

SOCIAL LEARNING FOR THE MASSES

James Bullard

Federal Reserve Bank of St. Louis

Computational & Experimental Economics Workshop

Simon Fraser University

Feb. 4, 2023

Vancouver, British Columbia

Any opinions expressed here are our own and do not necessarily reflect those of the FOMC.



FEDERAL
RESERVE
BANK *of*
ST. LOUIS

Introduction

THE INTELLECTUAL LEGACY OF JASMINA ARIFOVIC

- Jasmina Arifovic (JA) was a pioneer in the application of artificial intelligence to macroeconomics, helping us to gain insight into the question, “How is equilibrium achieved?”
- JA ideas will be even more important in the decades ahead as macroeconomists work with more and more granular models: more agents, more details, more shocks, more frictions.
- This paper: How is equilibrium achieved in these more complex environments?
 - An earlier and more preliminary version of this talk was given under the title “[Conjectures on Learning in Krusell-Smith-type Economies](#)” at the 2021 Bank of Canada Annual Economic Conference on Nov. 10, 2021.

A MORE COMPLEX ECONOMY

- I study a stylized DSGE heterogeneous agent life cycle model with a known competitive equilibrium featuring Gini coefficients close to those in the U.S. data.
- The model features three aggregate shocks as well as idiosyncratic risk, but also features policies that can mitigate both the aggregate risk and the idiosyncratic risk.
- A welfare theorem states the sense in which these policies can achieve an optimal allocation of resources.
 - The subtext in this talk is that models in this class represent, broadly, the current and future direction of macroeconomics, and that the learning literature will have to continue to refine methods to provide insight for these environments.

LEARNING

- I then turn to discuss how agents might learn in this relatively complex macroeconomic setting if agent behavior is at some point disturbed.
- I will conclude that social learning as promoted by Jasmina Arifovic is likely to provide the best path forward.
 - Unmodified concepts of econometric learning promoted and studied extensively in the existing literature are less likely to be appropriate in this environment.

Core argument

AN OLDER TRANSITION TIME RESULT

- Suppose the economy is initially on a balanced growth path but is suddenly disrupted by a “one-time, large, unanticipated shock.”
 - This shock is *above and beyond* the shocks envisioned within the ambient stochastic environment of the model.
 - To fix ideas, think of an unanticipated “financial crisis” or an unanticipated “pandemic.”
- In a related class of heterogeneous agent models, an earlier generation of quantitative study emphasized perfect foresight transition times following a disturbance of this type.
- That literature found that transition times are long—measured in years or decades—in this related class of models.

WHY SLOW CONVERGENCE?

- The slow convergence was because the disrupted agents experiencing the shock would have to complete their life cycle and exit the model before the long-run balanced growth path can be achieved.
- Taken literally, one might conclude that actual macroeconomies subject to occasional “large, unanticipated shocks” would nearly always be in transition, even if households had rational expectations following the large shock.
 - Examples: Auerbach and Kotlikoff (*Dynamic Fiscal Policy*, 1987); see also Cogley and Sargent (*JME*, 2008) in which the large shock twists the priors of a Bayesian learner and leads to slow learning over subsequent decades.

RELATIVELY FAST CONVERGENCE IN THE U.S. DATA

- I will calibrate the DSGE model used in this paper to U.S. data assuming U.S. macroeconomic policies are in fact the optimal ones the model requires.
- I will then provide *prima facie* evidence that actual convergence times in U.S. economic data following a “large, unanticipated shock” are an order of magnitude shorter than in the earlier literature.
- In particular, these transition times are measured in quarters rather than years or decades.

SOCIAL LEARNING FOR THE MASSES

- This suggests that in reality, the U.S. economy—despite its complexity—does not seem to follow the types of slow adjustment paths emphasized in some of the earlier literature.
- I will suggest that the rapid convergence observed in the U.S. data could occur if there is substantial communication across the society—social learning as implemented by JA.
- In the economy I describe, this can occur because there are many millions of agents that have already learned and retained the “DNA” of optimal decision rules for consumption, assets and hours worked before the shock occurred.
- Other agents that may not know these decision rules can learn relatively quickly from those that do.
- I call this phenomenon “social learning for the masses.”

Environment

ENVIRONMENT BASICS

- At each date t , a new continuum of households enters the economy, makes economic decisions over $T + 1 = 241$ dates, then exits the economy. (To fix ideas, think of $\approx 1m$ agents per quarterly cohort.)
- This corresponds to an agent entering the economy as a decision-maker at age 20 and exiting as a decision-maker at age 80, inclusive of end points, and making economic decisions at a quarterly frequency.
- Results are perfectly general for the choice of T , with higher values corresponding to decision-making at more frequent intervals.
- This class of models has a “paper-and-pencil” equilibrium solution, and so it provides a simple benchmark model for heterogeneous-agent macroeconomies with aggregate shocks.

RISKS FACED BY HOUSEHOLDS

- There are both aggregate risk and idiosyncratic risk.
- Idiosyncratic risk is borne as a productivity-profile scaling shock as the agent enters the economy, and also in the form of simple *i.i.d.* unemployment risk at each date.
- Monetary and fiscal policymakers provide a form of insurance against the aggregate risk, and a labor market authority provides unemployment insurance.
- The idiosyncratic risk borne as the agent enters the economy via the productivity-profile scaling shock is uninsurable.
- A welfare theorem describes the sense in which the equilibrium studied here represents a first-best allocation of resources.

ASSETS

- There are three *nominally denominated* assets: privately issued debt, publicly issued debt and capital.
- We think of these as representing U.S. data counterparts: (1) mortgage-backed securities (MBS), (2) federally issued debt and (3) physical capital, respectively.
- In the U.S. data, MBS net out, but federally issued debt and physical capital are in positive net supply and we target a value of the assets-to-GDP ratio equal to $1.23 + 3.32 = 4.55$.

NOMINAL CONTRACTING

- The credit market friction is non-state contingent nominal contracting (NSCNC): All debt contracts are stated in nominal terms, with a stated nominal interest rate, and repayment is not state-contingent.
- The role of monetary policy is to adjust the price level each period in order to convert these nominal, non-state contingent contracts into real, state-contingent contracts.

HOUSEHOLD TYPES

- Household types: “life cycle” (LC) and “hand-to-mouth” (HTM).
- The life-cycle households are assigned a hump-shaped productivity profile at the beginning of their life cycle. Accordingly, they need to use credit markets (hold assets) to smooth life-cycle consumption.
- The hand-to-mouth households are assigned a perfectly flat productivity profile as they enter the economy. Accordingly, they never need to use credit markets and instead consume their labor income each period.

ATTAINING THE CORRECT ASSET LEVEL

- The economy with only LC households wants to hold assets equal to $A/4Y = 5.71$, a value which is considerably higher than the value observed in the U.S. data, which is 4.55.
- The economy with only HTM households would be “Spartan,” and would hold no assets at all.
- We will adjust the fraction of HTM households in order to match the assets-to-GDP ratio in the U.S. data.

PREFERENCES

- Each household $i \in (0, 1)$ entering the economy at date t has preferences (the same for both LC and HTM types)

$$U_{t,i} = \sum_{s=0}^T [\eta \ln \tilde{c}_{t,i}(t+s) + (1-\eta) \ln \ell_{t,i}(t+s)].$$

- We define $\tilde{c}_{t,i}(t+s) = D(t+s) c_{t,i}(t+s)$, where $D(t+s)$ is the state of aggregate demand at date $t+s$. The state of demand evolves as

$$D_t = \delta(t-1, t) D_{t-1},$$

where $\delta(t-1, t)$ is the gross growth rate of demand, which follows an appropriate stochastic process that keeps $D(t) > 0 \forall t$.

PRODUCTIVITY PROFILES

- Agents entering the economy draw a scaling factor x from a lognormal distribution and receive a productivity profile that is a scaled version of a baseline profile, e_s :

$$e_{s,i} = x \cdot e_s,$$

where for LC agents $e_s^{LC} = 1 + p_1 \exp \left[- \left(\frac{s-p_2}{p_3} \right)^4 \right]$, and where p_1, p_2 and p_3 are chosen to match calibration targets given below, and for HTM agents $e_s^{HTM} = h (1/T) \sum_{s=0}^T e_s^{LC}$ where $h \in (0, 1)$.

- Huggett, Ventura and Yaron ([AER, 2011](#)) argue that differences in initial conditions are more important than differences in shocks for lifetime earnings.
- We think of all endowments at each date as containing linear labor tax factor $(1 - \tau^u)$, with τ^u set for all households in each cohort by the labor authority to fund unemployment insurance. This type of tax will not distort labor supply in this model.

TECHNOLOGY

- Aggregate real output $Y(t)$ is given by

$$Y(t) = [D(t) Q(t) N(t)]^{1-\alpha} K(t)^\alpha [L(t)]^{1-\alpha}, \quad (1)$$

where $K(t)$ is the real value of the physical capital stock, $L(t)$ is the aggregate effective human capital supply (hours \times productivity of various households), $Q(t)$ is a productivity index, $N(t)$ indexes the size of the labor force, and $D(t)$ is the state of aggregate demand.

- Q , N and D grow at stochastic gross rates λ , ν and δ respectively.
- These assumptions mean that real output grows at the stochastic rate $\lambda\nu\delta$ each period.
 - The aggregate demand assumption is a simple version of Bai, Ríos-Rull and Storesletten ([unpublished, 2019](#)).
 - The labor force growth assumption affects all cohorts proportionately and can be interpreted as “immigration.”

NOMINAL CONTRACTING AND TIMING PROTOCOL

- Under the assumptions outlined, the contract nominal interest rate is given by

$$R^n(t, t+1)^{-1} = E_t \left[\frac{\tilde{c}_{t,i}(t)}{\tilde{c}_{t,i}(t+1)} \frac{P(t)}{P(t+1)} \right]. \quad (2)$$

- The timing protocol is: (1) Nature assigns new entrant productivity profiles and also draws aggregate shocks; (2) The fiscal authority issues nominal debt; (3) The monetary authority sets the price level; (4) Households choose date t consumption, hours worked and net asset holding.
- Households will be able to make date t decisions without reference to future uncertainty, as the monetary policymaker is providing a type of perfect insurance.

THE FISCAL AUTHORITY

- The fully credible nominal debt issuance process is given by

$$B(t) = R^n(t-1, t) B(t-1), \quad (3)$$

where $B(t)$ is the total level of nominal debt and $B(0) > 0$.

- The fiscal authority is issuing enough new debt to maintain the level of assets in the economy at the appropriate level.

THE MONETARY AUTHORITY

- The monetary authority controls the price level directly and implements a price-level path criterion

$$P(t) = \frac{R^n(t-1, t)}{\delta(t-1, t) \lambda(t-1, t) \nu(t-1, t)} P(t-1). \quad (4)$$

- This criterion implements countercyclical price-level movements relative to the expectation embodied in the contract rate $R^n(t-1, t)$.
 - See Koenig (*IJCB*, 2013) and Sheedy (*BPEA*, 2014) on NGDP targeting.
 - See Andolfatto, et al. (unpublished, 2021, p. 14) for a discussion of how this criterion relates to a similar New Keynesian “targeting criterion” developed by Giannoni and Woodford (2004, pp. 101-2).

COMPETITIVE EQUILIBRIUM AND SOCIAL WELFARE

- Solution: Guess and verify that there is a competitive equilibrium in which the real rate of interest is always equal to the stochastic rate of real output growth.
 - The “Wicksellian natural real rate of interest” for this economy.
- A social planner would conclude that the allocation of resources is a social optimum provided (i) the planner places equal weight on all households for all time, (ii) the planner discounts backward and forward in time at the stochastic real rate of interest, (iii) the planner cannot alter the distribution of productivity profiles within the cohort, which are decided by nature at the beginning of the life cycle, (iv) the planner cannot alter the tax rate τ^u .

Calibration

MAPPING TO THE DATA

- Adjust cohort size based on data from the U.S. Census Bureau.
- Set the baseline hump-shaped life-cycle productivity profile such that households endogenously choose to work the hours worked by age in the U.S. data.
- Choose η to match average time devoted to market work across the economy.
- Set the fraction of HTM households (who do not hold assets) such that the aggregate level of assets to output, $A / (4Y)$, matches the U.S. data (4.55), with net assets defined as capital, $K / (4Y) = 3.32$, plus government issued debt, $B / (4Y) = 1.23$.
- Choose the within-cohort standard deviations of productivity for life-cycle and hand-to-mouth households, σ_{lc} and σ_{htm} , respectively, to approach the pre-taxes-and-transfers Gini coefficients for income and financial wealth in the U.S. data.

BASELINE LIFE-CYCLE PRODUCTIVITY




FIGURE: Baseline endowment profile of life-cycle agents.

THE MASS OF LIFE-CYCLE PRODUCTIVITY




FIGURE: The mass of endowment profiles: life-cycle agents (blue) and hand-to-mouth agents for $h = 0.5$ (red). The dashed lines denote the 25th and the 75th percentile of the endowment distributions.

HOURS WORKED BY AGE




FIGURE: Hours worked by age for life-cycle households: U.S. data (blue) and calibrated model (red).

POPULATION WEIGHTS

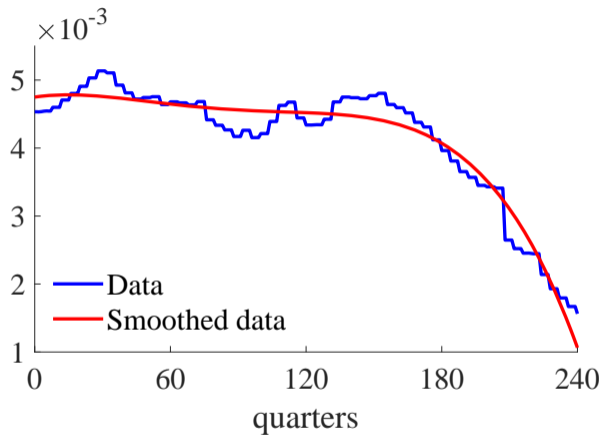


FIGURE: Population weights: U.S. data (blue) and 4th degree polynomial smoothed (red).

ASSETS AND GINI COEFFICIENTS

	Model	U.S. data
h	0.50	—
σ_{lc}	1.24	—
σ_{htm}	1.03	—
$A / (4Y)$	4.55	4.55
G_W	0.74	0.78
G_Y	0.66	0.63
G_C	0.62 [†]	0.32 [‡]

TABLE: Parameter values and associated assets-to-output ratio and Gini coefficients in the model equilibrium vs. the U.S. data.

[†] The consumption Gini in the model is based on a pre-taxes-and-transfers income concept.

[‡] The consumption Gini in the data is based on a post-taxes-and-transfers income concept.

THE CONSUMPTION GINI

- The consumption (out of *pre*-taxes-and-transfers income) Gini in the model equilibrium is $G_c = 0.62$.
- In the U.S. data, the consumption (out of *post*-taxes-and-transfers income) Gini is 0.32, about half as large.
- The model is saying that the net effect of taxes and transfers in the U.S. data is enough to reduce consumption inequality by half.
 - Some evidence: Using German data, Haan, Kemptner, and Prowse ([working paper, 2018](#)) use a life-cycle model to estimate that the tax-and-transfer system is sufficient to offset 54% of the inequality in lifetime earnings.

MARGINAL PROPENSITIES TO CONSUME SCHEMATIC




FIGURE: Cross section: Schematic of the marginal propensity to consume out of labor income by cohort for life-cycle agents. The MPC does not depend on the endowment scaling factor. Hand-to-mouth agents have a MPC of one.

MORE ASPECTS OF EQUILIBRIUM FIT TO U.S. DATA

- The model is calibrated to match hours worked by cohort for life-cycle households.
- Heckman: Appropriately specified “wage regressions” will suggest hours changes are independent of real wage changes.
- The model can be calibrated to fit U.S. real output growth exactly, attributing the growth in part to technological improvement, labor force growth uniform across cohorts, and the state of aggregate demand.
- The model predicts that consumption growth will be equalized across households at different ages and different income levels: economic growth gets “shared out” appropriately.
- The income, wealth and consumption distributions are maintained by a smoothly operating credit market with the correct interest rate.

Learning

NOMINAL RETURNS

- The model equilibrium predicts equalized nominal and real returns for three assets under optimal monetary policy: capital, MBS and Treasuries.
- These assets are not further differentiated inside the model.
- To compare with the data, we need an asset representing a return to capital in a format with risk characteristics similar to MBS and Treasuries.
- One candidate is a high-quality corporate bond.
- I will use a seven-year nominal investment-grade corporate bond metric. In the model and the data, this type of bond has a seven-year horizon but can be refinanced each period.

NOMINAL RETURNS VERSUS NOMINAL OUTPUT GROWTH

- The model equilibrium predicts that the nominal return on the assets should be equal to the nominal consumption growth rate, or, equivalently in the model, the nominal output growth rate.
- This prediction holds in periods of relative stability with optimal monetary, fiscal, and labor market policy.
- In these circumstances the private sector is able to set nominal debt contracts relying on the monetary authority to set the price level that ratifies those debt contracts ex-post.

LARGE UNANTICIPATED SHOCKS

- I will argue that the U.S. economy has been disturbed by two large unanticipated shocks since 2005: (1) the global financial crisis (GFC), and (2) the global pandemic.
- For my purposes, these events are simply “large disturbances” outside the scope of this model.
- The interim period, 2011-2019, fits the model assumptions better and we may expect the model to provide a better fit to the data during this time frame.
- It does not take long for the equilibrium conditions to be met after the GFC.

MODEL VERSUS U.S. DATA

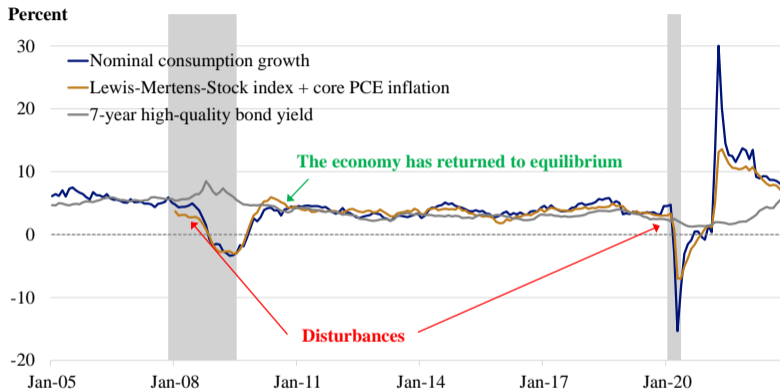


FIGURE: In "normal times," nominal consumption growth and nominal yields are close, as predicted by the model.

WHAT THE CHART SHOWS

- Measures of U.S. nominal consumption growth and nominal GDP growth on a 12-month basis are approximately equal to the nominal return on a 7-year high quality corporate bond between 2011 and 2019, as predicted by the model equilibrium.
- However, nominal growth rates and interest rates are considerably different during large, unanticipated shocks like the GFC and the pandemic.
- The chart suggests that the conditions of macroeconomic equilibrium were re-established relatively quickly after the GFC, and also appear to be close to being re-established following the pandemic.

ECONOMETRIC LEARNING

- The standard approach to learning—replace agents in the model with econometricians as in Cogley and Sargent (*JME*, 2008)—might interpret the large shocks as moments where rational expectations were badly disturbed across all agents in the economy: young and old, rich and poor.
- Forecasts that placed considerable weight on the chaotic observations from the crisis could lead to important changes in economic behavior, which could then feed back and continue to keep the economy away from its long-run equilibrium for some time.
- This vision of learning seems to be at odds with the data in the figure.

SOCIAL LEARNING

- This model has decision rules for *LC* households $i \in (0, 1)$:

$$\tilde{c}_{t-s,i}(t) = x_{lc} \eta \bar{e} w(t), \quad (5)$$

$$\ell_{t-s,i}(t) = (1 - \eta) \frac{\bar{e}}{e_s}, \quad (6)$$

$$\frac{a_{t-s,i}(t)}{P(t)} = x_{lc} w(t) \left\{ \left[\sum_{j=0}^s e_j \right] - \left(\frac{s+1}{T+1} \right) \sum_{j=0}^T e_j \right\}, \quad (7)$$

for $s = 0, \dots, T$, where \bar{e} is the average baseline endowment for *LC* agents and x_{lc} is the scale factor for agent i within the cohort, and for *HTM* agents $s = 0, \dots, T$:

$$\tilde{c}_{t-s,i}^{htm}(t) = x_{htm} \eta h \bar{e} w(t), \quad (8)$$

$$\ell_{t-s,i}^{htm}(t) = 1 - \eta, \quad (9)$$

$$a_{t-s,i}^{htm}(t) = 0. \quad (10)$$

SEEDED WITH DNA

- Households in the equilibrium of this model have very different incomes, levels of consumption, and assets.
- Nevertheless, they can learn from each other due to the fact that these optimal decision rules are transferable across agents because they adjust for age and productivity in the appropriate way.
- Furthermore, most agents would have had to learn these decision rules before the large, unanticipated shock occurred.
- The economy is in effect seeded with a sort of “DNA”—known, previously learned decision rules—even after the large shock occurs and begins to dissipate.

SOCIAL LEARNING

- After the GFC, for instance, there would be some cohorts whose only experience as decision-makers in the economy was during the crisis.
- However, there would be many more agents in the society, 95% or more, that would have knowledge of optimal decision-making in normal times.
- These known decision rules are relatively simple and can propagate exponentially quickly through the population following the large shock, returning the economy to equilibrium in short order.
- This “social learning for the masses” is more likely to be the successful learning concept in large heterogeneous agent economies.

Conclusions

JASMINA ARIFOVIC'S CONTRIBUTIONS

- JA was a pioneer in the application of methods from artificial intelligence to macroeconomics to try to help answer the question, “How is equilibrium achieved?”
- I have argued here that the combination of her insights and the likely future direction of macroeconomic research suggests that this work will be even more important in the decades ahead.