude de physique sur les élections*. 2014.
- [90] Serge Galam. 'Rational Group Decision Making A random field Ising model at $T = 0$ '. In: (1997).
- [91] Jean-Pierre Nadal et al. 'Monopoly Market with Externality: an Analysis with Statistical Physics and Agent Based Computational Economics'. In: (2003).
- [92] Daniel S. Seara et al. *Sociohydrodynamics: data-driven modelling of social behavior*. 2025. URL: <https://arxiv.org/abs/2312.17627>.
- [93] Haborld Hotblling. 'Stability in competition'. In: *The economic journal* 39.153 (1929), pp. 41–57.
- [94] Arthur Pigou. *The economics of welfare*. Routledge, 2017.
- [95] Dietrich Braess. 'Über ein Paradoxon aus der Verkehrsplanung'. In: *Unternehmensforschung* 12 (1968), pp. 258–268.
- [96] R. Hegselmann. 'Thomas C. Schelling and James M. Sakoda: The history of a model'. In: *Journal of Artificial Societies and Social Simulation* 20.3 (2017), p. 15.
- [97] L. P. Boustan. *Racial Residential Segregation in American Cities*. Tech. rep. National Bureau of Economic Research, 2013.
- [98] J. Trounstine. *Segregation by Design: Local Politics and Inequality in American Cities*. Cambridge University Press, 2018.

- [99] D. Vinković and A. Kirman. ‘A physical analogue of the Schelling model’. In: *Proceedings of the National Academy of Sciences* 103.51 (2006), pp. 19261–19265. doi: [10.1073/pnas.06093711103](https://doi.org/10.1073/pnas.06093711103).
- [100] L. Dall’Asta, C. Castellano, and M. Marsili. ‘Statistical physics of the Schelling model of segregation’. In: *Journal of Statistical Mechanics: Theory and Experiment* 2008.07 (2008), p. L07002. doi: [10.1088/1742-5468/2008/07/L07002](https://doi.org/10.1088/1742-5468/2008/07/L07002).
- [101] L. Gauvin, J. Vannimenus, and J.-P. Nadal. ‘Phase diagram of a Schelling segregation model’. In: *European Physical Journal B* 70.2 (2009), pp. 293–304. doi: [10.1140/epjb/e2009-00234-0](https://doi.org/10.1140/epjb/e2009-00234-0).
- [102] Tim Rogers and Alan J McKane. ‘A unified framework for Schelling’s model of segregation’. In: *Journal of Statistical Mechanics: Theory and Experiment* 2011.07 (July 2011), P07006. doi: [10.1088/1742-5468/2011/07/P07006](https://doi.org/10.1088/1742-5468/2011/07/P07006). (Visited on 07/21/2023).
- [103] L. Gauvin, J.-P. Nadal, and J. Vannimenus. ‘Schelling segregation in an open city: A kinetically constrained Blume-Emery-Griffiths spin-1 system’. In: *Physical Review E* 81.6 (2010), p. 066120. doi: [10.1103/PhysRevE.81.066120](https://doi.org/10.1103/PhysRevE.81.066120).
- [104] Pablo Jensen et al. ‘Giant Catalytic Effect of Altruists in Schelling’s Segregation Model’. In: *Physical Review Letters* 120.20 (May 2018), p. 208301. doi: [10.1103/PhysRevLett.120.208301](https://doi.org/10.1103/PhysRevLett.120.208301). (Visited on 02/09/2023).
- [105] George Barmpalias, Richard Elwes, and Andy Lewis-Pye. ‘Tipping points in 1-dimensional Schelling models with switching agents’. In: *Journal of Statistical Physics* 158 (2015), pp. 806–852.
- [106] Diego Ortega, Javier Rodríguez-Laguna, and Elka Korutcheva. ‘A Schelling model with a variable threshold in a closed city segregation model. Analysis of the universality classes’. In: *Physica A: Statistical Mechanics and its Applications* 574 (2021), p. 126010.
- [107] Diego Ortega, Javier Rodríguez-Laguna, and Elka Korutcheva. ‘Avalanches in an extended Schelling model: An explanation of urban gentrification’. In: *Physica A: Statistical Mechanics and its Applications* 573 (2021), p. 125943.
- [108] David Abella, Maxi San Miguel, and José J Ramasco. ‘Aging effects in Schelling segregation model’. In: *Scientific Reports* 12.1 (2022), p. 19376.
- [109] Jerome Garnier-Brun, Jean-Philippe Bouchaud, and Michael Benzaquen. ‘Bounded Rationality and Animal Spirits: A Fluctuation-Response Approach to Slutsky Matrices’. In: *Journal of Physics: Complexity* 4.015004 (2023).
- [110] R. Wittkowski et al. ‘Scalar φ^4 field theory for active-particle phase separation’. In: *Nature Communications* 5 (2014), p. 4351. doi: [10.1038/ncomms5351](https://doi.org/10.1038/ncomms5351).
- [111] John W. Cahn and John E. Hilliard. ‘Free Energy of a Nonuniform System. I. Interfacial Free Energy’. In: *Journal of Chemical Physics* 28.2 (1958), pp. 258–267. doi: [10.1063/1.1744102](https://doi.org/10.1063/1.1744102).
- [112] R. D. Luce. ‘The choice axiom after twenty years’. In: *Journal of Mathematical Psychology* 15.3 (1977), pp. 215–233.

- [113] J.-P. Nadal et al. 'A formal approach to market organization: choice functions, mean field approximation and maximum entropy principle'. In: *Advances in Self-Organization and Evolutionary Economics* (1998), pp. 149–159.
- [114] M. Marsili. 'On the multinomial logit model'. In: *Physica A* 269.1 (1999), pp. 9–15.
- [115] William A Brock and Steven N Durlauf. 'Discrete choice with social interactions'. In: *The Review of Economic Studies* 68.2 (2001), pp. 235–260.
- [116] A. Kolmogoroff. 'Zur Theorie der Markoffschen Ketten'. In: *Mathematische Annalen* 112.1 (1936), pp. 155–160. doi: [10.1007/BF01565412](https://doi.org/10.1007/BF01565412).
- [117] P. Tamayo and W. Klein. 'Critical dynamics and global conservation laws'. In: *Physical Review Letters* 63.25 (Dec. 1989), pp. 2757–2759. doi: [10.1103/PhysRevLett.63.2757](https://doi.org/10.1103/PhysRevLett.63.2757).
- [118] A. J. Bray. 'Comment on "Critical dynamics and global conservation laws"'. In: *Physical Review Letters* 66.15 (Apr. 1991), pp. 2048–2048. doi: [10.1103/PhysRevLett.66.2048](https://doi.org/10.1103/PhysRevLett.66.2048).
- [119] A. D. Rutenberg. 'Nonequilibrium phase ordering with a global conservation law'. In: *Physical Review E* 54.1 (July 1996), pp. 972–973. doi: [10.1103/PhysRevE.54.972](https://doi.org/10.1103/PhysRevE.54.972).
- [120] A. Tartaglia, L. F. Cugliandolo, and M. Picco. 'Coarsening and percolation in the kinetic 2d Ising model with spin exchange updates and the voter model'. In: *Journal of Statistical Mechanics: Theory and Experiment* 2018.8 (Aug. 2018), p. 083202. doi: [10.1088/1742-5468/aad366](https://doi.org/10.1088/1742-5468/aad366).
- [121] J. T. Siebert et al. 'Critical behavior of active Brownian particles'. In: *Physical Review E* 98.3 (2018), p. 030601. doi: [10.1103/PhysRevE.98.030601](https://doi.org/10.1103/PhysRevE.98.030601).
- [122] C. Maggi et al. 'Universality class of the motility-induced critical point in large scale off-lattice simulations of active particles'. In: *Soft Matter* 17.14 (2021), pp. 3807–3812. doi: [10.1039/D0SM02162H](https://doi.org/10.1039/D0SM02162H).
- [123] F. Dittrich, T. Speck, and P. Virnau. 'Critical behavior in active lattice models of motility-induced phase separation'. In: *European Physical Journal E* 44.4 (Apr. 2021), p. 53. doi: [10.1140/epje/s10189-021-00058-1](https://doi.org/10.1140/epje/s10189-021-00058-1).
- [124] M. Rovere, D. W. Heermann, and K. Binder. 'The gas-liquid transition of the two-dimensional Lennard-Jones fluid'. In: *Journal of Physics: Condensed Matter* 2.33 (Aug. 1990), pp. 7009–7032. doi: [10.1088/0953-8984/2/33/013](https://doi.org/10.1088/0953-8984/2/33/013).
- [125] M. Rovere, P. Nielaba, and K. Binder. 'Simulation studies of gas-liquid transitions in two dimensions via a subsystem-block-density distribution analysis'. In: *Zeitschrift für Physik B Condensed Matter* 90.2 (June 1993), pp. 215–228. doi: [10.1007/BF02198158](https://doi.org/10.1007/BF02198158).
- [126] K. Binder. 'Applications of Monte Carlo methods to statistical physics'. In: *Reports on Progress in Physics* 60.5 (May 1997), p. 487. doi: [10.1088/0034-4885/60/5/001](https://doi.org/10.1088/0034-4885/60/5/001).

- [127] N. Gnan and C. Maggi. ‘Critical behavior of quorum-sensing active particles’. In: *Soft Matter* 18.39 (2022), pp. 7654–7661. doi: [10.1039/D2SM00654E](https://doi.org/10.1039/D2SM00654E).
- [128] Hiroyoshi Nakano and Kyosuke Adachi. ‘Universal properties of repulsive self-propelled particles and attractive driven particles’. In: *arXiv preprint arXiv:2306.17517* (2023).
- [129] Fernando Caballero et al. ‘Strong Coupling in Conserved Surface Roughening: A New Universality Class?’ In: *Physical Review Letters* 121.2 (July 2018), p. 020601. doi: [10.1103/PhysRevLett.121.020601](https://doi.org/10.1103/PhysRevLett.121.020601). (Visited on 09/15/2023).
- [130] Thomas Speck. ‘Critical behavior of active Brownian particles: Connection to field theories’. In: *Physical Review E* 105.6 (June 2022), p. 064601. doi: [10.1103/PhysRevE.105.064601](https://doi.org/10.1103/PhysRevE.105.064601). (Visited on 05/09/2023).
- [131] A. I. Curatolo et al. ‘Cooperative pattern formation in multi-component bacterial systems through reciprocal motility regulation’. In: *Nature Physics* 16.11 (Nov. 2020), pp. 1152–1157. doi: [10.1038/s41567-020-0964-z](https://doi.org/10.1038/s41567-020-0964-z).
- [132] A. I. Curatolo et al. ‘Multilane driven diffusive systems’. In: *Journal of Physics A: Mathematical and Theoretical* 49.9 (2016), p. 095601. doi: [10.1088/1751-8113/49/9/095601](https://doi.org/10.1088/1751-8113/49/9/095601).
- [133] Samuel M. Allen and John W. Cahn. ‘A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening’. In: *Acta Metallurgica* 27.6 (1979), pp. 1085–1095.
- [134] John W. Cahn and John E. Hilliard. ‘Free Energy of a Nonuniform System. I. Interfacial Free Energy’. In: *The Journal of Chemical Physics* 28.2 (Aug. 2004), pp. 258–267. doi: [10.1063/1.1744102](https://doi.org/10.1063/1.1744102).
- [135] T. Grafke, M. E. Cates, and E. Vanden-Eijnden. ‘Spatiotemporal Self-Organization of Fluctuating Bacterial Colonies’. In: *Physical Review Letters* 119.18 (Nov. 2017), p. 188003. doi: [10.1103/PhysRevLett.119.188003](https://doi.org/10.1103/PhysRevLett.119.188003).
- [136] J. O’Byrne and J. Tailleur. ‘Lamellar to Micellar Phases and Beyond: When Tactic Active Systems Admit Free Energy Functionals’. In: *Physical Review Letters* 125.20 (Nov. 2020), p. 208003. doi: [10.1103/PhysRevLett.125.208003](https://doi.org/10.1103/PhysRevLett.125.208003).
- [137] J. O’Byrne. ‘Nonequilibrium currents in stochastic field theories: A geometric insight’. In: *Physical Review E* 107.5 (May 2023), p. 054105. doi: [10.1103/PhysRevE.107.054105](https://doi.org/10.1103/PhysRevE.107.054105).
- [138] Felix Otto. ‘The Geometry of Dissipative Evolution Equations: The Porous Medium Equation’. In: *Communications in Partial Differential Equations* 26.1-2 (Jan. 2001). Publisher: Taylor & Francis _eprint: <https://doi.org/10.1081/PDE-100002243>, pp. 101–174. doi: [10.1081/PDE-100002243](https://doi.org/10.1081/PDE-100002243). (Visited on 02/29/2024).
- [139] Alexandre Lefevre and Giulio Biroli. ‘Dynamics of interacting particle systems: stochastic process and field theory’. In: *Journal of Statistical Mechanics: Theory and Experiment* 2007.07 (2007), P07024.

- [140] J. Tailleur and M. E. Cates. ‘Statistical Mechanics of Interacting Run-and-Tumble Bacteria’. In: *Physical Review Letters* 100.21 (May 2008), p. 218103. doi: [10.1103/PhysRevLett.100.218103](https://doi.org/10.1103/PhysRevLett.100.218103).
- [141] Julien Tailleur, Jorge Kurchan, and Vivien Lecomte. ‘Mapping out-of-equilibrium into equilibrium in one-dimensional transport models’. In: *Journal of Physics A: Mathematical and Theoretical* 41.50 (Dec. 2008), p. 505001. doi: [10.1088/1751-8113/41/50/505001](https://doi.org/10.1088/1751-8113/41/50/505001).
- [142] Lorenzo Bertini et al. ‘Macroscopic fluctuation theory’. In: *Reviews of Modern Physics* 87.2 (June 2015), pp. 593–636. doi: [10.1103/RevModPhys.87.593](https://doi.org/10.1103/RevModPhys.87.593).
- [143] Yongjoo Baek, Yariv Kafri, and Vivien Lecomte. ‘Dynamical phase transitions in the current distribution of driven diffusive channels’. In: *Journal of Physics A: Mathematical and Theoretical* 51.10 (Mar. 2018), p. 105001. doi: [10.1088/1751-8121/aaa8f9](https://doi.org/10.1088/1751-8121/aaa8f9).
- [144] R. Zakine and E. Vanden-Eijnden. ‘Minimum action method for nonequilibrium phase transitions’. In: *arXiv preprint arXiv:2202.06936* (2022).
- [145] Wei Wu and Jin Wang. ‘Potential and flux field landscape theory. I. Global stability and dynamics of spatially dependent non-equilibrium systems’. In: *The Journal of chemical physics* 139.12 (2013).
- [146] Wei Wu and Jin Wang. ‘Potential and flux field landscape theory. II. Non-equilibrium thermodynamics of spatially inhomogeneous stochastic dynamical systems’. In: *The Journal of Chemical Physics* 141.10 (2014).
- [147] A. P. Solon et al. ‘Generalized thermodynamics of phase equilibria in scalar active matter’. In: *Physical Review E* 97.2 (Feb. 2018), p. 020602. doi: [10.1103/PhysRevE.97.020602](https://doi.org/10.1103/PhysRevE.97.020602).
- [148] M. E. Cates. ‘Active Field Theories’. In: *arXiv:1904.01330 [cond-mat]* (Apr. 2019). 36 pages, 7 figures.
- [149] A. P. Solon et al. ‘Generalized thermodynamics of motility-induced phase separation: phase equilibria, Laplace pressure, and change of ensembles’. In: *New Journal of Physics* 20.7 (July 2018), p. 075001. doi: [10.1088/1367-2630/aacdd](https://doi.org/10.1088/1367-2630/aacdd).
- [150] Reynolds Farley et al. ‘Continued Racial Residential Segregation in Detroit: “Chocolate City, Vanilla Suburbs” Revisited’. In: *Journal of Housing Research* 4.1 (1993). Publisher: American Real Estate Society, pp. 1–38. (Visited on 03/01/2024).
- [151] William A. V. Clark. ‘Ethnic Preferences and Ethnic Perceptions in Multi-Ethnic Settings’. In: *Urban Geography* 23.3 (May 2002), pp. 237–256. doi: [10.2747/0272-3638.23.3.237](https://doi.org/10.2747/0272-3638.23.3.237). (Visited on 03/01/2024).
- [152] Lawrence Bobo and Camille L. Zubrinsky. ‘Attitudes on Residential Integration: Perceived Status Differences, Mere In-Group Preference, or Racial Prejudice?’. In: *Social Forces* 74.3 (Mar. 1996), pp. 883–909. doi: [10.1093/sf/74.3.883](https://doi.org/10.1093/sf/74.3.883). (Visited on 03/01/2024).

- [153] S. Saha, J. Agudo-Canalejo, and R. Golestanian. ‘Scalar Active Mixtures: The Nonreciprocal Cahn-Hilliard Model’. In: *Physical Review X* 10.4 (Oct. 2020), p. 041009. doi: [10.1103/PhysRevX.10.041009](https://doi.org/10.1103/PhysRevX.10.041009).
- [154] A. Dinelli et al. ‘Non-reciprocity across scales in active mixtures’. In: *arXiv preprint arXiv:2203.07757* (2022).
- [155] <https://github.com/rzakine/sakoda-schelling>.
- [156] Lawrence E. Blume. ‘The Statistical Mechanics of Strategic Interaction’. In: *Games and Economic Behavior* 5 (3 1993). doi: [10.1006/game.1993.1023](https://doi.org/10.1006/game.1993.1023).
- [157] William A Brock. ‘Pathways to randomness in the economy: emergent nonlinearity and chaos in economics and finance’. In: *Estudios Economicos* (1993), pp. 3–55.
- [158] Christian Borghesi and Jean-Philippe Bouchaud. ‘Of songs and men: a model for multiple choice with herding’. In: *Quality & quantity* 41.4 (2007), pp. 557–568.
- [159] José Moran et al. ‘By force of habit: Self-trapping in a dynamical utility landscape’. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30.5 (2020).
- [160] A.-L. Barabási and H. E. Stanley. *Fractal Concepts in Surface Growth*. Cambridge University Press, 1995.
- [161] U. Frisch. *Turbulence: The Legacy of A. N. Kolmogorov*. Cambridge University Press, 1995.
- [162] C. Borghesi, J.-C. Raynal, and J.-P. Bouchaud. ‘Election turnout statistics in many countries: A diffusive field model for decision-making’. In: *PLoS ONE* 7.5 (2012), e36289.
- [163] J. Fernández-Gracia, K. Suchecki, J. J. Ramasco, et al. ‘Is the voter model a model for voters?’ In: *Physical Review Letters* 112.15 (2014), p. 158701.
- [164] J. Cagé and T. Piketty. *Une histoire du conflit politique*. Le Seuil, 2023.
- [165] J. Geanakoplos, R. Axtell, J. D. Farmer, et al. ‘Getting at systemic risk via an agent-based model of the housing market’. In: *American Economic Review* 102.3 (2012), pp. 53–58.
- [166] S. Rosen. ‘Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition’. In: *Journal of Political Economy* 82.1 (1974), pp. 34–55.
- [167] T. Besley and H. Mueller. ‘Estimating the Peace Dividend: The Impact of Violence on House Prices in Northern Ireland’. In: *American Economic Review* 102.2 (2012), pp. 810–833.
- [168] D. N. Figlio and M. E. Lucas. ‘What’s in a Grade? School Report Cards and the Housing Market’. In: *American Economic Review* 94.3 (2004), pp. 591–604.
- [169] J. M. Clapp and D. Tirtiroglu. ‘Positive feedback trading and diffusion of asset price changes: Evidence from housing transactions’. In: *Journal of Economic Behavior & Organization* 24.3 (1994), pp. 337–355.

- [170] P. A. Samuelson. 'Proof that properly anticipated prices fluctuate randomly'. In: *The World Scientific Handbook of Futures Markets*. World Scientific, 2016.
- [171] S. Basu and T. Thibodeau. 'Analysis of Spatial Autocorrelation in House Prices'. In: *Journal of Real Estate Finance and Economics* 17.1 (1998), pp. 61–85.
- [172] T. G. Conley and M. Kelly. 'The Standard Errors of Persistence'. In: *Journal of International Economics* 153 (2025).
- [173] F. Schweitzer and J. A. Holyst. 'Modelling collective opinion formation by means of active Brownian particles'. In: *European Physical Journal B* 15 (2000), pp. 723–732.
- [174] F. Schweitzer. 'Coordination of decisions in a spatial model of Brownian agents'. In: *The Complex Dynamics of Economic Interaction*. Springer, 2004.
- [175] J.-P. Bouchaud, C. Borghesi, and P. Jensen. 'On the emergence of an 'intention field' for socially cohesive agents'. In: *Journal of Statistical Mechanics: Theory and Experiment* 2014.3 (2014), P03010.
- [176] M. B. Short et al. 'Dissipation and displacement of hotspots in reaction-diffusion models of crime'. In: *Proceedings of the National Academy of Sciences* 107.9 (2010), pp. 3961–3965.
- [177] S. F. Edwards and D. R. Wilkinson. 'The surface statistics of a granular aggregate'. In: *Proceedings of the Royal Society of London A* 381.1780 (1982), pp. 17–31.
- [178] S. Bach, A. Thiemann, and A. Zucco. *The Top Tail of the Wealth Distribution in Germany, France, Spain, and Greece*. Discussion Paper. DIW Berlin, 2015.
- [179] J. O'Byrne. *Non-Equilibrium Currents in Stochastic Field Theories: a Geometric Insight*. Tech. rep. arXiv:2108.13535. arXiv, 2022.
- [180] Margaret Stevens. 'New microfoundations for the aggregate matching function'. In: *International Economic Review* 48.3 (2007), pp. 847–868.
- [181] Pascal Michailat and Emmanuel Saez. $u^* = \sqrt{uv}$. 2022.
- [182] Lorenz Goette, Uwe Sunde, and Thomas Bauer. 'Wage Rigidity: Measurement, Causes and Consequences'. In: *The Economic Journal* 117.524 (2007), F499–F507. (Visited on 06/12/2024).
- [183] Franklin M. Fisher. *Disequilibrium Foundations of Equilibrium Economics*. Online publication: January 2013. Cambridge, UK: Cambridge University Press, 1983.
- [184] Anton Pichler et al. 'Forecasting the propagation of pandemic shocks with a dynamic input-output model'. In: *Journal of Economic Dynamics and Control* 144 (2022), p. 104527. doi: [10.1016/j.jedc.2022.104527](https://doi.org/10.1016/j.jedc.2022.104527).
- [185] Sebastian Poledna et al. 'Economic forecasting with an agent-based model'. In: *European Economic Review* 151 (2023), p. 104306. doi: [10.1016/j.euroecorev.2022.104306](https://doi.org/10.1016/j.euroecorev.2022.104306).
- [186] J. Doynne Farmer. 'Chaotic attractors of an infinite-dimensional dynamical system'. In: *Physica D: Nonlinear Phenomena* 4.3 (1982), pp. 366–393. doi: [10.1016/0167-2789\(82\)90042-2](https://doi.org/10.1016/0167-2789(82)90042-2).

- [187] Bart Hobijn and Ayşegül Şahin. 'Job-finding and separation rates in the OECD'. In: *Economics Letters* 104.3 (2009), pp. 107–111. doi: [10.1016/j.econlet.2009.04.013](https://doi.org/10.1016/j.econlet.2009.04.013).
- [188] Gerd Gigerenzer and Florian M. Artinger. *Heuristic Pricing in an Uncertain Market: Ecological and Constructivist Rationality*. Report 2938702. SSRN, 2016.

APPENDIX

Part I: An out-of-equilibrium Sakoda-Schelling model

A.

A.1. Coarsening exponent

We start from a homogeneous system of size $L_x \times L_y$ with $L_x = L_y$, and we quench it below the critical temperature in the spinodal region. The system undergoes a spinodal decomposition where dense domains coarsen until forming one single large cluster. The typical size of the domains, denoted L_d grows with time as $\sim t^{1/z}$, where the growth exponent z indicates the physics at play. To measure the typical domain size, we compute first the structure factor, given by

$$S(\mathbf{k}, t) = \left| \sum_{\mathbf{r}} e^{-i\mathbf{k} \cdot \mathbf{r}} \phi(\mathbf{r}, t) \right|^2. \quad (\text{A.1})$$

Using isotropy of the system, we average the structure factor over given shells $q = (k_x^2 + k_y^2)^{1/2}$ and we obtain the radial structure factor $s(q, t) = \int_{[0, 2\pi]} S(q, \theta, t) d\theta$. The typical domain size is given by

$$L_d(t) = 2\pi \frac{\int_{k_1}^{\Lambda} s(q, t) dq}{\int_{k_1}^{\Lambda} q s(q, t) dq}, \quad (\text{A.2})$$

with Λ the ultraviolet cutoff and $k_1 = 2\pi/L_x$ the infrared cutoff. On our finite grid, the integral takes the form of a discrete sum, the wavenumber q ranges from $2\pi/N_x$ to $2\pi(N_x - 1)/N_x$, and the increment dq is replaced by $2\pi/N_x$.




Figure A.1: Snapshot of a Monte Carlo simulation. Green sites are occupied, black sites are empty. We draw the boxes of size $\ell = L_x/6$ that are used to measure the liquid and the gas densities. Here, the system size is $L_x = 200, L_y = 66$. Parameters: $\alpha = 3/2, \rho_0 = 0.35, T = 0.05$.

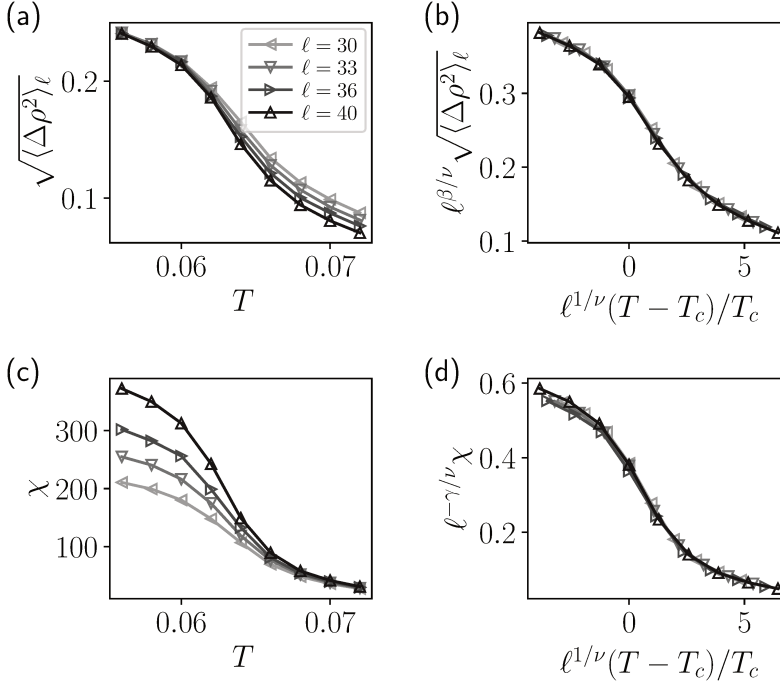


Figure A.2.: Binder cumulant, order parameter and compressibility close to the critical point $(\rho_c, T_c) = (0.271, 0.0620)$ computed for $\alpha = 3/2$ and $\sigma = 1$, as a function of the temperature T .

A.2. Monte Carlo simulations for 2nd-order phase transitions

The Monte Carlo simulation setup that we used is shown in Fig. A.1. As introduced in recent works on MIPS, we study the phase transition using four boxes of size $\ell \times \ell$ located in the bulks of the dense and “gas” phases. The initial condition for the simulations is a fully separated state, where a slab of density 1 coexists with a slab of density 0. We measure the system decorrelation time τ_d and we start recording data after $\sim 1.5\tau_d$. Each simulation is run for a time $\geq 5\tau_d$, and each symbol in Fig. A.2 aggregates the data of 80 independent simulations.

The collapse of the different observable with the 2D Ising critical exponents is displayed in Fig. A.2.

A.3. Local mean field description and LSA

In this section, we consider a modified dynamics where agents are allowed to relocate on neighboring site only. For simplicity, we also consider that the system is one dimensional. It is thus possible to perform a Taylor expansion of the different fields assuming that all fields are smooth in the mean-field limit. The jump probability between two neighboring sites becomes

$$f_{\Gamma}[u(x+a) - u(x)] = f_{\Gamma} \left(a \partial_x u + \frac{a^2}{2} \partial_x^2 u \right), \quad (\text{A.3})$$

where a is the lattice size, and u is the utility on position x . The evolution of the density (for non-overlapping agents) is thus given by

$$\begin{aligned}\partial_t \rho = & \rho(x+a)[1-\rho(x)]f_\Gamma(-a\partial_x u - \frac{a^2}{2}\partial_x^2 u) \\ & + \rho(x-a)[1-\rho(x)]f_\Gamma(a\partial_x u - \frac{a^2}{2}\partial_x^2 u) \\ & - \rho(x)[1-\rho(x+a)]f_\Gamma(a\partial_x u + \frac{a^2}{2}\partial_x^2 u) \\ & - \rho(x)[1-\rho(x-a)]f_\Gamma(-a\partial_x u + \frac{a^2}{2}\partial_x^2 u).\end{aligned}\quad (\text{A.4})$$

After Taylor expansion up to $O(a^2)$ and time rescaling, it turns out that the evolution equation simplifies into

$$\partial_t \rho = f_\Gamma(0)\partial_x^2 \rho - 2f'_\Gamma(0)\partial_x[\rho(1-\rho)\partial_x u]. \quad (\text{A.5})$$

Then, expanding around an homogeneous state, we write $\rho = \rho_0 + \rho_1(x, t)$, $\phi = \rho_0 + \phi_1(x, t)$, and we obtain to leading order in the perturbation:

$$\partial_t \rho_1 = f_\Gamma(0)\partial_x^2 \rho_1 - 2f'_\Gamma(0)\rho_0(1-\rho_0)u'(\rho_0)\partial_x^2 \phi_1. \quad (\text{A.6})$$

In Fourier space the evolution of the mode k is given by $\partial_t \hat{\rho}_1 = \Lambda(k)\hat{\rho}_1$, with

$$\Lambda(k) = -k^2 f_\Gamma(0) \left(1 - 2 \frac{f'_\Gamma(0)}{f_\Gamma(0)} \rho_0(1-\rho_0)u'(\rho_0)\hat{G}_\sigma(k) \right). \quad (\text{A.7})$$

From this, we deduce that the homogeneous system is unstable if there exists a mode k^* such that

$$1 < 2 \frac{f'_\Gamma(0)}{f_\Gamma(0)} \rho_0(1-\rho_0)u'(\rho_0)\hat{G}_\sigma(k^*). \quad (\text{A.8})$$

This criterion is exactly the same as the one found for the non-local move dynamics.

A.4. Local versus non-local PDEs

To illustrate the effectiveness of our local-move approximation in the description of the steady state of the system, we have solved numerically both the local and the non-local mean-field PDEs for the same parameters. The resulting density profiles, displayed in Fig. A.3, appear to be strictly identical up to numerical errors.




Figure A.3. 1D steady-state profiles of the density ρ computed by solving the mean-field dynamics with local moves (solid line), or non-local moves (dashed line). The two density profiles superimpose almost exactly, independently of the various parameter values. It was checked that different initial configurations lead to the same final state. a) Parameters: utility exponent $\alpha = 1$, $\rho_0 = 0.4$, $\sigma = 12$, $L = 300$, $= 9$. b) $\alpha = 3/2$, $\rho_0 = 0.3$, $\sigma = 10$, $L = 300$, $= 15$.

A.5. Linear stability for two coupled populations

We consider the evolution of a perturbation of the homogeneous state in Eq. (3.36) (and in its coupled analogue for the field ρ_B). Close to the homogeneous state $\rho_A(x) \equiv \bar{\rho}_A$, $\rho_B(x) \equiv \bar{\rho}_B$, with $\rho_0 = \bar{\rho}_A + \bar{\rho}_B$, we expand the fields $\rho_Z(x, t) = \bar{\rho}_Z + \rho_{Z,1}(x, t)$, with $Z = A$ or B , and the perturbation fields are denoted with index 1. One also has $\rho(x, t) = \rho_0 + \rho_1(x, t)$, and $\phi_Z(x, t) = \bar{\rho}_Z + \phi_{Z,1}(x, t)$. Keeping leading order terms in Eq. (3.36) yields

$$\partial_t \rho_{A,1} = \Omega \omega_A \left[-\bar{\rho}_A f_{\Gamma_A}(0) \rho_1(x, t) + 2(1 - \rho_0) \bar{\rho}_A f'_{\Gamma_A}(0) [\phi_{A,1}(x, t) \partial_1 u_A + \phi_{B,1}(x, t) \partial_2 u_A] - (1 - \rho_0) f_{\Gamma_A}(0) \rho_{A,1}(x, t) \right], \quad (\text{A.9})$$

where $\partial_1 u_A$ is a shorthand notation for $\frac{\partial u_A}{\partial \bar{\rho}_A} [\bar{\rho}_A, \bar{\rho}_B]$. Taking the logistic function $f_{\Gamma_A}(0) = \frac{1}{2}$, $f'_{\Gamma_A}(0) = \frac{\Gamma_A}{4}$, the linear evolution simplifies into

$$\partial_t \rho_{A,1}(x, t) = \frac{\Omega \omega_A}{2} \left[-\bar{\rho}_A \rho_1(x, t) + (1 - \rho_0) \bar{\rho}_A \Gamma_A [\phi_{A,1}(x, t) \partial_1 u_A + \phi_{B,1}(x, t) \partial_2 u_A] - (1 - \rho_0) \rho_{A,1}(x, t) \right]. \quad (\text{A.10})$$

Similarly, we obtain for the evolution of B :

$$\partial_t \rho_{B,1}(x, t) = \frac{\Omega \omega_B}{2} \left[-\bar{\rho}_B \rho_1(x, t) + (1 - \rho_0) \bar{\rho}_B \Gamma_B [\phi_{A,1}(x, t) \partial_1 u_B + \phi_{B,1}(x, t) \partial_2 u_B] - (1 - \rho_0) \rho_{B,1}(x, t) \right]. \quad (\text{A.11})$$

Denoting $\hat{\rho}_Z(k, t)$ the Fourier transform of $\rho_{Z,1}(x, t)$, the evolution equation can be cast in Fourier space into

$$\partial_t \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix} = L \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix}, \quad (\text{A.12})$$

with

$$L = \frac{\Omega}{2} \begin{pmatrix} \omega_A (\bar{\rho}_B - 1 + (1 - \rho_0) \bar{\rho}_A \Gamma_A \hat{G}_\sigma(k) \partial_1 u_A) & \omega_A (-\bar{\rho}_A + (1 - \rho_0) \bar{\rho}_A \Gamma_A \hat{G}_\sigma(k) \partial_2 u_A) \\ \omega_B (-\bar{\rho}_B + (1 - \rho_0) \bar{\rho}_B \Gamma_B \hat{G}_\sigma(k) \partial_1 u_B) & \omega_B (\bar{\rho}_A - 1 + (1 - \rho_0) \bar{\rho}_B \Gamma_B \hat{G}_\sigma(k) \partial_2 u_B) \end{pmatrix}. \quad (\text{A.13})$$

For simplicity, we will consider that agents are equally rational ($\Gamma_A = \Gamma_B = \Gamma$) and that their moving rates are also identical ($\omega_A = \omega_B = \omega$).

We are looking for conditions to observe dynamical patterns and/or static phase separation. Notably, the homogeneous state is linearly unstable if one eigenvalue of L has a positive real part. It is important to stress that the linear stability analysis is unable to predict the dynamic behavior when nonlinear terms become relevant. Whether the eigenvalues display an imaginary part or not *does not* bring any information on the final dynamics of the system. For the sake of completeness, we explicitate the criteria to have eigenvalues with positive real part and zero imaginary part, referred to as case (i), and eigenvalues with positive real part and nonzero imaginary part, referred to as case (ii). We lie in case (i) if

$$\begin{cases} \text{tr} L > 0 \\ (\text{tr} L)^2 - 4 \det L > 0, \end{cases} \quad \text{or} \quad \begin{cases} \text{tr} L < 0 \\ \det L < 0. \end{cases} \quad (\text{A.14})$$

Case (ii) is obtained if

$$\begin{cases} \text{tr} L > 0 \\ (\text{tr} L)^2 - 4 \det L < 0. \end{cases} \quad (\text{A.15})$$

The criterion $\text{tr} L > 0$ notably simplifies into

$$\bar{\rho}_A \partial_1 u_A + \bar{\rho}_B \partial_2 u_B > \frac{1}{\Gamma \hat{G}_\sigma(k)} \left(\frac{2 - \rho_0}{1 - \rho_0} \right). \quad (\text{A.16})$$

In the main text we have come up with utility functions that lead to eigenvalues with positive real parts and non zero imaginary parts, thus suggesting chasing instability. In some cases, oscillations were observed close to the homogeneous state but they eventually vanished at late times. Whether or not the chasing instability or oscillations are sustained cannot be predicted from the simple linear stability analysis but would require to perform a weakly non-linear analysis which is beyond the scope of this present thesis.

A.6. LSA for two populations with local moves

We start from the local jump approximation of the mean-field equation for the coupled fields. We find that the dynamics can be cast into

$$\partial_t \rho_A = \partial_x [\rho_A (1 - \rho_A - \rho_B) \partial_x \mu([\rho_{A,B}], x)], \quad (\text{A.17})$$

with $\mu = \mu_{ent.} + \mu_{util.}$,

$$\mu_{ent.} = w_A(0) \log \left(\frac{\rho_A}{1 - \rho_A - \rho_B} \right), \quad (\text{A.18})$$

$$\mu_{util.} = -2w'_A(0)u_A([\rho], x), \quad (\text{A.19})$$

and likewise for ρ_B . One can look into the stability of an homogeneous state with densities $\bar{\rho}_A$ and $\bar{\rho}_B$, expanding around this state with a utility $u(\phi_A, \phi_B)$ for agents A and $v(\phi_A, \phi_B)$ for agents B . For convenience, we will take $w_A(0) = w_B(0) = \omega f_\Gamma(0) = \omega/2$ and

$w'_A(0) = w'_B(0) = \omega f(0) = \omega/4$. Expanding around the homogeneous state $(\bar{\rho}_A, \bar{\rho}_B)$ leads to

$$\begin{cases} \partial_t \rho_{A,1} = \frac{\omega}{2} [(1 - \bar{\rho}_B) \partial_x^2 \rho_{A,1} + \bar{\rho}_A \partial_x^2 \rho_{B,1} - \partial_x (\Gamma \bar{\rho}_A (1 - \rho_0) \partial_x \phi_{A,1} \partial_1 u + \partial_x \phi_{B,1} \partial_2 u)] \\ \partial_t \rho_{B,1} = \frac{\omega}{2} [(1 - \bar{\rho}_A) \partial_x^2 \rho_{B,1} + \bar{\rho}_B \partial_x^2 \rho_{A,1} - \partial_x (\Gamma \bar{\rho}_B (1 - \rho_0) \partial_x \phi_{A,1} \partial_1 v + \partial_x \phi_{B,1} \partial_2 v)] \end{cases}, \quad (\text{A.20})$$

Hence, in Fourier space, the linear system can be cast into

$$\partial_t \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix} = K \begin{pmatrix} \hat{\rho}_A(k, t) \\ \hat{\rho}_B(k, t) \end{pmatrix}, \quad (\text{A.21})$$

with

$$K = \frac{\omega k^2}{2} \begin{pmatrix} \bar{\rho}_B - 1 + \Gamma \bar{\rho}_A (1 - \rho_0) \hat{G}_\sigma(k) \partial_1 u & -\bar{\rho}_A + \Gamma \bar{\rho}_A (1 - \rho_0) \hat{G}_\sigma(k) \partial_2 u \\ -\bar{\rho}_B + (1 - \rho_0) \bar{\rho}_B \Gamma \hat{G}_\sigma(k) \partial_1 v & \bar{\rho}_A - 1 + (1 - \rho_0) \bar{\rho}_B \Gamma \hat{G}_\sigma(k) \partial_2 v \end{pmatrix}. \quad (\text{A.22})$$

It is interesting to note that the evolution matrix K is directly proportional to L and, as a consequence, the stability criterion of the homogeneous state with local moves is exactly the same as the one found for non-local moves.

A.7. Introducing a non-linear price field

For completeness, we can also consider a non-monotonic price field (see Fig. A.4(a))

$$\psi = \rho_p^* - |\phi - \rho_p^*|^\alpha. \quad (\text{A.23})$$

The intuition behind this relation is that prices can be lower in overcrowded neighborhoods as well as in empty neighborhoods, and are maximized for a given density ρ_p^* . Assuming again that the price-adjusted utility has the form $\bar{u}(\phi) = u(\phi) - u$

$p[\psi(\phi)]$, with u

$p[\psi] = \lambda \psi$ and $\lambda > 0$, the total utility for the agents is then given by

$$\bar{u}(\phi) = -|\phi - \rho_p^*|^\alpha + \lambda |\phi - \rho_p^*|^\alpha + cst, \quad (\text{A.24})$$

We can inject this expression into the linear-stability condition, see Eq. (3.19), to pinpoint the condensation. In Fig. A.4(b) we take $\alpha = 3$, $\lambda = 5$, $\alpha = 2$, and we interestingly observe that some spinodal curves display several re-entrance points.




Figure A.4.: a) Non-monotonic price field $\psi(\phi)$ as a function of the locally perceived density ϕ , given by Eq. (A.23). Parameters: $\rho_p^* = 0.25$ and $\alpha = 2$. b) Spinodal curves for parameters $\lambda = 5$, $\alpha = 2$, $\alpha = 3$ and densities $[\rho^* = 0.2, \rho_p^* = 0.4]$ (solid line), $[\rho^* = 0.45, \rho_p^* = 0.2]$ (dotted line) and $[\rho^* = 0.5, \rho_p^* = 0.1]$ (dashed line).

Part I: The diffusive nature of housing prices

B.

B.1. Analytical derivation of the diffusive term

We assume that the diffusive term in the price field evolves through a mechanism of supply and demand such that the time evolution of the field ψ depends on the difference of the field between two locations $\psi(R_\alpha) - \psi(R_\beta)$ where R_α and R_β refer to the considered locations. We then propose the following generic equation to describe the propagation of the field with respect to its surrounding influences:

$$\partial_t \psi(R_\alpha, t) = \sum_{\beta} \Gamma_{\alpha, \beta} \psi(R_\beta) - \sum_{\beta} \Gamma_{\beta, \alpha} \psi(R_\alpha), \quad (\text{B.1})$$

where Γ is a symmetric influence matrix such that:

$$\Gamma_{\alpha, \beta} = \Gamma(R_\alpha | R_\beta) = t(R_\alpha - R_\beta | R_\beta). \quad (\text{B.2})$$

Hence, in the continuous limit and in one dimension for simplicity, it comes:

$$\partial_t \psi(x, t) = \int t(x - x' | x') \psi(x', t) dx' - \int t(x' - x | x) \psi(x, t) dx', \quad (\text{B.3})$$

which we can re write as:

$$\partial_t \psi(x, t) = \int t(y | x - y) \psi(x - y, t) dy - \int t(y | x) \psi(x, t) dy, \quad (\text{B.4})$$

changing variables to $y = x - x'$. The Kramers-Moyal expansion of (B.4) up to the order 2 in y then gives:

$$\partial_t \psi(x, t) = -\partial_x [R_1(x) \psi(x)] + \frac{1}{2} \partial_x^2 [R_2(x) \psi(x)], \quad (\text{B.5})$$

where:

$$R_1(x) = \int y t(y, x) dy; \quad (\text{B.6})$$

$$R_2(x) = \int y^2 t(y, x) dy. \quad (\text{B.7})$$

Moreover, the influence matrix is symmetric, hence the drift term $R_1(x)$ is set to zero and we retrieve the one dimensional diffusion equation:

$$\partial_t \psi(x, t) = \partial_x^2 [D(x) \psi(x)] \quad (\text{B.8})$$

with $D(x) = \frac{1}{2} \int y^2 t(y, x) dy$. Note that we retrieve here a non-uniform diffusion coefficient, but we assume in the rest of the study that we can take $D(x) = D$.

B.2. Computation of the temporal variogram

Let us consider the following stochastic partial differential equation:

$$\frac{\partial \psi(\mathbf{r}, t)}{\partial t} = D \Delta \psi(\mathbf{r}, t) - \kappa \psi(\mathbf{r}, t) + \eta(\mathbf{r}, t) + \xi(\mathbf{r}), \quad (\text{B.9})$$

where Δ is the Laplacian operator, D a diffusion coefficient, κ a mean-reversion coefficient, $\eta(\mathbf{r}, t)$ a Langevin noise with zero mean and short range time and space correlations, and $\xi(\mathbf{r})$ a static random field with zero mean and short range correlations. The correlators of these terms are assumed to be of the following type:

$$\begin{aligned} \langle \eta(\mathbf{r}, t) \eta(\mathbf{r}', t') \rangle &= \frac{A^2}{T a^2} e^{-|t-t'|/T} g_a(|\mathbf{r} - \mathbf{r}'|); \\ \langle \xi(\mathbf{r}) \xi(\mathbf{r}') \rangle &= \frac{\Sigma^2}{a^2} g_a(|\mathbf{r} - \mathbf{r}'|), \end{aligned} \quad (\text{B.10})$$

where $g_a(r)$ is a bell-shaped function that decays over length scale a , such that $2\pi \int_{r \geq 0} g_a(r) r dr = a^2$. For the rest of the calculations, we consider the regime where $|\mathbf{r} - \mathbf{r}'| = \ell \gg a$ which leads to $\frac{1}{a^2} g_a(|\mathbf{r} - \mathbf{r}'|) \approx \delta(|\mathbf{r} - \mathbf{r}'|)$.

We have shown in the main text that the contribution to the space time correlation function coming from field η reads:

$$\frac{A^2 (2\pi)^2}{2T((D\mathbf{k}^2 + \kappa)^2 - \frac{1}{T^2})} \left[e^{-\frac{|t-t'|}{T}} - \frac{e^{-(D\mathbf{k}^2 + \kappa)|t-t'|}}{T(D\mathbf{k}^2 + \kappa)} \right]. \quad (\text{B.11})$$

As we are now interested in the temporal variation of the same point in space, we will neglect the random static field $\xi(\mathbf{r})$ in the computation which will only yield constant terms. Moreover, we will again neglect the contribution κ in the calculations as the integration back to real space will impose $Dk^2 \gg \kappa$, as seen in the previous section.

The computation will yield different results depending on the relative values of $\tau = |t - t'|$ and T .

When $\tau = |t - t'| \gg T$, we can set $e^{-\frac{|t-t'|}{T}}$ to zero. Coming back in real space yields:

$$\mathbb{C}(0, |t - t'|) = -\frac{A^2}{T^2 (2\pi)^2} \int d\mathbf{k} \frac{e^{-D\mathbf{k}^2 |t-t'|}}{2D\mathbf{k}^2 (D^2 \mathbf{k}^4 - \frac{1}{T^2})}, \quad (\text{B.12})$$

which gives in polar coordinates:

$$\mathbb{C}(0, \tau) = -\frac{A^2}{T^2 (2\pi)^2} \int dk \int d\theta \frac{k e^{-Dk^2 \tau}}{2Dk^2 (D^2 k^4 - \frac{1}{T^2})}. \quad (\text{B.13})$$

It comes:

$$\mathbb{C}(0, \tau) = -\frac{A^2}{8\pi D T^2} \int_{\frac{D\tau}{a^2}}^{\frac{D\tau}{a^2}} du \frac{e^{-u}}{u(\frac{u^2}{\tau^2} - \frac{1}{T^2})}. \quad (\text{B.14})$$

Moreover, $\frac{u}{\tau} < \frac{1}{T}$ if $S = \frac{a^2}{D} > T$, which allows us to neglect this term,

leading to:

$$\mathbb{C}(0, \tau) \approx \frac{A^2}{8\pi D} \int_{\frac{D\tau}{\ell^2}}^{\frac{D\tau}{a^2}} du \frac{e^{-u}}{u}. \quad (\text{B.15})$$

Hence, in the regime where $T < S \ll \tau \ll \kappa^{-1} = \frac{\ell^2}{D}$:

$$\mathbb{C}(0, \tau) \approx -\frac{A^2}{8\pi D} \log \tau, \quad (\text{B.16})$$

up to constant terms. This finally yields:

$$\mathbb{V}(0, \tau) \approx \frac{A^2}{4\pi D} \log \tau. \quad (\text{B.17})$$

When $S \ll T \ll \tau \ll \kappa^{-1}$, logarithmic contributions can once again be obtained by performing a partial fraction decomposition in (B.14) prior to integration. For completeness, in the regime where $\tau \gg \kappa^{-1}, S, T$, the computation yields a constant value.

When $\tau = |t - t'| \ll T$, we come back to:

$$\frac{A^2(2\pi)^2}{2T(D^2\mathbf{k}^4 - \frac{1}{T^2})} \left[e^{-\frac{|t-t'|}{T}} - \frac{e^{-D\mathbf{k}^2|t'-t|}}{TD\mathbf{k}^2} \right]. \quad (\text{B.18})$$

If $\tau \ll S$, we can expand up to the order two in the exponentials for $D\mathbf{k}^2\tau \rightarrow 0$, in addition to the expansion for $\frac{\tau}{T} \rightarrow 0$, leading to:

$$\frac{A^2(2\pi)^2}{2T(D^2\mathbf{k}^4 - \frac{1}{T^2})} \left[\frac{TD\mathbf{k}^2 - 1}{TD\mathbf{k}^2} + \frac{1}{2}(1 - TD\mathbf{k}^2)\frac{\tau^2}{T^2} \right]. \quad (\text{B.19})$$

Hence, the temporal contribution in the correlation function, coming back to real space, is:

$$\mathbb{C}(0, \tau) = \frac{1}{2\pi} \int_{\frac{1}{\ell^2}}^{\frac{1}{a^2}} dk \frac{A^2 k}{2T(D^2k^4 - \frac{1}{T^2})} \frac{1}{2}(1 - TDk^2)\frac{\tau^2}{T^2}. \quad (\text{B.20})$$

This yields, after integration and up to constant terms:

$$\mathbb{C}(0, \tau) \approx \frac{A^2}{32\pi D} \log \left(\frac{\frac{TD}{\ell^2} + 1}{TD/a^2 + 1} \right) \frac{\tau^2}{T^2}, \quad (\text{B.21})$$

which we can re write as:

$$\mathbb{C}(0, \tau) \approx -\frac{A^2}{32\pi D} \log \left(\frac{\frac{T}{S} + 1}{\kappa T + 1} \right) \frac{\tau^2}{T^2}. \quad (\text{B.22})$$

This finally yields:

$$\mathbb{V}(0, \tau) \approx \frac{A^2}{16\pi D} \log \left(\frac{\frac{T}{S} + 1}{\kappa T + 1} \right) \frac{\tau^2}{T^2}. \quad (\text{B.23})$$

If $\tau \geq S$, we cannot expand in the second exponential term of (B.11). This leads us to study separately both terms. The first one will give,

after expanding up to the second order in $\frac{\tau}{T}$:

$$\frac{1}{2\pi} \int k dk \frac{A^2}{2T(D^2 k^4 - \frac{1}{T^2})} \left(1 - \frac{\tau}{T} + \frac{\tau^2}{2T^2}\right), \quad (\text{B.24})$$

which yields:

$$\frac{A^2}{32\pi D} \log \left(\frac{\left|\frac{T}{S} - 1\right|(\kappa T + 1)}{\left(\frac{T}{S} + 1\right)|\kappa T - 1|} \right) \left(1 - \frac{\tau}{T} + \frac{\tau^2}{2T^2}\right). \quad (\text{B.25})$$

The second term:

$$-\frac{1}{2\pi} \frac{A^2}{2T(D^2 \mathbf{k}^4 - \frac{1}{T^2})} \frac{e^{-D\mathbf{k}^2|t'-t|}}{TD\mathbf{k}^2} \quad (\text{B.26})$$

will give:

$$-\frac{A^2}{8\pi T^2 D} \int_{\frac{1}{T}}^{\frac{1}{\tau}} dk \frac{e^{-Dk^2\tau}}{k(Dk^2 - \frac{1}{T})(Dk^2 + \frac{1}{T})}. \quad (\text{B.27})$$

Changing variables to $u = Dk^2\tau$ yields, after a few integration steps:

$$-\frac{A^2}{32\pi D} \left[e^{\tau/T} \log \left(\frac{\left|\frac{T}{S} - 1\right|}{|\kappa T - 1|} \right) + e^{-\tau/T} \log \left(\frac{\frac{T}{S} + 1}{\kappa T + 1} \right) - 2 \log(\kappa S) \right], \quad (\text{B.28})$$

which gives, after expanding the two exponentials $e^{\tau/T}$ and $e^{-\tau/T}$ up to the order two in $\frac{\tau}{T}$:

$$-\frac{A^2}{32\pi D} \log \left(\frac{\left|\frac{T}{S} - 1\right|(\kappa T + 1)}{\left(\frac{T}{S} + 1\right)|\kappa T - 1|} \right) \frac{\tau}{T} - \frac{A^2}{64\pi D} \log \left(\frac{\left|\frac{T^2}{S^2} - 1\right|}{|\kappa^2 T^2 - 1|} \right) \frac{\tau^2}{T^2}. \quad (\text{B.29})$$

This finally yields, after adding the first and second term contribution from (B.11):

$$\mathbb{V}(0, \tau) \approx \frac{A^2}{8\pi D} \log \left(\frac{\left|\frac{T}{S} - 1\right|(\kappa T + 1)}{\left(\frac{T}{S} + 1\right)|\kappa T - 1|} \right) \frac{\tau}{T} + \frac{A^2}{16\pi D} \log \left(\frac{\frac{T}{S} + 1}{\kappa T + 1} \right) \frac{\tau^2}{T^2}. \quad (\text{B.30})$$

We hence lose the quadratic behavior for the variogram when $S \leq \tau \ll T$ and the dominant behavior becomes linear.

B.3. Spatial variogram in the ballistic case

Starting from the ballistic equation

$$\partial_t \psi(\mathbf{r}, t) + \widehat{\mathbf{v}} \cdot \nabla \psi(\mathbf{r}, t) = -\kappa \psi + \eta, \quad \|\widehat{\mathbf{v}}\| = v,$$

the stationary solution for a fixed direction $\widehat{\mathbf{v}}$ is

$$\psi_{\mathbf{k}}(t) = \int_0^t e^{-\kappa(t-\tau)} e^{-i v \widehat{\mathbf{v}} \cdot \mathbf{k} (t-\tau)} \eta_{\mathbf{k}}(\tau) d\tau, \quad (\text{B.31})$$

where, again, the Ornstein-Uhlenbeck noise obeys

$$\langle \eta_{\mathbf{k}}(t_1) \eta_{\mathbf{k}'}(t_2) \rangle = (2\pi)^2 \frac{A^2}{T} e^{-|t_1-t_2|/T} \delta(\mathbf{k} + \mathbf{k}'), \quad (\text{B.32})$$

and $\xi(\mathbf{r})$ a static random field with zero mean and short range correlations. Then, we insert (B.31) into $\langle \psi_{\mathbf{k}}(t) \psi_{\mathbf{k}'}(t') \rangle$ and set $\mathbf{k}' = -\mathbf{k}$ through the δ in (B.32). One obtains, for the contribution coming from noise η :

$$C_{\eta}^{\widehat{\mathbf{v}}}(\mathbf{k}; t, t') = \frac{A^2(2\pi)^2}{T} \int_0^t d\tau_1 \int_0^{t'} d\tau_2 e^{-\kappa(t-\tau_1)} e^{-\kappa(t'-\tau_2)} e^{-|\tau_1-\tau_2|/T} e^{-i v \widehat{\mathbf{v}} \cdot \mathbf{k} (t-\tau_1-t'+\tau_2)}. \quad (\text{B.33})$$

We then shift integration variables to lag coordinates

$$s_1 = t - \tau_1, \quad s_2 = t' - \tau_2,$$

so $s_1, s_2 \in [0, \infty)$ when $t, t' \rightarrow \infty$. Moreover, $t - \tau_1 - t' + \tau_2 = s_2 - s_1$ and $|\tau_1 - \tau_2| = |s_1 - s_2|$. We thus get

$$C_{\eta}^{\widehat{\mathbf{v}}}(k; \Delta t) = \frac{A^2(2\pi)^2}{T} \int_0^{\infty} ds_1 \int_0^{\infty} ds_2 e^{-\kappa(s_1+s_2)} e^{-|s_1-s_2-\Delta t|/T} e^{-i v \widehat{\mathbf{v}} \cdot \mathbf{k} (s_1-s_2-\Delta t)}, \quad (\text{B.34})$$

with $\Delta t = t - t'$. For the equal-time spatial correlation function we set $\Delta t = 0$. We now set $s = s_1 - s_2$ and perform the angular average over all directions $\widehat{\mathbf{v}}$, yielding:

$$\frac{1}{2\pi} \int_0^{2\pi} e^{-i v k \cos \theta s} d\theta = J_0(vk|s|).$$

A further change of variables to (u, s) with $u = s_1 + s_2$ ($|s| \leq u$) gives

$$\langle \psi_{\mathbf{k}} \psi_{-\mathbf{k}} \rangle_{\eta} = \frac{A^2(2\pi)^2}{2T} \int_{-\infty}^{\infty} ds e^{-|s|/T} J_0(vk|s|) \int_{|s|}^{\infty} du e^{-\kappa u}.$$

Carrying out the u -integral yields $e^{-\kappa|s|}/\kappa$, leaving the Laplace-Bessel integral

$$\int_0^{\infty} e^{-\alpha s} J_0(\beta s) ds = \frac{1}{\sqrt{\alpha^2 + \beta^2}},$$

with $\alpha = \kappa + 1/T$ and $\beta = vk$. One thus finds

$$\langle \psi_{\mathbf{k}} \psi_{-\mathbf{k}} \rangle_{\eta} = \frac{A^2(2\pi)^2}{2T} \frac{1}{\sqrt{(\kappa + 1/T)^2 + (vk)^2}}. \quad (\text{B.35})$$

Transforming (B.35) back to real space via

$$e^{i\mathbf{k}\cdot\mathbf{r}} = e^{i k \ell \cos \phi}, \quad \frac{1}{2\pi} \int_0^{2\pi} e^{i k \ell \cos \phi} d\phi = J_0(k\ell),$$

gives:

$$\mathbb{C}_\eta(\ell, 0) = \frac{A^2}{2} \int_{1/\ell^*}^{1/a} \frac{k J_0(k\ell)}{\sqrt{1 + (vTk)^2}} dk, \quad (\text{B.36})$$

where we have assumed that $\kappa T \ll 1$. In the regime where $a \ll \ell \ll \ell^*$, it should be justified to replace the bounds of the integral above by 0 and ∞ , in which case the integral can be performed exactly (see Gradshteyn & Ryzhik, Eq. 6.554-1):

$$\mathbb{C}_\eta(\ell, 0) \approx \frac{A^2}{2vT} \int_0^\infty \frac{k J_0(k\ell)}{\sqrt{(1/vT)^2 + k^2}} dk = \frac{A^2}{2vT\ell} e^{-\ell/vT}. \quad (\text{B.37})$$

Using $v \approx 1$ km/yr and $T \approx 3$ years, we find $vT \approx 3$ km, so that for $\ell \gtrsim a = 5$ km we are always in the regime where this correlation decays exponentially fast with ℓ , a behaviour very different from the logarithmic decay observed empirically. Note that the contribution from the static field ξ gives

$$\langle \psi_{\mathbf{k}} \psi_{-\mathbf{k}} \rangle_\xi = (2\pi)^2 \Sigma^2 \frac{1}{\kappa^2 + (v k)^2}. \quad (\text{B.38})$$

Interestingly, this would contribute again to a logarithmic behaviour of the spatial correlator. However, using the value for the upper bound of idiosyncratic effects proposed in the main text, we have $\Sigma^2/v^2 \lesssim 10^{-3}$. Hence such static-field contribution with ballistic propagation does not seem plausible to explain our empirical finding.

B.4. Additional Plots




Figure B.1: Distribution of all transaction log-prices $p := \log P$, for the 5 years of DVF data. Note the double hump shape, reflecting a mixture of two distributions, corresponding to prices in cities and prices in the countryside. The right tail, for property prices above 500,000 Euros, corresponds to a power-law tail for prices as $P^{-1-\mu}$ with $\mu \approx 1.5$.




Figure B.2: Cross-sectional variability of the spatial variograms for the log price field $p(\mathbf{r})$ averaged over the period 2018-2022 for France's *régions*, *départements* and cities. Shaded bands show one standard deviation from the mean, with $\bar{V}(\ell) = \frac{1}{N} \sum_{i=1}^N V_i(\ell, 0)$, and σ

$$V(\ell) = \sqrt{\frac{1}{N-1} \sum_{i=1}^N [V_i(\ell, 0) - \bar{V}(\ell)]^2},$$
 where $V_i(\ell, 0)$ is the variogram of unit i (region, département or city) and N the number of units at that scale. Furthermore, the variogram slopes across units are remarkably consistent; the observed dispersion is driven almost entirely by differences in their baseline levels. The black dashed lines have a slope equal to 0.19 for all scales, corresponding to $A^2/D = 1.2$. We display selected examples of individual variograms at each scales, for completeness.




Figure B.3.: Spatial variogram for the log-price field per squared meter $\hat{p}(\mathbf{r}) := \log(\hat{P})$, where the notation \hat{P} indicates the prices per squared meter, averaged over the period 2018-2022 for France as a whole, its *régions*, *départements* and cities, with their respective cross-sectional variability highlighted in shaded colors and the averages for each scale as filled circles. The different off-sets in the y direction corresponds to the measurement noise contribution to the empirical field $\hat{p}(\mathbf{r})$. The observed empirical behavior is once again logarithmic, with slopes ranging between 0.03 and 0.06, with an average value approximately half of that found in [88].




Figure B.4.: Distribution of all transaction log-prices $p := \log P$, averaged for the 5 years of DVF data, both for the *département* of la Creuse and for Paris. These locations were chosen as typical examples of both the countryside and cities, showing clearly two different shapes. This explains the double-hump nature of the global log-price distribution for the whole of France, discussed in the main text.




Figure B.5.: Distribution of all transaction log-prices per squared meter \hat{p} , for the 5 years of DVF data. The right tail corresponds to a power-law tail for prices per squared meter as $\hat{P}^{-1-\mu}$ with $\mu \approx 1.4$, close to 1.5, as found for the log-prices above.




Figure B.6.: Spatial variograms for the log price field $p(\mathbf{r})$ for every year between 1970 to 2022, using the data from [164]. We see that the slope of these variograms is only weakly time-dependent, and that the logarithmic behavior is robust in time. The variogram saturates for $\ell \approx 70$ km in 1970 and for $\ell \approx 300$ km in 2022.

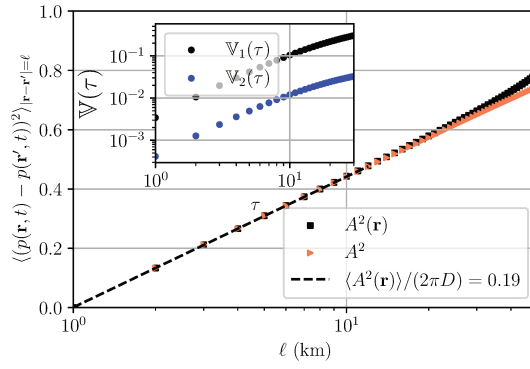


Figure B.7.: Theoretical predictions for the spatial variogram, computed when the noise amplitude is uniform, equal to A^2 , and when the noise amplitude $A^2(\mathbf{r})$ is strongly heterogeneous, with $\langle A^2(\mathbf{r}) \rangle = A^2 = 2\pi D \times 0.19$. We obtain a similar logarithmic behavior in both cases. The inset shows a comparison between $V_1(\tau)$ and $V_2(\tau)$ computed for data simulated on a lattice with the same strongly heterogeneous noise amplitude $A^2(\mathbf{r})$. We hence qualitatively retrieve the observed empirical temporal behavior.

B.5. Estimating the diffusion constant D

To estimate the order of magnitude of the diffusion constant D in France, we examine the propagation of price “shocks” induced by the opening of a TGV (*Train à Grande Vitesse*) station in several cities (Lyon, Bordeaux, and Tours) and their surrounding areas. For each area, we compute

$$\sigma^2(t) = \frac{\sum R(\mathbf{r}, t)(\mathbf{r} - \bar{\mathbf{r}})^2}{\sum R(\mathbf{r}, t)},$$

where $R(\mathbf{r}, t) = p(\mathbf{r}, t) - p(\mathbf{r}, t_0)$, with the summation taken over all *communes* within the considered areas. Our findings indicate that, for these three regions (indexed by i), the relation

$$\sigma_i^2(t) \approx D_i t + C_i$$

holds, allowing us to estimate an order of magnitude for the diffusion constant from the slopes of the corresponding curves, as illustrated in Fig. 4 of the main text.

B.6. Deriving the coupled linear stability criterion

We start again from:

$$\begin{cases} \partial_t \psi_{\mathbf{k}} = -(D|\mathbf{k}|^2 + \kappa) \psi_{\mathbf{k}} + \alpha \rho_{\mathbf{k}}, \\ \partial_t \rho_{\mathbf{k}} = -f_{\beta}'(0) |\mathbf{k}|^2 \rho_{\mathbf{k}} - 2f_{\beta}'(0) \rho_0(1 - \rho_0) u'(\rho_0) |\mathbf{k}|^2 G_{\sigma}(\mathbf{k}) \rho_{\mathbf{k}} + 2\lambda f_{\beta}'(0) \rho_0(1 - \rho_0) |\mathbf{k}|^2 \psi_{\mathbf{k}}. \end{cases} \quad (\text{B.39})$$

We begin by writing the linearized dynamics in Fourier space as

$$\partial_t \begin{pmatrix} \psi_{\mathbf{k}} \\ \rho_{\mathbf{k}} \end{pmatrix} = L(\mathbf{k}) \begin{pmatrix} \psi_{\mathbf{k}} \\ \rho_{\mathbf{k}} \end{pmatrix}, \quad (\text{B.40})$$

with the operator

$$L(\mathbf{k}) = \begin{pmatrix} -(D|\mathbf{k}|^2 + \kappa) & \alpha \\ -K_2 |\mathbf{k}|^2 & -[f_{\beta}'(0) - K_1(\mathbf{k})] |\mathbf{k}|^2 \end{pmatrix}. \quad (\text{B.41})$$

Here we define

$$K_1(\mathbf{k}) = 2f_{\beta}'(0) \rho_0(1 - \rho_0) u'(\rho_0) G_{\sigma}(\mathbf{k}), \quad G_{\sigma}(\mathbf{k}) = e^{-\sigma^2 \mathbf{k}^2 / 2}, \quad (\text{B.42})$$

$$K_2 = 2\lambda f_{\beta}'(0) \rho_0(1 - \rho_0). \quad (\text{B.43})$$

Note that with the logit rule, $f_{\beta}(0) = 1/2$ and $f_{\beta}'(0) = \beta/4$.

The eigenvalues of L are determined by the characteristic equation

$$\lambda^2 - \text{Tr } L \lambda + \det L = 0,$$

where

$$\text{Tr } L = -[D + f_\beta(0) - K_1(\mathbf{k})] \mathbf{k}^2 - \kappa, \quad (\text{B.44})$$

$$\det L = (D \mathbf{k}^2 + \kappa) [f_\beta(0) - K_1(\mathbf{k})] \mathbf{k}^2 + \alpha K_2 \mathbf{k}^2. \quad (\text{B.45})$$

Introducing

$$A(\mathbf{k}) \equiv f_\beta(0) - K_1(\mathbf{k})$$

and notation $|\mathbf{k}| = k$, the larger eigenvalue $X(\mathbf{k})$ satisfies

$$2 X(\mathbf{k}) = -[D + A(\mathbf{k})] k^2 - \kappa + \sqrt{\Delta(k)}, \quad (\text{B.46})$$

with the discriminant

$$\Delta(k) = [D - A(\mathbf{k})]^2 k^4 + [2 \kappa (D - A(\mathbf{k})) - 4 \alpha K_2] k^2 + \kappa^2. \quad (\text{B.47})$$

In the long-wavelength limit ($k \ll 1$), one finds $X > 0$ when

$$-4 D A(\mathbf{k}) k^4 - 4 [\kappa A(\mathbf{k}) + \alpha K_2] k^2 > 0. \quad (\text{B.48})$$

Expanding $A(\mathbf{k}) = f_\beta(0) - B e^{-\sigma^2 k^2/2}$ with

$$B = 2 f'_\beta(0) \rho_0 (1 - \rho_0) u'(\rho_0),$$

and retaining terms up to $O(k^2)$, this yields the stability criterion

$$\beta < \frac{2 f_\beta(0)}{\rho_0 (1 - \rho_0) [u'(\rho_0) - \alpha \lambda / \kappa]}. \quad (\text{B.49})$$

Part II: A zero intelligence agent-based model of the labor market

C.

C.1. Additional Plots




Figure C.1: D_∞ against D_0^\dagger for the fully connected network (in black) and the occupational mobility network (in cyan). This confirms the linear relationship between the stationary realised demand and the desired demand.




Figure C.2: Convergence time τ against the network spectral gap $\Delta\lambda$.




Figure C.3: Simulated Beveridge curve with $\delta_u < \delta_v$, which leads to a clockwise cycle in the unemployment-vacancy plane. We kept the same target demand as in Figure (5.2)




Figure C.4: Behavior of $D(t) - D_\infty$ when introducing constant demand shocks, compared with the exponential decay behavior fitted on the demand D (in red). We find $\tau \approx 9$ months, with time steps of approximately 6.75 weeks and parameters fitted on US labor market data.




Figure C.5: Behavior of simulated $D(t)$ (black line) compared to the functional form for $D(t)$ introduced in this study (in cyan). We also plotted $D^+(t)$ (black line).

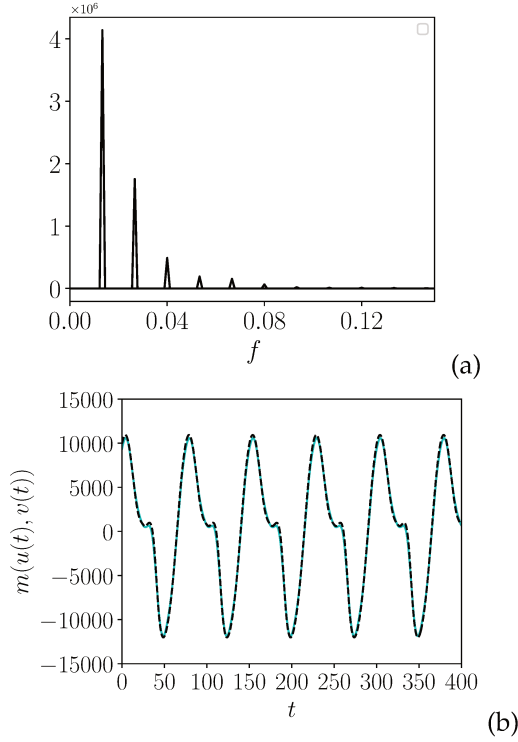


Figure C.6: a) Fourier spectrum of $m(u(t), v(t))$ (only the first harmonics are displayed). b) Behavior of simulated $m(u(t), v(t))$ (black line) compared to the functional form for $m(u(t), v(t))$ introduced in this study (in cyan).




Figure C.7: Comparison between the Beveridge curve obtained when simulating the complete labor market model (in black) with the Beveridge curve obtained when simulating the reduced model with the occupational mobility network matching process (in cyan). The global behavior of the curve is well preserved.

C.2. Stationary point and equilibrium in the ABM

Assuming that all the target demands reach a stationary value, $d_{i,t}^+ \xrightarrow{t \rightarrow \infty} d_i^{+*}$, the values of the other quantities in the economy can be found by setting the right hand side equations (5.4) to 0, which is precisely the same as writing

$$\frac{de(t)}{dt} = 0 \quad (\text{C.1})$$

in the DMP framework.

This leads to the conditions

$$\omega_i^* = v_i^* = \sum_j f_{ij}^* = \sum_j f_{ji}^*, \quad (\text{C.2})$$

which simply means that when the economy reaches a stationary state, the number of workers that are separated must match the incoming flow of workers, leading to a stationary unemployment rate. Similarly, at the stationary point the number of vacancies that are created matches the number of workers that are hired at every time-step.

Because of the very complicated nature of the f_{ji} coefficients finding an analytical expression for the equilibrium unemployment and vacancy rates is not possible. We can however give some intuition into the nature of the stationary point. Assume for example that $\delta_u > \delta_v$ and that $d_i^{+*} > d_i^*$, $\forall i$, then we can equate $\omega_i^* = v_i^*$, obtaining that

$$e_i^* = \frac{1}{1 + \kappa} (d_i^{+*} - v_i^*); \quad \kappa = \frac{\delta_u - \delta_v}{(1 - \delta_v)\gamma_v}, \quad (\text{C.3})$$

which is a rewriting of the expression $d_i^{+*} = d_i^* + \frac{\delta_u - \delta_v}{(1 - \delta_v)\gamma_v} e_i^*$ obtained in [82]. In this case, we see that $\kappa > 0$ and therefore $e_i^* < (d_i^{+*} - v_i^*)$. This means that, at equilibrium, firms are not able to hire as many people as they require. A strategy to overcome this would require bringing κ closer to 0, either by bringing δ_u closer to δ_v (that is, equating the rate at which people are hired with the rate at which vacancies are posted) or by increasing γ_v to make firms more reactive to labor supply-demand imbalances.

We can sum over i and note that the total unemployment rate at equilibrium U^* is equal to $L - \sum_i e_i^*$, where L is the amount of people in the economy. We therefore obtain a relation of the form

$$1 - \frac{1}{1 + \kappa} D^{+*} + \frac{1}{1 + \kappa} V^* = U^*, \quad (\text{C.4})$$

after imposing the normalisation $E + U = 1$, implying that at equilibrium the relationship between vacancies and unemployment is linear. We note that, formally, the stationary quantities depend on the target demands on the network and on all the other parameters, and therefore $v_i^* = v_i^*(\{d_j^{+*}\}, \gamma, \delta)$ and $u_i^* = u_i^*(\{d_j^{+*}\}, \gamma, \delta)$, with $\gamma = (\gamma_u, \gamma_v)$ and $\delta = (\delta_u, \delta_v)$.

C.3. Additional information on time scales in the ABM

In equation (5.14), the time-scale τ represents the time it takes for the labor market to reach the steady-state. We naturally expect it to depend on the different quantities at play in our equations. To understand this, note that the right-hand-side of Equations (5.4) must be homogeneous to an inverse time: the rate ω_i (resp. v_i) represents the number of separated workers (resp. created vacancies) *per unit time*, while f_{ij} represents also the flow of workers from an occupation to another *per unit time*.

This means that e.g. each worker has a probability $\delta_u dt$ of being separated spontaneously between t and $t + dt$, and a probability $\gamma_u [d_i(t) - d_i^+(t)]^+ dt$ of being separated to balance supply and demand.¹

Switching to a continuous-time formulation, we replace quantities such as $x_{i,t+1} - x_{i,t}$ by $\frac{dx_i(t)}{dt}$, and the corresponding right hand side by

$$\begin{aligned}\omega_{i,t+1} &\longrightarrow \omega_i(t) = \delta_u e_i(t) + \gamma_u [d_i(t) - d_i^+(t)]^+ \\ v_{i,t+1} &\longrightarrow v_i(t) = \delta_v e_i(t) + \gamma_v [d_i^+(t) - d_i(t)]^+.\end{aligned}\quad (\text{C.5})$$

1: Note that the term $\delta_u dt \gamma_u dt$ is of order dt^2 and therefore negligible in the continuous formulation.

C.4. Asymmetric behavior

When $\omega\tau \lesssim 1$, that is when the shock is very quick in relation to the adjustment time of the economy, the behavior of the economy during the crisis is not the same as the behavior during recovery and the asymmetric behavior arises from the model itself and not from ad hoc arguments. This asymmetric behavior with respect to positive or negative output shocks is a result of the DMP framework, so it is interesting to see that our less-constrained agent-based model manages to reproduce this fact.

To see this, note that it's possible to write

$$D^+(t) - D(t) = A^*(\omega) \cos(\omega t + \psi(\omega)), \quad (\text{C.6})$$

with $A^*(\omega) = \sqrt{A_2^2 \left(1 + \frac{\alpha^2}{1+\omega^2\tau^2}\right) - 2\alpha \frac{A_2^2}{\sqrt{1+\omega^2\tau^2}} \cos(\varphi(\omega))}$ and $\psi(\omega) = \arctan\left(\frac{\alpha \sin(\varphi(\omega))}{\alpha \sin(\varphi(\omega)) - \sqrt{1+\omega^2\tau^2}}\right)$. Assuming then that all demands are homogeneous, i.e. $d_i(t) = d(t) \quad \forall i$,

then we can write a mean-field version of equations Eq. (5.4) by summing over i , yielding

$$\begin{aligned}\frac{de}{dt} &= \delta_u e(t) + \gamma_u [D(t) - D^+(t)]^+ + F(t) \\ \frac{dv}{dt} &= \delta_v e(t) + \gamma_v [D^+(t) - D(t)]^+ - F(t),\end{aligned}\quad (\text{C.7})$$

with $F(t) = \sum_{ij} f_{ij}(t)$.

When $\omega\tau \rightarrow +\infty$, that is for very fast shocks, we will get $D^+ - D \approx A_2 \cos(\omega t)$: the economy never manages to match the global target demand and there is an oscillating mismatch between target and realised demand. In contrast, for slow shocks when $\omega\tau = 0$ the amplitude of $D^+ - D$ is equal to $A_2(1 - \alpha) \approx 0$ (if $\alpha \approx 1$), and also $\psi(\omega) \approx 0$: the economy is always close to its target demand.

In the second case, the γ terms in the equation above will not play a role in the dynamics, and the only source of non-linearities will be $F(t)$. Thus, equations Eq. (C.7) are symmetric in time, and the values taken by $e(t)$ and $v(t)$ during the crisis period $t \in [0; T/2]$ will follow the same path as those taken during the recovery period $t \in [T/2; T]$, but in reverse. This means that there is *no Beveridge cycle* for very slow shocks. This is confirmed by simulating Beveridge curves for different shock frequency as seen on Figure (5.2)b), where we see that the area of the Beveridge curve cycle is an increasing function of the shock frequency ω .

In the first case however, when $D^+ - D \approx A_2 \cos(\omega t)$, we will see a clear asymmetric behavior. As the crisis unfolds starting at $t = 0$ with a drop in the target demand, we will start in a setting where $D^+ - D > 0$, where the economy keeps creating vacancies at a rate proportional to γ_v . However, at $t = T/4$ we will enter a phase where $D^+ - D < 0$, as the current demand is higher than the target demand, leading to a stop in the creation of demand-induced vacancies and starting instead a phase of excess-labor-induced firings at a rate proportional to γ_u . There will therefore be a 'switch' in (C.7) as the demand-related term will go to 0 in the vacancies equation and will be simultaneously activated in the unemployment equation. During the recovery, after $t = T/2$, the behavior will therefore *not* be the same as during recession.

Part II: A simple yet effective disequilibrium model of the Beveridge curve

D.

D.1. Data analysis

We compare three search-and-matching labor market models, all written in terms of labour-force *fractions* ($0 < u, v < 1$) but displayed in this thesis in percentage points for readability. Throughout we work with monthly U.S. data¹ spanning January 1951 – December 2023.

- Vacancies are floored at 10^{-4} to avoid division by zero.
- GDP is filtered using a HP filter allowing to extract the cycle component from the trend and then *standardized with a 120-month rolling standard deviation*, then lagged one months (g_{t-1} drives \dot{u}_t).²
- We assume that, for Model 1, the spontaneous separation probability s is equal to a fraction of the empirical average separation probability measured by Hobijn and Sahin ($s_0 = 0.0106$), found through the following process: we compute

$$s' = \frac{1}{N} \sum_{t=1}^N \left(s_0 + C \left[\frac{dg}{dt} \right]_- \right),$$

where N is the number of observations and $\left[\frac{dg}{dt} \right]_-$ denotes the negative part of the derivative of GDP. We then set $s \leftarrow s'$ and recompute the average to obtain

$$s'' = \frac{1}{N} \sum_{t=1}^N \left(s' + C \left[\frac{dg}{dt} \right]_- \right).$$

Finally, we update the target value with

$$s \leftarrow s'' - s'.$$

This process converges to $s = 0.00265$. Models 2–3 treat s as free.

D.2. Fitting procedure and error evaluation for unemployment

We estimate each model by non linear least squares (NLLS), holding the separation rate fixed at $s = 0.00265$. For Models 1 and 3, which feature state-dependent dynamics, we define a residual vector

$$r_i(\theta) = u_{\text{sim},i}(\theta) - u_i, \quad i = 1, \dots, N,$$

where $u_{\text{sim},i}(\theta)$ is obtained by iterating $\Delta u_t = f(u_{t-1}, v_{t-1}, \Delta g_{t-1}; \theta)$. For Model 2, the static root form $u_t = s/(s + K v_t)$ is likewise recast as $r_i(K) = s/(s + K v_i) - u_i$. We then call `scipy.optimize.least_squares` with a finite-difference Jacobian. At the optimum $\hat{\theta}$, the estimated co-

1: JOLTS vacancy rate from 2001 aggregated with composite Help Wanted Index data spanning from 1951, BLS unemployment rate, and FRED real GDP converted to quarterly cyclical component by one-sided HP filter, then linearly interpolated to months.

2: We empirically find that a lag of one months gives the best evaluation metric when fitting the data.

variance matrix is

$$\widehat{\text{Var}}(\hat{\theta}) = \hat{\sigma}^2 [J(\hat{\theta})^\top J(\hat{\theta})]^{-1}, \quad \hat{\sigma}^2 = \frac{\text{SSE}(\hat{\theta})}{N - p},$$

where J is the $N \times p$ Jacobian of residuals, p is the number of parameters and SSE denotes the sum of squared errors. The diagonal entries of $\widehat{\text{Var}}(\hat{\theta})$ yield standard errors, and we form two-sided 95 % confidence intervals via $\hat{\theta}_j \pm t_{0.975, N-p} \text{se}(\hat{\theta}_j)$.

To test whether Model 1's "shock" parameter C is significantly different from zero, we note that Model 3 is nested within Model 1 via $C = 0$. Denoting SSE_{full} and SSE_{rest} as the residual sums of squares of Models 1 and 3 respectively, the classical F -statistic

$$F = \frac{(\text{SSE}_{\text{rest}} - \text{SSE}_{\text{full}})/1}{\text{SSE}_{\text{full}}/(N - 2)} \sim F(1, N - 2)$$

evaluates the null $C = 0$. Finally, because Model 2 and Model 1 are not nested, we carry out a residual bootstrap with $B = 500$ replications. For each draw, we resample the original model's residuals with replacement, re-fit both Model 1 and Model 2 to the pseudo-data, and record the bootstrap distributions of R^2 . We then compute 95 % bootstrap confidence intervals for $R^2_{(1)}$ and $R^2_{(2)}$, as well as for $\Delta R^2 = R^2_{(1)} - R^2_{(2)}$. Because the 95 % CI for ΔR^2 does not cross zero, we conclude that Model 1's in-sample fit is significantly better than Model 2's at the 5 % level.

Model	Param	Estimate	Std. Err.	95% CI Lower	95% CI Upper
1	C	0.011998	0.000631	0.010760	0.013237
	K	1.730775	0.025917	1.679908	1.781643
2	K	1.397351	0.014049	1.369778	1.424924
3	K	1.343137	0.014484	1.314710	1.371563

Table D.1.: Estimated Parameters (s fixed at 0.00265)

Model	R^2	MAE	RMS	Corr.
1	0.3733	0.01113	0.01362	0.6537
2	0.1265	0.01245	0.01608	0.5282
3	0.0600	0.01343	0.01668	0.3895

Table D.2.: Fit-Quality Metrics

Quantity	Value
SSE(Model 3, restricted)	0.243587
SSE(Model 1, full)	0.162389
$F(1, 874)$	437.017
p -value	0.0000
Decision	Reject $H_0: C = 0$. Model 1 fits significantly better than Model 3.

Table D.3.: F-Test for Model 1 vs Model 3 ($C = 0$ Restriction)

Description	Value
Bootstrap replicates (B)	500
Elapsed time (sec)	7.6
Model 1: R^2	0.3733 (95% CI: [0.3253, 0.4280])
Model 2: R^2	0.1265 (95% CI: [0.0326, 0.2371])
$\Delta R^2 = R^2_{(1)} - R^2_{(2)}$	0.2469 (95% CI: [0.1517, 0.3390])
Fraction of $\Delta R^2_{\text{boot}} > 0$	1.0000
Inference	Model 1's R^2 is significantly greater than Model 2's at the 5% level (no overlap).

Table D.4.: Bootstrap Inference for R^2 (Model 1 vs Model 2)

Results for vacancies

The results for fitting vacancies, when letting all parameters be free, can be summarized in the following tables:

Parameter	Estimate
v_0	0.005
C_v	0.008
K	2.65

Table D.5.: Estimated parameters for the vacancy equation

Metric	Value
R^2	0.09
MAE	0.007
RMS	0.01
Pearson correlation	0.52

Table D.6.: Goodness-of-fit metrics

D.3. Additional plots




Figure D.1: *Asymmetric lags for several periods recession.* Smoothed monthly unemployment (blue) and vacancies (orange) for the United States for several time periods, with the characteristic matching time $1/(Kuv)$ (black, right-hand scale) plotted in units of months required to achieve a one percent change based on matching alone. Vertical dashed lines represent key events. We overlay dashed colored lines to distinguish the mean slopes of unemployment and vacancy rates during recessionary and recovery phases. Our analysis shows that, for both series and all periods, the average slope during recessions is steeper than the subsequent recovery slopes.

Titre : Modèles minimaux hors-équilibre de systèmes socio-économiques : du marché immobilier français au marché du travail américain.

Mots clés : Physique statistique, hors équilibre, systèmes socio-économiques, modèles d'agents.

Cette thèse étudie la dynamique complexe et souvent contre-intuitive des systèmes socio-économiques en mariant la rigueur mathématique de la physique statistique à la flexibilité adaptative des simulations de modèles d'agents. À travers deux volets complémentaires, nous montrons que des interactions purement locales — qu'il s'agisse de champs continus comme les prix de l'immobilier ou d'agents individuels suivant des règles de décision simples et limitées — donnent naissance à des phénomènes agrégés riches que les théories économiques classiques ne peuvent pas expliquer. Dans la première partie, nous revisitons le modèle de ségrégation de Sakoda-Schelling sur un réseau. À partir d'un modèle d'agents qui évaluent localement la densité perçue, nous réalisons une analyse de stabilité linéaire et des simulations à grande échelle pour en déduire des équations hydrodynamiques stochastiques, nous permettant de lier le comportement hors d'équilibre du modèle à la théorie de la matière active. Nous démontrons que des fonctions d'utilité non linéaires suffisent à engendrer une transition de phase, pour une large gamme de paramètres, appartenant à la classe d'universalité du modèle d'Ising en deux dimensions. Nous étudions ensuite le marché immobilier français via une équation de diffusion stochastique micro-fondée heuristiquement à travers un mécanisme d'ajustement local des prix. Calibrée sur un demi-siècle de données, cette équation reproduit les régularités spatiales et temporelles observées — variogrammes spatiaux log-croissants, autocorrélations longue portée, persistance des chocs sur plusieurs

années —, et une analyse par réponse impulsionnelle fournit une constante de diffusion vraisemblable, confirmant des conjectures de longue date et soulignant le caractère non efficient de ce marché.

La seconde partie étudie un modèle d'agents zéro-intelligence dans lequel des demandeurs d'emploi postulent aléatoirement à des offres compatibles avec leurs compétences, tandis que les entreprises ajustent leurs annonces selon des règles heuristiques simples. De ces simulations discrètes, nous établissons rigoureusement des équations différentielles qui modélisent les délais de coordination, l'asymétrie entre embauches et licenciements, et le caractère non instantané du processus de mise en relation entre demandeurs d'emploi et entreprises. Ces lois continues comblent le fossé micro-macro et montrent comment des déphasages purement mécaniques génèrent l'hystérésis observée sur la courbe de Beveridge, sans faire appel aux hypothèses classiques d'agents rationnels. Nous en extrayons enfin un modèle dynamique minimal hors d'équilibre, calibré sur les données américaines de 1951 à 2023. À partir d'une trajectoire de taux de place vacantes donnée, ce modèle parcimonieux reproduit fidèlement le taux de chômage observé. Son caractère analytique le rend particulièrement utile pour l'analyse de politiques publiques, en capturant la mécanique hors-équilibre essentielle des fluctuations du marché du travail.

Title : Minimal Non-Equilibrium Models of Socio-Economic Systems: From the French Housing Market to the US Labor Market.

Keywords : Statistical Physics, Out-of-Equilibrium, Socio-Economic Systems, Agent-Based Models.

This thesis investigates the complex, often counter-intuitive dynamics of socio-economic systems by fusing the mathematical rigor of statistical physics with the adaptive flexibility of agent-based simulations. Across two complementary strands, we show that purely local interactions—whether governing continuous fields such as housing prices or guiding individual agents through simple, boundedly rational decision rules—generate rich aggregate phenomena that classical equilibrium paradigms cannot explain. In the first part of this thesis, we revisit the Sakoda-Schelling segregation model on a “neighborhood-less” lattice. From an agent-based model where each agent evaluates only the density in the vicinity of its own site, we perform linear-stability analysis and large-scale simulations to derive stochastic hydrodynamic equations mapping its out-of-equilibrium dynamics onto active-matter form. We demonstrate that nonlinear utilities alone drive phase-separation—segregation patterns—over a broad parameter range, and we identify the transition as belonging to the 2D Ising universality class. We then turn to the French housing market. We propose a stochastic diffusion equation micro-founded on simple, locally applied price-adjustment heuristics and microscopic noise. Calibrated on fifty years of regional price data, this equation reproduces key spatial and tempo-

ral regularities—variograms growing logarithmically with distance, long-range autocorrelations and multi-year shock persistence. An impulse-response analysis yields a diffusion constant of plausible magnitude, confirming longstanding conjectures about price diffusion and highlighting deviations from efficient-market behavior.

The second part of this thesis studies a zero-intelligence agent-based model in which unemployed workers apply randomly to skill-compatible vacancies and firms adjust postings via simple heuristics. From these discrete simulations we rigorously derive continuous disequilibrium equations isolating coordination lags, asymmetric hiring versus firing speeds and non-instantaneous matching. These laws bridge the micro-macro gap and show how purely mechanical delays generate the observed Beveridge-curve hysteresis, without rational-expectations assumptions. We finally distill these insights into a minimal dynamical disequilibrium model: a small system of coupled differential equations calibrated on US data (1951–2023). Given an observed vacancy trajectory, this parsimonious model reproduces the unemployment path with high fidelity. Its analytical tractability makes it a powerful tool for policy analysis, capturing the essential out-of-steady-state mechanics of labor-market fluctuations.