A Neoclassical Theory of the Keynesian Multiplier*

George-Marios Angeletos† **Chen Lian**‡

December 31, 2017

Abstract

This paper develops a novel theory of how a drop in consumer spending, or aggregate demand, can trigger a series of feedback loops between spending, employment, and income, ultimately leading to a sizable recession. Unlike the one embedded in the New Keynesian framework, our theory does not hinge on nominal rigidities and on the failure of monetary policy to replicate flexible prices. Instead, it is based on the idea that firms and consumers alike are unable to disentangle idiosyncratic from aggregate shocks and to reach common knowledge of the latter. This in turn could be either because of an information friction or because of bounded rationality. As a result, our theory bypasses the empirical failings of old and new Philips curves. It also allows for sizable fiscal multipliers without commensurate inflationary pressures.

*An earlier version was entitled "A (Real) Theory of the Keynesian Multiplier". For helpful comments and suggestions, we thank seminar participants at NYU and MIT.

†MIT and NBER; angelet@mit.edu.

‡MIT; lianchen@mit.edu.

1 Introduction

Can a drop in consumer spending, such as that triggered by a collapse in housing wealth, a credit crunch, or pessimistic beliefs about the future, trigger a recession? A popular narrative suggests yes: the drop in consumer spending is said to cause a fall in aggregate employment, output and income, which in turn feed back to a further drop in consumer spending, and so on.

This mechanism, which is known as the Keynesian multiplier, had a core position in the old IS-LM framework, but finds no place in the context of elementary, micro-founded, rational-expectations settings such as the baseline RBC model. In that context, countervailing mechanisms regarding the general-equilibrium adjustment of real wages and real interest rates turn the narrative on its head.

The New Keynesian framework fixes this problem, and makes room for the aforementioned narrative, by introducing nominal rigidities, preventing monetary policy from replicating the underlying flexible-price allocations, and equating a demand-driven recession with a monetary contraction.[1] This works, but raises two challenges. One is conceptual: presumably, one would like to make sense of the Great Recession and, more generally, of the idea that shifts in aggregate demand can trigger business cycles even when nominal rigidities are negligible or non-binding. The other is empirical: in the absence of appropriate bells and whistles, the New Keynesian model predicts, counterfactually, that the Great Recession should have been the Great Deflation or, more generally, that inflation should co-move strongly with output and employment.[2]

In this paper, we develop a formalization of the Keynesian multiplier that does not hinge on nominal rigidities and that bypasses these challenges. At the core of this theory is a friction, relative to the Arrow-Debreu and RBC frameworks, in the ability of economic agents to disentangle the forces that drive the variation in the behavior of others and in their own fortunes. Formally, we depart from the baseline RBC model by removing common knowledge of the underlying aggregate demand shocks, allowing them to be confounded with idiosyncratic shocks, and preventing the agents from reaching common knowledge about one another's behavior. This friction helps generate realistic business cycles. It also rationalize sizable fiscal multipliers without commensurate inflationary pressures.

Preview. The backbone of our model is a minimalistic general-equilibrium economy of the kind that are in the heart of the Arrow-Debrue and RBC frameworks. We simplify the exposition by assuming two periods ("today" and "tomorrow") and abstracting from investment. We abstract from any kind of nominal rigidity. We finally let an exogenous shock to the discount factor of the households proxy for shifts in consumer spending (or "aggregate demand").

[1] A detailed exposition of these points can be found in ???.

[2] Think of the empirical failings of old and new Philips curves and of the delicate assumptions that the DSGE literature has made in order to bypass these failings.

The representative-agent version of this model provides a stark benchmark with which we can compare our main results. In this benchmark, the aforementioned shock is predicted to have no effect on employment, output, and consumption. The reason is that, in the absence of investment, the general-equilibrium adjustment in the real interest rate and the real wage perfectly offsets the direct, or partial-equilibrium, effect of the shock on consumer spending and labor supply.[3]

We depart from this benchmark by populating the economy with large number of consumers and producers ("farmers"), who face a variety of realistic idiosyncratic shocks in addition to the aforementioned aggregate shock, and who interact in decentralized markets ("islands"). Each household ("family") contains a single farmer and multiple consumers. The farmer-member of any given family is matched to a particular island and produces only the good that is specific to that island, whereas the consumer-members are spread across many islands and consume many goods. This captures the fact that real-world households are specialists in production but generalists in consumption. It also lets an island represent a market between a firm and a set of consumers that are neither the owners nor the employees of that firm.

Firms have monopoly power, but prices are flexible: firms observe the realized demand in their market prior to setting their prices. To guarantee monetary neutrality and distinguish our contribution from the older literature on nominal confusion (**??**), we let all prices be denominated in real terms, relative to a composite of the goods produced and consumed in the second period.[4] We finally introduce the friction of interest by letting markets be informationally segmented in the following sense: the firms (respectively, the consumers) have perfect knowledge of the demand (respectively, the prices) in their own islands, but not necessarily of the relevant outcomes in other islands. And because the decisions and the trades that take place in one island depend on beliefs of the outcomes in other islands, the assumed friction amounts to introducing an imperfection in the coordination of economic decisions across markets.

Consider now an aggregate discount-rate shock that, other things equal, causes the consumers in some or all the islands to cut down on their spending. When the aforementioned friction is switched off, the shock fails to generate a drop in employment and output, exactly as in the representative-agent benchmark. But once the friction is switched on, this prediction is overturned and the Keynesian narrative starts making sense.

Holding constant the outcomes in other islands, the underlying shock triggers a drop in the demand faced by each farmer (or firm) in her own island. Because information is incomplete, some

[3] Adding investment makes things worse for the narrative under consideration: in general equilibrium, the decrease in consumption is accompanied by an increase in employment, output, and investment.

[4] An extension, however, introduces a nominal unit of account in order to offer a more realistic interpretation of our mechanism and to explore some interactions with monetary policy.

farmers fail to understand that the drop in their demand is due to aggregate rather than idiosyncratic reasons. As a result, these farmers fail to anticipate the adjustments that are likely to take place in other islands, find it optimal to work less, and instruct their consumer-siblings to spend less. But as these consumers spend less, other farmers experience a further drop in their demand. These farmers may now find it optimal to work less and to instruct their own siblings to spend less, even if they themselves are fully aware that the initial trigger was an aggregate discount-rate shock. An extra round of reduction in production, employment, and consumption therefore takes place. And so on.

A more realistic version of our theory replaces the farmers with collections of firms and workers (and adds labor markets). In response to the aggregate discount-rate shock, some firms see a decrease in the demand for their products and start hiring less. Some workers see wages go down, or unemployment go up, and start spending less. Additional firms then see their demand go down and respond by contributing to even less hiring. And so on. Our main analysis abstracts from investment; the logic, however, directly extends from the firms' incentives to hire to the firms' incentives to invest.

Another extension introduces government spending in order to let our theory speak to fiscal policy. By accommodating a mechanism akin to the Keynesian multiplier, our paper allows fiscal stimuli to have large effects. But unlike the New Keynesian model, the fiscal multipliers in our setting does not hinge either on monetary policy or inflationary pressures. Furthermore, whereas the New Keynesian model suggests that policy makers can leverage on these pressures and achieve a higher effect on current activity by backloading the fiscal stimuli, our theory favors the front-loading of fiscal stimuli. This is because the effects of fiscal stimuli in our setting depend on the firms experiencing an increase in their sales and the consumers experiencing an increase in their income, as opposed to them expecting high inflation and low real rates in the future.

Discussion. Our theory offers a simple account of two salient feature of the data. First, the bulk of the business cycle does not appear to be driven either by supply shocks, as proxied by TFP and labor productivity. Second, the bulk of the fluctuations in inflation are disconnected from the business cycle in the real quantities, which is at odds with the Keynesian mechanism.

These facts are illustrated Figure 1, which reports the scatterplots of the business-cycle component of output (on the horizontal axis) against the business-cycle components of hours worked, investment, consumption, TFP, labor productivity, and inflation (on the vertical axis).[5] The top three panels reveal the strong co-movement of the real quantities; the bottom three panels reveal the absence of

[5]The data are in quarterly frequency and cover the 1960-2015 period. Output is measured by GDP; hours worked by the hours of all persons in the non-farm business sector; consumption by the sum of personal consumption expenditures in nondurables goods and services; investment by the sum of personal consumption expenditures on durables goods, fixed private investment and changes in inventories; TFP by the utilization-adjustment measure provided in **?**; labor productivity by the ratio of GDP to total hours; and inflation by the change in the CPI index. The business-cycle components are obtained by applying the Band-Pass filter and isolating the frequencies corresponding to 6-32 quarters, as in **?**. The results are nearly identical if the HP filter is applied instead.

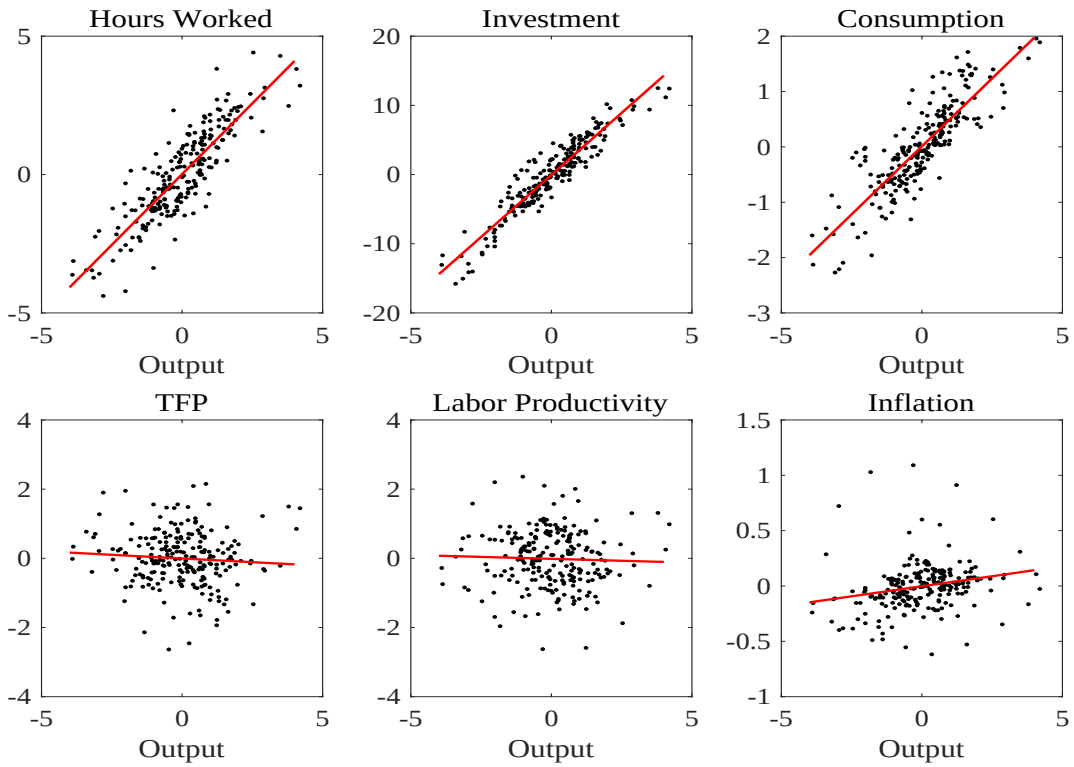

Figure 1: Business-Cycle Comovement with Output (or lack thereof).

commensurate co-movement with either TFP and labor productivity, or inflation. A similar picture is painted in Angeletos, Collard and Dellas (2015) by running a VAR on the key macroeconomic time series and by identifying a shock that accounts for up to 70% of the variation in hours or unemployment at business cycles frequencies, but less than 10% of the variation in TFP, labor productivity, and inflation at any frequency. And while it is certainly possible to accommodate these patterns in DSGE models featuring a multitude of shocks and propagation mechanisms, in our eyes these patterns indicate the value of theories that let the business cycle be disconnected from *both* productivity and inflation—which is precisely what our paper delivers.

To accomplish this, we assume that firms (or farmers) are unable to tell apart the sources of variation in their demand and similarly that consumers are unable to tell apart the sources of variation in their employment opportunities and income. This ingredient of our theory is reminiscent of **?**. There are, however, three important differences.

First, while the agents in that earlier work confuse nominal shocks for real shocks, the agents in our model confuse one kind of real shocks for another kind of real shocks. Our paper therefore has something entirely novel to say about the response of non-monetary economies to real shocks.

Second, the aforementioned decision-theoretic confusion is only part of the story in our paper;

the other part, which is absent from **?**, is the feedback loops between the firms and the consumers described above. Because these feedback loops are akin to strategic complementarity in "beauty contests," the confusion of *some* agents rationalize a similar behavior by *other* agents regardless of whether the latter are also confused or not. It is this part of our theory that helps formalizes the Keynesian multiplier and that lets small exogenous disturbances have large endogenous effects.

Last but not least, the confusion can but does not have to be the product of segmented market interactions and of missing public signals about the state of the economy. It can also be the product of rational inattention or, erratically but perhaps realistically, the product of bounded rationality.

Let us expand on this last point and on the broader empirical plausibility of our theory. In the real world, the bulk of the variation that the typical firm faces in her sales and profitability, or that the typical consumer faces in her income and employment opportunities, is driven by idiosyncratic forces. Through experience, or even selection, agents may have learned how to optimize their behavior in response to idiosyncratic forces. The main idea behind our contribution is that agents extrapolate from this kind of behavior when responding to aggregate shocks.

In our model, this kind of extrapolation is reconciled with rational expectations by allowing idiosyncratic shocks to be confounded with aggregate shocks and by removing common knowledge of the latter. This seems conceptually appealing, not only because of the methodological advantages of retaining the rational-expectations solution concept, but also because it helps us illustrate how the mistakes of one agent can rationalize the mistakes of another agent: when some agents engage in the aforementioned kind of extrapolation because they are either naive or informationally-constrained, other agents have the incentive to do the same even if they are themselves unboundedly rational and perfectly informed. That said, we invite the reader to adopt a flexible interpretation of the proposed mechanism, either in terms of informational frictions and higher-order uncertainty, or in terms of bounded rationality.

[TBC]