Department of Economics - FEA/USP

Endogenously time-varying heterogeneity in unemployment expectations: A discrete choice approach in an agent-based modeling

Julia de Sousa Pinheiro
Jaylson Jair da Silveira
Gilberto Tadeu Lima

Working Paper Series Nº 2026-11

DEPARTMENT OF ECONOMICS, FEA-USP
WORKING PAPER Nº 2026-11

Endogenously time-varying heterogeneity in unemployment expectations: A discrete choice approach in an agent-based modeling

Julia de Sousa Pinheiro (julia.sp215@gmail.com)

Jaylson Jair da Silveira (jaylson.silveira@ufsc.br)

Gilberto Tadeu Lima (giltadeu@usp.br)

Abstract:

This paper investigates the complex interplay between the dynamics of heterogeneous unemployment expectations across working households and observed unemployment dynamics. Drawing on considerable survey evidence that documents the persistent heterogeneity in households` unemployment expectations and its coevolution with observed unemployment, we develop an agent-based modeling (ABM) that incorporates a heterogeneous expectations-augmented efficiency wage mechanism. The time evolution of unemployment expectations is endogenized by means of a discrete choice protocol, allowing their frequency distribution across households to change endogenously over time and coevolve with observed unemployment. By integrating economic dynamics modeling with survey-based data, the simulated model reproduces qualitatively salient empirical regularities, including self-sustaining and coevolving cyclical fluctuations in both the heterogeneity in working households` expectations regarding unemployment and observed unemployment.

Keywords: Heterogeneous expectations; unemployment expectations; unemployment rate; discrete choice; agent-based modeling.

JEL Codes: C02; D84; E24; E70; J64.

Endogenously time-varying heterogeneity in unemployment expectations: A discrete choice approach in an agent-based modeling*

Júlia de Sousa Pinheiro[a], Jaylson Jair da Silveira[b], Gilberto Tadeu Lima[c]

[a] *Graduate Program in Economics, Federal University of Santa Catarina, Brazil*
[b] *Department of Economics and International Relations, Federal University of Santa Catarina, Brazil*
[c] *Department of Economics, University of São Paulo, Brazil (corresponding author)*

Abstract

This paper investigates the complex interplay between the dynamics of heterogeneous unemployment expectations across working households and observed unemployment dynamics. Drawing on considerable survey evidence that documents the persistent heterogeneity in households' unemployment expectations and its coevolution with observed unemployment, we develop an agent-based modeling (ABM) that incorporates a heterogeneous expectations-augmented efficiency wage mechanism. The time evolution of unemployment expectations is endogenized by means of a discrete choice protocol, allowing their frequency distribution across households to change endogenously over time and coevolve with observed unemployment. By integrating economic

*We are quite grateful to Andrea Roventini, Eva Catela and Roberto Meurer for helpful comments. Any remaining errors are our own. This research was supported by CNPq (the Brazilian National Council of Scientific and Technological Development) [grant numbers 311811/2018-3 and 316740/2023-3 (GTL); and 313628/2021-1 and 310013/2025-9 (JJS)]. This research was also financed by CAPES (Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil) [Finance Code 001 and grant number 88882.438558/2019-01 (JSP)], and in part by FAPESP (the São Paulo Research Foundation) [grant number 2025/27201-3 (GTL, JJS)]. The preparation of this paper benefited from a fruitful research visit by the first author (JSP) at the Sant'Anna School of Advanced Studies in Pisa, Italy, from September 2022 to April 2023. JSP gratefully acknowledges the kind hospitality and stimulating academic environment provided by the institution during this period.

Email addresses: `julia.sp215@gmail.com` (Júlia de Sousa Pinheiro), `jaylson.silveira@ufsc.br` (Jaylson Jair da Silveira), `giltadeu@usp.br` (Gilberto Tadeu Lima)

dynamics modeling with survey-based data, the simulated model reproduces qualitatively salient empirical regularities, including self-sustaining and coevolving cyclical fluctuations in both the heterogeneity in working households' expectations regarding unemployment and observed unemployment.

Keywords: Heterogeneous expectations; unemployment expectations; unemployment rate; discrete choice; agent-based modeling.

---

1. Introduction

Individuals' perceptions and expectations regarding key macroeconomic indicators can exert a significant influence on aggregate economic outcomes by shaping their decision-making processes (Curtin, 2019a). In particular, expectations regarding labor market conditions are especially salient for working-age individuals, as job loss represents one of the most consequential risks they face—often leading to reduced consumption and substantial welfare losses (Hendren, 2017). Households that overestimate current or future unemployment may perceive themselves to be in a weaker bargaining position during wage negotiations and, consequently, accept lower wages or jobs of lower quality than warranted (Orland, 2017).

Given the importance of individual and household unemployment expectations for fluctuations in key macroeconomic variables, several surveys have been conducted across different countries to gather data on these expectations. Typically, these surveys capture unemployment expectations along a spectrum ranging from (using this paper's terminology) pessimism to optimism, including neutral expectations that unemployment will remain about the same in the near future. The resulting survey data have provided robust evidence of the existence and persistence of heterogeneity in unemployment expectations across individuals or households (Curtin (2003); Carroll (2003); Fullone et al. (2007); Blanchflower and Kelly (2008); Malgarini and Margani (2008); Orland (2017); Kuchler and Zafar (2019)).

Several studies point to differences in information sets and in the cognitive ability to acquire and process information as potential drivers of the observed heterogeneity in expectations. However, Curtin (2019a) argues that consumer expectations data on unemployment in most developed economies provide compelling evidence against the presence of high costs associated with collecting, processing, and interpreting information on unemployment rate trends. In a related contribution, Curtin (2019b) contends that, in forming expectations, most households do not rely exclusively on national economic statistics but instead place greater weight on the labor market conditions they directly experience. Accordingly, one of the hypotheses underlying the agent-based modeling developed in this paper is that heterogeneity in expectations arises, *inter alia*, from differences in the general economic conditions faced by households in the labor market.

Meanwhile, other studies have presented robust econometric evidence of correlation between the state of unemployment expectations and the observed level or rate of unemployment (Curtin (2003); Dickerson and Green (2012); Leduc and Sill (2013); Girardi (2014); Lehmann and Weyh (2016); Curtin (2019a)). Drawing on different US and European surveys, these studies have found that households' unemployment expectations are an important driver of observed unemployment, with an increase in pessimism (optimism) leading to a rise (fall) in observed unemployment. This empirical finding is a second motivating factor considered in developing the model in this paper. In fact, one of the channels through which individual unemployment expectations may affect the observed unemployment rate can be intuitively understood within a heterogeneous expectations-augmented efficiency wage framework of the labor market, which is incorporated into the analysis in this paper.

In order to explore the dynamics of heterogeneity in unemployment expectations and its implications for the evolution of observed unemployment, we develop an agent-based modeling (ABM) that builds on the heterogeneous expectations–augmented efficiency wage framework proposed by Silveira and Lima (2021). In that framework, the authors show that the significant positive correlation between pessimistic unemployment expectations and the observed unemployment rate documented in the empirical literature can arise through a composition effect within a heterogeneous expectations–augmented efficiency wage setting.

The short-run equilibrium solution derived in Silveira and Lima (2021) is parameterized by the frequency distribution of unemployment expectations across working households. In the agent-based modeling developed in this paper, we endogenize the frequency distribution of unemployment expectations across working households within a dynamic version of the heterogeneous expectations–augmented efficiency wage framework of Silveira and Lima (2021). In this setting, workers' periodic revision of unemployment expectations is formally modeled through a discrete choice protocol, following the well-known framework of Brock and Hommes (1997). As a result, the frequency distribution of unemployment expectations across workers is endogenously time-varying and jointly determined with the dynamics of observed unemployment, with both exhibiting self-sustained cyclical fluctuations that qualitatively replicate several well-established empirical regularities.

A qualifying methodological remark is warranted. As is standard in formal modeling contributions, we abstract from several other determinants of observed unemployment dynamics to focus on—and gain insight into—the effects of time-varying heterogeneity in household expectations about near-future unemployment, for which there is extensive empirical evidence. We restrict attention to a limited set of deterministic and random transmission channels governing the discrete-choice process of unemployment expectation formation, allowing us to derive rationalizable results. In this respect, our approach is analogous to controlling for covariates when empirically testing for causal effects. Additionally, our purpose is to qualitatively, rather than quantitatively, replicate some related well-established empirical regularities.

In addition to this introduction, the paper is structured as follows. Section 2 draws on relevant literature that documents persistent heterogeneity in unemployment expectations and its interaction with observed unemployment, thereby setting the stage for the agent-based model developed in the next section. Section 3 details the formal structure of the model, while Section 4 presents and analyzes the main simulation results. Finally, Section 5 summarizes the key findings and offers concluding remarks.

2. Related literature on heterogeneous unemployment expectations and their coevolution with observed unemployment

To set the stage, we briefly survey key empirical contributions documenting the existence and persistence of heterogeneous unemployment expectations, which are central to motivating the modeling choices underlying the analysis developed herein. We also briefly report representative empirical evidence on the correlation between labor market pessimism or optimism and the observed unemployment rate, thereby underscoring the empirical relevance of the mechanisms and transmission channels incorporated into the modeling framework that follows.[1]

[1] The U.S. Michigan Survey of Consumers asks households each month: "How about people out of work during the coming 12 months—do you think there will be more unemployment than now, about the same, or less?" Responses of "don't know" and "no answer" appear only in very small proportions (https://data.sca.isr.umich.edu/) and

Unemployment expectations vary systematically across groups with different socioeconomic characteristics (Fullone et al. (2007); Blanchflower and Kelly (2008); Malgarini and Margani (2008); Orland (2017)). Using survey data, Kuchler and Zafar (2019) show that personal unemployment experience makes respondents significantly more pessimistic about the future national unemployment rate. Upon becoming unemployed, individuals assess the probability of rising U.S. unemployment over the next twelve months as 4–5 percentage points higher than when they were employed. Using survey data and a binary indicator of whether individuals expect unemployment to rise over the next twelve months, Blanchflower and Kelly (2008) find that this expectation is positively influenced by the current unemployment rate. Using Italian survey data, Malgarini and Margani (2008) find significant heterogeneity in unemployment expectations across socioeconomic groups, with some individuals relying on personal experiences rather than broader labor-market trends due to the cost of information gathering.

There are also studies that provide evidence of the link between expectations of unemployment and the observed unemployment rate, showing that a rise in negative beliefs about job prospects can cause a corresponding increase in observed unemployment. The study by Curtin (2003) examines U.S. unemployment expectations using data from the Michigan Survey spanning 1961 to 2002. It finds that expectations are positively correlated with future changes in unemployment, negatively correlated with past unemployment, and also reflect contemporaneous labor market conditions. The study by Curtin (2019a) employs Granger causality tests to examine the temporal relationship between unemployment expectations and the observed unemployment rate, using data from 15 countries between 1978 and 2016. The results indicate that in only seven countries do past unemployment changes predict expectations, whereas in twelve countries, expectations predict subsequent unemployment changes. For the U.S., relying on Michigan Survey data, past unemployment Granger-causes expectations only for the tertiary-education group, while expectations predict unemployment across all education levels.

are excluded when computing the proportions used in this paper. The main EU survey on unemployment expectations similarly asks households, on a monthly basis, how they expect the number of unemployed people in their country to change over the next 12 months. The response options are: increase sharply, increase slightly, remain the same, fall slightly, fall sharply, and don't know (`https://economy-finance.ec.europa.eu/economic-forecast-and-surveys`).

Tortorice (2012) uses microdata from the Michigan Survey to show that households form unemployment expectations in a systematically biased and heterogeneous way, especially around business cycle turning points. Rather than aligning with statistically predictable movements in unemployment, a large share of individuals continues to expect rising unemployment even at the onset of recoveries, revealing persistent pessimism and significant forecast errors. The evidence indicates that expectations are strongly backward-looking, driven by recent changes in unemployment rather than forward-looking information, with extrapolative behavior outperforming rational or learning-based models in explaining the data.

Leduc and Sill (2013) use data from the Livingston Survey and the Survey of Professional Forecasters to examine how changes in expectations affect aggregate macroeconomic variables. Using a VAR model, they find that declines in expected unemployment—reflecting more optimistic economic expectations—are associated with higher economic activity and lower observed unemployment. Impulse-response analyses indicate that a negative shock to expected unemployment six months ahead reduces the current unemployment rate. The results are robust when using Michigan Survey data as an alternative measure of unemployment expectations. Drawing on European survey data, Girardi (2014) shows that changes in household unemployment expectations are linked to fluctuations in the observed unemployment rate. VAR analysis indicates that, on average, a negative shock reducing expectations leads to a decrease in the current unemployment rate.

Lehmann and Weyh (2016) examine the forecasting power of employment expectations on employment growth in 15 European countries (1998–2014) using Granger causality tests and out-of-sample models, finding that three-month-ahead survey expectations efficiently anticipate employment growth in most countries. Meanwhile, Dickerson and Green (2012) examine the relationship between labor-market expectations and observed employment in Germany and Australia, relying on data on expected job loss and reemployment probabilities. They find that job loss expectations strongly predict subsequent job loss, and expectations of finding a good replacement correlate with actual reemployment success, though workers in both countries exhibit a tendency toward pessimism in their job-loss expectations.

3. General structure of the model

The research issues outlined in the Introduction are examined by developing an agent-based modeling that draws on the discrete choice framework of Brock and Hommes (1997) to extend the short-run efficiency wage model with heterogeneous expectations proposed in Silveira and Lima (2021) beyond the short run, thus rendering it dynamic.[2] The empirical contributions reported in the preceding section, documenting the existence and persistence of heterogeneous unemployment expectations as well as the correlation between labor market pessimism or optimism and the observed unemployment rate, discipline and guide the mechanisms and transmission channels incorporated into the modeling framework that follows.

The agent-based modeling approach offers several advantages that make it a powerful tool for studying complex economic dynamics. First, it provides the computational capacity to analyze social interactions involving heterogeneous agents on a large scale, enabling researchers to examine how individual behaviors aggregate into macro-level outcomes. Second, agent-based modeling allows behavior to be represented through decision rules or behavioral heuristics rather than strict mathematical optimization, making it possible to model more realistic forms of bounded rationality. Third, the growing prominence of behavioral research in economics supports and enriches the development of agent-based modeling by offering empirically grounded insights into how agents actually behave. Fourth, rapid advances in network theory within the social sciences enhance the formalization of interactions among agents, allowing for more accurate representations of social and economic linkages. Finally, improvements in the estimation and calibration of agent-based modeling increasingly facilitate rigorous assessments of how well these models align with empirical data, strengthening their credibility and usefulness in applied research (Dawid and Delli Gatti, 2018; Steinbacher et al., 2021).

Silveira and Lima (2021) propose a novel efficiency wage model augmented

[2] Evidence that firms pay efficiency wages for various reasons is documented using different methodologies by Krueger and Summers (1988), Dickens and Katz (1987), Cappelli and Chauvin (1991), and Goldsmith et al. (2000). More recent firm-level evidence shows that productivity growth and performance-linked pay are positively associated with higher wages and remuneration incentives, consistent with efficiency-wage mechanisms (Wallskog, 2026). Admittedly, measuring effort in real-world data is inherently difficult, which limits clean empirical tests.

with exogenously given, constant heterogeneous unemployment expectations. As this short-run model provides the foundation for wage-setting behavior in our analytical framework, its main structure will be summarized in Subsection 3.1.

3.1. Heterogeneous expectations-augmented efficiency wage setting in the short run

In Silveira and Lima (2021), building on the Michigan Survey, workers' unemployment expectations are classified as neutral (n), optimistic (o), or pessimistic (p), using the adjective labels introduced in that study, since the original survey does not assign such labels. Neutral workers expect unemployment to remain the same and exert more effort on the job than optimistic workers, pessimistic workers expect it to rise and exert the most effort, while optimistic workers expect it to fall and exert the least effort. Firms cannot perfectly observe individual effort, and due to potential heterogeneity in expectations across hired workers, they set a uniform wage to maximize profits and minimize labor cost per unit of average effort. Consequently, optimistic workers are more costly per unit of effort than neutral workers, and neutral workers are more costly than pessimistic workers, since all receive the same wage while exerting different levels of effort.

A worker's effort depends positively on the wage received and negatively on an indicator of wage compensation reflecting the worker's expected labor market conditions. Formally, the respective effort function is given by:

$$\varepsilon_\tau = \begin{cases} \left(\dfrac{w_\tau - \mu_\tau}{\mu_\tau}\right)^{\gamma}, & \text{if } w_\tau > \mu_\tau, \\ 0, & \text{otherwise}, \end{cases} \tag{1}$$

where ε_τ denotes the level of effort exerted by a worker of type $\tau \in \mathcal{T} = \{n, o, p\}$, $w_\tau \in \mathbb{R}_{++}$ is the wage received by the worker, and $\mu_\tau \in \mathbb{R}_{++}$ is the wage compensation indicator reflecting the expected labor market conditions for a worker of type $\tau \in \mathcal{T}$. The parameter $\gamma \in (0, 1) \subset \mathbb{R}$ captures the effect of increasing effort when paying a worker a wage above the level associated with their expected labor market conditions.

The functional form of μ_τ is given by:

$$\mu_\tau = (1 - u^e_\tau) w_{a,t}, \tag{2}$$

where $u^e_\tau \in [0,1] \subset \mathbb{R}$ denotes the expected unemployment rate of a worker of type τ. Drawing on the empirical literature on heterogeneity in unemployment expectations, Silveira and Lima (2021) postulate that the expected unemployment rate varies across agents and can be plausibly ordered according to worker type as follows:

$$0 < u^e_o < u^e_n = u < u^e_p < 1, \tag{3}$$

where u is the current rate of unemployment.

Each firm, unable to perfectly identify worker types, sets a uniform wage w that minimizes the cost of labor per unit of average effort ε. More specifically, each firm chooses the labor quantity L and the wage w that maximize its profit:

$$\pi = F(\varepsilon L) - wL, \tag{4}$$

where $F(\cdot)$ is the production function, with $F'(\cdot) > 0$ and $F''(\cdot) < 0$ in $\mathbb{R}_+$, and the average effort is defined as:

$$\varepsilon = \varepsilon_n^\eta \varepsilon_o^\theta \varepsilon_p^\rho, \tag{5}$$

where η, θ, and ρ denote the proportions of neutral, optimistic, and pessimistic workers, respectively, with $\eta + \theta + \rho = 1$.

Assuming that $w > \mu_\tau$, the first-order conditions for an interior solution to the maximization problem in (4) can be expressed as:[3]

$$\frac{\partial \pi}{\partial w} = F'(\varepsilon L) L \frac{\partial \varepsilon}{\partial w} - L = 0, \tag{6}$$

$$\frac{\partial \pi}{\partial L} = F'(\varepsilon L)\varepsilon - w = 0. \tag{7}$$

Substituting (6) into (7) yields the Solow condition, according to which the profit-maximizing pair (w, L) implies a unitary wage-effort elasticity:

$$\frac{\partial \varepsilon}{\partial w}\frac{w}{\varepsilon} = 1. \tag{8}$$

Taking into account the heterogeneity in unemployment expectations, the average effort level in (8) can be replaced by the definition of ε given in (5). With this substitution, the condition in (8) can be rewritten as:

[3]The first-order conditions are sufficient to guarantee a maximum, as the production function is assumed to be strictly concave.

$$\eta \frac{\partial \varepsilon_n}{\partial w} \frac{w}{\varepsilon_n} + \theta \frac{\partial \varepsilon_o}{\partial w} \frac{w}{\varepsilon_o} + \rho \frac{\partial \varepsilon_p}{\partial w} \frac{w}{\varepsilon_p} = 1, \tag{9}$$

which Silveira and Lima (2021) dub the weighted Solow condition. Using (1) and (2), the above condition can be rewritten as:

$$\left(\frac{\eta}{u_n^e} + \frac{\theta}{u_o^e} + \frac{\rho}{u_p^e} \right) \gamma = 1. \tag{10}$$

Silveira and Lima (2021) assume the following specific form for the well-defined ordering for the unemployment expectations of employed workers of type $\tau = o, n, p$ in 3:

$$u_\tau^e = \begin{cases} (1-\delta)u, & \text{if } \tau = o, \\ u, & \text{if } \tau = n, \\ (1+\delta)u & \text{if } \tau = p, . \end{cases} \tag{11}$$

where $\delta \in (0, 1-\gamma) \subset \mathbb{R}$ is a dispersion factor.[4]

Substituting (11) into (10), the short-run equilibrium unemployment rate of the model—prevailing for a predetermined distribution of unemployment expectations across workers—is given by:

$$u^* = \left[1 + \left(\frac{\delta}{1-\delta} \right) \theta - \left(\frac{\delta}{1+\delta} \right) \rho \right] \gamma. \tag{12}$$

Once the equilibrium unemployment rate is determined, we can use (1), (2), (5), and (12) to derive the short-run equilibrium average effort level, which is given by:

$$\varepsilon^* = \left[\left(\frac{u^*}{1-u^*} \right)^\gamma \right]^{1-(\theta+\rho)} \left[\left(\frac{(1-\delta)u^*}{1-(1-\delta)-u^*} \right)^\gamma \right]^\theta \left[\left(\frac{(1+\delta)u^*}{1-(1+\delta)-u^*} \right)^\gamma \right]^\rho . \tag{13}$$

The equilibrium effort of each agent of type $\tau \in \mathcal{T}$ can be deduced as follows. First, by substituting (11) into (2) and considering the symmetric

[4] As noted in footnote 1, Michigan Survey respondents do not report precise numerical values for expected unemployment, so the cardinal dispersion across worker types is treated as exogenous, with equal gaps between optimistic–neutral and neutral–pessimistic workers in the model developed here.

Nash equilibrium condition ($w_{a,\tau} = w$, $\forall \tau \in \mathcal{T}$), we obtain the expected labor market conditions at this equilibrium:

$$\mu^*_\tau = \begin{cases} [1-(1-\delta)u^*]w^*, & \text{if } \tau = o, \\ (1-u^*)w^*, & \text{if } \tau = n, \\ [1+(1+\delta)u^*]w^* & \text{if } \tau = p. \end{cases} \tag{14}$$

Next, by substituting (14) into (1), we derive the short-run equilibrium effort for each type of worker:

$$\varepsilon^*_o = \left\{ \frac{w^* - [1-(1-\delta)u^*]w^*}{[1-(1-\delta)u^*]w^*} \right\}^\gamma = \left[\frac{(1-\delta)u^*}{1-(1-\delta)u^*} \right]^\gamma \tag{15}$$

$$\varepsilon^*_n = \left[\frac{w^* - (1-u^*)w^*}{(1-u^*)w^*} \right]^\gamma = \left(\frac{u^*}{1-u^*} \right)^\gamma, \tag{16}$$

$$\varepsilon^*_p = \left\{ \frac{w^* - [1-(1+\delta)u^*]w^*}{[1-(1+\delta)u^*]w^*} \right\}^\gamma = \left[\frac{(1+\delta)u^*}{1-(1+\delta)u^*} \right]^\gamma. \tag{17}$$

It is noteworthy that, as expected, in the short-run equilibrium, workers holding pessimistic expectations regarding the unemployment rate exert a relatively higher level of effort, reflecting a higher expected cost of job loss. Conversely, workers with optimistic expectations exert a relatively lower level of effort due to a lower expected cost of job loss, while neutral workers exert an intermediate level of effort, falling between optimistic and pessimistic workers. Accordingly, $\varepsilon^*_p > \varepsilon^*_n > \varepsilon^*_o$.

3.2. Unemployment expectations formation as a discrete choice process

This subsection develops a protocol in which near-future unemployment expectations are formed through a discrete choice mechanism, following the Adaptively Rational Equilibrium Dynamics (ARED) framework of Brock and Hommes (1997). In the ARED setting, agents form predictions of future endogenous variables using equilibrium relationships and update their beliefs over time by selecting from a finite set of predictors. Each predictor maps past observations into a forecast, and its performance metric is publicly observable. Agents choose among predictors according to their relative past performance, and this responsiveness is governed by a parameter measuring the intensity of choice, which plays a central role in the framework.

Inspired by the ARED framework, the agent-based modeling developed here assumes that, in each period, workers form optimistic, pessimistic, or neutral expectations about near-future unemployment. Pessimistic workers exert more effort than neutral workers, who in turn exert more than optimistic workers, although firms cannot perfectly observe individual expectation and hence effort. When the unemployment rate is relatively low, workers tend to form more favorable expectations about labor market conditions and exert less effort, as the expected cost of job loss is lower. Thus, according to the equilibrium unemployment rate in (12), a higher share of optimistic workers relative to pessimistic or neutral ones increases the unemployment rate, whereas a lower share reduces it. A low unemployment rate therefore reinforces optimism, reduces average effort, and pushes the unemployment rate upward, while a high unemployment rate induces pessimism, strengthens effort incentives, and lowers the equilibrium unemployment rate. This feedback mechanism potentially generates fluctuations in both the unemployment rate and the composition of workers' expectations about it, as suggested by empirical evidence. Workers' expectations and associated effort levels depend on the equilibrium unemployment rate, which itself depends on the frequency distribution of optimistic, neutral, and pessimistic workers. Therefore, as in the example of two predictors for price expectations presented by Brock and Hommes (1997) to illustrate the logic of the ARED mechanism, our model also features two opposing forces governing the co-evolution of the observed unemployment rate and the frequency distribution of workers' expectations about near-future unemployment.

3.2.1. The discrete choice protocol

Consider a worker (agent) $i \in \{1, 2, 3, \ldots, A\}$, where A denotes the total number of workers. Using the Michigan Survey as a reference, in each period t worker i is characterized by a type $\tau_i \in \mathcal{T} = \{n, o, p\}$, where the respective subscripts denote neutrality, optimism, and pessimism regarding the unemployment rate in a future period $t + \Delta$. Here, $\Delta \in \mathbb{N}_{++}$ represents a given and finite forecasting horizon. In the Michigan Survey, for example, $\Delta = 12$ months. Henceforth, when we refer to the expected (near-future) unemployment rate, or to the unemployment rate in the near future, we mean the expectation formed in a given period t regarding the unemployment rate in period $t + \Delta$. More precisely, if a worker is of type $\tau_i = o$ (type $\tau_i = p$), she holds an optimistic (pessimistic) expectation about unemployment and believes that the unemployment rate in the near future will be lower (higher) than the current rate. By contrast, a worker of type $\tau_i = n$ expects the unemployment rate in the relevant future to remain equal to the current rate. In the environment under analysis, the next 12 months constitute the relevant near-future horizon for forming unemployment expectations.

Based on Train (2009), the utility (or payoff) function of a worker i can be additively decomposed into a deterministic component, denoted by $U^d(\tau_i)$, associated with her observable motivations, and a random component, denoted by $\varsigma(\tau_i)$, capturing her unobservable motivations, as follows:

$$U(\tau_i) = U^d(\tau_i) + \varsigma(\tau_i). \tag{18}$$

We adopt the standard assumption in the discrete choice literature that agents of the same type share the same deterministic component in their utility function. That is, for any two workers i and j, if $\tau_i = \tau_j$, then $U^d(\tau_i) = U^d(\tau_j)$. By contrast, the unobserved component of the utility (or payoff) may be heterogeneous across workers: for any two workers i and j, even if $\tau_i = \tau_j$, it may still be the case that $\varsigma(\tau_i) \neq \varsigma(\tau_j)$. In the model set forth herein, this unobserved component reflects idiosyncratic, exogeneously given characteristics that influence the formation of unemployment expectations.

At the beginning of each period t, an individual worker forms her unemployment expectation (which we interpret as her choice of type τ_i in that period) by selecting the available alternative that provides her with the highest utility. Consequently, an individual worker will be of the type $\tau_i \in \mathcal{T}$ whose associated utility satisfies:

$$U(\tau_i) \geq U(\tau_i^{'}), \forall \tau_i^{'} \in \mathcal{T}. \tag{19}$$

Considering the utility function in (18), the discrete choice rule in (19) can be equivalently expressed as:

$$U^d(\tau_i) - U^d(\tau_i^{'}) \geq \varsigma(\tau_i^{'}) - \varsigma(\tau_i), \forall \tau_i^{'} \in \mathcal{T}. \tag{20}$$

Note that the inequality in (20) implies that even if the deterministic component of the utility associated with the unemployment expectation of type τ_i is strictly greater than that associated with an alternative unemployment expectation of type τ_i'—that is, even if $U^d(\tau_i) - U^d(\tau_i') > 0$—the unemployment expectation of type τ_i will not necessarily be the one chosen by an individual worker i in a given period t.

Given the presence of a random component affecting the choice of unemployment expectation type by the i-th worker, we can only characterize her statistical propensity to adopt each type of expectation in $\mathcal{T}$. Using

(19) and (20), the probability that the i-th worker selects the unemployment expectation of type τ_i can be formally expressed as:

$$\begin{aligned} Prob(\tau_i) &= Prob\big(U(\tau_i) \geq U(\tau_i^{'}), \forall \tau_i^{'} \in \mathcal{T}\big) \\ &= Prob\big(\varsigma(\tau_i^{'}) - \varsigma(\tau_i) \leq U^d(\tau_i) - U^d(\tau_i^{'}), \forall \tau_i^{'} \in \mathcal{T}\big) \\ &= \int_{-\infty}^{\infty} I[\varsigma(\tau_i^{'}) - \varsigma(\tau_i) \leq U^d(\tau_i) - U^d(\tau_i^{'}), \forall \tau_i^{'} \in \mathcal{T}] f(\vec{\varsigma_i})(d\vec{\varsigma_i}), \end{aligned} \tag{21}$$

where $f(\vec{\varsigma_i})$ is the joint probability density function of the vector of random variables $\vec{\varsigma_i}$, which is composed of the random variables $\varsigma(\tau_i)$, with $\tau_i \in \mathcal{T}$, and $I[\cdot]$ is the indicator function that takes the value 1 if the condition within the brackets is satisfied and 0 otherwise. Importantly, the expression in (21) implies that the propensity of worker i to select the strategy τ_i increases as the differential in observable incentives favoring this alternative becomes larger.

As noted by Train (2009), different specifications of $f(\vec{\varsigma_i})$ give rise to various discrete choice models. In particular, a specification that is often considered convenient—because it yields a closed-form expression for the integral in (21)—is the logit specification. This version is obtained by assuming that the random components of the utility functions in (18) are independent and identically distributed according to the type I extreme value (Gumbel) distribution, whose probability density function is given by:

$$f(\varsigma(\tau_i)) = \beta e^{-\beta\varsigma(\tau_i)} e^{-e^{-\beta\varsigma(\tau_i)}}, \tag{22}$$

where $\beta \in \mathbb{R}_+$ is a parametric constant, referred to as the intensity of choice in Brock and Hommes (1997). The cumulative distribution function associated with the function in (22) is given by:

$$F(\varsigma(\tau_i)) = e^{-e^{-\beta\varsigma(\tau_i)}}. \tag{23}$$

Using the functions in (22) and (23), we can solve the integral in (21) to obtain the choice probability of the i-th worker, which is given by the logistic cumulative distribution function:

$$Prob(\tau_i) = \frac{e^{\beta U^d(\tau_i)}}{\sum\limits_{\tau_i^{\prime} \in \mathcal{T}} e^{\beta U^d(\tau_i^{\prime})}} = \frac{1}{1 + \sum\limits_{\tau_i^{\prime} \in \mathcal{T}, \tau_i^{\prime} \neq \tau_i} e^{-\beta[U^d(\tau_i) - U^d(\tau_i^{\prime})]}}. \tag{24}$$

It follows from (24) that, *ceteris paribus*, the higher the value of the intensity-of-choice parameter β, the greater the relative influence of the deterministic component (and thus of observable motivations) compared with the random component (and hence idiosyncratic motivations) in determining the propensity of worker i to adopt the unemployment expectation of type τ_i. Put differently, the intensity-of-choice parameter measures how strongly workers respond to differences in deterministic attractiveness when forming unemployment expectations. As the intensity of choice, β, approaches $+\infty$, the discrete choice protocol in (24) becomes fully deterministic, so that the alternative with the highest deterministic utility component is chosen with probability one. Conversely, as $\beta \to 0$, the protocol becomes entirely random, and the choice probabilities converge to a uniform distribution (one-third in the present case of three unemployment expectation types), rendering the deterministic utility components irrelevant.

The discrete choice protocol outlined above will be employed in the next subsection to model unemployment expectation formation within an agent-based modeling framework.

3.2.2. An agent-based modeling of unemployment expectation formation as a discrete choice process

In the environment of unemployment expectation formation considered herein, each worker i, in each period t, forms either an optimistic, pessimistic, or neutral unemployment expectation depending on her private (or individual) incentive to obtain the highest wage per unit of effort and her social incentive to conform, to some extent, to the predominant type of unemployment expectation in the population of workers. The latter incentive reflects the presence of a popularity (or diffusion) effect, whereby the unemployment expectation adopted by worker i in period t is also influenced by the expectations prevailing in the population of workers.

Following Brock and Durlauf (2001), we decompose the total deterministic utility into two components: the private deterministic utility, associated with individual incentives, and the social deterministic utility, associated with social incentives. Specifically, the total deterministic utility $U^d(\tau_i)$ in (18), in each period t, is expressed as the sum of the private deterministic utility, denoted by $V(\cdot)$, and the social deterministic utility, denoted by $S(\cdot)$, with the latter weighted by a positive parametric constant $\psi \in \mathbb{R}_{++}$, as follows:

$$U^d(\tau_{i,t}) = V(\tau_{i,t}) + \psi S(\tau_{i,t}). \tag{25}$$

We assume that workers derive higher private utility from higher wages per unit of effort. However, as the unemployment rate increases, the likelihood that a given worker has earned that wage per unit of effort decreases. More formally, the private utility of the i-th worker in a given period t, with $i \in \{1, 2, \dots, A\}$, is represented by her wage per unit of effort weighted by the level of employment, as follows:

$$V(\tau_{i,t}) = (1 - u_t^*) \frac{w_{i,t}^*}{\varepsilon_{\tau_{i,t}}^*}, \tag{26}$$

where $\varepsilon_{\tau_{i,t}}^*$ denotes the short-run equilibrium level of effort supplied by a worker holding an unemployment expectation of type $\tau_{i,t} \in \mathcal{T}$, while u_t^* and $w_{i,t}^*$ represent, respectively, the unemployment and wage rates prevailing in a given short-run equilibrium in period t.[5]

Let η_t, θ_t, and ρ_t denote the respective proportions of neutral, optimistic, and pessimistic workers in their population during period t, with $(\eta_t, \theta_t, \rho_t) \in \Delta$, where $\Delta \equiv \{(\eta, \theta, \rho) \in \mathbb{R}_+^3 : \eta + \theta + \rho = 1\}$ represents the unit simplex in $\mathbb{R}_+^3$, comprising all feasible frequency distributions of unemployment expectations across workers.

The short-run equilibrium unemployment rate in a given period t, taking into account (12), can be expressed as:

$$u_t^* = \left[1 + \left(\frac{\delta}{1-\delta}\right)\theta_t - \left(\frac{\delta}{1+\delta}\right)\rho_t\right]\gamma, \tag{27}$$

recalling that $\delta \in (0, 1-\gamma) \subset \mathbb{R}$ is a parametric constant capturing the dispersion among the three types of unemployment expectations, and $\gamma \in (0,1) \subset \mathbb{R}$ is a parametric constant measuring the effort-enhancing effect of paying a worker a wage above the compensation associated with her expected labor market conditions.

Meanwhile, using (1) and (7), and assuming a Cobb-Douglas production function of the form $F(\cdot) = (\epsilon L)^\alpha$, the wage rate in the short-run equilibrium in a given period t, which is assumed homogeneous across workers, is given by:

[5] That is, the values of the variables $\varepsilon_{\tau_{i,t}}^*$, u_t^*, and $w_{i,t}^*$, determined for a given frequency distribution of unemployment expectations across workers (and thus parameterized by it) in period t, as derived in the short-run heterogeneous expectations-augmented efficiency wage model developed in Silveira and Lima (2021).

$$w_{i,t}^* = \alpha(\varepsilon_t^*)^\alpha(1-u_t^*)^{\alpha-1} \equiv w_t^*, \tag{28}$$

where ε_t^* denotes the average effort level in the short-run equilibrium in a given period t, which, taking into account (13), is given by:

$$\varepsilon_t^* = \left[\left(\frac{u_t^*}{1-u_t^*}\right)^\gamma\right]^{1-(\theta_t+\rho_t)} \left[\left(\frac{(1-\delta)u_t^*}{1-(1-\delta)u_t^*}\right)^\gamma\right]^{\theta_t} \left[\left(\frac{(1+\delta)u_t^*}{1-(1+\delta)u_t^*}\right)^\gamma\right]^{\rho_t}. \tag{29}$$

Finally, taking into account (15)–(17), the effort level exerted by a worker i holding an unemployment expectation of type $\tau_{i,t} \in \mathcal{T}$ in the short-run equilibrium in a given period t is given by:

$$\varepsilon_{\tau_{i,t}}^* = \begin{cases} \left(\frac{u_t^*}{1-u_t^*}\right), & \text{if } \tau_{i,t} = n, \\ \left[\frac{(1-\delta)u_t^*}{1-(1-\delta)u_t^*}\right]^\gamma, & \text{if } \tau_{i,t} = o, \\ \left[\frac{(1+\delta)u_t^*}{1-(1+\delta)u_t^*}\right]^\gamma, & \text{if } \tau_{i,t} = p. \end{cases} \tag{30}$$

Meanwhile, the social utility in (25) of a worker i holding an unemployment expectation of type τ_i is determined by the popularity (or diffusion) of that expectation within the population of agents. Specifically, we assume that the social utility associated with a worker's unemployment expectation increases with the proportion of workers who held the same expectation in the previous period:

$$S_{\tau_{i,t}} = \begin{cases} \theta_t, & \text{if } \tau_{i,t} = o, \\ \eta_t, & \text{if } \tau_{i,t} = n, \\ \rho_t, & \text{if } \tau_{i,t} = p. \end{cases} \tag{31}$$

Having specified the private and social components of the utility function in (25), then (24) can be rewritten as follows:

$$Prob(\tau_{i,t}) = \frac{1}{1 + \sum\limits_{\tau'_{i,t-1} \in \mathcal{T}, \tau'_{i,t-1} \neq \tau_{i,t-1}} e^{-\beta\left\{[V(\tau'_{i,t-1})+\psi(S(\tau'_{i,t-1}))]-[V(\tau_{i,t-1})+\psi(S(\tau_{i,t-1}))]\right\}}}. \tag{32}$$

3.3. Computational implementation and calibration of the proposed agent-based model

The number of agents (workers) used to implement the proposed agent-based modeling was chosen based on the household survey serving as a reference for model calibration. Specifically, we used data from the Michigan Survey, ensuring the model comprises 501 agents.[6] In the initial period, each type of unemployment expectation—neutral, optimistic, and pessimistic—is held by one-third of the population of agents, so that 167 agents held each type. Recall that at the beginning of each period t, a worker i forms either a neutral ($\tau_i = n$), optimistic ($\tau_i = o$), or pessimistic ($\tau_i = p$) expectation about the unemployment rate in that period. After establishing the initial conditions, we compute the short-run equilibrium values of the unemployment rate, the wage rate, and the average and individual effort levels in period 1 using (27) first, followed by (28), (29), and (30). Based on the short-run equilibrium unemployment rate, we then compute the corresponding wage and individual effort levels for period 1. These values are subsequently employed to calculate the private utility in this period via (26), and the social utility via (31). After computing the private and social deterministic utilities for period 1, we calculate the total deterministic utility in this period using (25). Given the short-run equilibrium value of total private utility in period 1, we then compute the choice probabilities specified in (32), which play a central role in determining the frequency distribution of unemployment expectations across workers in the subsequent period.

Following the literature suggesting that agents' unemployment expectations reflect their personal unemployment experiences (see, e.g., Campbell et al. (2007); Kassenboehmer and Schatz (2017); Curtin (2019a)), we reasonably specify that workers who experienced unemployment in any of the $t - l$ periods, with $l \in \{1, 2, \ldots, L\} \subset \mathbb{N}$, where L denotes the number of past periods during which unemployment experience continues to increase the propensity to hold a pessimistic unemployment expectation in period t, are more likely to form pessimistic expectations about their near-future employment prospects than workers who were employed in all of the L past periods. However, it does not seem plausible to specify that the effect of

[6] The Michigan Survey includes approximately 500 respondents; we consider 501 agents here to allow an equal distribution of agents among the three alternative types of unemployment expectation in the initial period.

the most recent unemployment experience on the propensity for pessimism is the same as that of earlier unemployment states. Accordingly, we assume that the probability that a worker who was unemployed in period $t-l$ holds a pessimistic unemployment expectation in period t decreases with l. With this in mind, we introduce the following indicator function:

$$I^i_{t-l} = \begin{cases} 1, & \text{if the worker } i \text{ was unemployed in period } t-l, \\ 0, & \text{otherwise.} \end{cases} \tag{33}$$

Assuming that the weight of past unemployment on the propensity to hold pessimistic expectations (ι) decays geometrically with l at a ratio $q \in (0,1) \subset \mathbb{R}$, we have:

$$\iota_{t-l} = q\iota_{t-l-1}. \tag{34}$$

Considering that $\sum_{l=1}^{L} \iota_{t-l} = 1$ and (34), we can rewrite ι_{t-1} as:

$$\iota_{t-1} = \frac{q-1}{q^L - 1}. \tag{35}$$

Let F^i_t denote the unemployment factor of worker i in period t, which we define as:

$$F^i_t = \iota_{t-1}I^i_{t-1} + \iota_{t-2}I^i_{t-2} + ... + \iota_{t-l}I^i_{t-L}. \tag{36}$$

Using (33) and (34) in (36), the unemployment factor can be expressed as:

$$F^i_t = \iota_{t-1}\sum_{l=1}^{L} q^{l-1}I^i_{t-l} = \frac{q-1}{q^L-1}\sum_{l=1}^{L} q^{l-1}I^i_{t-l}. \tag{37}$$

At the end of each period, given that all workers face the same probability of obtaining a job, we randomly select the workers who were not hired by any firm. For example, if the unemployment rate in a given period is 5%, we randomly select 5% of the worker population (or the nearest integer) to experience unemployment in that period.

For workers who have been unemployed in at least one of the previous L periods leading up to period t, we assume a bias toward pessimism in period t, using the following indicator as a reference:

$$\mathcal{Q}_t = min\{(1 + F^i_t)Prob(\tau_{i,t} = \rho), 1\}, \tag{38}$$

where $Prob(\tau_{i,t} = \rho)$ is calculated using (32). Given the value of $\mathcal{Q}_t$, we draw a random number $r \in [0,1] \subset \mathbb{R}$ from a uniform distribution. If $r_{i,t} < \mathcal{Q}_t$, then worker i in period t is assigned a pessimistic unemployment expectation.

In turn, for workers who were employed in all of the L periods prior to period t, the type of unemployment expectation they will hold in the next period is determined solely by the choice probabilities specified in (32). Given these probabilities, we draw a random number $r \in [0, 1] \subset \mathbb{R}$ from a uniform distribution to assign the unemployment expectation of such workers in any period $t \geq 2$, following the rules outlined in Table 1.

Table 1: Algorithm for unemployment expectation formation in period $t \geq 2$ by a worker i who does not have a pessimistic bias.

Possible cases	**Worker i's unemployment expectation at period $t \geq 2$**
$r \leq \text{Prob}(\tau_{i,t} = p)$	Pessimistic
$\text{Prob}(\tau_{i,t} = p) < r \leq \text{Prob}(\tau_{i,t} = p) + \text{Prob}(\tau_{i,t} = n)$	Neutral
$r > \text{Prob}(\tau_{i,t} = p) + \text{Prob}(\tau_{i,t} = n)$	Optimistic

The values of the following parameters were obtained through calibration to identify the combination that provides the best fit of the model to the empirical data: β (the intensity of choice in (22)–(24)), ψ (the weight of the social component in the total deterministic utility in (25)), δ (the measure of dispersion among the three types of unemployment expectations in (27)), q (the decay rate of the unemployment weight on the bias to pessimism in (34)), and L (the number of lags during which past unemployment experience continues to influence the bias to pessimism, as formalized in (33)–(38)).

Our calibration strategy consisted of identifying the combination of parameter values that minimizes the sum of squared deviations between the simulated data and the corresponding observed data. Two empirical time series were used for calibration. The first was the Balance Score (BS) monthly time series, calculated from the proportion of optimistic and pessimistic respondents in the Michigan survey. As noted by Curtin (2019a), the BS is a standard method for reporting survey-based expectations data and is defined as the percentage of respondents who anticipate an increase in the unemployment rate minus the percentage who anticipate a decrease, plus 100. The second time series used for calibration was the U.S. monthly unemployment rate.[7] The series employed in our calibration procedure spans January 1978 to December 2019, totaling 504 months. The parameter combination selected

[7] Available at `https://fred.stlouisfed.org/series/UNRATE`.

was the one that minimized the following objective function:

$$\sum_{t=1}^{T}[(u_t^{obs} - u_t^{sim})^2 + (BS_t^{obs} - BS_t^{sim})^2], \tag{39}$$

where $T = 504$ denotes the total number of periods, u^{obs} is the observed unemployment rate, u^{sim} is the simulated unemployment rate, BS^{obs} is the observed BS, and BS^{sim} is the simulated BS. The simulated BS was calculated in the same manner as the observed BS in the Michigan Survey, as follows:

$$BS_t = 100(\rho_t - \theta_t + 1). \tag{40}$$

The MATLAB function *fminsearchbnd* was employed to solve this minimization problem. The function explores different combinations of parameter values, subject to pre-specified bounds. For an initial randomly selected parameter vector, the algorithm computes the simulated values of the BS and the unemployment rate at each simulation step. These simulated values, together with the corresponding empirical observations, are then inserted into the objective function in (39) to obtain the function value associated with the given parameter vector. The same procedure is repeated for subsequent randomly selected parameter vectors. If a new combination yields higher accuracy—i.e., a smaller distance between the simulated and empirical series—it is retained, while the previous combination is discarded by the *fminsearchbnd* search algorithm. This iterative process continues for up to 1,000 different parameter combinations or until the algorithm identifies a set that reduces the value of (39) by less than 0.01.

We defined the following plausible ranges within which the *fminsearchbnd* algorithm searched for the parameter values that provide the best fit of the model to the empirical data: $0 \leq \beta \leq 10$, $0 \leq \delta \leq 0.97$, $0 \leq \psi \leq 2$, $0 \leq q \leq 1$, and $1 \leq L \leq 12$. To mitigate the risk of convergence to local minima, we sought an initial configuration of parameters that would ultimately generate the smallest value of the objective function in (39). For this purpose, 400 initial parameter vectors were randomly sampled within the previously specified admissible ranges. Among these 400 initial configurations, the combination yielding the best preliminary fit was $\beta = 7.72$, $\delta = 0.91$, $\psi = 0.33$, $L = 9.09$, and $q = 0.11$. Using this configuration as the initial input, the *fminsearchbnd* algorithm subsequently produced the calibrated parameter values reported in Table 2.

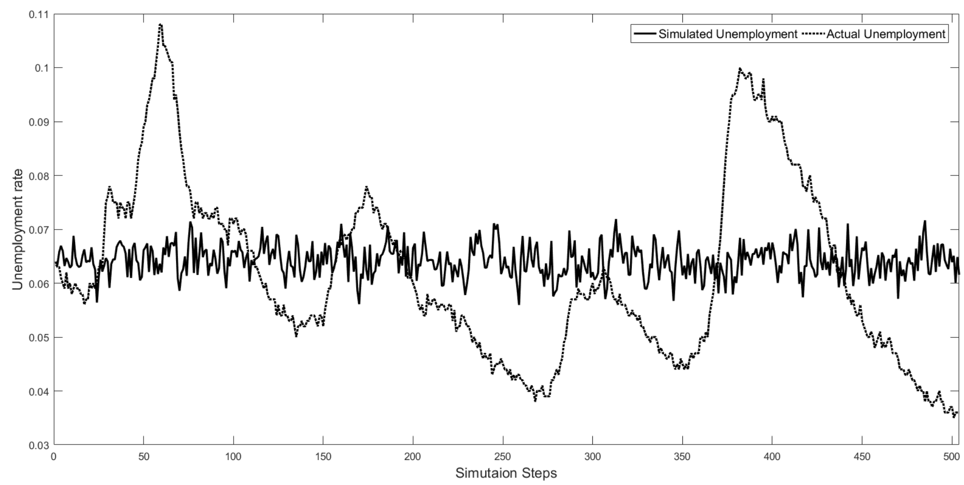

Figure 1: Actual (observed) and simulated unemployment rate.

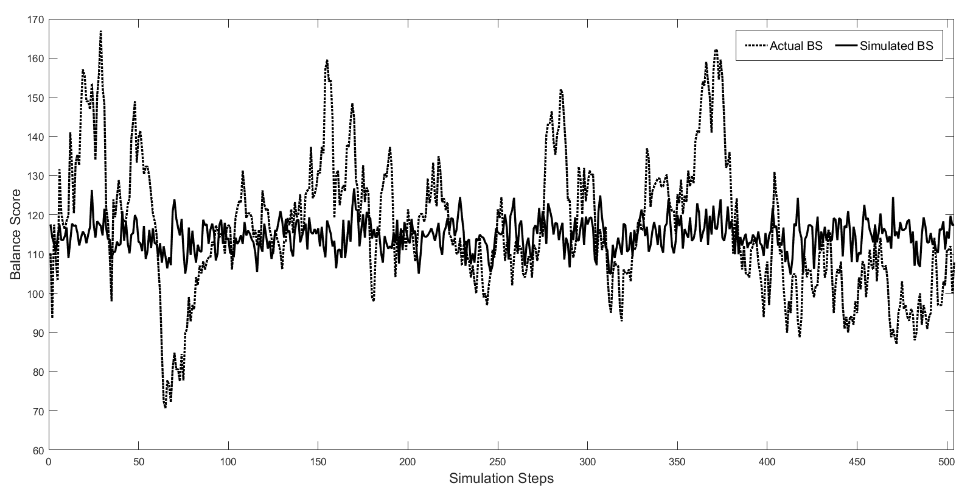

Figure 2: Actual (observed) and simulated balance-score (BS) unemployment expectations.

Table 2: Calibrated parameter values.

Parameter	**Value**
Intensity of choice (β)	5.44
Dispersion in unemployment expectations (δ)	0.81
Weight of the social component in deterministic utility (ψ)	0.17
Decay rate of unemployment's influence on pessimism bias (q)	0.26
Number of lags considered in the pessimism bias (L)	12

In addition to the parameters whose calibrated values are reported in Table 2, the value of the parameter γ in the effort function in (1)—which influences the unemployment rate in the short-run equilibrium in a given period t as shown in (27)—is set following Romer (2012, p.465), with $\gamma = 0.03$. This parameter captures the effort-enhancing effect of paying a worker a wage compensation that exceeds the wage associated with her expected labor market conditions.

Figures 1 and 2 present the observed and simulated time series for the unemployment rate and the BS, respectively. The simulated data were generated using the parameter values reported in Table 2 and assuming a uniform frequency distribution of the three types of unemployment expectations across workers in the initial period, as described above. Overall, the simulated series closely reproduce the qualitative patterns observed in the Michigan survey data, although the empirical series exhibits substantially greater volatility. It should be noted that the empirical series is affected by additional intervening factors that are abstracted in this agent-based modeling analysis.

4. Emergent properties of the agent-based modeling

In this section, we present a set of emergent properties generated by the agent-based modeling proposed in this paper. The simulation results are based on the calibrated parameter values reported in Table 2, combined with an initial uniform frequency distribution of workers across the three types of unemployment expectations. Each simulation was conducted over 566 periods, and all reported emergent properties are derived from simulations initialized with an identical random seed. The first 62 periods are treated as a transient phase and are therefore discarded; accordingly, only the dynamics observed over the subsequent 504 periods are presented in the figures that follow. Recall that the time series used in the calibration procedure spans January 1978 to December 2019, comprising a total of 504 monthly observations.

The first examined concerns the heterogeneity in the choice of unemployment expectations. Figure 3 presents the dynamics of the proportion of each type of unemployment expectation in the population of workers over the last 504 steps of the simulation. The small dotted curve represents the proportion of pessimistic workers. The large dotted curve represents the cumulative proportion of pessimistic and optimistic workers; hence, the difference between the large and the small dotted curves measures the proportion of optimistic workers. The proportion of neutral workers is given by the difference between the bold solid line parallel to the x-axis that intersects the y-axis at 1 and the large dotted curve.

Consistent with the empirical evidence, heterogeneity in the frequency distribution of unemployment expectations across workers is persistent, with no single expectation type ever being adopted by the entire population of them. Although this heterogeneity varies endogenously over time, the associated proportions—and thus the degree of heterogeneity—do not display any discernible long-run trend; rather, they exhibit dynamics consistent with mean reversion.

A similar pattern can be observed when examining the empirical dynamics of the actual proportions of each type of unemployment expectation in the Michigan survey. Figure 4 presents the corresponding time series. As in Figure 3, the small dotted curve represents the proportion of pessimistic workers, while the large dotted curve represents the cumulative proportion of pessimistic and optimistic workers; hence, the difference between the large and the small dotted curves measures the proportion of optimistic workers. The proportion of neutral workers is given by the difference between the bold solid line parallel to the x-axis that intersects the y-axis at 1 and the large dotted curve.

As in the simulated series, heterogeneity in the frequency distribution of unemployment expectations across workers is persistent, and no single type of expectation dominates at any point in time. However, the volatility of the proportions appears to be higher in the empirical series derived from

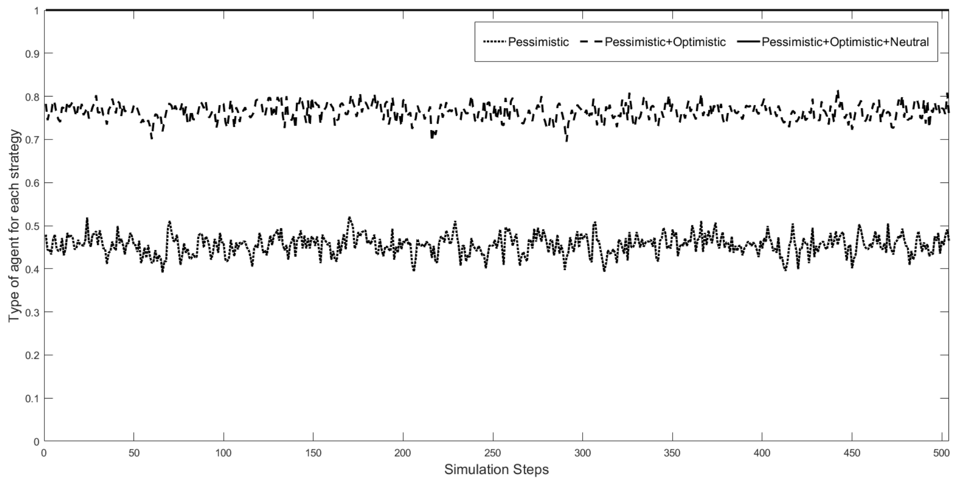

Figure 3: Simulated proportion of neutral, optimistic, and pessimistic workers.

the Michigan survey, which is influenced by several other intervening factors from which we abstract in this agent-based modeling analysis. As in Figure 3, the heterogeneity does not exhibit any apparent long-run trend—instead, it exhibits dynamics consistent with mean reversion.

Another relevant emergent property concerns the dynamics of the simulated unemployment rate and the Balance Score (BS) monthly time series described in the preceding section, shown in Figures 5 and 6. To ensure comparability across variables, the BS expression specified in (40) is normalized by dividing it by 200. Consequently, the normalized BS takes a value of 1 when all workers hold pessimistic expectations, 0 when all workers hold optimistic expectations, and 0.5 when the proportions of pessimistic and optimistic workers are equal. According to the specification in (40), the BS value, and hence its normalized counterpart, depends critically on the proportion of pessimistic workers relative to the proportion of optimistic workers. Furthermore, an increase in the proportion of pessimistic workers can lead to an increase, a decrease, or no change in the BS, depending on whether the proportion of optimistic workers remains constant (or increases less than proportionally, or decreases), increases more than proportionally, or increases proportionally, respectively.

A key analytical result derived in Silveira and Lima (2021) is that changes in the frequency distribution of unemployment expectations across workers

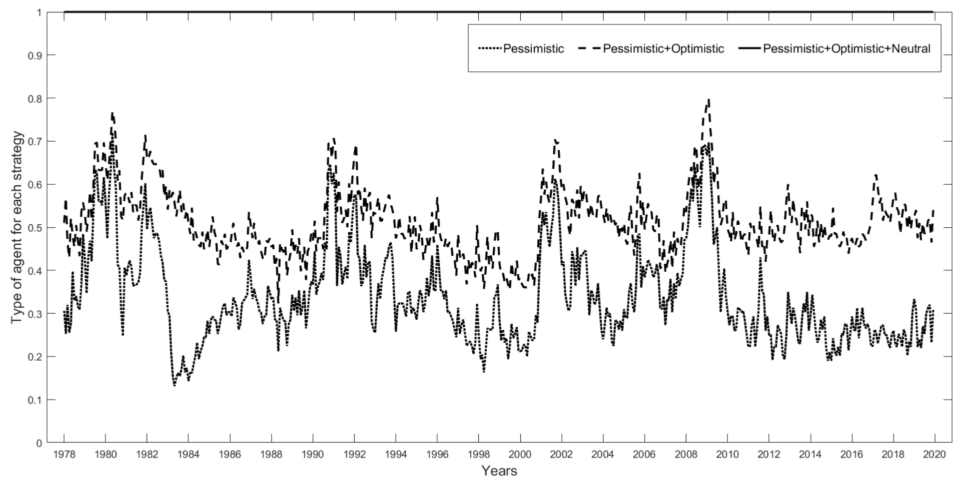

Figure 4: Actual (observed) proportion of neutral, optimistic, and pessimistic workers.

affect observed unemployment in a non-linear manner. Specifically, whether an increase in the proportion of workers holding pessimistic unemployment expectations—and thus facing a higher expected cost of job loss—leads to a lower or higher unemployment rate depends on the prevailing configuration of heterogeneity in unemployment expectations across workers. Indeed, a visual inspection of Figure 5, given the parameter values used in these simulations, reveals short-run episodes in which higher values of the BS—which, according to (40), correspond to elevated pessimistic expectations about near-future unemployment—are followed by periods of either higher or lower unemployment.

To further examine the joint dynamics of the normalized BS and the unemployment rate, Figure 6 presents these series over a 50-step window, from step 90 to step 140. The dynamics of trend reversals are irregular: some upward trends are reversed to downward trends within one or two periods, whereas other reversals unfold more gradually over multiple periods. As in Figure 5, there are short-run episodes in which higher values of the normalized BS are followed by periods of either higher or lower unemployment. Figure 7 presents the observed unemployment rate alongside the normalized BS. Compared with the simulated series, the observed unemployment exhibits longer cycles of growth and decline, rendering the correspondence between higher (lower) normalized BS and lower (higher) unemployment less

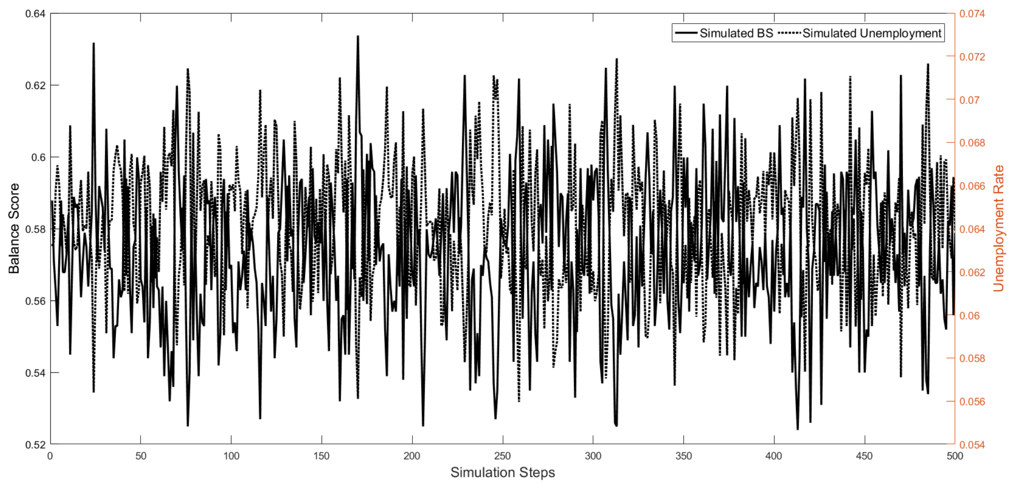

Figure 5: Normalized BS and unemployment rate simulated over the last 504 simulation steps.

pronounced. Nonetheless, there are periods in which this pattern is clearly observable. Again, short-run episodes exist in which higher values of the normalized BS are followed by periods of either higher or lower unemployment, underscoring the non-linear and variable nature of this relationship at a high frequency.

As reported in Section 2, substantial empirical evidence indicates a correlation between unemployment expectations and the observed unemployment rate. For example, Curtin (2019a) employs a Granger causality test to examine the temporal relationship between these variables. This test evaluates whether one variable can predict future changes in another beyond what is explained by the latter's own lagged values. Using data spanning 1978 to 2016, Curtin (2019a) find that the unemployment rate Granger-causes the BS and, conversely, that the BS Granger-causes the unemployment rate.

To determine whether the agent-based modeling developed here replicates this empirical relationship, we conducted a Granger causality test using the simulated series of the BS and the unemployment rate. Prior to this test, we verified that both series could be analyzed in levels by applying the Augmented Dickey-Fuller (ADF) unit root test, which rejected the null hypothesis of a unit root at the 5% significance level for both variables. We then

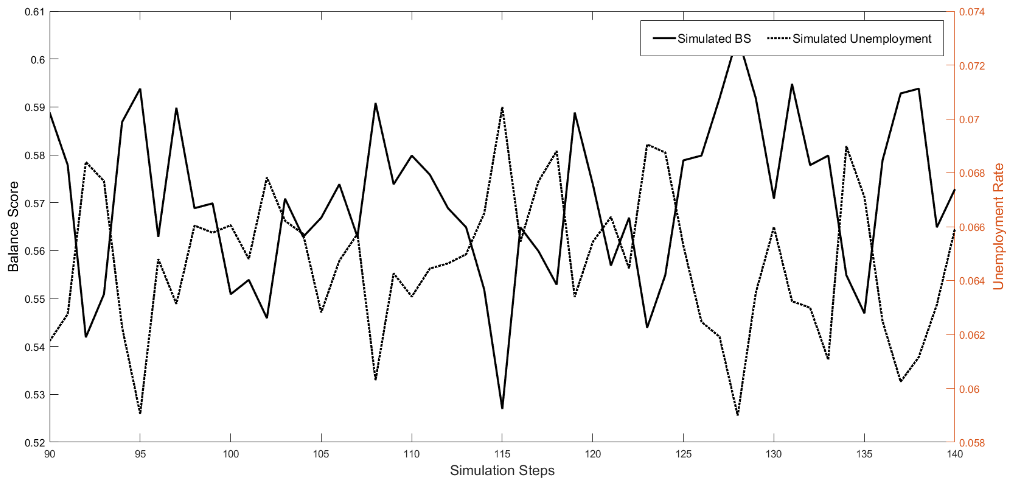

Figure 6: Simulated normalized balance-score (BS) unemployment expectations and unemployment rate between simulation steps 90 and 140.

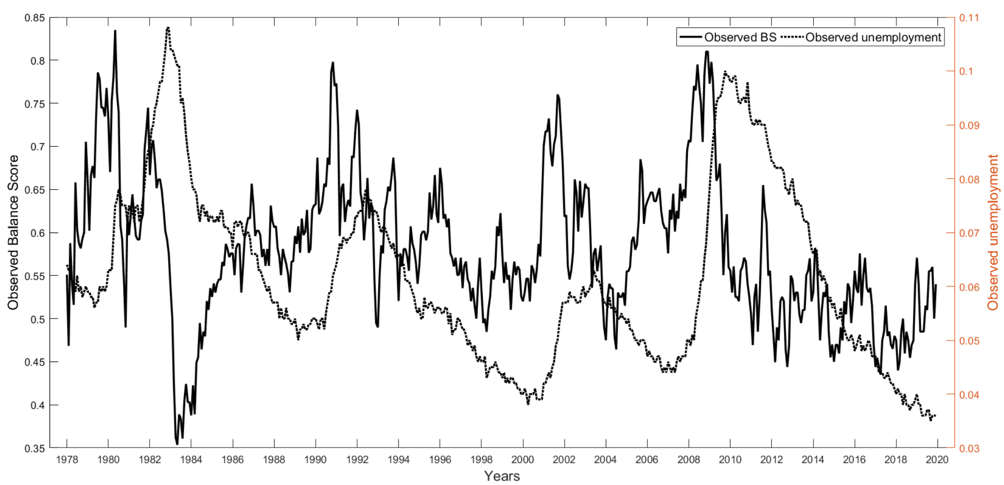

Figure 7: Observed normalized balance-score (BS) unemployment expectations and unemployment rate.

examined three information criteria to select the optimal number of lags for the model: the Akaike Information Criterion (AIC), the Hannan-Quinn information criterion (HQIC), and the Schwarz information criterion (SIC). All three criteria indicated that a single lag was sufficient. Table 3 reports the Granger causality statistics. Both the null hypothesis that the unemployment rate does not Granger-cause the BS and the null hypothesis that the BS does not Granger-cause the unemployment rate are rejected at the 5% significance level.

Similar results were obtained when performing the same test using the empirical data described in Subsection 3.3, which were used for model calibration. The test statistics are reported in Table Appendix A. Specifically, we find evidence that the observed unemployment rate Granger-causes the observed BS and, conversely, that the observed BS Granger-causes the observed unemployment rate over the period studied. This finding indicates that the simulated model is able to replicate the temporal precedence relationships observed in the empirical data. It should be noted, however, that the information criteria for the empirical data suggested models with either two or six lags.

Table 3: Granger causality test: simulated unemployment and normalized balance-score (BS) unemployment expectations.

Null Hypothesis	p-value
Unemployment rate does not Granger-cause BS	0.02
BS does not Granger-cause unemployment rate	0.00

One issue that warrants further investigation is whether the temporal causality persists for different values of the intensity of choice (β) and the social influence weight (ψ) in the probability-of-choice specification given in (32). To address this question, Figures 8 and 9 present heat maps of the p-values from the Granger causality tests applied to the BS and unemployment rate outputs of the model for various combinations of the parameters β and ψ. Specifically, Figure 8 illustrates the Granger causality test for the unemployment rate causing the BS, whereas Figure 9 illustrates the test for the BS causing the unemployment rate.

Both Figures 8 and 9 contain cells shaded in gray, representing parameter combinations for which the BS and/or unemployment rate remain constant. For these parameter values, causality testing is not required. Causality is

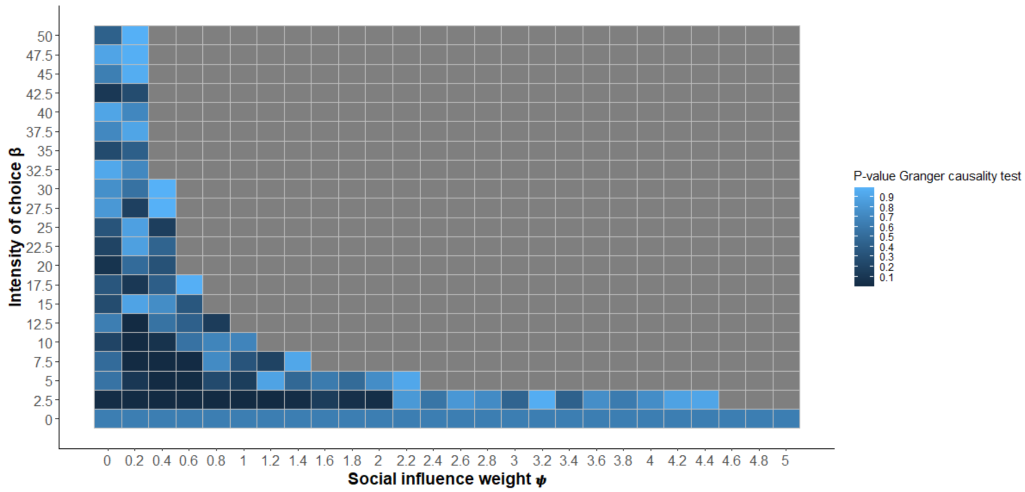

Figure 8: Granger causality test of normalized balance-score (BS) unemployment expectations on the unemployment rate across varying parameter combinations of social influence weight (ψ) and intensity of choice (β).

considered significant only if the p-value is less than 0.1. It is noteworthy that when the intensity of choice is zero, no significant temporal causality exists between the BS and the unemployment rate in either direction. Recall that the intensity-of-choice parameter measures the extent to which workers respond to differences in deterministic attractiveness, which includes both private and social components as specified in (25), when forming unemployment expectations. Therefore, as $\beta \to 0$, the discrete choice protocol becomes entirely random, and the choice probabilities converge to a uniform distribution (one-third in the present case of three unemployment expectation types), rendering the deterministic utility components irrelevant. Furthermore, for high values of ψ, causality is also absent. Specifically, the maximum value of ψ that generates causality from the BS to the unemployment rate is $\psi = 2$ when combined with $\beta = 2.5$. Finally, too high values of β also preclude causality, with the highest value of β that still produces significant causality from the BS to the unemployment rate being $\beta = 20$ when $\psi = 0$—recall that as the intensity of choice, β, approaches $+\infty$, the discrete choice protocol becomes fully deterministic, so that the alternative with the highest deterministic utility component is chosen with probability one.

An alternative approach to analyzing the temporal relationship between

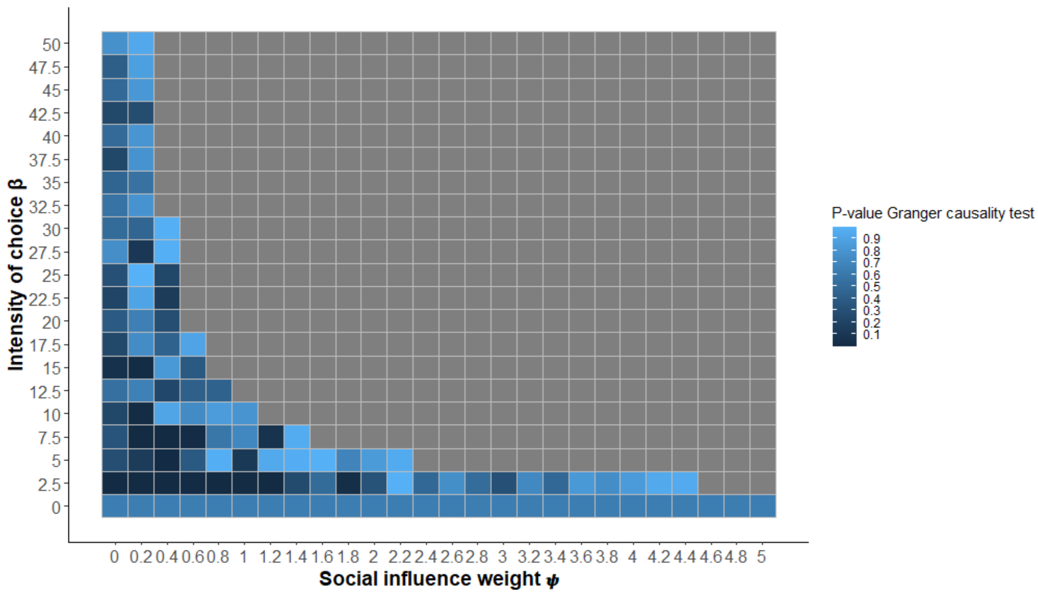

Figure 9: Granger causality test of the unemployment rate on normalized balance-score (BS) unemployment expectations across varying parameter combinations of social influence weight (ψ) and intensity of choice (β).

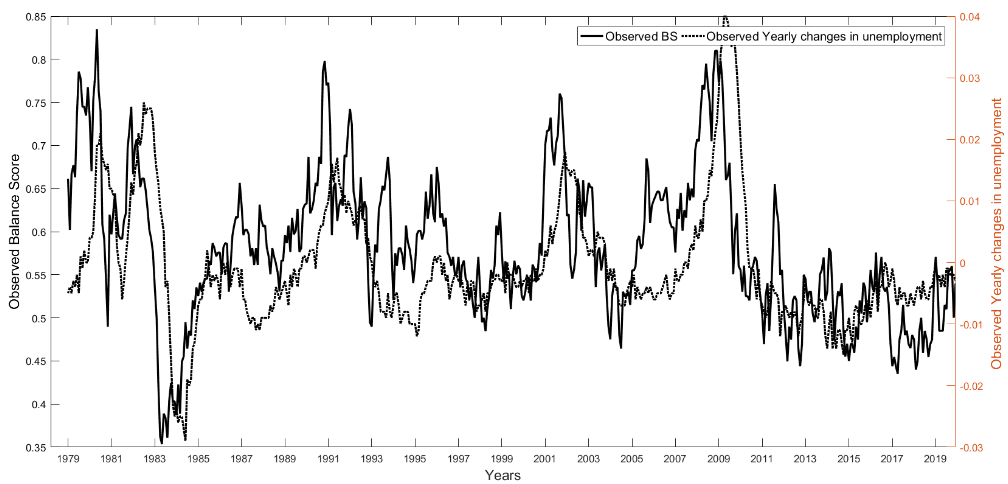

Figure 10: Observed normalized balance-score (BS) unemployment expectations and yearly change in the unemployment rate.

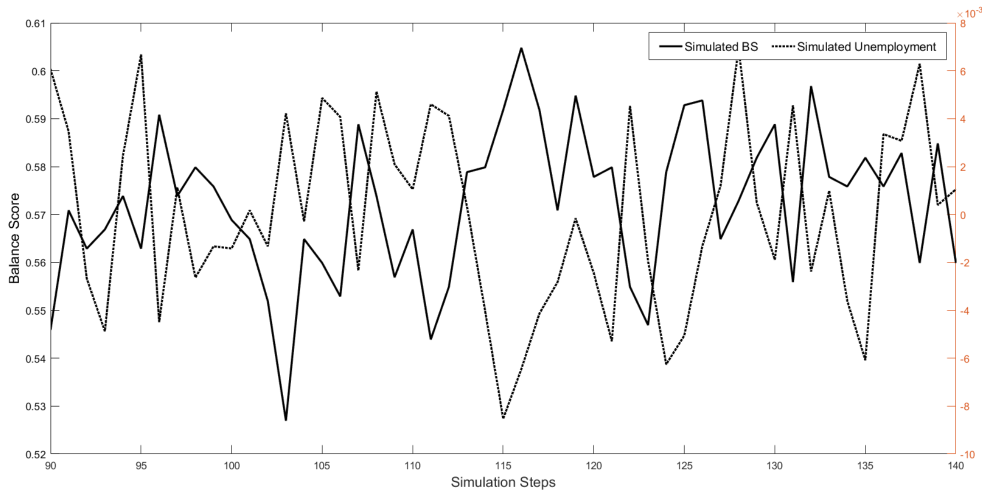

Figure 11: Simulated normalized balance-score (BS) unemployment expectations and yearly change in the unemployment rate between simulation steps 90 and 140.

unemployment expectations and the actual unemployment rate is proposed by Leduc and Sill (2013) and Girardi (2014). Leduc and Sill (2013) investigate the dynamic interaction between unemployment expectations and the unemployment rate within a vector autoregressive (VAR) framework using U.S. survey data. Their expectations indicator can be interpreted as a balance-score measure, summarizing the net difference between respondents anticipating an increase and those anticipating a decrease in unemployment. In contrast, Girardi (2014) employ European survey data to examine the relationship between unemployment expectations—also measured as balance scores—and year-over-year changes in the unemployment rate, defined as the difference between the unemployment rate in a given month and that observed in the same month of the preceding year. While the VAR framework in Leduc and Sill (2013) assesses causality via dynamic interactions within a multivariate system, Girardi (2014) evaluates whether expectations contain predictive information regarding subsequent changes in unemployment.

Building on the work of Leduc and Sill (2013) and Girardi (2014), we conducted a comparable analysis for the U.S. unemployment rate and expectations data from the Michigan Survey employed in this paper. Figure 10 presents the observed variables, while Figure 11 illustrates the coevolution

Table 4: Granger causality between normalized balance-score (BS) unemployment expectations and yearly changes in the unemployment rate.

Null Hypothesis	**p-value**
BS does not Granger-cause yearly change in simulated unemployment rate	0.00
BS does not Granger-cause yearly change in observed unemployment rate	0.00
Yearly change in observed unemployment rate does not Granger-cause BS	0.00
Yearly change in simulated unemployment rate does not Granger-cause BS	0.22

of the simulated time series between simulation steps 90 and 140. Visual inspection reveals that the same pattern is present in the U.S. data. To obtain more precise insights, we performed the Granger causality test. The results indicate that the normalized BS measure Granger-causes the yearly change in the unemployment rate in both the simulated and observed data, as reported in Table 4. The only discrepancy concerns the best VAR model selected according to the AIC information criterion, which corresponds to 13 lags for the observed data and 12 lags for the simulated data. Furthermore, while the yearly change in the unemployment rate also Granger-causes the normalized BS in the empirical data, there is insufficient statistical evidence to support this effect in the simulated data at the 1% significance level, where the p-value is 0.22.

5. Conclusion

This paper investigates the intricate dynamic interaction between heterogeneous unemployment expectations and observed unemployment dynamics. Motivated by considerable survey evidence documenting persistent heterogeneity in households' unemployment expectations and their coevolution with observed unemployment, we develop an agent-based analytical modeling framework of the labor market that incorporates a fitting heterogeneous expectations-augmented efficiency wage mechanism calibrated to observed data. Within this framework, heterogeneity in working households' unemployment expectations readily translates into corresponding heterogeneity in their perceived cost of job loss and, consequently, in their effort provision on the job.

Working households' near-future unemployment expectations evolve endogenously through a discrete choice protocol, allowing the frequency distribution of these expectations to coevolve with observed unemployment. To maintain the requisite analytical focus on the core research questions, however, the modeling developed in this paper necessarily abstracts from several related issues.

Calibrated with data from the U.S. Michigan Survey of Consumers and related labor market data, the simulated agent-based modeling captures key empirical patterns, including persistent heterogeneity in unemployment expectations and self-sustained cyclical fluctuations in both unemployment expectations and the observed unemployment rate. Short-run episodes reveal that higher optimistic expectations can be followed by immediate periods of either higher or lower observed unemployment, reflecting nuanced high-frequency dynamics. Granger causality tests further confirm a bidirectional temporal relationship between a balance-score measure of the state of unemployment expectations and observed unemployment. This temporal relationship is crucially shaped by the intensity-of-choice and social-influence-weight parameters, which play a central role in the discrete choice protocol.

Overall, these findings indicate that properly accounting for the complex co-evolution between the frequency distribution of unemployment expectations across working households and observed unemployment is relevant for better understanding the dynamics of observed unemployment and, more broadly, labor market fluctuations.

Appendix A. Granger causality test statistics for the relationship between the unemployment rate and normalized balance-score (BS) unemployment expectations using U.S. data

Table A.5: Results of the Augmented Dickey-Fuller (ADF) unit root test for U.S. empirical data

Variable	p-value (ADF test)
BS	0.01
Unemployment rate	0.07

Table A.6: Model selection criteria for empirical U.S. data, including the Akaike Information Criterion (AIC), Hannan-Quinn Information Criterion (HQIC), and Schwarz Information Criterion (SIC)

Criterion	AIC	HQIC	SIC
Number of lags	6	2	2

Table A.7: Results of the Granger causality test with two lags

Null hypothesis	p-value
Unemployment rate does not Granger-cause BS	0.00
BS does not Granger-cause unemployment rate	0.00

Appendix B. Granger causality test statistics for the relationship between balance-score (BS) unemployment expectations and yearly changes in the unemployment rate

Table A.8: Results of the Granger causality test with six lags

Null hypothesis	p-value
Unemployment rate does not Granger-cause BS	0.00
BS does not Granger-cause unemployment rate	0.00

Table B.9: ADF test results for yearly data

Variable	P-value (ADF test)
BS	0.01
Unemployment rate	0.07

Table B.10: Information criteria for yearly data

Criterion	AIC	HQ	SC
Number of lags	6	2	2

Table B.11: Granger causality test results with 2 lags (yearly data)

Null hypothesis	p-value
Unemployment rate does not Granger-cause BS	0.00
BS does not Granger-cause unemployment rate	0.00

Table B.12: Granger causality test results with 6 lags (yearly data)

Null hypothesis	p-value
Unemployment rate does not Granger-cause BS	0.00
BS does not Granger-cause unemployment rate	0.00

References

Blanchflower, D., Kelly, R., 2008. Macroeconomic literacy, numeracy and the implications for monetary policy. Technical Report. Working Paper Bank of England.

Brock, W., Hommes, C., 1997. A rational route to randomness. Econometrica 65, 1059–1095.

Brock, W.A., Durlauf, S.N., 2001. Discrete choice with social interactions. Review of Economic Studies , 235–260.

Campbell, D., Carruth, A., Dickerson, A., Green, F., 2007. Job insecurity and wages. The Economic Journal 117, 544–566.

Cappelli, P., Chauvin, K., 1991. An interplant test of the efficiency wage hypothesis. Quarterly Journal of Economics 106, 769–787.

Carroll, C., 2003. Macroeconomic expectations of households and professional forecasters. Quarterly Journal of Economics 118, 269–298.

Curtin, R., 2003. Unemployment expectations: The impact of private information on income uncertainty. Review of Income and Wealth 49, 539–554.

Curtin, R., 2019a. Consumer expectations: Micro foundations and macro impact. Cambridge University Press.

Curtin, R., 2019b. Consumer expectations: A new paradigm. Business Economics 54, 199–210.

Dawid, H., Delli Gatti, D., 2018. Agent-based macroeconomics. Handbook of Computational Economics 4, 63–156.

Dickens, W.T., Katz, L.F., 1987. Inter-industry wage differences and industry characteristics. NBER Working Paper .

Dickerson, A., Green, F., 2012. Fears and realisations of employment insecurity. Labour Economics 19, 198–210.

Fullone, F., Gamba, M., Giovannini, E., Malgarini, M., 2007. What do citizens know about statistics? the results of an oecd/isae survey on italian consumers. Statistics, Knowledge and Policy , 197–212.

Girardi, A., 2014. Expectations and macroeconomic fluctuations in the euro area. Economics Letters 125, 315–318.

Goldsmith, A.H., Veum, J.R., Darity, W.J., 2000. Working hard for the money? Efficiency wages and worker effort. Journal of Economic Psychology 21, 351–385.

Hendren, N., 2017. Knowledge of future job loss and implications for unemployment insurance. American Economic Review 107, 1778–1823.

Kassenboehmer, S.C., Schatz, S.G., 2017. Re-employment expectations and realisations: Prediction errors and behavioural responses. Labour Economics 44, 161–176.

Krueger, A.B., Summers, L.H., 1988. Efficiency wages and the inter-industry wage structure. Econometrica 56, 259–293.

Kuchler, T., Zafar, B., 2019. Personal experiences and expectations about aggregate outcomes. The Journal of Finance 74, 2491–2542.

Leduc, S., Sill, K., 2013. Expectations and economic fluctuations: An analysis using survey data. Review of Economics and Statistics 95, 1352–1367.

Lehmann, R., Weyh, A., 2016. Forecasting employment in Europe: Are survey results helpful? Journal of Business Cycle Research 12, 81–117.

Malgarini, M., Margani, P., 2008. Unemployment expectations across heterogeneous groups. Economics Bulletin 4, 1–6.

Orland, A., 2017. Personality traits and the perception of macroeconomic indicators: Survey evidence. Bulletin of Economic Research 69, E150–E172.

Romer, D., 2012. Advanced Macroeconomics, 4e. New York: McGraw-Hill.

Silveira, J.J., Lima, G.T., 2021. Can workers' increased pessimism about the labor market conditions raise unemployment? International Review of Economics & Finance 72, 125–134.

Steinbacher, M., Raddant, M., Karimi, F., Camacho Cuena, E., Alfarano, S., Iori, G., Lux, T., 2021. Advances in the agent-based modeling of economic and social behavior. SN Business & Economics 1, 1–24.

Tortorice, D.L., 2012. Unemployment expectations and the business cycle. The B.E. Journal of Macroeconomics 12.

Train, K.E., 2009. Discrete choice methods with simulation. Cambridge University Press.

Wallskog, M., 2026. Within-firm pay inequality and productivity: Evidence from U.S. census data. The Economic Journal (Forthcoming).