Exploring the Social Architecture Model

Alan G. Isaac
American University*

April 9, 2017

Abstract

Microfoundations proposed for macroeconomic models often involve strong counterfactual assumptions about the knowledge and foresight of agents and about the pervasiveness of equilibrium exchange. We explore an agent-based model of a monetary exchange economy that discards such assumptions. Heavily influenced by econophysics, the social architecture model proposes implicit microfoundations, grounded in Keynes's principle of indifference. The model emphasizes stochastic processes, disequilibrium exchange, and unpredictable individual behavioral. Econophysicists have demonstrated that models with implicit microfoundations can reproduce important stylized facts of real economic systems. This paper re-examines the GDP and unemployment distributions produced by the social architecture model. We offer general support for previous findings, subject to a modest model reinterpretation and minor model modifications.

Keywords: agent-based modeling, econophysics, microfoundations, macroeconomics
JEL: E03, C6, B5, B4

*Correspondence: Alan G. Isaac, Department of Economics, American University, Washington, DC 20016

1 Background: The Social Architecture Model

The social architecture model is an agent-based macromodel that is heavily influenced by the methods of econophysics.[1] These methods emphasize the role of constraints over the role of choice, in alignment with the zero-intelligence-trader models of Becker (1962) and Gode and Sunder (1993). Models influenced by econophysics strive to minimize the dependence of model outcomes on the peculiarities of specific theories of economic behavior, favoring instead mild distributional assumptions on behavior subject to uncontroversial constraints (such as budget constraints).

Wright (2009) presents the social architecture model as a marriage of the methods of econophysics to the concerns of macroeconomics. It is an agent-based macromodel of monetary exchange and labor market allocation. Drawing on the research on zero-intelligence agents, the social architecture model generates macroeconomic outcomes from the interactions of many weakly coordinated and unpredictable individuals. These agents are constrained by the institutional structure in which they interact, including labor markets and the institution of monetary exchange. However, the model does not impose an equilibrium-trading constraint on these agents, and the agents do not embody the unbounded cognition typically assumed in mainstream macromodels.[2]

In this paper, we explore several aspects of the social architecture model. We uncover tensions between the methodology of econophysics and the concerns of macroeconomics, which leads us to questions about the methodological justification of some of the model's key assumptions and simplifications—particularly justifications rooted in what we will call *behavioral agnosticism* (as explained below). We also explore some difficulties with the existing

[1]Econophysics is the economic study of the statistical equilibria that emerge from the interactions of large numbers of heterogeneous agents (Yakovenko and Rosser, 2009).

[2]This paper uses the modifier *mainstream* very broadly to denote macromodels that attempt to characterize all behaviorally relevant preferences, constraints, and information of cognitively superhuman optimizing agents who face equilibrium-trading constraints. We believe that this remains a convenient shorthand terminology, even though the scope of the mainstream is continually evolving. (As an example of this evolution, macroeconomists have recently shown an increasing willingness to incorporate results from behavioral economics (Driscoll and Holden, 2014).)

formulation of the social-architecture model, and we suggest ways to begin addressing these difficulties. In particular, we discuss the unemployment problem in the social architecture model, and we propose a partial solution.

1.1 Implicit Microfoundations and Social Architecture

In agent-based macromodels, the evolution of the macroeconomy emerges from the interactions of many heterogeneous agents. Agent heterogeneity is a fundamental feature of these models; there is no "representative" agent. For example, in any period some agents may be unemployed, some may be employed, and some may be employers. Individual-level (microeconomic) interactions and the global (macroeconomic) state of the economy influence the evolution of the employment status of of each agent. Since individual-level interactions affect individual outcomes, the history of each agent is fundamentally idiosyncratic. Nevertheless, we may detect persistent structure in the resulting macroeconomy, often in the form of stylized facts about statistical equilibria.

Econophysicists have repeatedly demonstrated that unpredictable micro-level interactions can produce predictable macro-level regularities. In the 1990s, Gene Stanley and other physicists applied the tools of statistical mechanics to economics, achieving clear recognition for their efforts by the mainstream economics profession just after the turn of the century (Gabaix et al., 2006). Econophysics has two core methodological emphases: an empirical search for stylized economic facts (in the form of stable statistical distributions), and related theoretical attempts to produce similar market-wide or economy-wide outcomes from the interactions of many heterogeneous and individually unpredictable agents. Models in the econophysics tradition eschew the standard theoretical toolkit of mainstream economists. In particular, they discard the assumption of essentially homogeneous rational agents in continuous equilibrium.[3]

Around the same time, economists displayed increasing interest in agent-based econo-

[3]More precisely, the mainstream reliance on imposed mechanical equilibrium is discarded in favor of emergent statistical equilibrium (Yakovenko and Rosser, 2009).

mics and behavioral economics, which also discard such assumptions. Agent-based social science blossomed with the publication of the monographs of Epstein and Axtell (1996) and Axelrod (1984, 1997). As in the econophysics literature, the agent-based literature evinced relatively little interest in characterizing the psychological state of individual agents. In striking contrast, behavioral economists—heavily influenced by the work of Kahneman and Tversky—were intensely interested in achieving a more realistic understanding of individual agents.

Despite these deep differences, there are also important similarities. Agent-based economics and behavioral economics both reject the thin rationality and equilibrium-trading constraints imposed on the agents in typical mainstream macroeconomic models. By *thinly rational* agents, we mean agents that are computationally unlimited but act exclusively in response to the constraints and payoffs that can be explicitly characterized by modelers. In saying that agents face *equilibrium-trading constraints*, we mean that they are unable to act in response to disequilibrium signals in the economy. (For example, all trades must take place at market-clearing prices.)

The econophysics and agent-based research threads intertwined to produce simulation studies that were heavily influenced by econophysics but which included some explicit behaviorally informed and institutionally constrained agent-based features. Literature in this vein includes the research on zero-intelligence agents, which traces back to Becker (1962) but became truly influential with the work of Gode and Sunder (1993).[4] In such models, aggregate outcomes derive from budget constraints and institutional features more than from individual perceptions and goals. Indeed, researchers influenced by the econophysics literature often model agents as choosing randomly among their feasible behaviors. Underpinning this approach is a suggestion that fundamental heterogeneity in the environments and constraints faced by individuals imply that systemic outcomes diverge from those implied by a representative agent (Aoki, 2002). Researchers found that interactions between hetero-

[4]For further examples and discussion, see Farmer et al. (2005), Wright (2008), or Cottrell et al. (2009).

geneous, weakly coordinated, unpredictable individuals can produce stochastic equilibria that display macro-level regularities in the form of distributions over possible macro-level outcomes. Economists calling for more attention to the characteristics of such stochastic equilibria include Steindl (1965), Foley (1994), and Aoki (1996, 2002).

Attempts to apply methods of econophysics specifically to the concerns of macroeconomics is often traced to Drăgulescu and Yakovenko (2000). Researchers propose that a statistical treatment of monetary exchange can improve our understanding of actual macroeconomic distributions. In the present paper, we explore a more recent, well-known example from this literature: the social architecture model of Wright (2009). Wright argues that zero-intelligence agents can provide useful *implicit* microfoundations for macroeconomic modeling. In stark contrast to the methods of mainstream macroeconomic theory, implicit microfoundations impose neither rationality nor equilibrium. Wright (2009) finds that the social architecture model produces distributions over a variety of macroeconomic variables that bear tantalizing resemblances to real economies. (We explore some examples below.)

Instead of characterizing individual behavior in terms of thinly rational optimization, models in this tradition typically turn to simple behavioral rules and allow for substantial behavioral randomness. This modeling strategy is meant to acknowledge the modeler's ignorance of the circumstantial details and underlying motivations of individual agents. It is *not* meant to constitute a claim that real individuals ignore goals or fail to formulate them. Rather this modeling strategy reflects a belief that the modeler's knowledge of the individuals in a macroeconomy is too vague and incomplete to support substantive speculation as to their goals and particular circumstances.[5] In contrast, mainstream macroeconomic models assume cognitively superhuman optimizing agents and attempt to characterize all their relevant preferences and constraints.

The turn toward implicit microfoundations is an attempt by researchers to acknowledge their ignorance of the circumstances, goals, and capacities of individual agents. The met-

[5] While it is far simpler to model a role of a single die in carefully controlled circumstances, even in this instance a simple stochastic model is usually the first choice.

hodological motivation is that we know too little about agents to accurately model their incentives and opportunities (and thereby their behavior) at the individual level. In this sense, reliance on implicit microfoundations is intended to support a behavioral agnosticism, eschewing the "pretence of knowledge" embodied in characterizing individuals as essentially homogeneous and readily comprehensible. Additionally, models drawing on the econophysics tradition do not impose market equilibrium. As in the real world, exchange takes place in real-time by a process of discovery that involves out-of-equilibrium trades. In contrast, exchange in mainstream macroeconomic models is equilibrium-constrained, effectively treating discovery as transpiring in imaginary time.

1.2 Model Overview

For the purposes of the present paper, the social architecture model has certain crucial features: replication is ensured by the availability of the source code, and its choice of methods allows us to explore the compatibility between the concerns of macroeconomists and the core methods of econophysics. Since Wright (2009) discusses model details and provides an implementation in code, we will only briefly review the details of the model.

The social architecture model is an agent-based model (Wilensky and Rand, 2015). There is a single type of agent, although different agents may be in fundamentally different states. An agent is essentially characterized by its attribute values and a few behaviors that are defined in the context of an economy. We assign each agent an index on the closed integer interval $[1..N]$, where N is the number of agents. An agent's state comprises the values of five core attributes: money (the transactions medium and unique store of value), an employer index in $[0..N]$ (where 0 indicates the agent has no employer), the value of the last wage received, a reservation wage (which determines which job offers will be accepted), and a set of employees. The set of employees is empty if the agent is not an employer; an agent is unemployed if its employer index is 0 and its employee set is empty.

At each point in time, an agent can be in one of three different employment states:

unemployed, employed, or employing. We will refer to an agent as a worker, a jobless agent, or an employer based on this employment state. Employment matches are determined by random search, which connects the social architecture model to the literature on search unemployment (Mortensen, 1970). Agents search for for positions by sampling employers. Employed agents may search in an attempt to improve their position; a worker need not leave a job in order to search for a new one. Jobless agents search to find a position. In any period, a jobless agent may fail to find a match, thereby remaining jobless for another period. The reservation wage falls during jobless spells, in accord with the empirical evidence (Addison et al., 2013; Brown and Taylor, 2013).

An interesting feature of the model is that the class structure of the economy is endogenous. This connects it to the work of Roemer (1982) and its ramifications such as Eswaran and Kotwal (1986). However, class in the social architecture model evolves dynamically: an agent may transition from being unemployed to being a worker or even to being an employer.

The state of an economy comprises its microstate (i.e., the state of each agent) and its macrostate. In the social architecture model, there is a single macroeconomic attribute: the level of *latent demand* in the economy. (See section 1.2.2.) All other macro characteristics are determined by aggregation over the microstate. The time scale of simulation is one month per iteration. Each month we produce a new state of the economy by applying a simple evolution rule to a list of agents. Each of these agents sequentially applies the following subrules (if applicable): job search, consumer spending (which adds to latent demand), firm revenue generation (which depends on latent demand), and firm management (comprising wage payment and separation decisions). Processing an agent naturally changes the agent, but it also affects other agents and the macrostate. The outcome is a new state of the economy (including a new state for that agent and a new level of latent demand).

Table 1: Reservation Wage Adjustment

	employed	unemployed
$w_o \leq w_d$	no change	w_d falls
$w_o > w_d$	switch firms; w_d rises	join firm; w_d rises

1.2.1 Job Search

Job-search is a behavior of workers and the unemployed, who search for a potential match and accept a position whenever the wage offered exceeds the reservation wage.(Employers do not seek to become employees.) An agent may search for a better (i.e., higher paying) match at any firm, even the current employer. Whenever a jobless agent's search proves fruitless, that agent's reservation wage falls. If a worker finds a position that offers a better match (i.e., a higher wage), the agent's reservation wage rises to the new wage level.

Job search behavior produces an updated economy. This new economy can include adjustments to wage expectations: an agent's reservation wage rises with an accepted offer or falls during unemployment spells. This means that a worker's reservation wage will ratchet upwards until the worker faces unemployment, at which point it starts to decline.[6] No workers are fired at this stage—that happens during the firm-management stage—but an unemployed worker may remain unemployed. We may summarize job-search behavior as follows:

- if the agent is an employer, do nothing; otherwise:
- randomly pick a potential employer; if there are no potential employers, do nothing; otherwise:
- negotiate a wage offer with the potential employer
- compare the wage offer (w_o) with the reservation wage (w_d) and accept the wage offer if it is adequate
- adjust the reservation wage (according to Table 1)

[6] An implication is that wages are sticky downwards, but as it turns out, not very much so.

Even though the social architecture model emphasizes random behavior subject to constraints, job search includes non-random behaviors. Most obviously, employed and unemployed agents engage in job search, and they only accept offers that exceed their reservation wages. This means that the model incorporates wage-ladder effects, where an agent who remains employed experiences upwards wage ratcheting. The role of the reservation wage is a virtue of the model, suggesting that interesting macromodels cannot reasonably adhere to strict behavioral agnosticism. Even in the social architecture model, the microfoundations are not entirely implicit.

Each recourse to an economic specification of behavior (beyond raw randomness subject to constraints) may be considered to be a methodological deviation from pure econophysics. Our exploration of job search behavior suggests that such deviations will often be desirable. As a practical matter, modelers will find that an emphasis on constraints and institutions does not entail complete liberation from the need to think economically about behavior. Even for researchers inclined towards behavioral agnosticism, modeling decisions must be judged pragmatically: which aspects of behavior can we usefully neglect, and which must we attend to in some detail? Once we raise the issue of usefulness, macroeconomists drawing on econophysics must debate their methods on the same pragmatic terrain as other macroeconomists. While it is a commonplace that macroeconomic models must be unrealistic, we must nevertheless ask whether any particular lapse of realism promotes our research goals or proves fatal to them. ("All models are wrong but some are useful.")

The job-search behavior involves subroutines for employer selection, wage negotiation, reservation-wage adjustment, and firm affiliation.[7] The first three of these subroutines specify key behavioral assumptions of the model. In particular, we are confronted with the following questions. How do agents search for positions? What does wage negotiation look like, and what influences an agent's reservation wage? Although our answers must in some sense be transparent when implemented in code, the social architecture model intentionally leaves the

[7]Additional subroutines perform bookkeeping duties: we need to update the state of the economy whenever a worker joins or leaves a firm.

answers to such questions relatively opaque at the level of individual goals and constraints.

When plausible, the social architecture model relies heavily on randomness in the description of behavior. In the pursuit of maximal behavioral agnosticism, the model repeatedly turns to the standard uniform distribution—often characterized as a reliance on Bernoulli's principle of insufficient reason. Even so, and even within its uncontroversial constraints (e.g., budget constraints), behavior in the model is not purely random. Consider a few examples of behavioral assumptions of the model that raise theoretical questions. A firm with more working capital is more likely to be chosen as a potential employer. Wage negotiations lead to offers that are (uniformly) between one and two times the job seeker's reservation wage. The reservation wage rises in response to a good offer but falls in response to unemployment. Pointing out that such assumptions are behavioral in no way implies that they are poor choices, but in each case there is a drift away from purely statistical reasoning towards economic reasoning. To put it another way, these assumptions reflect a compromise between the goals of parsimony, plausibility, and performance. But such compromises are coextensive with the compromises of any economic theorizing, not just that inspired by econophysics.

1.2.2 Spending Behavior and Firm Revenues

One virtue of agent-based modeling is that maintaining stock-flow consistency becomes almost trivial, since model transactions are generally explicit. In the social-architecture model, all transactions involve the unique medium of exchange. All expenditures require money, and all firm revenues are the receipt of money. Macroeconomic models in the econophysics tradition often adopt such a cash-in-advance constraint on transactions, and this is also common in the mainstream macroeconomics literature (Clower, 1967; Lucas and Stokey, 1987; Drăgulescu and Yakovenko, 2000). In an interesting deviation from the standard theoretical treatments, however, the social-architecture model allows some temporal slippage between expenditures and receipts. These are mediated by a macroeconomic state variable, which we call *latent demand*. As a result, expenditure decisions may not immediately produce receipts for firms: expenditure contributes directly to latent demand, but firm revenues accrue from effective demand.[8]

The attraction of this treatment of latent demand lies not in any particular story but rather in in its support for motivational agnosticism. Specifically, models hewing to the methods of econophysics tend to avoid speculations about how a consumer will chose a particular firm for particular consumption expenditures. This allows us to discard the counterfactual assumption that a consumer faced with unchanging circumstances (including relative prices) will continually purchase an unchanging bundle of goods and services. The social architecture model goes further: it does not characterize each consumer's diachronic allocation of expenditure among available options. Contrast this with complex shopping models and evolving consumer-firm networks, as found in Edoardo Gaffeo and Gallegati (2008), Neveu (2013), or Ashraf et al. (2017). Realistically, a consumer's discretionary expenditures may be big or small from month to month, and may be allocated to a few firms one month and many firms the next. As we should expect from our previous methodological discussion, the social architecture model shies away from specifying the details of how each consumer allocates expenditures among firms. The resulting modeling strategy proves simple and attractive. At the level of the individual consumer, only total spending is explicitly characterized, not its detailed allocation among goods and services. We may summarize consumer expenditure behavior as follows.

- set expenditure to a random fraction of the consumer's wealth
- decrement the consumer's wealth by the amount of the expenditure
- increment latent demand by the amount of the expenditure

Next, a firm-revenue rule determines how firms accrue revenues based on latent demand. Some portion of latent demand is realized as effective demand. Effective demand is received

[8] This slippage is perhaps easiest to understand from the consumer side. For example, a consumer might allocate cash for a purchase but then take a while to complete the purchase. An explicit escrow account may even be involved. On the firm side, the employees of a firm may accumulate consumer payments but not recognize these as increments to working capital until the payments are centrally collected and counted. In the absence of explicit institutional detail, we are free to speculate about the background story.

by firms (i.e., by employers). Workers are agents of firms; they generate the firm's revenue (implicitly, via sales effort). Except for the unemployed, each agent has a shot at garnering a portion of latent demand for a firm. The resulting allocation of latent demand favors bigger firms (i.e., firms with more employees), which implies an employment-weighted allocation of revenue. Since large firms have proportionally greater revenue generating capacity, the social architecture model imposes no ex ante size limitation on firms. The firm-revenue rule can be summarized as follows.

- if the agent is unemployed, do nothing (no revenue is generated); otherwise:
- set effective demand to a random fraction of latent demand
- augment the wealth of the firm owner by effective demand
- decrement latent demand by the allocated effective demand

1.2.3 Employer Behavior

The final agent-processing stage characterizes firm-management behavior: wage payment, or employee firing. Employers pay employees in the order hired; employees are let go if there is not enough money to pay them. (They are paid nothing at all in this case.) If all employees depart, the firm dies, and the employer enters the ranks of the unemployed.

This firm-management behavior brings us again to the tension between behavioral agnosticism and economic theorizing. How agnostic is it to pay employees in the order hired? To reiterate, such a question need not imply a criticism of the assumption, which seems reasonably descriptive (Foulkes, 1980). Rather, such questions highlight the tension between implicit theorizing and the desire for behavioral agnosticism. Would it be more behaviorally agnostic to pay employees in random order? Would it be less plausible to first fire the most expensive of these equally qualified (in the expected revenues sense) individuals? Or should workers be let go without regard to their previous tenure, as in Neveu (2013)? We find that the goal of behavioral agnosticism is simply too vague to provide substantial guidance at crucial points in model formation. It is at such points that the methodological guidance of mainstream microfoundations—which discard relatively vague notions of behavioral plausibility and conformance to stylized facts in favor of the implausibly optimal pursuit of counterfactually stable goals subject to an unrealistically precise characterization of current conditions—may seem a relief. Setting aside these concerns for the moment, we summarize firm management behavior in the social architecture model as follows.

- if the agent is not an employer, do nothing; otherwise:
- iterate through employee list (in order hired) and pay wages owed as long as working capital is not exhausted[9]
- employees that the firm cannot pay become unemployed

1.2.4 Updating Agents and Economies

These four agent-processing stages—job-search behavior, spending behavior, revenue generation, and firm-management behavior—determine all changes in the state of the economy. Call this sequence of actions the *agent schedule*. Agents are processed sequentially, and an agent is always processed by the same schedule. From a functional perspective, the agent schedule is the composition of the four agent-processing stages.

- job search
- consumer spending
- firm revenue determination
- firm management

Recall that the time scale of our simulation is one month per iteration. We therefore call the evolution rule for the economy the *one-month rule*. We implement one-month rule as follows. The agent-schedule is effectively a binary function: taking as arguments an economy and an agent, it returns an updated economy. We can therefore conceptualize the evolution rule for the economy as a binary fold.[10] That is, given an economy and a list of agents, we

[9] Actual wage payment is passed off to a subroutine. This subroutine primarily performs accounting, but in doing so it alters one agent attribute: the employee's last wage payment received.

[10] A binary fold, also called a reduction, involves a repeated application of a binary operation. For example,

can fold the processing rule over the list of agents to produce the new end-of-month economy. Each month, we update the economy by applying the agent schedule to N agents.[11]

Since agents are interdependent, they must be processed sequentially, each in the context of an updated economy. For example, the consumption expenditure of an employee depends on the wage payment of the employer. Similarly, processing an employee will generally change the state of the employer.

The one-month rule is our core state-transition rule for the economy, from which we can build up trajectories of arbitrary length. For example, given an initial state of the economy, we can produce a trajectory for one year by means of 12 repeated applications of the one-month rule. Correspondingly, one simulation run constructs a trajectory of economies from a given initial economy. We specify the length of the simulation in years; a year is of course twelve months, where one month is the time scale of a single iteration. When we want an annual description of a trajectories, we use the last month of each year for stock variables, and changes from end-of-year to end-of-year for flow variables.

2 Macroeconomic Considerations

An economy comprises a microstate and a macrostate. The microstate is just the state of all the agents in the economy. In the social architecture model, macroeconomic data are for the most part produced by aggregating over the microstate. However, the model also includes a purely macroeconomic variable that cannot be produced by aggregation of the microstate: latent demand.

A complete description of an economy therefore comprises two elements: a description of all the agents (its microstate), and the level of latent demand (its macrostate). We can correspondingly construct an initial economy by constructing a list of agents and assigning a value to latent demand.

addition can be folded over a list of numbers to produce their sum.

[11] Each month, agents are in a new random order. Wright (2009) samples with replacement, justifying this as producing better calibration of the model to his data.

Our next task is to explore some of macroeconomic properties of the model. A core message of this section is that macroeconomists who turn to the methods of econophysics must temper them with explicit economic reasoning, especially macroeconomic reasoning.

2.1 GDP in the Baseline Model

Researchers in the econophysics tradition tend to emphasize model outputs over model structure. The idea is roughly the following: if the simulation results resemble the stylized facts inspiring the model, then the model is achieving its goals. In this restricted sense, the econophysics approach to model assessment resembles mainstream calibration approaches. At the level of the formal model, detailed realism is a secondary consideration; qualitative matching of selected views of the data has priority.

In order to produce analytically tractable models, mainstream microfoundations include extremely strong assumptions about individual circumstances and motivations. In contrast, models influenced by econophysics tend towards extreme behavioral agnosticism. In this paper we have repeatedly proposed that macroeconomists need more than behavioral agnosticism for model construction. We can pursue this observation again by asking a question: to maximize behavioral agnosticism, why should we not just select our stylized facts and then search for an appropriate joint distribution that we can sample from? For example, if we simply care about the distribution of GDP, why do we need more than a documentation of its empirical distribution, which we can sample from whenever we wish to reproduce our stylized GDP facts?

For economists drawn to macroeconomic modeling, the answer is obvious: we want macoreconomic models to usefully embody persistent real-world economic structure. For example, we would like to augment our understanding of GDP and unemployment. Ideally, we would like to improve our predictions about how these respond to exogenous changes, especially policy changes. This means we care about how the distributions of macroeconomic variables depend on the parameters of our model.

It is therefore natural, as illustrated by the social architecture model, that the influence of econophysics on macroeconomic modeling has not led to a true behavioral agnosticism. Rather it has provided support for efforts to free macroeconomic theory from the contemporary mainstream tropes of explicit microfoundations: essentially homogeneous agents, unbounded cognition, thin rationality, and equilibrium-trading constraints. Although these tropes achieved normative dominance in the 1970s and 1980s, the history of macroeconomics is filled with models that have not been bound by these norms. Ultimately, researchers must choose among models based on their usefulness (for understanding, prediction, or control). Like other theoretical decisions, the choice of more or less explicit microfoundations is ultimately justified by its contributions towards achieving these goals.

Once we acknowledge that model building is fundamentally a pragmatic endeavor, our discussions of macroeconomic modeling are quickly drawn into very subtle questions in the philosophy of science, which lie beyond the scope of the present paper. For example, we are drawn into questions about the role of ever-present simplifying assumptions in macroeconomic theory, along with questions of whether the assumptions of a model count among the predictions of the model. In contrast, models influenced by econophysics have been more concerned with simply demonstrating that some empirical distributions can be readily reproduced in the stochastic steady state of some relatively theory-free simulation model.

2.2 Roughly Lognormal GDP

Evaluation of the social architecture model has focused on its ability to roughly reproduce certain macroeconomic stylized facts.[12] Lee et al. (1998) use data from the Penn World Table to argue that detrended GDP data is roughly lognormally distributed. Taking this as a stylized fact about the distribution of GDP, Wright (2009) shows that the distribution

[12] Attempts to empirically evaluate macroeconomic models face deep methodological problems, many of which have been widely recognized (if habitually discarded) since Keynes (1939) offered his critique of Tinbergen's methods. As in the calibration literature, the econophysics literature tends to deal in practitioner-specific notions of what constitutes a resemblance of simulated outcomes to real-world data. Any effort to explore what might plausibly be meant by *resemblance* would require delving into the philosophy of science, which again lies beyond the scope of this paper.

of GDP generated by the social architecture model is similarly crudely lognormal. To this extent, we have conformance of simulated GDP to the empirical data. This kind of comparison of simulated stochastic equilibria to stylized facts typifies models that are influenced by econophysics.

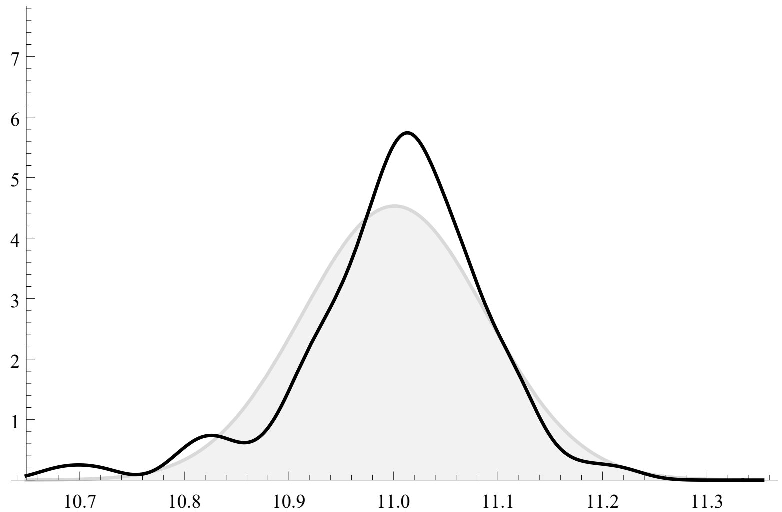

Logarithm of GDP; baseline parameterization. The solid line represents a kernel density estimate. The filled gray area represents a fitted normal distribution.

Figure 1: Distribution of Log GDP

We begin this section by confirming these GDP findings under a baseline parameterization of the model. This parameterization is chosen to match Wright (2009): 1000 agents, each with money of 10.[13] We measure annual GDP on the expenditure side, computed as the total revenues during a year by all firms. This means that we need to sum revenues for the year across all agents who spent time as employers. To represent the distribution of these GDPs, we produce a kernel density estimate (using a Gaussian kernel). For comparison, we superimpose this estimate on a fitted lognormal distribution (i.e., a normal distribution fitted to log GDP). Although the tails of log GDP appear a bit fatter than normal (but not as fat as Cauchy), the lognormal distribution provides a rough visual fit to the simulation data.[14]

We next consider a problem that arises from the construction of the simulated GDP data.

[13]For full details, including initialization and burn in, see the appendix. The scaling effects of money are reported below. Absence of important scaling effects in the number of agents is reported by Wright (2009,

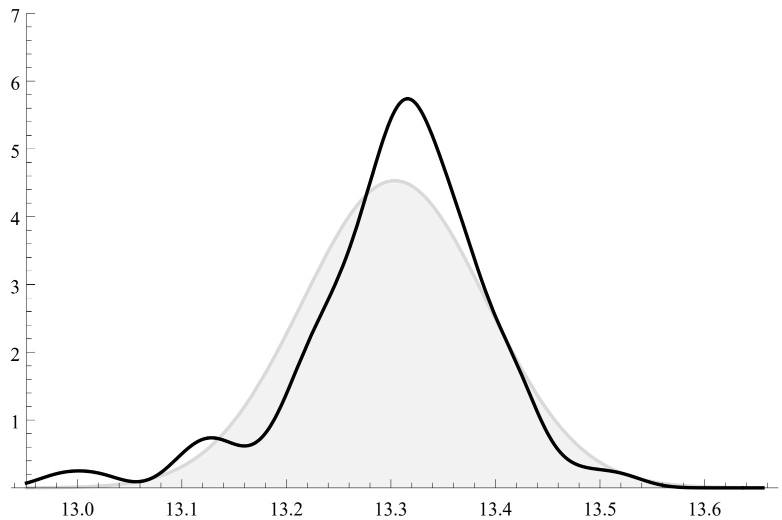

Figure 2: Monetary Expansion Increases Measured GDP
Logarithm of GDP; baseline parameterization, except for money supply increase. The solid line represents a kernel density estimate; the filled gray area represents a fitted normal distribution.

Total annual revenues of the economy is a measure of *nominal* GDP. But Lee et al. (1998) naturally analyze the behavior of real GDP. One might imagine attempting a reconciliation of the nominal and real data by invoking a constant price level, but this would create more problems than it would solve. In particular, the social architecture model implies a clear prediction: in order to permanently double mean GDP, we need only permanently double the money supply. Figure 2 illustrates the results of increasing the money supply by an order of magnitude, without making any other changes in the model. Comparison with our previous Figure makes it clear that we have shifted mean GDP to the right by about an order of magnitude.

We have found that on average a monetary expansion simply shifts measured GDP. Few economists will be surprised by this result: the independence of average real output from the level of the nominal money supply has been repeatedly supported theoretically and

p.8), and we explored smaller economies down to 100 agents to confirm this finding. (Not shown.)

[14]Lee et al. (1998) and Wright (2009) report only visual evidence of this fit, comparing a histogram to a fitted lognormal distribution. Our use of a kernel density estimate instead of a histogram makes it easier to judge how the data diverge from the fitted model, and it is clear in this case that the data have modest excess positive kurtosis. For this reason, we emphasize that the data are *roughly* lognormal, in the sense of approximate visual fit. In the discussion of real GDP below, we will additionally see some positive skew in the data.

verified empirically (McCandless and Weber, 1995; Bullard, 1999; Lucas, 2014).[15] Therefore, any reasonable reconciliation of the behavior of GDP in the social architecture model with the available macroeconomic facts must be compatible with such independence. In the next section, we propose a simple method to extract a real GDP measure from the social architecture model.

2.3 Nominal vs Real GDP

Production is left entirely in the background of the basic social architecture model. In order to extract a measure of real GDP from this model, we must add some description of production. In this section, we offer a minimal supplement to the model that allows us to discuss real GDP. This supplement is no more than a simple reinterpretation of the model via an added background assumption that production is linear in inelastically supplied homogeneous labor. Linear production is common in agent-based macroeconomics (Delli Gatti et al., 2005; Russo et al., 2007; Edoardo Gaffeo and Gallegati, 2008; Neveu, 2013). Furthermore, it aligns the model with a large heterodox literature in macroeconomics, where (given labor productivity) fluctuations in real GDP correspond to fluctuations in employment. The resulting salience of labor inputs also creates an attractive link to the circuitist school, which seems particularly appropriate given the emphasis of the social architecture model on the monetary flows underpinning transactions.

We compute the total labor input for a year by summing total employment each month over the year. Since the basic social architecture model abstracts from technological change, we may arbitrarily normalize units of real GDP. For simplicity, we set one unit of GDP to one person month. A kernel density estimate of the resulting distribution of real GDP still appears crudely lognormal, as illustrated in Figure 3. In sum, switching to our more appropriate real GDP measure does not diminish the model's ability to accommodate the

[15]To avoid misunderstandings, note that these empirical observations do *not* imply that money is neutral in the senses attacked by Davidson (1987), that the money supply is is exogenous, or that monetary policy is ineffective for macroeconomic stabilization.

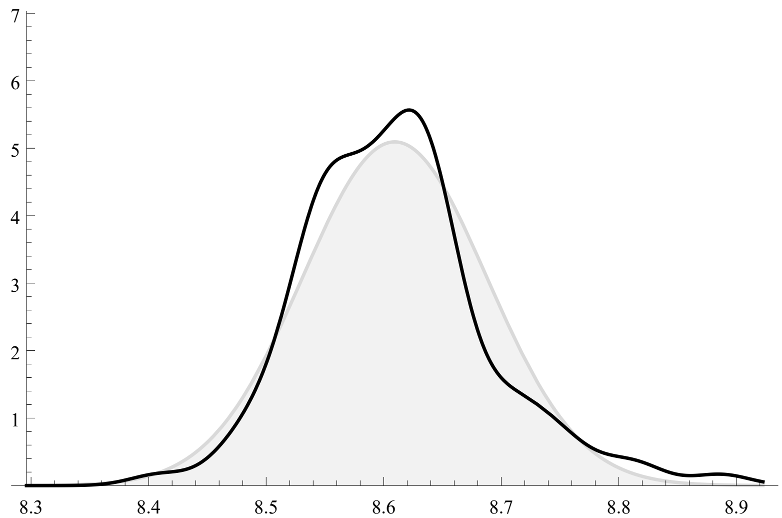

Logarithm of real GDP; baseline parameterization. The solid line represents a kernel density estimate; the filled gray area represents a fitted normal distribution.

Figure 3: Distribution of Real GDP

macroeconomic stylized fact of lognormal GDP.

3 Individual Behavior and Macroeconomic Outcomes

The social architecture model readily survived the challenges posed in section 2.3. The present section proposes a different kind of exploration. It asks how sensitive the model's results are to its particular form of behavioral agnosticism.

Recall that key behaviors in the model are randomized. Section 1.1 explored the justification for this while explicating the idea of implicit microfoundations. Reliance on implicit microfoundations links the social architecture model to the econophysics literature, so the explorations of the present section bear on this link.

Randomized behavior is often seen as an application of Keynes's principle of indifference, which Keynes in turn equated to Jacob Bernoulli's principle of insufficient reason (Keynes, 1921, ch.4). Roughly speaking, the principle of indifference proposes that ignorance justifies diffuse priors. Sinn (1980) discusses many difficulties and attractions of the principle of indifference; these lie beyond the scope of our paper. Instead, we focus on the observation of Becker (1962) and Gode and Sunder (1993) that aggregate outcomes may be relatively robust to behavioral specifications. For example, Becker argued that market demand curves tend to be negatively inclined whether consumers are neoclassically rational, irrationally impulsive, or irrationally inertial. In this section, we make a related point for the social architecture model.

3.1 Spending and Effective Demand

We first consider two places in the model where the specific modeling choices intuitively seem least likely to matter: consumer spending behavior, and the firm-revenue rule. With many consumers, the probability is high that aggregate spending per capita will approximate its expected value. For similar reasons, random conversion of latent demand to effective demand seems unlikely to matter much in the aggregate.

To explore the validity of these intuitions, we remove all randomness from consumer spending and from the firm-revenue rule, while matching the mean values. This means a consumer will always spend half of current wealth, and half of latent demand will always be converted to effective demand. Figure 4 illustrates the real-GDP consequences of independently making each of these changes. Comparison with Figure 3 reveals only small effects on the behavior of real GDP. There is no detectable effect on mean GDP. There is little effect on overall dispersion, but there is a slight skewness evident in the kernel density estimate. The same is true when we simultaneously implement both changes. (Not shown.)

3.2 Reservation Wage

We next reconsider the adjustment of the reservation wage. Recall that when the job search of an unemployed agent fails, the agent's reservation wage falls. Once again the social architecture model keeps microfoundations implicit: invoking our ignorance of the actual psychology of reservation wage adjustment, it simply treats the adjustment size as random. Yet once again, in the presence of many jobseekers, it is unclear that this randomness (rather

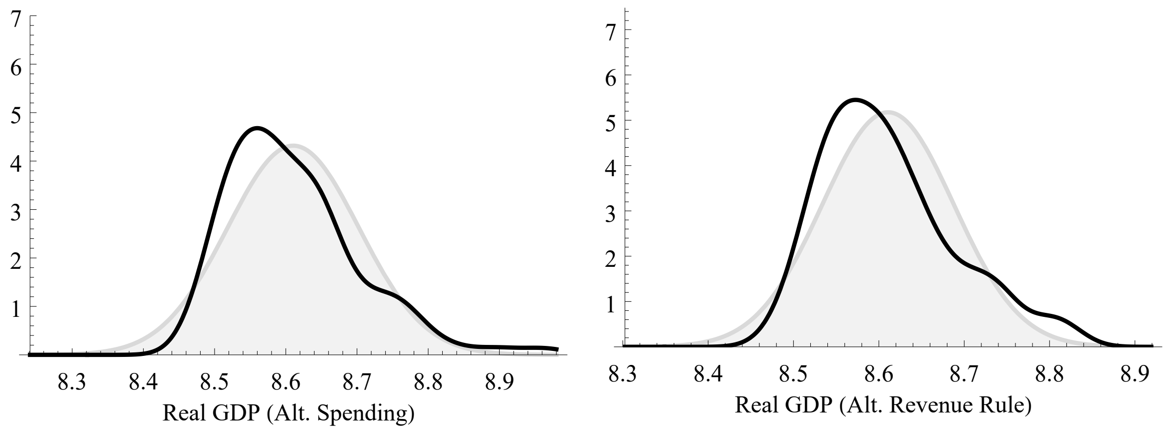

Logarithm of real GDP; baseline parameterization, except for the following changes. In the *Alt. Spending* subfigure, spending is deterministic; in the *Alt. Revenue Rule* subfigure, the conversion of latent demand into firm revenues is deterministic. A solid line represents a kernel density estimate. A filled gray area represents a fitted normal distribution.

Figure 4: Effect of Spending and Revenue Rules on Real GDP

than the average downward adjustment) should matter much for macroeconomic outcomes.

We therefore take the same approach as in the preceding subsection: we replace the random adjustment process with a deterministic process, preserving the mean. The resulting real GDP outcomes are illustrated in Figure 5, which also shows the outcomes from combining this change with the previous two changes. The two subfigures are very similar, although we can see a slightly reduced dispersion in the second.

By comparing Figure 3 and Figure 4 with Figure 5, we see that changing the reservation wage behavior has bigger effects than the previous changes. While the mean of the distribution is roughly the same, the variance is somewhat larger. This result may at first appear surprising: despite the shared mean adjustment, we find that less behavioral uncertainty in the model produces a *greater* variation in real activity. Yet this reflects the simple fact that with behavioral uncertainty most agents adjust their wage demand more than the average decline.[16] To put it another way, with uncertainty in the wage adjustment, the median

[16]This is not a quirk of the model but a general observation about uniformly distributed proportional declines. To see this, consider simple discrete distribution: the wage demand is reduced by either $1/4$ or

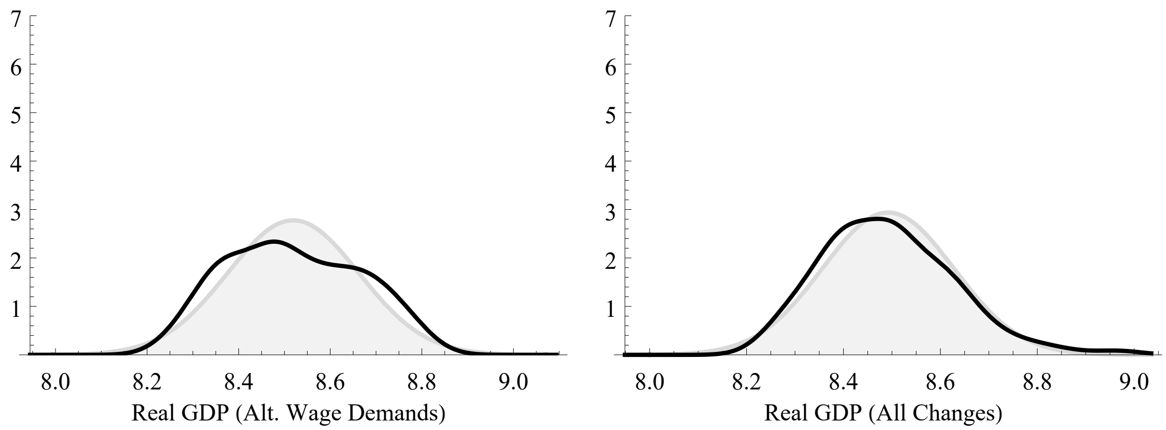

Logarithm of real GDP; baseline parameterization, except as follows. In the *Alt. Wage Demands* subfigure, we render deterministic the reservation wage adjustment of the unemployed; in the *All Changes* subfigure, we addtionally include the changes illustrated in Figure 4. A solid line represents a kernel density estimate; a filled gray area represents a fitted normal distribution.

Figure 5: Effects of Reservation Wage Behavior on Real GDP

reservation wage falls more quickly than the mean reservation wage.

3.3 Wage Negotiation and the Unemployment Problem

Although the effects on real GDP were still fairly modest, the previous subsection suggested that labor market considerations loom relatively large in the social architecture model. We now pursue this observation more vigorously, and in the process we explore the unemployment problem in the social architecture model.

Wright (2009, p.9) finds that the unemployment rate in the social architecture model averages 18.5%, "higher than is usually reported in modern economies". Unemployment measures accounting for discouraged and marginally attached workers sometimes reach such levels, and Wright suggests that this offers the best comparison to the social architecture

3/4, each with probability 1/2. This produces a mean reduction in wage demands of 1/2. Two declines by half would of course produce a total decline of 3/4. However, given two rounds of declines, the wage demand falls by 7/16 with probability 1/4, by 15/16 with probability 1/4, and by 13/16 with probability 1/2. In most cases, the decline is more than the deterministic decline.

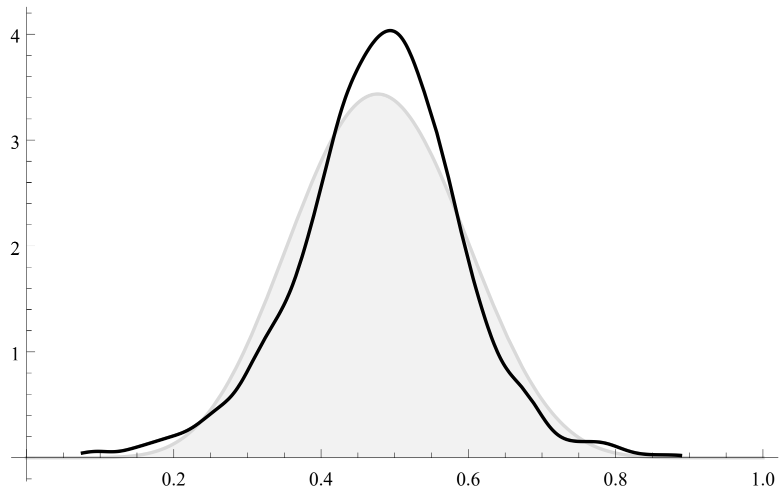

Unemployment rate; baseline parameterization. The solid line represents a kernel density estimate. The filled gray area represents a fitted beta distribution.

Figure 6: The Unemployment Problem

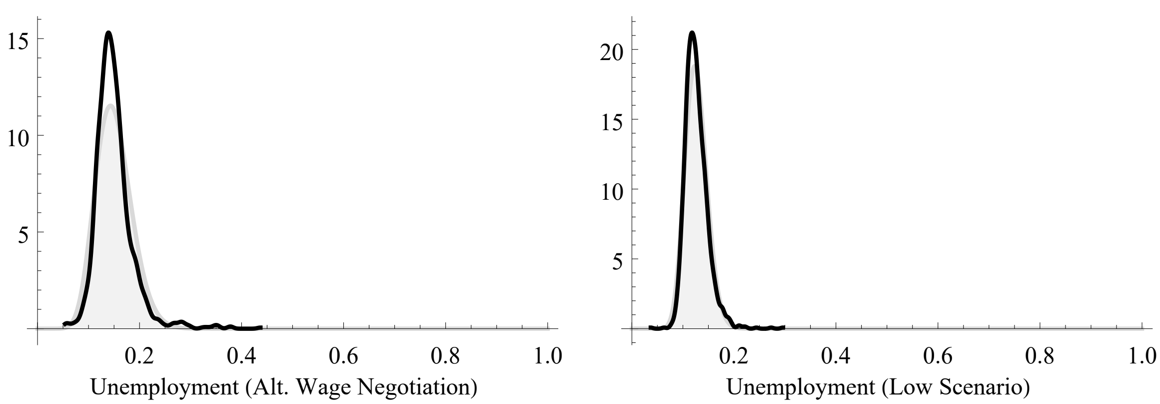

Unemployment rate; Baseline parameterization except for the following changes. The *Alt. Wage Negotiation* subfigure includes the new wage negotiation behavior; the *Low Scenario* subfigure additionally the spending and latent demand changes discussed above. The solid line represents a kernel density estimate. The filled gray area represents a fitted beta distribution.

Figure 7: Effect of Wage Negotiation on Unemployment

model.[17] Nevertheless, these high unemployment rates are *prima facie* a problem for the social architecture model.

In fact, the unemployment problem in the social architecture model is substantially worse than reported by Wright. Recall that in this model, a fired employee adopts a wage demand equal to the last wage payment received. Wright (2009, p.27) has a coding error in the `WagePaymentAndFiring` subroutine: instead of resetting the wage received of the employee, the published code resets this attribute of the employer. Our reported simulations include a correction of this small error. We illustrate the consequences in Figure 6, which plots a kernel density estimate for the unemployment distribution in the baseline model. For comparison, we subimpose a fitted beta distribution.[18]

It may seem surprising that this change produces such a sizable effect on the unemployment distribution, particularly since the reservation wage is reset each time an employee takes a new position. However, when an employee is downsized, the reservation wage is reset to the last wage received. In the erroneous code, the last wage received will never be updated during the employment period. The wage laddering promised by the model description is thereby missing from its implementation in code. With this understanding, it becomes less surprising that the corrected implementation produces substantially higher unemployment rates.

This implied unemployment distribution poses a serious challenge to the social architecture model. To address this, we re-examine the labor market, focusing on the wage negotiation process. Recall that when an employee approaches a potential employer, wage negotiations lead to an offer between the one and two times the employee's reservation wage. Since the time scale of the simulation is one month per iteration, this characterization of wage negotiation can produce highly unrealistic earnings acceleration. To produce more realistic wage behavior, we will reset the mean increase to 1% per negotiation. This would still

[17] For example, in the US the U6 unemployment rate briefly exceeded 17% after the Great Recession.

[18] The switch to a beta distribution accommodates the restricted range of possible unemployment values. Fitting a normal distribution produces very similar results. (Not shown.)

allows increases of more than 25% per year for a very lucky worker, but in keeping with the explorations of this paper, we also remove the uncertainty in the negotiation. The result can be seen in the first subfigure in Figure 7. As a comparison of possible interest, the second subfigure additionally incorporates the spending and latent demand changes discussed above. This low-unemployment scenario pushes mean unemployment down to about 12.5%. This is still a bit high when compared with typical U6 figures for the US, but it is a plausible value for many other countries (e.g., Japan) and even low for others (including much of Europe).

We have demonstrated that the characterization of wage negotiation has large effects on the unemployment outcomes of the social architecture model. We have also shown that a small, plausible change in the average payoff to job finding produces a marked change in the plausibility of the unemployment outcomes. This does not mean that we have completely eliminated the unemployment problem in the social architecture model, but we have taken a substantial step in that direction.

4 Conclusion

Agent-based macroeconomics often discards the mainstream assumptions of thin rationality and equilibrium-trading constraints. The social-architecture model takes this further than most, often discarding behavioral speculations altogether in favor of behavioral agnosticism. The model strives in this way to provide *implicit microfoundations* for a monetary macroeconomy. A core motivation for implicit microfoundations is to minimize the "pretence of knowledge" while still allowing the construction of formal macroeconomic models. Implicit microfoundations replace the counterfactual mainstream assumptions about the knowledge and foresight of agents with constrained randomness.

Implicit microfoundations in this sense may be seen as an application the Keynes's principle of indifference (Keynes, 1921, ch.4). However, in light of the results of Becker (1962) and Gode and Sunder (1993), a natural question arises: once we consider large numbers of heterogeneous agents facing basic market constraints, might macro-outcomes prove relatively robust to our treatment of micro-behaviors? For the specific behaviors we consider, this paper provides some support for that speculation.

We nevertheless argue that microfoundations cannot be as implicit as some researchers have hoped. In this sense, the methodological compromises of the social architecture model are appropriate, and the model offers an interesting marriage of the methods of econophysics to the concerns of macroeconomics. Previous research found that the this model can closely match many key macroeconomic outcomes. We re-examine the correspondence of the model outcomes with some stylized facts about real GDP and unemployment, expose some problems in the social architecture model, suggest ways to mitigate those problems without radically changing the model, and finally reach mostly favorable conclusions.

References

Addison, John T., José A. F. Machado, and Pedro Portugal (2013). "The Reservation Wage Unemployment Duration Nexus." *Oxford Bulletin of Economics and Statistics* 75(6), 980–987.

Aoki, Masanao (1996). *New Approaches to Macroeconomic Modeling: Evolutionary Stochastic Dynamics, Multiple Equilibria, and Externalities as Field Effects.* Cambridge University Press.

Aoki, Masanao (2002). *Modeling Aggregate Behavior and Fluctuations in Economics.* Cambridge, UK: Cambridge University Press.

Ashraf, Quamrul, Boris Gershman, and Peter Howitt (2017, March). "Banks, Market Organization, And Macroeconomic Performance: An Agent-based Computational Analysis." *Journal of Economic Behavior and Organization* 135, 143–180.

Axelrod, Robert (1984). *The Evolution of Cooperation.* New York, NY: Basic Books.

Axelrod, Robert (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.

Becker, Gary S. (1962, February). "Irrational Behavior and Economic Theory." *Journal of Political Economy* 70(1), 1–13.

Brown, Sarah and Karl Taylor (2013). "Reservation Wages, Expected Wages and Unemployment." *Economics Letters* 119(3), 276 – 279.

Bullard (1999, November/December). "Testing Long-Run Monetary Neutrality Propositions: Lessons from the Recent Research." Federal Reserve Bank of St. Louis *Review* 81(6), 57–77.

Clower, Robert (1967, December). "A Reconsideration of the Microfoundations of Monetary Theory." *Western Economic Journal* 6(1), 1–8.

Cottrell, Alan F., Paul Cockshott, Gregory John Michaelson, Ian P. Wright, and Victor M. Yakovenko (2009). *Classical Econophysics*. Routledge Advances in Experimental and Computable Economics. London, UK: Routledge.

Davidson, Paul (1987). "Sensible Expectations and the Long-Run Non-Neutrality of Money." *Journal of Post Keynesian Economics* 10(1), 146–153.

Delli Gatti, Domenico, Edoardo Gaffeo, Mauro Gallegati, and Antonio Palestrini (2005). "The Apprentice Wizard: Monetary Policy, Complexity and Learning." *New Mathematics & Natural Computation* 1(1), 109–128.

Driscoll, John C. and Steinar Holden (2014, September). "Behavioral Economics and Macroeconomic Models." *Journal of Macroeconomics* 41, 133–147.

Drăgulescu, Adrian A. and Victor M. Yakovenko (2000, October). "Statistical Mechanics of Money." *The European Physical Journal B* 17(4), 723–729.

Edoardo Gaffeo, Domenico Delli Gatti, Saul Desiderio and Mauro Gallegati (2008, Fall). "Adaptive Microfoundations for Emergent Macroeconomics." *Eastern Economic Journal* 34, 441–463.

Epstein, Joshua M. and Robert L. Axtell (1996). *Growing Artificial Societies: Social Science from the Bottom Up.* Washington, DC and Cambridge, MA: Brookings Institution Press and MIT Press.

Eswaran, Mukesh and Ashok Kotwal (1986, June). "Access to Capital and Agrarian Production Organisation." *The Economic Journal* 96(382), 482–498.

Farmer, J. Doyne, Paolo Patelli, and Ilija I Zovko (2005). "The Predictive Power Of Zero Intelligence In Financial Markets." *Proceedings of the National Academy of Sciences of the United States of America* 102(6), 2254–2259.

Foley, Duncan K. (1994). "A Statistical Equilibrium Theory of Markets." *Journal of Economic Theory* 62(2), 321–345.

Foulkes, Fred K. (1980). *Personnel Policies in Large Nonunion Companies.* Prentice Hall.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley (2006, May). "Institutional Investors and Stock Market Volatility." *Quarterly Journal of Economics* 121(2), 461–504.

Gode, D.K. and S. Sunder (1993, February). "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality." *Journal of Political Economy* 101(1), 119–137.

Keynes, John Maynard (1921). *A Treatise on Probability.* London: MacMillan and Co., Limited.

Keynes, John Maynard (1939, November). "Professor Tinbergen's Method." *Economic Journal* 49(195), 558–577.

Lee, Youngki, Lus A. Nunes Amaral, David Canning, Martin Meyer, and H. Eugene Stanley (1998). “Universal Features in the Growth Dynamics of Complex Organizations.” *Physical Review Letters* 81(15), 3275–3278.

Lucas, Robert E. and Nancy L. Stokey (1987). “Money and Interest in a Cash-in-Advance Economy.” *Econometrica* 55(3), 491–513.

Lucas, Jr., R. E. (2014, Third Quarter). “Liquidity: Meaning, Measurement, Management.” *Federal Reserve of St. Louis Review* 96(3), 199–212.

McCandless, George T., Jr. and Warren E. Weber (1995, Summer). “Some Monetary Facts.” Federal Reserve Bank of Minneapolis *Quarterly Review* 19(3), 2–11.

Mortensen, Dale (1970, December). “Job Search, the Duration of Unemployment and the Phillips Curve.” *American Economic Review* 60(5), 847–862.

Neveu, Andre R. (2013). “Fiscal Policy and Business Cycle Characteristics in a Heterogeneous Agent Macro Model.” *Journal of Economic Behavior and Organization* 92, 224–240.

Roemer, John E. (1982). *A General Theory of Exploitation and Class.* Cambridge, MA: Harvard University Press.

Russo, A., M. Catalano, E. Gaffeo, Mauro Gallegati, and M. Napoletano (2007). “Industrial Dynamics, Fiscal Policy and R&D: Evidence from a Computational Experiment.” *Journal of Economic Behavior and Organization* 64(3–4), 426–447.

Sinn, Hans-Werner (1980). “A Rehabilitation of the Principle of Insufficient Reason.” *The Quarterly Journal of Economics* 94(3), 493–506.

Steindl, Josef (1965). *Random Processes and the Growth of Firms.* London: Charles Griffin and Company.

Wilensky, Uri and William Rand (2015). *An Introduction to Agent-Based Modeling: Modeling Natural, Social, and Engineered Complex Systems with NetLogo*. Cambridge, MA: MIT Press.

Wright, Ian (2008). "The Emergence Of The Law Of Value In a Dynamic Simple Commodity Economy." *Review of Political Economy* 20(3), 267–391.

Wright, Ian (2009, May). "Implicit Microfoundations for Macroeconomics." *Economics* 3, Article 2009–19.

Yakovenko, Victor M. and J. Barkley Rosser (2009, December). "*Colloquium* : Statistical Mechanics of Money, Wealth, and Income." *Reviews of Modern Physics* 81, 1703–1725.

A Parameterization, Initialization, and Source Code

Our parameterization is chosen to match the description in Wright (2009): 1000 agents, with initial money holdings of 10 per agent. (These are the values used for the reported GDP and unemployment results.) We arbitrarily initialize the macrostate with a latent demand of zero, but this arbitrariness is offset by a substantial burn-in period for the simulation. Initialization of agents is also somewhat arbitrary (but again, this is offset by the simulation's burn-in period). Every agent gets the same initial wealth (in the form of money), which removes a scaling consideration.

We run each simulation for 150 years. We treat the first 50 years as burn in, so we produce 100 years of monthly data for each simulation. For reasons elaborated in Wright (2009, p.9) we report results for a single replicate of each scenario. (Briefly, the model quickly reaches a stochastic steady state.)

Wright (2009, pp.26–27) includes full Mathematica source code for the social architecture model. The changes described in the present paper can be implemented readily by anyone familiar with the Wolfram Language.[19] (For example, to implement the deterministic spending rule, simply change `expenditure = RandomReal[]*agent[[agentMoney]]` to `expenditure = 0.5*agent[[agentMoney]]`.)

[19]We ran the simulations with Mathematica 11.0.1 on Windows 10. Prior to each run, we set the random seed to 314, ensuring exact replicability. (The seed is arbitrary; Wright (2009, p.20) shows that, as one expects, the model results are robust to the choice of seed.) Note that we use the term *latent demand* for what Wright calls *effective demand*.