The Macroeconomics of the Financial Crisis and Its Aftermath: The Role of the Collapse and Restoration of Confidence*

Robert E. Hall
Hoover Institution and Department of Economics
Stanford University
rehall@stanford.edu; stanford.edu/∼rehall

July 1, 2018

Abstract

The notion that confidence varies over the business cycle has an important and growing role in macroeconomic theory. The volatility of the stock market, investment, and unemployment seems hard to understand without a powerful force that affects the willingness of investors to defer consumption to build plants, equip them, and create jobs. Though Keynes argued that confidence had a key role in the Great Depression and other manifestations of the business cycle, the idea that confidence mattered had received diminished attention in modern macroeconomics until the financial crisis of 2008. I develop a model suitable for understanding how the collapse of confidence influences key macro variables.The collapse enters the model as a large increase in the utility discounts of investor-consumers. The model matches three important events following the crisis: the sharp rise in unemployment, the large decline in investment, and the almost perfect stability of nondurables and services consumption, 61 percent of GDP.

*This research was supported by the Hoover Institution. It is part of the Economic Fluctuations and Growth Program of the National Bureau of Economic Research. Complete backup for all of the calculations and data sources will be available from my website, stanford.edu/∼rehall

1 Introduction

Big events provide the best evidence about the structure of highly interactive systems. The financial crisis of 2008 was a big event. Many macroeconomic theorists have proposed models of the crisis focusing on some notion of confidence as the centerpiece. Models view investors and consumers as losing confidence because their perceptions of the likelihood of future terrible outcomes jump (disaster models), because their beliefs about the beliefs of others shift downward, or because they become impatient. This paper is a contribution to the literature on the magnitudes of the driving forces for the crisis and their contributions to the large movements in key macro variables during and after the crisis, including consumption, investment, employment, stock-market values, and financial returns. Success in this exercise means that the model hypothesizes plausible driving forces, economic structure, and parameter values, and demonstrates that the model's description of the resulting movements of the key variables roughly matches the observed movements over the period from 2007 to 2017.

Parts of the model are shared with modern general-equilibrium modeling, derived from growth models. The labor market uses a simple version of the Diamond-Mortensen-Pissarides matching and bargaining model, while dropping the Nash-bargaining protocol in favor of one that can match the observed volatility of unemployment. I follow modern strands of financial-market economics by recognizing that large changes in financial discounts occur in general, and that, in particular, discounts rose to extreme levels during the financial crisis. These high discounts resulted in dramatic declines in investment in plant, equipment, and intellectual property and large declines in investment in job creation.

The model follows a number of others in characterizing the driving force as a substantial rise in the discount that consumer-investors apply to near-future utility. Thus, the paper does not contribute to the interesting and growing literature on the more fundamental mechanisms dealing the loss of confidence that obviously occurred in the fall of 2008. Rather, the paper posits a jump in the utility discount and maps it into the resulting financial discount in a general-equilibrium setting.

With one important exception, the model succeeds in this endeavor. It matches the rise in unemployment immediately after the crisis, along with the collapse of investment, including investment in consumer durables. It almost matches the remarkable fact that consumption of non-durables and services—61 percent of GDP—hardly responded to the crisis. The exception is that the stock market in the model falls by considerably less than the actual decline of 50 percent.

Christiano, Eichenbaum and Trabandt (2015) is an earlier paper with a much wider and somewhat different scope from this one. Their paper inherited the complicated structure and scope of the first two authors' major contribution to the New Keynesian model, with its focus on monetary driving forces. Among the many driving forces in the model, one that turns out to be particularly important—identified as a financial wedge—is related to the utility discount and its downstream counterpart, the financial discount, in this paper. However, the model does not include a stock market. The financial wedge is derived from the pricing of the yield spread between risky and riskless bonds. Philippon (2009) observes that bond spreads help measure valuations of rare events, because the payoffs of risky and safe bonds do not differ much except in fairly seriously adverse conditions, when the risky bonds default. But the authors do not interpret the yield spread in that way.

Farmer (2012) emphasizes the correlation between unemployment and the stock market. He lays out a model of the labor market in which demand determines employment and household accept any level of employment passively. His model differs in almost all respects from the one developed here.

Angeletos and Huo (2018) observe that agents with imperfect information who need to infer the future behavior of other agents will make decisions comparable to those of fully informed agents who have higher discount rates. They cite Farhi and Werning (2017) and Gabaix (2018) for similar theoretical reasoning. Angeletos and Lian (2016) survey a large recent literature on expectations with limited information, in which variations in sentiment or confidence generate fluctuations in economic activity resembling the business cycle.

This paper focuses on confidence as the central driving force of economic activity in the US in the decade following the financial crisis of 2008. For simplicity, it omits other driving forces. One is productivity. Fluctuations in productivity growth and in the size of the labor force are important contributors to overall movements in output at low frequencies but not at cyclical frequencies. Another is shifts in markups of prices over marginal cost, which are at the heart of the New Keynesian model of cyclical movements—see Rotemberg and Woodford (1999). But measured shifts, though large, are at frequencies well below the business cycle. The attention of the macro profession has shifted to driving forces characterized as financial frictions, variable discounts, or fluctuations in disaster prospects.

2 Financial Driving Forces

The model developed in this paper takes the discount embedded in investors' utility functions as the driving force accounting for the collapse of investment and job creation following the financial crisis. Another financial driving force considered in recent macro-finance models is the perceived probability or magnitude of a disaster—an event that results in a substantial drop in consumption. This section explores these alternatives in a simple model, derived from Lucas (1978).

The economy lasts for two periods. An investor has a consumption endowment of 1 unit in the first period. In the second, the endowment is random—1 unit in state N (normal) and $1-q$ units in state B (bad). State N arises with probability $1-\pi$ and state B is a bad state that arises with probability π. Equilibrium establishes *state prices*, p_N and p_B, denominated in terms of first-period output, for the receipt of one unit of value in each of the future states. Utility is

$$u(1) + \frac{1}{1+\rho}[(1-\pi)u(1) + \pi u(1-q)]; \quad (1)$$

ρ is the *utility discount rate*. The first-order conditions for the household's portfolio are:

$$\frac{1}{1+\rho}(1-\pi) = p_N \quad (2)$$

and

$$\frac{1}{1+\rho}\pi\mu = p_B; \quad (3)$$

$\mu > 1$ is the ratio of marginal utility in state B, $u'(1-q)$, to marginal utility in period 1, $u'(1)$, The ratio exceeds 1 because state B is a bad state with lower consumption than in period 1. Because the endowments are predetermined, the first-order conditions determine p_N and p_B.

Although investment is not feasible in this endowment economy, the state prices measure the payoff to a hypothetical investment. Specifically, the payoff, measured as of period 1, to a risky payoff that delivers one unit of output in period 2 under state N and $1-q$ units of output under state B, is

$$\frac{1}{1+\rho}[1-\pi+\pi\mu\cdot(1-q)]. \quad (4)$$

The factor in brackets is positive, so the payoff is unambiguously a decreasing function of ρ. The relation of the payoff to the disaster probability π is ambiguous, because $\mu\cdot(1-q)$ could exceed one. And μ is an increasing function of q. With a constant coefficient of relative risk aversion of γ, $\mu = (1-q)^{-\gamma}$, so $\mu\,(1-q) = (1-q)^{-(\gamma-1)}$. Under the reasonable assumption that $\gamma > 1$, $\mu \cdot (1-q)$ is an increasing function of q and the payoff to the investment is correspondingly a an increasing function of q. Because $q > 0$, $(1-q)^{-(\gamma-1)} > 1$ and the payoff is an increasing function of the disaster probability π.

To summarize in terms of a driving force that would result in a collapse of investment, under the assumption that $\gamma > 1$, the incentive to invest will decline if

- the utility discount rate ρ rises
- the disaster shortfall of consumption, q, falls
- the disaster probability π falls

The model's approach to risk is to consider the discounts applied to risky payoffs. Thus the model does not have an explicit stochastic discounter, but rather applies the principle embodied in the CAPM and related financial models, where each stochastic payoff has its own discount rate reflecting its covariance with the stochastic discounter. A basic principle of finance holds that, in financial equilibrium, the discount rate for a risky investment equals the expected return to that investment. Note that the financial discount is a feature of general equilibrium and is not necessarily the same as the utility discount.

In the two-period, two-future-state example, the expected return ratio for an investment that pays off one unit of value in the second period in state 1 and $1-q$ units in state 2 is

$$R = \frac{1-\pi+\pi(1-q)}{p_N + p_B(1-q)}. \tag{5}$$

Under the assumption from before that marginal utility is one in state N and $(1-q)^{-\gamma}$ in state B, the expected return ratio is

$$R = (1+\rho)\,\frac{1-\pi+\pi(1-q)}{1-\pi+\pi(1-q)^{-(\gamma-1)}}. \tag{6}$$

The financial discount rate is

$$r = R - 1, \tag{7}$$

so, if $\pi = 0$ or $q = 0$, the financial rate would be the utility discount rate, ρ. One of the reasons that the financial discount rate differs from the utility discount rate is risk, which exists if $q > 0$ and $\pi > 0$. Another reason, not brought out in this example, is expected growth in consumption with consequent expected shrinkage in marginal utility.

In equation (6), three financial determinants could be driving forces determining the financial discount rate: the utility discount rate, ρ, the disaster probability π, and the disaster consumption shortfall, q. The first two are taken as exogenous in the general-equilibrium model developed later in this paper, while the consumption level is an equilibrium outcome in that model. Exogenous fluctuations in the utility discount rate ρ and in the perceived probability of a disaster have both appeared in GE models as driving forces. The model in this paper considers ρ as the exclusive driving force. Note that increases in the disaster probability π and magnitude q *lower* the financial discount rate because—under the assumption that $\gamma > 1$—the numerator in equation (6) falls with π or q while the denominator rises, so R and thus r declines as π or q rise.

The model portrays the utility discount factor as fluctuating smoothly in continuous time. Its change is governed by the instantaneous discount rate ρ—the discount applied to time-t utility as of time zero is

$$\beta(t) = \exp\left(-\int_0^t \rho(\tau)d\tau\right). \tag{8}$$

The discount is not influenced by risk in the current version of the model.

3 Data

I consider five basic time series describing the evolution of the US economy from the end of 2007 through the end of 2017:

3.1 Consumption

I take real consumption of nondurables and services from the US national income and product accounts as the measure of consumption; I treat consumer durables as part of investment. I normalize and detrend the series as an index beginning and ending at one, to compare to the output of the model which has no trend. The top line in Figure 1 shows the result. Consumption of nondurables and services responded only slightly to the financial crisis. Those categories of consumption accounted for 61 percent of GDP in 2017.

3.2 Investment

The middle line in Figure 1 shows real fixed investment, including investment in intellectual property and housing, plus consumer expenditures on durables, detrended and normalized.

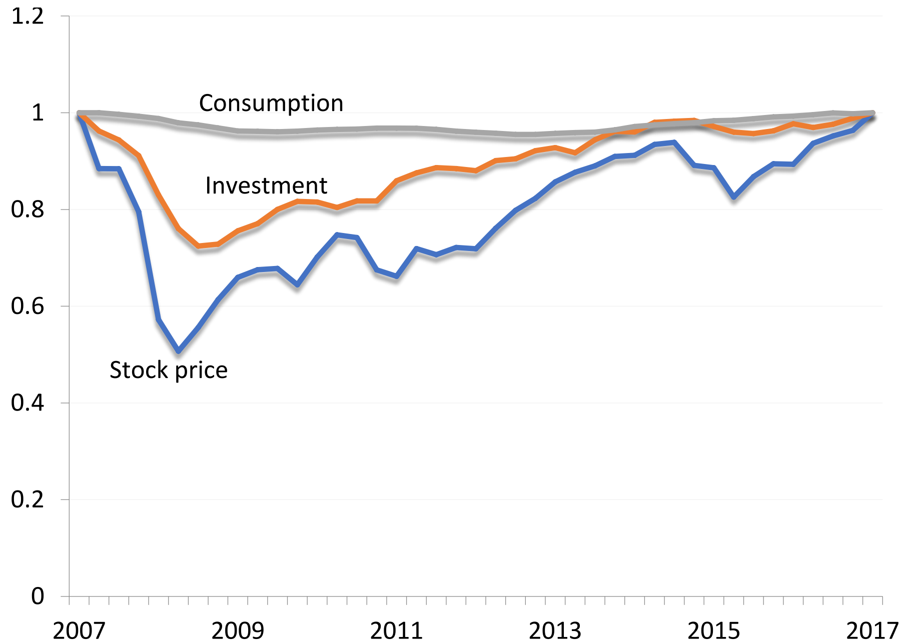

Figure 1: Normalized and Detrended Data for Consumption, Investment, and Stock Price, 2007 through 2017

Investment declined by about 25 percent immediately after the crisis and then rose gradually, relative to trend.

3.3 The stock market

The bottom line in Figure 1 shows the value, excluding compounded dividends, of the Wilshire 5000 index of the value of the US equity market, deflated by the NIPA price index for consumption. The stock market fell about 50 percent and gradually recovered, relative to trend. Again, the trend is measured by the growth of the market in real terms from 2007 through 2017.

3.4 The employment rate

The upper line inn Figure 3 shows the employment rate, defined as the fraction of the labor force who are employed. It is measured as one minus the unemployment rate. The line is not normalized or detrended, as unemployment is a remarkably stationary series. Employment fell rapidly by 5 percentage points, then slowly grew back to normal.

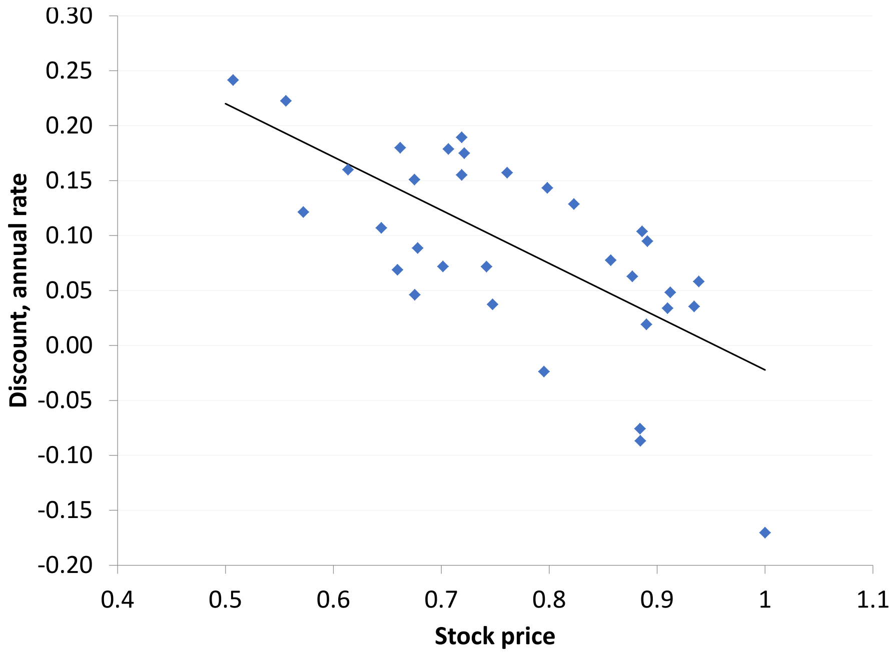

Figure 2: Campbell-Shiller Analysis of Expected Returns in the Stock Market

3.5 The discount rate implicit in the stock market

Campbell and Shiller (1988) and Fama and French (1988) pioneered the extraction of the implicit discount rate in the stock market. Their idea started from the observation that the discount is also the expected return. The level of the stock market is a powerful indicator of the expected return. Figure 2 shows an application of the principle using data from the crisis and its aftermath. The horizontal axis is the stock-price index shown earlier. The vertical axis is the 2-year annualized realized real return rate $r_t = (\tilde{p}_{t+2}/\tilde{p}_t)^{1/2} - 1$, where $\tilde{p}_t$ is the Wilshire 5000 price index including compounded dividends, deflated by the consumption price index. The downward sloping straight line is the regression line, $d_t = 0.46 - 0.48 p_t$. The standard error of the slope is 0.09. I calculate the discount d_t from this equation and show it as the lower line in Figure 3. The discount rate peaked just after the crisis and then gradually declined.

4 Model

The model describes a standard one-sector economy in continuous time with capital adjustment costs and labor-market equilibrium according to the principles of Diamond, Mortensen, and Pissarides. Everybody is in the labor market—the only variation in employment, n,

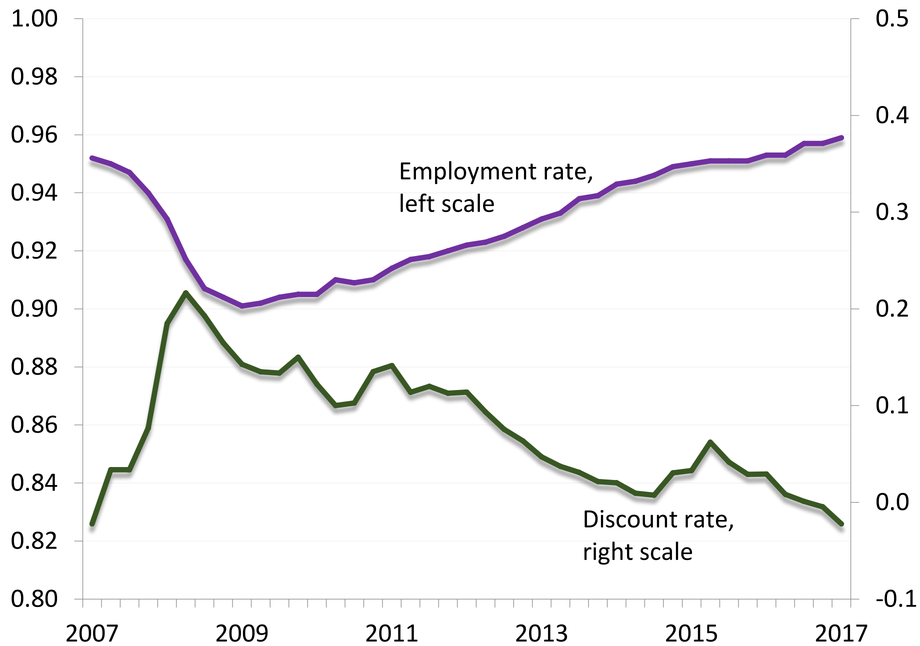

Figure 3: Data for the Employment Rate and the Discount Rate for the Stock Market, 2007 through 2017

arises from variations in unemployment. Production is Cobb-Douglas with capital elasticity α. Capital depreciates at rate δ. Consumption is c. The law of motion of the capital stock, k is

$$\dot{k} = k^{\alpha} n^{1-\alpha} - \delta k - c. \tag{9}$$

Investment incurs quadratic adjustment costs. The supply of installed capital relates to Tobin's q, with cost κ, as

$$\dot{k} = \frac{k}{\kappa}(q - 1). \tag{10}$$

The demand for installed capital satisfies

$$\dot{q} = q(r + \delta) - \alpha \left(\frac{n}{k}\right)^{1-\alpha}. \tag{11}$$

Infinitely-lived households spread consumption over time according to an Euler equation with an elasticity of intertemporal substitution σ and a time-varying utility discount rate ρ:

$$\dot{c} = \sigma(r - \rho)c \tag{12}$$

Here r is the economy's time-varying financial discount.

With respect to the labor market, employers place a value J on an employee, on account of the recruiting cost involved in adding one to the employer's work force. J is the present value of the future stream of value the worker will contribute, which is the difference between the worker's marginal product and the worker's wage. I normalize that difference at one:

$$\dot{J} = (r + s)J - 1 \tag{13}$$

According to standard DMP principles, J determines the flow rate of vacancy-filling, which determines the tightness of the market, which in turn determines the job-finding rate, which determines the unemployment rate, and the employment rate is the complement of the unemployment rate, as noted earlier. I summarize this chain by a linear relation between J and n, with intercept $\bar{n}$ and non-negative slope ϕ:

$$n = \bar{n} + \phi J. \tag{14}$$

In the DMP model, the vacancy/unemployment ratio θ serves as the index of labor-market tightness. The vacancy-filling rate for employers, $q(\theta)$, is a decreasing function of tightness and the job-finding rate for the unemployed, $f(\theta)$, is an increasing function of tightness. Potential employers expand vacancy postings to the point that the flow value from a vacancy, $q(\theta)J$, equals the flow cost of maintaining a vacancy, taken to be constant. Consequently, a higher job value results in a tighter labor market: Tightness θ is an increasing function of the job value J. Thus the job-finding rate $f(\theta)$ increases in tightness—jobs are easier to find in a tighter market. The inflow to employment is the product of the job-finding rate and the volume of unemployment, $f(\theta)(1 - n)$, and the outflow is the product of the separation rate and employment, $s\,n$. In stochastic equilibrium, the inflow matches the outflow, so the stochastic equilibrium employment rate is

$$n = \frac{f(\theta)}{s + f(\theta)}, \tag{15}$$

an increasing function of θ and thus of the job value J.

The single driving force in the model is ρ, the utility discount. It begins at a high level, ρ_0, then returns exponentially at rate ω to its normal level, ρ^*:

$$\dot{\rho} = -\omega(\rho - \rho^*). \tag{16}$$

The only place that ρ appears in the model is in the Euler equation (12). Variations in ρ are much like variations in a financial friction or wedge that limits the return received by savers relative to the return earned by capital. Thus the qualitative results in this paper are similar to those from a financial friction. To my knowledge, work in the financial-friction setup has not taken advantage of the stock market as a direct source of information about expected returns.

The appendix describes solving the model. It reduces to a three-dimensional system of ordinary differential equations in q, k, and n, with three boundary conditions. The system can be solved easily with standard software. The solution is exact.

5 Parameters

Five of the parameters are uncontroversial. These are $\alpha = 0.4$ (boosted to account for the inclusion of consumer durables in capital), $\delta = 0.1$, $s = 0.18$, $\rho^* = 0.05$, and $\bar{n}$ chosen to equate the stationary value of employment, n, to its long-run average value of 0.945.

Three other parameters are important in the model and less well pinned down by data and research. I discuss results for a base case and for a variant, for each of the three.

The first is the adjustment cost parameter, κ. A standard value at an annual rate from research based on the first-order condition for optimal investment is 2, but the results in this paper point toward a considerably higher value, 8. Values this high have been found in research based on direct estimation of equation (10) using data on firms' market values. I use $\kappa = 2$ as a variant.

The second is the intertemporal elasticity of substitution, σ. The results here point toward a low value of 0.2. I use 0.5 as a variant.

The third is the sensitivity ϕ of the employment rate n to the job value J. The results here point toward a value of 0.05. Shimer (2005) found that the parameter values in Mortensen and Pissarides (1994) made unemployment insensitive to the driving force of productivity because the Nash bargain resulted in wage movements that almost fully offset productivity movements. In the standard DMP setup, this finding implies that J is insensitive to productivity. Hall (2017) showed that Shimer's finding carries over to discounts as a driving force. I use $\phi = 0$ as a variant, interpreted as a growth model without employment volatility.

6 Results

6.1 Financial Discount Rate

The model contains two distinct discount rates.

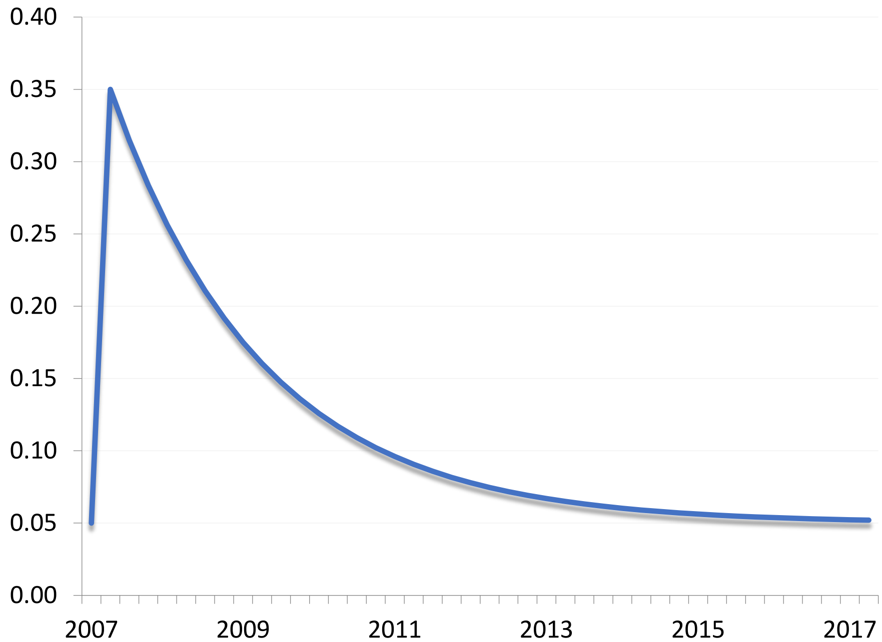

Figure 4: The Utility Discount Rate, Driving Force of the Model

I also use the term financial *discount rate* in the same way as in the title of Campbell and Shiller (1988). It is the asset-specific rate r that discounts future expected payoffs back to the present. It is also the expected return to the asset. I measure this financial discount as the expected return from a comprehensive portfolio of US traded equities. I assume that the discount rates are roughly the same for other assets—the marginal product of capital in the equation for Tobin's q and the discount for the job value J. Both of these payoffs are subject to random determinants similar to those perturbing the cash flows discounted in the stock market. The consumption Euler equation (12) shows that the utility discount and the financial discount are the same when consumption is constant over time. With consumption growth, the financial discount is higher than the utility discount.

Figure 4 shows the driving force of the model, the utility discount rate. The loss of confidence associated with the financial crisis results in a sharp spike of the utility discount, from its normal rate of 5 percent to a peak of 35 percent. Thereafter, the discount declines back to normal.

Figure 5 shows the response of the financial discount rate in the model and the values measured from the stock market. The financial rate spiked at the time of the crisis but by less than the utility discount rate, according to the model. During the period following the

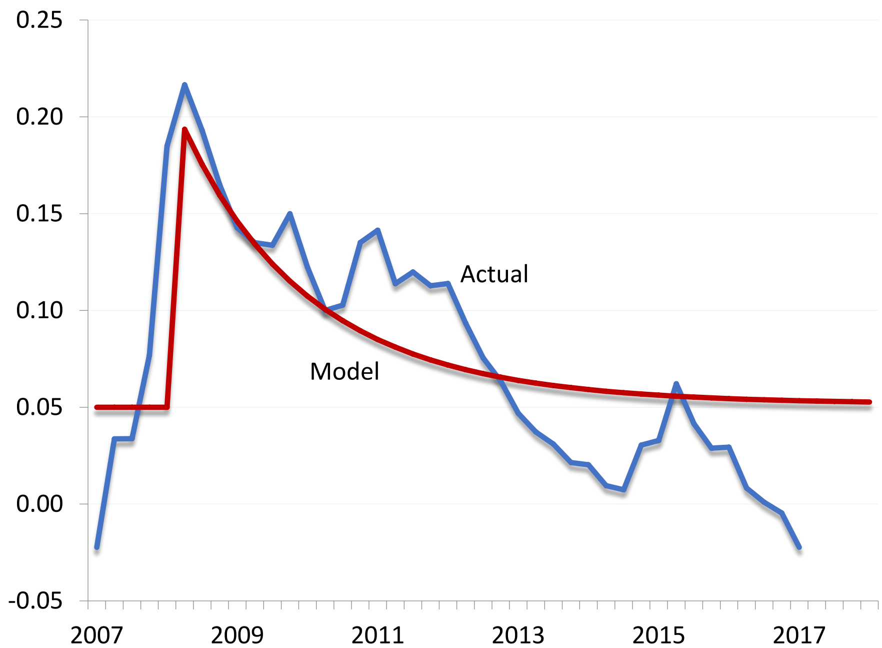

Figure 5: Financial Discount Rate

crisis, the financial rate fell short of the utility rate, so consumption declined during that period.

6.2 The stock price

The value of the stock market in the model is the product, $q\,k$, of the market price of installed capital and the quantity of capital (see Hall (2001) for a discussion of this view of the stock market) plus the value of the established relationships with workers, $n\,J$. Figure 6 shows the values from that calculation along with the detrended real stock price shown earlier. The model's version captures the rapid decline and gradual recovery of the price, but the decline is less than half of the decline that actually occurred. One important fact in interpreting this finding is that the level of the stock market in 2007, and even more in 2017, much exceeded empirical counterparts of $q\,k$—Tobin's q measured as the ratio of the market value of financial claims on traded corporations to the current market price of replacing their plant, equipment, and intellectual property, was around 2 in 2007 and higher in 2017. The stock market is considerably more volatile—in both directions—than any reasonable version of $q\,k$. In the 1970s, empirical q fell to about 0.4, inexplicably low by the standards of $q\,k$. The inclusion of intellectual property in the calculations only somewhat lessens the mystery

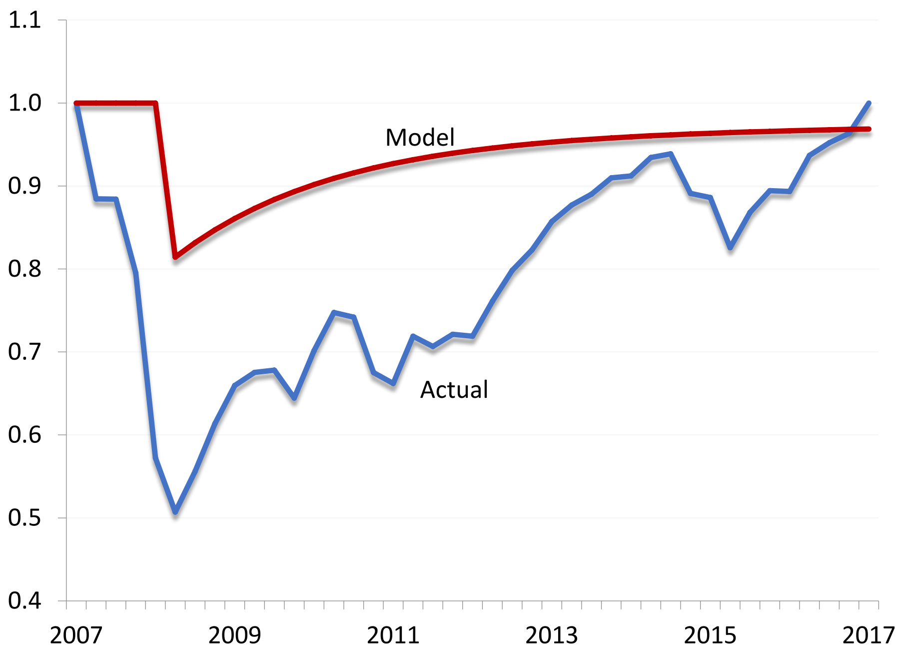

Figure 6: Stock Price

of periods of high empirical q and only deepens the mystery of extremely low values of q in other periods.

6.3 The Employment Rate

Figure 7 compares the model's implication for the employment rate (one minus the unemployment rate) to the actual employment rate. According to the model, the high financial discount rate eroded incentives for job creation sufficiently to drive the employment rate down to about 0.90, roughly its actual value in the worst year, 2009. The model does not capture the slow recovery of the employment rate in the years 2010 through 2012. With the sharp declines in the financial discount rate shown in Figure 5, the model infers a speedy recovery in employment during those years, but the actual speeding up occurred after 2012. One reason for the lag of actual employment behind the improvement in job-creation incentives is the neglect in my stripped-down version of the DMP model of its internal dynamics. Whereas market tightness is a jump variable in the model, unemployment is a state variable that responds to changes in tightness with a lag. Although the lag appears to be short in DMP models that consider only labor-market-wide job-finding rates, a disaggregation that recognizes groups with low job-finding rates has much longer lags.

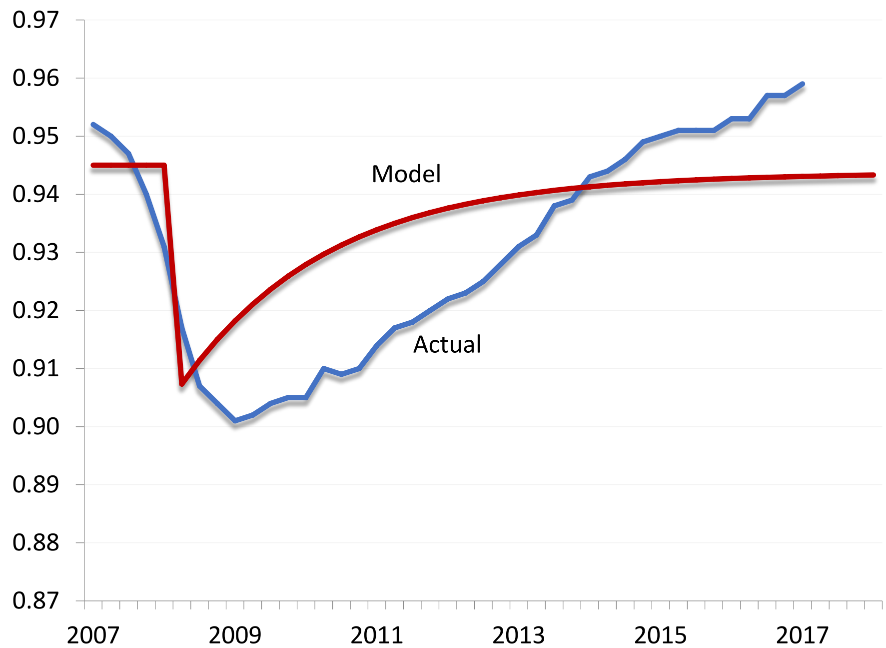

Figure 7: Employment Rate

6.4 Consumption

Figure 8 compares the paths of consumption over the crisis and aftermath in the model to the actual data. Recall that the data capture non-durables and services, a measure suitable for an Euler equation. The important fact about the data is that the crisis left the path of this measure of consumption essentially untouched. The model, on the other hands, suffers from a mild version of the problem identified in Barro and King (1984)—quantitative general-equilibrium models tend to infer an increase in consumption when an adverse shock hits. The reason is that the shock diminishes the incentive to invest, so it frees up resources for consumption. The jump in consumption is consistent with a subsequent decline in consumption associated with a financial discount rate that is below the utility discount rate.

The upper line in the figure shows that the model has a small version of the Barro-King problem. Consumption jumps by about 4 percent upon the crisis shock, then declines during the period when the financial discount is low relative to the utility discount. Some further fine-tuning may be able to eliminate the model's failure to replicate that fact that the crisis was a complete non-event for a category of spending that is more than half of GDP. More fundamentally, if the model recognized the role of households whose Euler equations

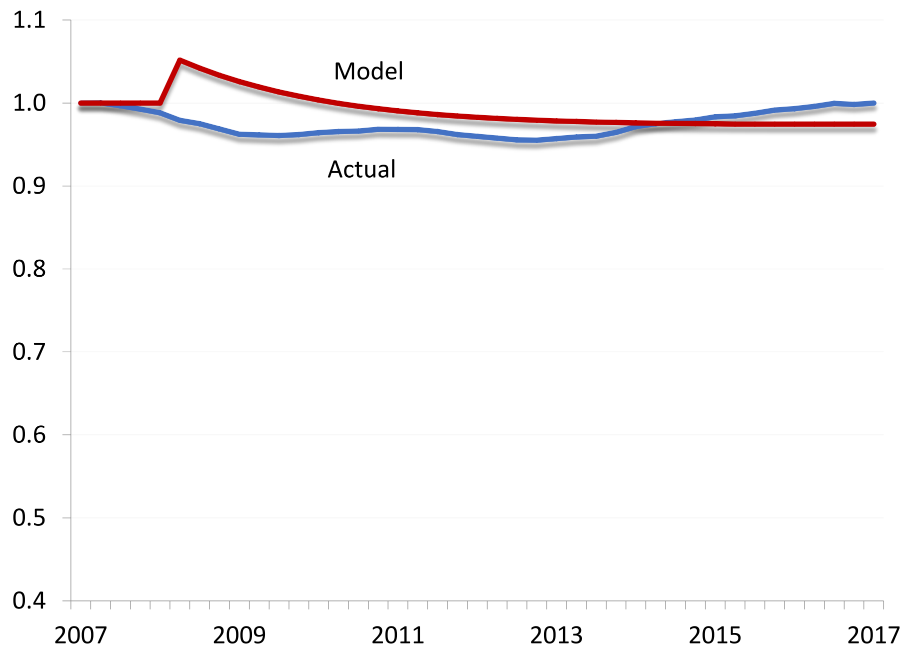

Figure 8: Consumption

have large and highly variable terms arising from borrowing constraints, the tightening of those constraints in the crisis would result in an offset to the bulge in the consumption of households with stable Euler equations.

6.5 Investment

Figure 9 shows that the model tracks the huge decline in investment, including consumer durables investment, that the crisis triggered. Two factors account for the decline, both visible in the capital-demand equation (11). One is the large increase in the financial discount, r. The other is the decline in the marginal product of labor resulting from the diminution of employment, n.

7 Alternative Values of Key Parameters

Table 1 summarizes the findings in variants of the base case. The figures are the values of the five key variables the moment after the crisis hits. The leftmost column shows that actual values. For the first three variables, the stock price, investment, and consumption, the values just prior to the cdrisis are normalized at one. For employment, the prior value

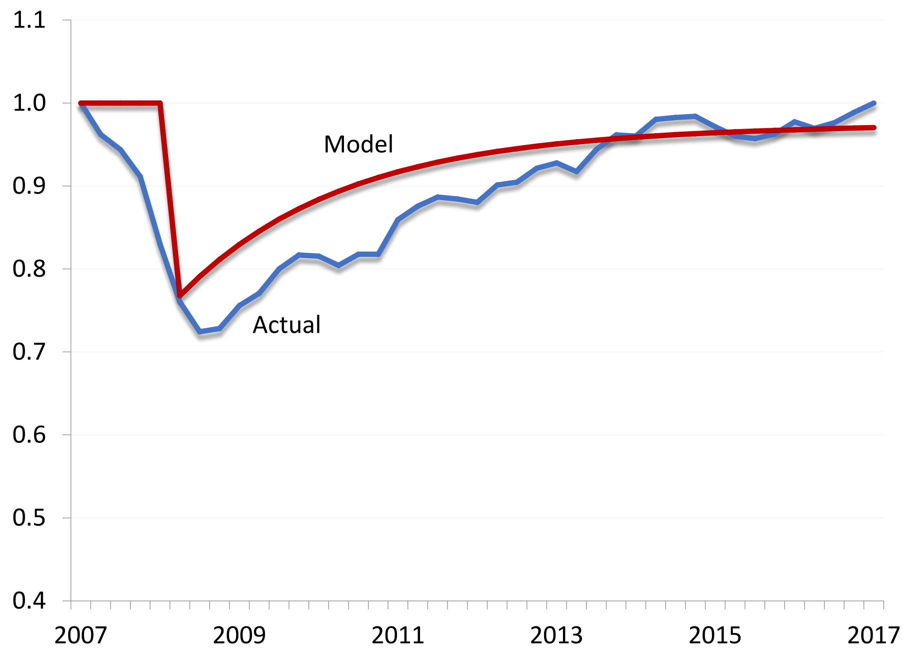

Figure 9: Investment

is 0.945, the long-run average of one minus the unemployment rate. For the discount rate, the prior value is 0.05.

7.1 Lower capital adjustment cost

The second column from the left gives the immediate effects for the base case already discussed. The next column describes an economy identical to the one in the base case except that the capital adjustment cost is 2 instead of 8. The resulting post-crisis stock price is even further from the actual value of 0.50 that in the base case, because Tobin's q is much

	Actual	*Base model*	*Lower adjustment cost*	*Growth model*	*Higher intertempral substitution*
			$\kappa = 2$	$\varphi = 0$	$\sigma = 0.5$
Stock price	0.50	0.82	0.90	0.85	0.74
Investment	0.73	0.77	0.74	0.81	0.67
Consumption	1.00	1.05	1.08	1.07	1.07
Employment rate	0.910	0.907	0.923	0.945	0.892
Discount rate	0.20	0.19	0.13	0.17	0.27

Table 1: Results for Variants of Parameter Values

less sensitive to capital demand than in the base case. Investment is correspondingly more sensitive, falling to 0.74 compared to 0.77 in the base case. The Barro-King problem is more pronounced with milder adjustment cost, because the larger downward movement in investment frees up more resources than in the base case. At 0.923, the employment rate falls much less than in the base case, because, as the bottom line shows, the effect of the crisis on the financial discount rate is quite a bit smaller than in the base case. As the larger initial increase in consumption indicates, expected declines in the financial discount are higher, and thus the discount increase is smaller, than in the base case. The overall conclusion from these results is that fairly high adjustment costs are central to the success of the base case in matching the actual behavior of the key variables, apart from the stock price.

7.2 Constant employment rate

The next column examines the importance of the induced movements of the employment rate, by setting the sensitivity parameter ϕ to zero, so the model becomes like a traditional growth model without fluctuations in employment. Again, the match to the actual data worsens relative to the base case. The response of the stock market is even smaller. Investment does not fall as much. The Barro-King bulge in consumption is bigger. By construction, there is no movement in employment, compared to a fall of 4.5 percentage points in the data and a bit more in the base case of the model. And the discount rate rises to 0.17, compared to 0.20 in the data and 0.19 in the base case. The conclusion is that an economic model of unemployment, relating the employment rate to the incentives to create jobs, makes an important contribution to matching the actual behavior of the economy during and after the crisis.

7.3 Higher intertemporal elasticity of substitution

The rightmost column of Table 1 considers a higher value of the elasticity of intertemporal substitution, σ, of 0.5 in place of 0.2 in the base case. This value improves the poor match to the contraction of the stock market in immediate response to the crisis, with a value of 0.76, compared to the actual value of 0.50 and a value of 0.83 in the base case. But higher intertemporal substitution overstates the decline in investment, at 0.67, compared to 0.73 in the data. And the Barro-King problem in consumption is more acute, at an increase resulting from the crisis to 1.06, compared to no increase in the data and an increase to 1.04

in the base case. Because, as shown in the bottom line, the rise in the financial discount to 0.27 overstates the actual increase to 0.20, this version of the model also overstates the depression of the employment rate on account of the crisis. The conclusion is that a relatively low elasticity of intertemporal substitution is important for matching the economic events triggered by the financial crisis.

8 Concluding Remark

The results in this paper support the idea that macroeconomics can understand the behavior of the US economy during and after the financial crisis. Keynes looms over macroeconomics for his belief that a collapse in confidence was an important part of the story of the Great Depression and other macro fluctuations—see the famous chapter 12 of Keynes (1936), the *General Theory*. In that respect, the ideas here are Keynesian, just not New Keynesian. Many recent papers try to capture the notion of confidence within modern theoretical frameworks. This paper does not advance that topic. Its reliance on movements of the utility discount as a way to capture fluctuations in confidence is not novel. The paper does advance quantitative understanding of how fluctuations in the utility discount make their way into the financial discount and thus into job creation and investment, but not into consumption, apart from purchases of durables.

References

Angeletos, G.-M. and C. Lian, "Chapter 14 - Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination," in John B. Taylor and Harald Uhlig, eds., *John B. Taylor and Harald Uhlig, eds.*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier, 2016, pp. 1065 – 1240.

Angeletos, George-Marios and Zhen Huo, "Myopia and Anchoring," Working Paper 24545, National Bureau of Economic Research April 2018.

Barro, Robert J. and Robert G. King, "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *The Quarterly Journal of Economics*, 1984, *99* (4), 817–839.

Campbell, John Y. and Robert J. Shiller, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, *1* (3), 195–228.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt, "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 110–167.

Fama, Eugene F and Kenneth R French, "Dividend yields and expected stock returns," *Journal of Financial Economics*, 1988, *22* (1), 3 – 25.

Farhi, Emmanuel and Iván Werning, "Monetary Policy, Bounded Rationality, and Incomplete Markets," Technical Report, National Bureau of Economic Research 2017.

Farmer, Roger EA, "The stock market crash of 2008 caused the Great Recession: Theory and evidence," *Journal of Economic Dynamics and Control*, 2012, *36* (5), 693–707.

Gabaix, Xavier, "A Behavioral New Keynesian Model," March 2018.

Hall, Robert E., "The Stock Market and Capital Accumulation," *American Econmic Review*, December 2001, *91* (5), 1185–1202.

——, "High Discounts and High Unemployment," *American Economic Review*, February 2017, *107* (2), 305–330.

Keynes, John Maynard, *The General Theory of Employment Interest and Money*, New York: Harcourt, Brace and Company, 1936.

Lucas, Robert E., "Asset Prices in an Exchange Economy," *Econometrica*, 1978, *46* (6), 1429–1445.

Mortensen, Dale T. and Christopher Pissarides, "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 1994, *61*, 397–415.

Philippon, Thomas, "The Bond Market's q," *Quarterly Journal of Economics*, August 2009, *124* (3), 1011–1056.

Rotemberg, Julio J. and Michael Woodford, "The Cyclical Behavior of Prices and Costs," 1999. in Taylor and Woodford, eds., *Handbook of Macroeconomics*, Chapter 16, pp. 1051–1135, 1999.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 2005, *95* (1), 24–49.

Appendix: Solving the Model

The canonical form for an ODE for a two-point boundary-value solver is

$$\dot{y} = f(y), \tag{17}$$

for a vector y. The model fails to adhere to the canonical form in two respects. First, it has two equations with $\dot{k}$ on the left side, which imposes a restriction on the joint levels of the variables. Second, it contains the latent variable r.

To deal with the first issue, I equate the right-hand sides of equation (9) and equation (10) to get

$$c = k^{\alpha} n^{1-\alpha} - \delta k - \frac{k}{\kappa}(q-1). \tag{18}$$

I refer to this formula as $\mathbb{C}(k, q)$. And I refer to the formula for its time derivative as $\dot{\mathbb{C}}(k, \dot{k}, q, \dot{q}, n, \dot{n})$.

For simplicity, I also combine equation (13) and equation (14) to form

$$\dot{n} = (r+s)(n-\bar{n}) - \phi. \tag{19}$$

The model becomes

$$\dot{k} = \frac{k}{\kappa}(q-1) \tag{20}$$

$$\dot{q} = q(r+\delta) - \alpha\left(\frac{n}{k}\right)^{1-\alpha} \tag{21}$$

$$\dot{\mathbb{C}} = \left[\sigma(r-\rho) + \gamma\frac{\dot{n}}{n}\right]\mathbb{C}(k, q) \tag{22}$$

$$\dot{n} = (r+s)(n-\bar{n}) - \phi. \tag{23}$$

These are four equations in the four variables $\dot{q}, \dot{k}, \dot{n}$, and r. I use an equation solver to find their values, given $y = [q, k, n]$. I deliver the first three values as $\dot{y}$ to the two-point boundary-value ODE solver.