Discounting: Investment Sensitivity and Aggregate Implications*

Russell Cooper†and Jonathan L. Willis‡

December 12, 2014

Abstract

This paper studies the effects of discounting on plant-level and aggregate investment. We study a number of date based processes to represent the state dependent discount factor. Empirically, the stochastic discount factor is procyclical. The investment decision at the plant level is sensitive to the specification of the stochastic discount factor. Non-convexities in adjustment costs at the plant level have aggregate implications: lumpy investment is not smoothed.

1   Motivation

This paper studies the dependence of plant-level and aggregate investment on the stochastic discount factor when adjustment costs are non-convex. It asks two questions. First, what are the effects of alternative representations of the stochastic discount factor on plant-level investment? Second, do non-convexities at the plant-level have aggregate implications?

The specification of the stochastic discount factor (SDF) is significant: plant-level investment behavior is sensitive to the SDF. Model-based specifications are at odds with the SDF inferred from data. Our analysis highlights these differences and studies their implications for plant-level and aggregate investment.

As is now understood from a number of studies, investment at the establishment level is characterized by periods of only minimal changes in the capital stock coupled with intermittent

periods of large capital adjustments. These patterns are difficult if not impossible to mimic in the standard quadratic adjustment cost model. Instead, these patterns are captured in models which rely on the presence of non-convex adjustment costs.[1]

The aggregate implications of these plant-level nonconvexities though remain in dispute. Following the lead of Thomas (2002), Khan and Thomas (2003) and Khan and Thomas (2008), one might conjecture that the non-convexities at the plant-level are not important for aggregate investment. In those papers, state dependent interest rates are determined in equilibrium. A striking result from this literature is that the absence of aggregate effects of lumpy investment: the aggregate model with non-convexities at the micro-level is essentially indistinguishable from an aggregate model without non-convexities. A key to the result, as noted by Thomas (2002), is the response of the interest rate to aggregate shocks and the evolution of the capital stock.[2]

Those results, however, do not address the question we raise here. The issue is the process for the stochastic discount factor (interest rate). Our focus is on the effects of the stochastic discount factor inferred from the data, rather than that created in the underlying equilibrium of a stochastic business cycle model. It is entirely conceivable that the equilibrium smoothing of plant-level investment through interest rate movements does not arise in a model economy with an empirically relevant stochastic discount factor.

The problem is that the interest rate process from the standard RBC model does not match the data well, as discussed, for example, in Beaudry and Guay (1996).[3] Thus, the smoothing of lumpy investment through interest rate movements produced in these models might be both theoretically of interest and model-consistent, but not empirically based.

Our analysis is conducted in a model with two key features. The first, as in Cooper and Haltiwanger (2006), is the presence of non-convexities in the capital adjustment process. The model of plant-level adjustment costs is taken from Cooper and Haltiwanger (2006). As discussed below, the specification of adjustment costs includes an opportunity cost which is independent of the magnitude of investment, a quadratic adjustment cost and a form of irreversibility.

The second feature is a SDF that is empirically based rather than the outcome of a particular stochastic equilibrium model. A number of empirically based representations of the SDF are considered.

One approach begins with a specification of household utility and to infer the stochastic discount factor from a household Euler equation. In this case, the SDF is mildly countercycli-

---

[1]See, for example, the results reported in Caballero and Engel (1999) and Cooper and Haltiwanger (2006).

[2]Cooper and Haltiwanger (2006) find some, but not complete, smoothing by aggregation arising from idiosyncratic shocks when interest rates are held constant.

[3]This point appears in Thomas (2002) as well: Table 5 indicates a correlation of -0.385 between the real interest rate and output in the data but a correlation of 0.889 in the benchmark model.

cal, with minimal variability.[4] The version of the SDF is similar to that obtained from the standard RBC model.

In contrast, the SDF can be inferred directly from asset prices from portfolio choice based upon standard orthogonality conditions. This version of the SDF is strongly procyclical and variable. The behavior of plant-level and aggregate investment is very different for this version of the SDF compared to the one inferred from the household Euler equation.

The results indicate that investment behavior at the plant-level is sensitive to the specification of the SDF. Further, non-convexities at the plant-level have aggregate implications for empirically consistent stochastic discount factors. These effects are particularly prominent when the SDF is inferred from portfolio choices.

2　Dynamic Capital Demand

Our approach is to begin with a dynamic capital demand problem at the plant level, following Cooper and Haltiwanger (2006) for the specification of the adjustment costs. But, in contrast to that analysis, we allow the discount factor to be state dependent rather than constant. The stochastic discount factor is inferred from data.

2.1　Plant Level: Dynamic Optimization

The dynamic programming problem is specified as:

$$V(\varepsilon, k, Z) = \max\{V^i(\varepsilon, k, Z), V^a(\varepsilon, k, Z)\}, \qquad \forall (\varepsilon, k, Z) \tag{1}$$

where ε is the idiosyncratic productivity shock, k represents the beginning of period capital stock at the plant, and Z is a vector of aggregate variables, including aggregate productivity, A. The superscripts refer to active investment "a," where the plant undertakes investment to obtain capital stock k' in the next period, and inactivity, "i," where no gross investment occurs.

The options in (4) are defined by:

$$V^i(\varepsilon, k, Z) = \Pi(\varepsilon, k, Z) + E_{\varepsilon', Z'|\varepsilon, Z}\left[\tilde{\beta}(Z, Z')V(\varepsilon', k(1-\delta), Z')\right] \tag{2}$$

and

$$V^a(\varepsilon, k, Z) \;=\; \max_{k'}\left\{\Pi(\varepsilon, k, Z) - C(\varepsilon, k, Z, k') + E_{\varepsilon', Z'|\varepsilon, Z}\left[\tilde{\beta}(Z, Z')V(\varepsilon', k', Z')\right]\right\} \tag{3}$$

---

[4]To be precise, the cyclical properties depend on the the specification of the utility function and the method of detrending as discussion below.

In this problem, the state variables (ε, k, Z) directly impact the profit flow of the plant in the current period. Other aggregate variables, included in Z, are in the state vector insofar as they provide information on the stochastic discount factor, denoted $\tilde{\beta}(Z, Z')$. In principle, higher order moments of the cross sectional distribution of the capital stock and profitability as well as measures of aggregate uncertainty could appear in Z.

The model includes three types of adjustment costs which, as reported in Cooper and Haltiwanger (2006), are the leading types of estimated adjustment costs:

$$
\begin{aligned}
C(\varepsilon, k, Z, k') \;=\; & \overbrace{(1 - \lambda)\,\Pi(\varepsilon, k, Z)}^{\text{disruption cost}} + p_b(I > 0)(k' - (1 - \delta)k) \\
& - \underbrace{p_s(I < 0)((1 - \delta)k - k')}_{\text{irreversibility}} + \underbrace{\frac{\nu}{2}\left(\frac{k' - (1 - \delta)k}{k}\right)^2 k}_{\text{convex cost}}
\end{aligned}
$$

The first is a disruption cost parameterized by Λ. If $\Lambda < 1$, then any level of gross investment implies that a fraction of revenues is lost. The second is the quadratic adjustment cost parameterized by ν. The third is a form of irreversibility in which there is a gap between the buying, p_b, and selling, p_s, prices of capital. These are included in (3) by the use of the indicator function for the buying $(I > 0)$ and selling of capital $(I < 0)$.

The profit function is

$$
\Pi(A, \varepsilon, k) = A\varepsilon k^{\theta}. \tag{4}
$$

Here A is an aggregate shock and ε is an idiosyncratic shock to plant-level profitability.

This is a reduced-form profit function which can be derived from an optimization problem over flexible factors of production (i.e. labor, materials, etc.). The parameter θ reflects factor shares as well as the elasticity of demand for the plant's output.

3 Stochastic Discount Factors

The optimization problem given in (4) includes a stochastic discount factor, $\tilde{\beta}_{t+1} \equiv \tilde{\beta}(Z_t, Z_{t+1})$. General equilibrium models tie the stochastic discount factor of the plant to the intertemporal preferences of the households. Empirically, it is natural to use asset pricing conditions to infer the stochastic discount factor.

The standard asset pricing formula of

$$
E_t\left[\tilde{\beta}_{t+1} R_{t+1}^j\right] = 1 \tag{5}
$$

for any asset j provides the framework. The idea is to first infer the stochastic discount factor

from observations on returns using (5). This representation of the household based SDF is mapped into the state space for (4) and the values of the various options for plant-level optimization are evaluated.

The analysis presents three different strategies for determining $\tilde{\beta}(Z_t, Z_{t+1})$. One is based upon a household Euler equation, thereby replacing $\tilde{\beta}_{t+1}$ with some version of an intertemporal marginal rate of substation for a representation household.

A second approach dispenses with the direct link between household utility and $\tilde{\beta}(Z_t, Z_{t+1})$. Instead, this approach estimates a parameterized model of the stochastic discount factor directly from asset pricing data. The final approach looks directly at interest rates.

For each, we are interested in a state space including aggregate capital, aggregate productivity and other elements of Z. This form of state dependence is immediate in versions of the stochastic growth model with capital and a technology shock in the state vector. Once there are nonlinearities in adjustment costs as well as plant-specific shocks, the state vector is more complex, potentially including measures of the cross sectional distribution of capital.

The appendix provides a detailed discussion of our data. For our analysis, we consider A_t to be total factor productivity in period t. With a focus on business cycle dynamics, we model a stationary specification of the shock process, which we parameterize from annual data after detrending using the H-P filter. As discussed in the appendix, our results depend on the choice of the parameter of the H-P filter, denoted λ. The choice of this parameter determines directly the serial correlation of the total factor productivity process and consequently the dependence of $\tilde{\beta}_{t+1}$ on the state vector (Z_t, Z_{t+1}). The mapping from the values of λ to the serial correlation of the aggregate shock, ρ_A, is given in Table 13 in the Appendix. We report results for two distinct values of λ: $\rho_A = 0.14$ results from the band-pass filter, approximated on annual data by setting $\lambda = 7$, while $\rho_A = 0.84$ comes from removing a linear trend, approximated by setting $\lambda = 100,000$.

For this analysis, we initially assume $Z_t = (A_t, K_t)$. Thus we exclude higher order moments of the cross sectional distribution of capital and plant-level profitability shocks. Under the null established by the results of Thomas (2002), the mean of the aggregate capital stock, K_t along with the aggregate profitability shock, A_t, suffice to characterize the current state of the system. This will be an issue we return to below.

3.1   Household Based SDF

As is customary within a stochastic equilibrium model in which households do not face costs of adjusting their portfolios, rates of return are linked to household preferences through an Euler equation. In that case, $\tilde{\beta}_{t+1}$ is the *ex post* marginal rate of substitution for consumption

across two periods. Hence (5) becomes

$$E_t \left[\frac{\beta u'(c_{t+1})}{u'(c_t)} R_{t+1}^j \right] = 1 \tag{6}$$

for asset $j = 1, 2, ...J$ where c_t is household consumption in period t and $u(\cdot)$ represents utility for a representative household.

Using data on consumption and some assumptions on preferences, (6) will generate a time series for $\tilde{\beta}_{t+1}$. For this exercise, assume $u(c) = \log(c)$ and set $\beta = 0.95$.[5]

With this parameterization it is easy to extract $\tilde{\beta}_{t+1} = \beta \frac{c_t}{c_{t+1}}$ from the data. To compute the empirical relationship between the stochastic discount factor and observables corresponding to key state variables in our model, we regress this stochastic discount factor on measures of the capital stock, current and future productivity: (A_t, A_{t+1}, K_t). The inclusion of both current period and future period values of the aggregate shock, A_t, are required since $\tilde{\beta}_{t+1}$ measures the realized real return between periods t and $t+1$. The future value of the capital stock is excluded as it is determined from (A_t, K_t) through the law of motion of the capital stock and is thus implicit in our approach.

Our results for this specification are presented in the first block of Table 1, labeled 'HE' (Household Euler). The columns indicate the regression coefficients and goodness of fit measures. The results ignore the capital stock since the inclusion of K_t is not statistically significant and adds very little to the empirical model.

Table 1: Stochastic Discount Factor: Household Euler Equation

SDF Specification	ρ_A	A_t	A_{t+1}	R^2	$\frac{dE_t[\tilde{\beta}(\cdot)]}{dA_t}$
HE	0.14	0.21	-0.61	0.46	0.12
		(0.09)	(0.09)		
	0.84	0.34	-0.46	0.56	-0.04
		(0.05)	(0.06)		

The table reports coefficients from a log-linear regression of the stochastic discount factor on aggregate productivity in the current year and next year using annual data from 1948 to 2008. The coefficient representing the constant is not reported. The final column reports the elasticity of the expected discount factor with respect to current productivity. Results are reported for two stationary specifications: 1) an approximation of the band-pass filter ($\rho_A = 0.14$) and 2) an approximation of linear detrending ($\rho_A = 0.84$).

These results indicate that this measure of the SDF varies positively with period t (lagged) aggregate productivity and negatively with period $t+1$ (current) productivity. As is made clear in the table, the magnitude of the coefficient depends on the choice of ρ_A. A lower value

---

[5]This is not the result of estimation of (6). Rather the log utility specification is common in the RBC literature and thus allows for comparison.

of this parameter leads to more dependence on A_t and less, in absolute value, dependence on A_{t+1} in the second block of results. This will be important later when we study how investment responds to productivity shocks.

3.2 Inferring the SDF from Asset Prices

Many studies have pointed to the difficulty in rationalizing asset prices within the standard specifications of household utility.[6] This has led to research on alternative representations of preferences, such as recursive utility models, as well as the inclusion of internal and external habits.

Instead of linking the stochastic discount factor to a richer model of household optimization, an alternative is to go directly to data on asset prices to infer the stochastic discount factor directly. The approach follows Zhang (2005), Bansal and Viswanathan (1993) Gallant and Hong (2007) and is based upon (5), $E_t\left[\tilde{\beta}_{t+1}R^j_{t+1}\right] = 1$.[7] Specifically, suppose that $\tilde{\beta}(A, K, A', K')$ is a logistic function, so that the discount factor lies in $(0, 1)$:

$$\tilde{\beta}(A_t, K_t, A_{t+1}, K_{t+1}) = \frac{e^{(\alpha_0 + \alpha_1 A_t + \alpha_2 A_{t+1} + \alpha_3 K_t + \alpha_4 K_{t+1})}}{1 + e^{(\alpha_0 + \alpha_1 A_t + \alpha_2 A_{t+1} + \alpha_3 K_t + \alpha_4 K_{t+1})}}. \tag{7}$$

The assets, indexed by j, represent six Fama-French portfolios adjusted for inflation and the 1-year Treasury rate adjusted for inflation.[8] Note that the portfolios include firm equity and hence include the variations similar to those produced in our estimated model.

The coefficients $\alpha \equiv (\alpha_0, \alpha_1, \alpha_2, \alpha_3, \alpha_4)$ are estimated from (5) using GMM to bring the sample analogue of (5) as close as possible to zero. For this estimation, the coefficients on the capital variables were initially included but, as in 'HE' model, were statistically insignificant.[9] Hence we report the coefficients only for current and lagged aggregate productivity. The instruments used were a constant along with period t consumption and capital. The parameter estimates are shown in Table 2 for the row 'Portfolio FF'.

Following Zhang (2005), the block labeled 'Portfolio Z' is from a specification of the SDF that highlights the countercyclical pricing of risk. Here the SDF is calibrated to match a number of cross-sectional asset pricing facts. It is modeled as

$$log\tilde{\beta}_{t+1} = log\beta + (logA_t - logA_{t+1})[\alpha_1 + \alpha_2(logA_t - \bar{A})] \tag{8}$$

---

[6]See, for example, the discussion and references in Cochrane (2011).

[7]Thanks to Ron Gallant, Monika Piazzesi and Lu Zhang for discussions on this approach.

[8]The data are taken from `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

[9]For the exercises with K_t in the state vector, the choice model was supplemented with a transition equation for aggregate capital taken from the data. In principle, other moments from the cross-sectional distribution of (k_i, ε_i) could be included as well.

Table 2: Stochastic Discount Factor: Nonlinear asset-based specifications

SDF Specification	ρ_A	α_1	α_2
Portfolio FF	0.14	72.48	10.94
		(15.35)	(16.83)
	0.84	66.11	-68.92
		(38.37)	(40.34)
Portfolio Z	0.14	4.84	-4.65
		(1.27)	(745.46)
	0.84	3.83	-2.70
		(9.63)	(390.03)

The table reports coefficients NL2SLS estimations of two portfolio-based specifications of the stochastic discount factor using annual data from 1948 to 2008. The specification for Portfolio FF is given by (7). The specification for Portfolio Z is from Zhang (2005) and given by (8). The coefficient representing the constant is not reported. Results are reported for two stationary specifications: 1) an approximation of the band-pass filter ($\rho_A = 0.14$) and 2) an approximation of linear detrending ($\rho_A = 0.84$).

where $\bar{A}$ is the mean of log(A). This specification of the SDF is used in (5) along with the Fama-French portfolios to estimate (α_1, α_2).

The results of the estimation are shown in the 'Portfolio Z' block of Table 2.[10] For both values of ρ_A, the SDF responds positively to variations in lagged productivity and negatively with current productivity.

The results in Table 1 related the stochastic discount factor from the household Euler equation to the state variables through a log-linear regression. The coefficients in Table 3 represent a log linear approximation of the two empirical SDFs from Table 2, making the results comparable to Table 1. The patterns from the 'HE' model appear in this representation as well: for $\rho_A = 0.84$, the SDF is procyclical with respect to lagged productivity and countercyclical with respect to current productivity.

3.3 Market Rates

Instead of working through an asset pricing equation, the last two blocks of Table 3 report a measure of $\tilde{\beta}_{t+1}$ based upon observed returns on different assets. In this way, we can look directly at interest rates that a firm might use to discount profits. These are *ex ante* rates and thus depend on current A_t but not A_{t+1}.

We study two real interest rates, the 30-day T bill and a long term AAA bond. These are *ex ante* rates and thus depend on the current state, A_t. The real rates are constructed by

---

[10]Zhang (2005) does not estimate this relationship but rather sets $\alpha_0 = 50$ and $\alpha_1 = -1000$ to match some moments.

Table 3: Stochastic Discount Factors: Log-linear representations of asset-based specifications

SDF Specification	ρ_A	A_t	A_{t+1}	R^2	$\frac{dE_t[\tilde{\beta}(\cdot)]}{dA_t}$
Portfolio FF	0.14	7.12	0.73	0.74	7.22
		0.63	0.63		
	0.84	4.60	-4.67	0.60	0.67
		0.57	0.60		
Portfolio Z	0.14	4.85	-4.82	1.00	4.18
		0.02	0.02		
	0.84	3.81	-3.80	1.00	0.62
		0.02	0.02		
30-day T-bill rate	0.14	-0.37		0.05	0.36
		(0.23)			
	0.84	-0.25		0.19	0.24
		(0.07)			
AAA LT bond yield	0.14	-0.27		0.03	0.25
		(0.21)			
	0.84	-0.40		0.57	0.37
		(0.05)			

The table reports coefficients from a log-linear regression of the stochastic discount factor for each specification on aggregate productivity in the current year and next year using annual data from 1948 to 2008. The coefficient representing the constant is not reported. The final column reports the elasticity of the expected discount factor with respect to current productivity. Results are reported for two stationary specifications: 1) an approximation of the band-pass filter ($\rho_A = 0.14$) and 2) an approximation of linear detrending ($\rho_A = 0.84$).

subtracting realized inflation for the 30-day T bill and long-term inflation expectations from the Survey of Professional Forecasters for the AAA bond from the respective nominal rate.

Importantly, the response of the return to a variation in A_t is generally negative and it is significant for the larger values of λ. Since the interest rate is inversely related to the stochastic discount factor, high realizations of productivity translate into higher discount factors. That is, once again the SDF is procyclical.

3.4 Response of the expected SDF to productivity

While the specifications differ, with the exception of the 'HE' case, the various estimates have a common feature: the SDF is procyclical. Though the plant-level optimization problem does include a covariance of future values with the stochastic discount factor, the response of the expected discount factor to variations in current productivity will be a key element in the results that follow.

To see this procyclicality, Tables 1 and 3, report the elasticity of $E_t[\tilde{\beta}(A_t, A_{t+1})]$ with

respect to A_t.[11] The results appear in the last column of those tables.

For the $\tilde{\beta}(\cdot)$ estimates using the household Euler equation, corresponding to $\rho_A = 0.14$, the expected stochastic discount factor is increasing in current productivity. The response is slightly negative at $\rho_A = 0.84$ as the negative coefficient on A_{t+1} is given more weight due to the higher serial correlation of the shock.

For all of the other specifications, the expected SDF is positively related to current pro-ductivity. This is true for both of the estimated models as well as the market interest rates reported in Table 3. Generally, the procyclicality is lower for $\rho_A = 0.84$ compared to $\rho_A = 0.14$. The pro cyclicality of the two market based SDFs is consistent with evidence of countercyclical real interest rates, as carefully documented in Beaudry and Guay (1996).

The procyclical SDF adds another dimension to the response of plant-level investment to an increase in A. First, as profitability increases, current profits increase as well. This alone has no affect on investment. Second, the conditional expectation of A' is increasing in A and this creates an incentive to invest more as A increases. Third, the expected SDF is increasing in A: as profitability increases, plant's discount less and invest more. Not surprisingly, the serial correlation of the profitability shock is central to the response of investment to variations in A.

Figure 1 summarizes these results for the $\rho_A = 0.84$ case. The top panel displays the SDF from the household Euler specification as well as the two market interest rates. The bottom panel displays the SDF for the two portfolio specifications. In both panels, the expected SDF, conditional on current productivity, is graphed. The solid line is aggregate productivity.[12]

Importantly, the SDF for the 'HE' case is relatively flat. As we shall see, this case creates investment behavior not too different from a specification with a fixed discount factor.

The other series are procyclical, peaking in the mid-1960s and showing some volatility over the sample. In fact, despite the differences in the estimated coefficients, the *ex ante* SDFs are very similar for these cases.[13]

4   Plant-level Implications

Using these processes for the stochastic discount factor, we study the response of investment to shocks. We do so first at the plant-level and then study aggregate effects in the next section.

To obtain these results, we solve (4) for different specifications of the stochastic discount factor reported in Tables 1 and 3.[14] We then study investment choices at the plant-level and

---

[11]Recall these are log-linear regressions.

[12]At $\rho_A = 0.84$, there is only one peak in the series and one trough.

[13]But they can have different implications on plant-level investment if their covariance with the value of capital differs. This will be included in the analysis of the plant-level dynamic optimization problem.

[14]These tables include a case labeled "Model-CC" which we explain below.

Figure 1: Stochastic Discount Factors.

Aggregate shock and expected discount factor (E[β]) when ρ_A = 0.84

Aggregate shock (A) and expected discount factor (E[β]) when ρ_A = 0.84

This figure shows the relationship between the various measures of the stochastic discount factor and aggregate productivity.

in aggregate. Thus the investment choices depend on empirically relevant representations of the stochastic discount factor, both through the conditional expectation emphasized in Tables 1 and 3 and the covariance between the aggregate state and the value of the plant, $V(\varepsilon, k, Z)$, from (4).

Table 4: Plant-level moments from simulated data

SDF Specification	ρ_A	mean$(\frac{I_i}{k_i})$	Frac Inactive	Frac -	Spike +	Spike -	mean$(\text{abs}(\frac{I_i}{k_i}))$	Corr$(I_i, I_{i,-1})$	Corr(I_i, ε_i)
Fixed β	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.13	0.20
HE	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.13	0.20
Portfolio FF	0.84	0.092	0.88	0.000	0.12	0.000	0.75	-0.09	0.13
Portfolio Z	0.84	0.090	0.87	0.000	0.13	0.000	0.71	-0.04	0.13
30-day T-bill rate	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.11	0.18
AAA LT bond yield	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.10	0.16
Model-CC	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.13	0.19
Model-KPR	0.84	0.090	0.88	0.000	0.12	0.000	0.72	-0.13	0.20

The table reports plant-level moments from a simulation of 1000 plants for 1000 periods under various specifications of the stochastic discount factor.

Table 5: Plant-level investment regression on simulated data: $I_i(A, K, \varepsilon_i, k_i)$

SDF Specification	ρ_A	A	K	ε_i	k_i	R^2
Fixed β	0.84	24.11		29.61	-0.43	0.36
		(1.21)		(0.27)	(0.00)	
HE	0.84	14.68		29.82	-0.43	0.36
		(1.21)		(0.27)	(0.00)	
Portfolio FF	0.84	65.88		7.90	-0.25	0.28
		(0.79)		(0.13)	(0.00)	
Portfolio Z	0.84	156.15		14.64	-0.30	0.34
		(1.51)		(0.22)	(0.00)	
30-day T-bill rate	0.84	90.47		24.95	-0.39	0.35
		(1.32)		(0.25)	(0.00)	
AAA LT bond yield	0.84	118.29		21.61	-0.36	0.35
		(1.39)		(0.24)	(0.00)	
Model-CC	0.84	14.57	0.38	28.94	-0.43	0.36
		(1.48)	(0.02)	(0.26)	(0.00)	
Model-KPR	0.84	13.41	0.41	28.87	-0.42	0.36
		(1.39)	(0.02)	(0.26)	(0.00)	

The table reports coefficients from an OLS regression of plant-level investment on the state variables of the plant's dynamic optimization decision. Data are produced from a simulation of 1000 plants for 1000 periods under various specifications of the stochastic discount factor. All variables are expressed in logs for the regression.

For these simulations, we follow Cooper and Haltiwanger (2006) and assume $\alpha = 0.58$, $\rho_\varepsilon = 0.885$ $\sigma_\varepsilon = 0.1$ for the idiosyncratic shock process.[15] The adjustment costs are given by $\Lambda = 0.8, \nu = 0.15, p_s = 0.98$. In the simulated data set, we follow 1000 plants for 500 periods.

Table 4 reports moments at the plant-level for different interest rate processes. These are the moments used in Cooper and Haltiwanger (2006) for the estimation of adjustment cost parameters.[16] There is an important point to gather from this table: these plant-level moments are essentially independent of the representation of the stochastic discount factor. Hence the parameter estimates from Cooper and Haltiwanger (2006), which assumes a constant discount factor, are robust to an analysis allowing a stochastic discount factor.

This does not imply though that the investment decision is independent of the specification of the stochastic discount factor. Table 5 reports results of plant-level regressions on current state variables for different measures of the SDF, with $\rho_A = 0.84$.

---

[15]In Cooper and Haltiwanger (2006) the estimates of the aggregate and idiosyncratic shock processes correspond to profitability shocks, as technology, cost and demand shocks cannot be separately identified.

[16]In that analysis, the correlation of productivity and investment was used rather than the correlation of plant-specific productivity and investment. In the data, these correlations are about the same.

Table 6: Plant-level investment regression (Adjusters only): $I_i(A, K, \varepsilon_i, k_i)$

SDF Specification	ρ_A	A	K	ε_i	k_i	R^2
Fixed β	0.84	25.48		28.79	0.04	0.98
		(0.51)		(0.22)	(0.00)	
Household Euler	0.84	18.94		28.77	0.05	0.98
		(0.44)		(0.21)	(0.00)	
Portfolio FF	0.84	77.46		12.35	0.19	0.94
		(0.89)		(0.13)	(0.01)	
Portfolio Z	0.84	209.15		25.44	-0.09	0.74
		(4.44)		(0.47)	(0.01)	
30-day T-bill rate	0.84	80.33		26.18	0.06	0.93
		(1.61)		(0.36)	(0.01)	
AAA LT bond yield	0.84	111.61		25.50	0.04	0.88
		(2.50)		(0.42)	(0.01)	
Model-CC	0.84	22.76	0.31	26.83	0.08	0.97
		(0.66)	(0.01)	(0.24)	(0.01)	
Model-KPR	0.84	20.82	0.30	26.80	0.07	0.97
		(0.66)	(0.01)	(0.25)	(0.01)	

The table reports coefficients from an OLS regression of plant-level investment on the state variables of the plant's dynamic optimization decision for all observations with nonzero. Data are produced from a simulation of 1000 plants for 1000 periods under various specifications of the stochastic discount factor. Only observations with nonzero investment are included in the regression. All variables are expressed in logs for the regression.

There are three important results here. First, investment increases in response to an increase in idiosyncratic profitability. Second, the response of investment to current aggregate productivity depends on the specification of the stochastic discount factor. Third, though not shown explicitly, this response depends on ρ_A.

Investment is less responsive to A in the 'HE' model compared to the fixed β case since the SDF is countercyclical in the 'HE' case. The response for these two specification is muted compared to the other SDF models. For the portfolio based SDF specifications as well as the market rates, the coefficient on aggregate productivity is considerably larger. This procyclical SDF amplifies the response of plant-level investment to A, relative to the fixed β and 'HE' cases.

These responses depend on $\rho_A = 0.84$. The persistence of the shock has two affects. First, it increases the conditional expectation of A', thus increasing the returns to investment. Second, from Table 3, it tends to reduce the procyclicality of the SDF, thus reducing the response to investment to A.

Tables 6 and 7 break the investment response into two components: the intensive margin

Table 7: Linear probability regression using plant-level simulated data (extensive margin)

SDF Specification	ρ_A	A	K	ε_i	k_i	R^2
Fixed β	0.84	1.04		1.27	-0.80	0.37
		(0.05)		(0.01)	(0.01)	
Household Euler	0.84	0.61		1.27	-0.81	0.37
		(0.05)		(0.01)	(0.01)	
Portfolio FF	0.84	4.98		0.58	-0.54	0.29
		(0.06)		(0.01)	(0.01)	
Portfolio Z	0.84	6.62		0.60	-0.53	0.35
		(0.07)		(0.01)	(0.01)	
30-day T-bill rate	0.84	3.98		1.07	-0.72	0.36
		(0.06)		(0.01)	(0.01)	
AAA LT bond yield	0.84	5.16		0.92	-0.66	0.36
		(0.06)		(0.01)	(0.01)	
Model-CC	0.84	0.55	0.65	1.25	-0.80	0.37
		(0.07)	(0.03)	(0.01)	(0.01)	
Model-KPR	0.84	0.48	0.73	1.25	-0.79	0.37
		(0.06)	(0.03)	(0.01)	(0.01)	

The table reports coefficients from an OLS regression of the plant-level extensive margin for investment on the state variables of the plant's dynamic optimization decision for all observations with nonzero. Data are produced from a simulation of 1000 plants for 1000 periods under various specifications of the stochastic discount factor. All explanatory variables are expressed in logs for the regression.

indicating the investment response of the adjusters and the extensive margin regarding the choice to invest or not.

Comparing the intensive margin regressions in Table 6 with the results in Table 5, there are a couple of points to note. For the fixed β case, the response to A is positive for the adjusters. Once the selection effect from the extensive margin is removed, investment is increasing in productivity.

For the intensive margin the response to k_i is almost zero. The explanation for the inverse relationship between investment rates and the stock of capital, shown in Table 5, must come from the extensive margin.

Comparing the extensive margin regressions in Table 7 from the linear probability model with results in Table 5, the adjustment probability is increasing in A_t for all specifications. To stress an important point, the response to variations in A in the fixed β and 'HE' cases are muted compared to those obtained from the portfolio and market based models of the SDF. The negative effects of high capital on the adjustment choice is very strong. The procyclical adjustment rate is consistent with evidence from Cooper, Haltiwanger and Power (1999).

5 Aggregate Implications

We aggregate our simulated data to study aggregate investment with different specifications of the stochastic discount factor. There are a couple of main findings. First, the stochastic discount factor matters for the behavior of aggregate investment. As noted earlier, the different models of the stochastic discount factor translate into different sensitivities of the expected discount factor to variations in current aggregate productivity. Through this mechanism, the response of aggregate investment to profitability shocks depends on the SDF representation.

Second, the non-convexities at the plant level have aggregate implications. By this we mean that the aggregate investment produced by the model with non-convexities is not the same as that produced by a model with quadratic adjustment costs.

5.1 Aggregate Moments

Table 8: Correlation of key variables with aggregate productivity (A)

Data or Model	ρ_A	Aggregate variables		Extensive margin Fract. of adjusters	Intensive margin mean($\frac{I_i}{k_i}\vert\frac{I_i}{k_i} > 0.2$)
		$\frac{I}{K}$	K'		
Data	0.84	0.37	0.13	0.19	NA
Fixed β	0.84	0.43	0.95	0.45	0.10
Household Euler	0.84	0.42	0.93	0.41	0.71
Portfolio FF	0.84	0.44	0.83	0.42	0.52
Portfolio Z	0.84	0.42	0.85	0.49	-0.01
30-day T-bill rate	0.84	0.43	0.94	0.48	-0.85
AAA LT bond yield	0.84	0.42	0.91	0.49	-0.87
Model-CC	0.84	0.55	0.45	0.52	0.48
Model-KPR	0.84	0.52	0.53	0.49	0.45

For the data results, the method used to construct the productivity data is described in the appendix. The other measures are constructed from the Census Bureau's Annual Survey of Manufacturers (ASM) as reported and described in Gourio and Kashyap (2007). The results in the first two rows are based on annual data from 1974 to 1998. The other results are based on simulated data from the model. The threshold used for the extensive and intensive margins is an investment rate (I/K) greater than 20 percent in absolute value.

Table 8 presents some basic correlations from the data as well as aggregated simulated data from the model.[17] A key point is how these aggregate moments depend on the SDF.

For the data, productivity is measured as describe in the appendix. The other measures are constructed from the Census Bureau's Annual Survey of Manufacturers (ASM) as reported and described in Gourio and Kashyap (2007). This data source is used because it provides

---

[17]Starting with this table, results are shown only for the $\rho_A = 0.84$ case.

important evidence for this analysis regarding the extensive margin investment decisions of manufacturing establishments. The reported statistics are based on annual data from 1974 to 1998.

In the data there are positive, but small, correlations between productivity and most of the aggregate series. The correlation between productivity and aggregate investment is modestly positive. The correlation between productivity and the fraction of establishments with investment rates greater than 20 percent is approximately 0.2. The correlation between productivity and the capital stock is negative.

The correlations between productivity and aggregate variables from the simulated data depend on the specification of the stochastic discount factor. For the investment rate, the correlations are surprisingly not too sensitive to the specification of the SDF.[18] The correlations of productivity with the adjustment rate are all positive, with a slight magnification for the portfolio based measures.

Table 9: Aggregate investment rate regression: OLS regression of $\frac{I_t}{K_t}$ on A_t and K_t

Model	ρ_A	A_t	K_t	R^2
Fixed β	0.84	0.63	-0.51	0.91
		(0.01)	(0.01)	
Household Euler	0.84	0.38	-0.47	0.82
		(0.01)	(0.01)	
Portfolio FF	0.84	3.45	-0.33	0.57
		(0.14)	(0.02)	
Portfolio Z	0.84	4.08	-0.37	0.60
		(0.15)	(0.02)	
30-day T-bill rate	0.84	2.46	-0.49	0.85
		(0.05)	(0.01)	
AAA LT bond yield	0.84	3.19	-0.45	0.76
		(0.08)	(0.01)	
Model-CC	0.84	0.37	-0.07	0.40
		(0.02)	(0.01)	
Model-KPR	0.84	0.31	-0.07	0.39
		(0.02)	(0.01)	

A different perspective appears when looking at aggregate regressions. Table 9 reports regressions of the aggregate investment rate, defined as the ratio of aggregate investment to the aggregate capital stock, on the current state for alternative specifications of the stochastic discount factor. Consistent with the correlations reported in Table 8, the coefficients on A_t

---

[18] At $\rho_A = 0.14$, more differences appear.

Figure 2: Aggregate Investment Rate with Non-convex Adjustment Costs.
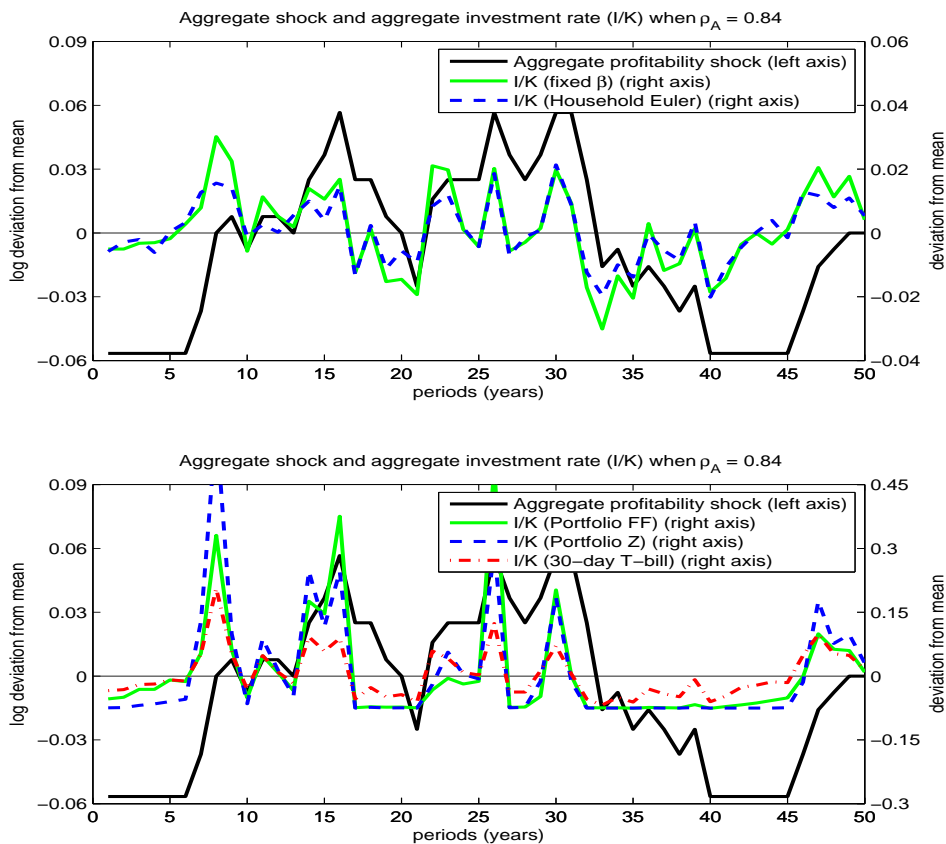
This figure shows aggregate investment rate, defined as aggregate investment divided by the aggregate capital stock, and the aggregate shock for the model of non-convex adjustment costs under different specifications of the stochastic discount factor for the case of $\rho_A = 0.84$.

are all positive. But the magnitudes vary considerably across the SDF specifications. The response of the aggregate investment rate is relatively muted in the fixed β and 'HE' models compared to the portfolio based specifications. The coefficient on capital is negative for all the models.

Figure 2 further highlights the importance of the stochastic discount factor for aggregate investment in the baseline non-convex adjustment cost case. The figure shows simulated aggregate investment rates (deviation from trend) series as well as simulated aggregate profitability. As suggested by Table 8, the aggregate investment rate series respond very differently to variations in aggregate productivity depending on the SDF.

From the top panel, the fixed β and HE cases are quite similar. This is to be expected given that these two representations of the SDF are very similar. The aggregate investment rate is procyclical for these cases.

The bottom panel shows the aggregate investment rates for the two portfolio and the T-bill rate versions of the SDF. These are strongly procyclical. These series are considerably more volatile than either the fixed β or HE induced series. This is consistent with the relative magnitude of the coefficients in Table 9.

5.2 Do Non-convexities Matter for Aggregate Investment?

In this section we explore the aggregate implications of a model with non-convex adjustment costs. The section addresses two related questions. First, are the non-convexities at the plant-level smoothed by variations in the empirically based SDF? Second, is there a specification of the SDF such that smoothing occurs?

Based upon the contributions of Thomas (2002), Khan and Thomas (2003) and Khan and Thomas (2008), the behavior of aggregate investment can be very well approximated by a standard real business cycle (RBC) model **even if** investment is lumpy at the plant level. Our interest here is to see whether those findings apply in our environment with: (i) a different model of non-convexity at the plant-level and (ii) an empirically based SDF.

For this exercise, we search over the parameters of a log-linear representation of $\tilde{\beta}(A, A', K)$, denoted Θ_B, and well as a quadratic adjustment cost parameter, ν, to minimize the distance between the aggregate investment behavior of two models. The first, denoted 'NC' for non-convex model, is the baseline model from Cooper and Haltiwanger (2006) solved with a stochastic discount factor parameterized by Θ_B. The second, denoted 'QAC', assumes plants are homogeneous and have **only** quadratic adjustment costs parameterized by ν. The homogeneous plants in the QAC case utilize the same stochastic discount factor parameterized by Θ_B.

Given (Θ_B, ν), each of the models can be solved and simulated using the same aggregate shocks, creating two times series for aggregate investment. A goodness of fit measure, R^2, is computed between these two series. The procedure finds (Θ_B, ν) to maximize the goodness of fit.

If the goodness of fit is near 1, we conclude that there does exist a stochastic discount factor that would make the NC model essentially identical to the QAC model. In that case, complete smoothing is feasible. If the goodness of fit is not near 1, then smoothing does not arise.[19] This addresses the second question of the existence of a SDF that smooths the non-convexities.

The first question is answered by comparing the parameters of Θ_B against those found in the data based models of the SDF. If these parameters are (close to) identical, then smoothing in fact will occur with that representation of the SDF.

---

[19]This statement, of course, is conditional on searching in the class of log linear models.

Table 10: Goodness of Fit: QAC vs Nonconvex Adjustment Costs

Model	$\Theta_B(A)$	$\Theta_B(A')$	$\Theta_B(K)$	ν	R^2
Best Fit	0.39	-0.41	na	0.15	0.832
Best Fit	0.38	-0.72	0.096	0.075	0.868
Fixed β	0	0	0	0.187	0.723
HE	0.34	-0.46	na	0.194	0.559
Portfolio FF	4.60	-4.67	na	0.202	0.460
Portfolio Z	3.81	-3.80	na	0.285	0.522
30-day T-bill	-0.25	na	na	0.185	0.770
Model-CC	0.43	-0.64	0.09	0.157	0.656
Model-KPR	0.37	-0.59	0.09	0.195	0.673

The results are reported in Table 10. The first row reports estimates of Θ_B imposing the restriction that the coefficient on K is zero. The estimate from this exercise is $\nu = 0.15$ with an $R^2 = 0.832$.[20] The parameter estimates of Θ_B are 0.39 for the coefficient on A_t and -0.41 for the coefficient on A_{t+1}. This implies a derivative of the expected SDF with respect to A of 0.046. In comparison to the parameters of Θ_B shown in Tables 1 and 3 for the HE and portfolio-based specifications, this stochastic discount factor with the best smoothing properties is more procyclical than the HE specification and less procyical than the other specifications.

The second row of the table adds K to the model. The fit improves slightly and the dependence on A_{t+1} is larger in absolute value. Also the estimated quadratic adjustment cost is lower.[21]

In both cases, the quadratic adjustment cost model does not fit the aggregate data produced by the model with non-convex adjustment costs very well. The R^2 are all well below unity. Perhaps a better fit would be obtained if the model of the SDF included moments of the cross sectional distribution of (k_i, ε_i). The significance of these moments would, by itself, indicate that the nonconvexities have aggregate implications.

This leads to two conclusions. First, none of the SDF representations lead to complete smoothing of aggregate investment: the best fitting Θ_B differs from these representations. Second, movements in the SDF do not smooth the lumpy investment: the fit of the best model is significantly less than unity.

Using this procedure, we study the goodness of fit between the 'QAC' and 'NC' models

---

[20]This estimated quadratic adjustment cost and goodness of fit are close to those reported Cooper and Haltiwanger (2006) for minimizing the distance between the NC model with a **fixed** β and the QAC model. In that paper, the reported estimate is $\nu = 0.195$ with an $R^2 = 0.859$.

[21]A word of caution is in order. In contrast to the specification excluding K, the objective function in this case is relatively insensitive to variations in parameters, indicating an identification problem.

for each of the other specifications of the stochastic discount factor. This gives us a sense of how much smoothing is induced by the various measure of the SDF, rather than the one that produces the most smoothing. For each specification of the SDF, we simulate to create an aggregate investment series and compare it against the series produced by the 'QAC' model using the **same SDF**.[22] The parameter ν is re-estimated for each of the SDF specifications. Our results are summarized in the bottom part of Table 10.

Here we find that the stochastic discount factors generate less smoothing than a fixed discount specification. This again reflects the procyclical nature of these representations of the SDF. So while the estimation above demonstrates that there exists a specification of the stochastic discount factor that generates more smoothing than in a model with a fixed discount specification, the results in the table show that the 'HE' , 'Portfolio FF' and 'Portfolio Z' specifications generate considerably less smoothing. The exception is the 30-day T-bill, which generates a better fit than the fixed β model.

6   A General Equilibrium Approach

Thus far our analysis has focused on empirically based representations of the stochastic discount factor, rather than those emerging from a stochastic general equilibrium model. That is, our model excludes the household optimization that would support the stochastic discount factor uncovered from the data. This does not imply that our approach is inconsistent with a general equilibrium structure. Rather, we focus on the response of plants to shocks given the SDF process inferred from data.

The results stand in contrast to those of Thomas (2002) and the literature that followed. The aggregate investment series produced by our model of non-convex adjustment costs and the empirical based SDF are not well approximated by a quadratic adjustment cost model.

This section contrasts the empirically based representations of the SDF with those coming from a RBC model, and does so for different models of adjustment costs. We first present a model-based version of the stochastic discount factor and then turn to whether non-convexities at the plant level are smoothed through aggregation and movements in the model-based stochastic discount factor. We then move away from the Cooper and Haltiwanger (2006) model of adjustment costs to study the aggregate implications of other specifications.

6.1   Model-Based SDF

We study three model based SDFs. Throughout the model-based approach, the utility function is $log(c)$. Table 11 reports results for three models. For each model, we use three different

---

[22]The entry in Table 10 for the 'Portfolio FF' case is the log-linear representation. The analysis uses the non-linear representation from Table 2.

values of ρ_A corresponding to those used in the empirical specification.

The first model, labeled 'RBC', is the standard real business cycle model, building from King, Plosser and Rebelo (1988). It forms the basis of comparison in Thomas (2002) and the work that followed.

But that model is based upon perfect competition while one interpretation of the curvature in the profit function is market power. Hence we study a second model, labeled 'Model-CC' and taken from Chatterjee and Cooper (1993), which studies a stochastic real business cycle model with monopolistic competition. This environment is closer to the underlying market structure assumed in Cooper and Haltiwanger (2006). The results for the Model-CC specification in Table 11 come from a model with no entry and exit and a markup of 25 percent.[23]

The third model, labeled 'Model-GK', represents the Gourio and Kashyap (2007) specification of the model proposed by Thomas (2002). The relationship in Table 11 is only an approximation of the true stochastic discount factor because the state space for the Gourio and Kashyap (2007) model includes the cross-sectional distribution of capital vintages as well as the capital stock and productivity shocks. Thus, the stochastic discount factor should, in principal, depend on the cross-sectional distribution.

However, given that the real allocations from the Thomas (2002) model with lumpy investment are so close to the stochastic growth model, one would conjecture that the process for the stochastic discount factor would be close to that of the RBC and/or Model-CC specifications. The coefficients shown in Table 11 for 'Model-GK', are estimated based on simulated data from the model using the Gourio-Kashyap specification of the underlying parameters. So while the model-based stochastic discount factor is a function of a larger set of state variables accounting for underlying heterogeneity, the simpler representation that includes only productivity and the capital stock captures 99.1% of the variation in the stochastic discount factor, as measured by the R^2 in the regression.

For the RBC, model the return responds positively to both the current productivity shock and the capital stock and negatively with future productivity for all values of ρ_A. The Model-CC results are quite similar to those in the standard RBC model. The Model-GK results are close to the RBC model, though the SDF is a bit more responsive to A_t.

As discussed earlier, a key issue is the sensitivity of the expected SDF to variations in current productivity. From the model based results, the expected stochastic discount factor is **countercyclical** for the $\rho_A = 0.84$ case and **procyclical** for $\rho_A = 0.14$. This response is more countercyclical than in the empirical based results. Recall that for the portfolio based SDF, the expected SDF was always procyclical and for the household Euler equation approach, the SDF was countercyclical with a response of -0.04.

---

[23]The markup is based on a CES specification where the elasticity of substitution is set to 5 for both consumption and capital goods.

Table 11: Model-based relationship between state variables and discount factor

Model	ρ_A	A_t	A_{t+1}	K_t	$\frac{dE_t[\tilde{\beta}(A_t,A_{t+1},K_t)]}{dA_t}$
RBC	0.14	0.08	-0.39	0.09	0.031
	0.84	0.37	-0.59	0.09	-0.121
Model-CC	0.14	0.09	-0.41	0.09	0.032
	0.84	0.43	-0.64	0.09	-0.110
Model-GK	0.14	0.11	-0.35	0.10	0.062
	0.84	0.43	-0.59	0.10	-0.066

These differences between data and model are consistent with earlier findings, such as Beaudry and Guay (1996), on the inability of the standard RBC model to match interest rate movements. Our interest here is not in these deviations between data and models *per se* but rather their implications for aggregate investment.

The current capital stock enters into the SDF as well. For all models, the coefficient on K_t is positive and about 0.1.

Figure 3 complements Figure 1 by showing the comovement of a model-based SDF, Model-CC and two of the empirically based SDFs, HE and Portfolio FF, over the sample period. Note that the model based SDF and that obtained from the household Euler equation are quite similar and, as noted earlier, quite smooth and less procyclical than the Portfolio FF specification.

6.2   Moments

It is useful to compare the moments from a model-based SDF with those from the empirically-based representations. We use Model-CC and KPR for this comparison.
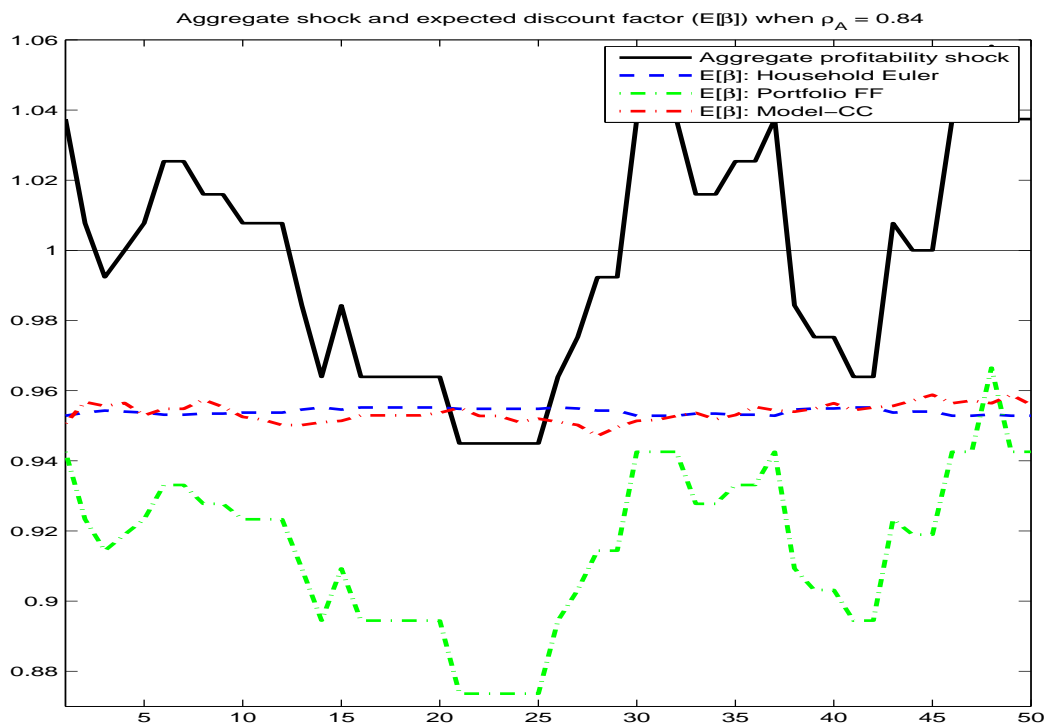
Given that the stochastic discount factor depends, in these models, on the current capital stock, a law of motion for the aggregate capital stock is needed so that the optimizing plant can predict future capital and thus discount appropriately in future periods. Denote this function as $K' = \Gamma(A, K)$. For the experiments using Model-CC, for example, the log linear representation of $\Gamma(A, K)$ from that solution is used to frame the expectations at the plant-level:

$$\hat{K}_{t+1} = 0.840\hat{A}_t + 0.366\hat{K}_t. \tag{9}$$

A comparable equation is used for the KPR based simulations.

The plant-level moments, as noted earlier, are largely independent of the SDF. This holds for the model-based representations as well as shown in the last block of Table 4. Plant-level

Figure 3: Model Based Stochastic Discount Factors

This figure shows the relationship between the various measures of the stochastic discount factor and aggregate productivity for two model based representations of the SDF as well as the Portfolio FF case..

investment is increasing in A, with procyclical adjustments on both the extensive and intensive margins. Further, with K in the state space for the SDF, investment is increasing in K as well.

The investment regressions, reported in Table 5, are similar to those produced by the HE specification. This the case for the response on the extensive margin as well, as in Table 7. The magnitude of the response to variations in A, particularly on the extensive margin, is much lower for the Model-CC and KPR cases compared to the other SDF representations.

From Table 8, the investment rate, as well as the extensive and intensive margins, are correlated with aggregate productivity for the model-based SDF. The correlations are in line with the other specifications of the SDF.

Figure 4 shows the aggregate investment rate series created by Model-CC, compared to the 'HE' and 'Portfolio FF' specifications. The investment rate predictions from Model-CC and the 'HE' cases are quite similar. In contrast, there is considerable volatility in investment produced by the 'Portfolio FF' case because of the more procyclical SDF in this empirically based specification.

Figure 4: Aggregate Investment Rate with Non-convex Adjustment Costs.
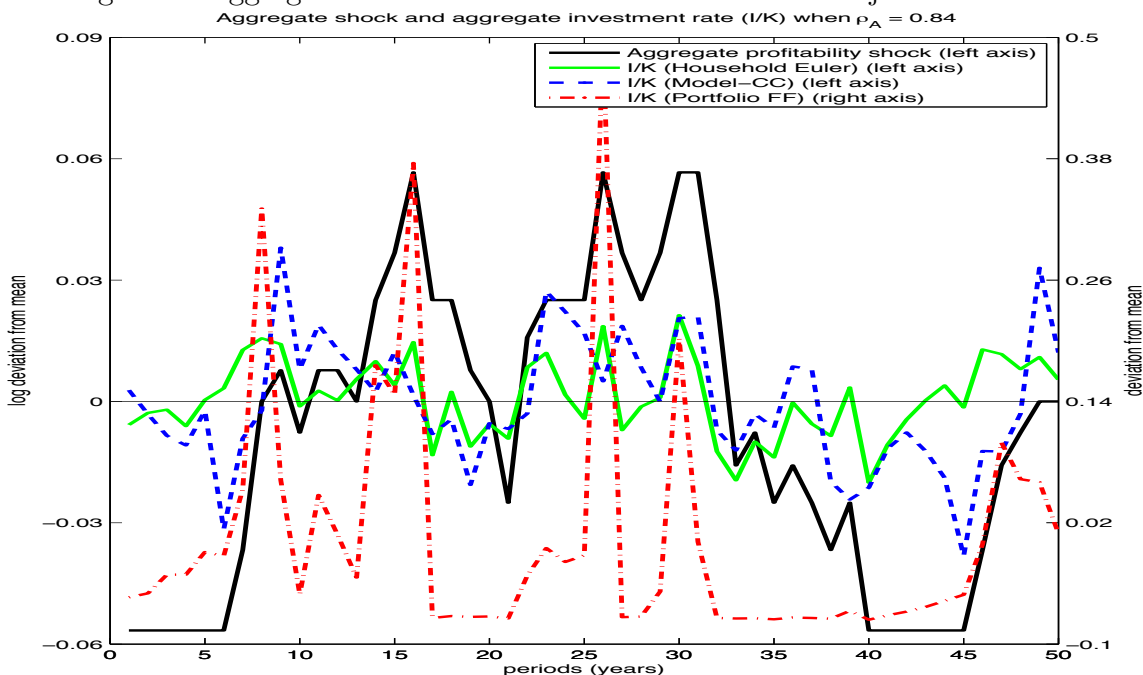
This figure shows aggregate investment rate, defined as aggregate investment divided by the aggregate capital stock, and the aggregate shock for the model of non-convex adjustment costs under different specifications of the stochastic discount factor for the case of $\rho_A = 0.84$.

6.3 Does the Model Based SDF Smooth Plant-level Non-convex Adjustment Costs?

In this section we return to the aggregate implications of a model with non-convex adjustment costs to see how much smoothing is obtained from the model based SDF. The results are shown in the last block of Table 10.

For this exercise, we find that the model-based stochastic discount factor **does not** smooth aggregate investment. The goodness of fit is significantly below unity. That is, using the SDF from either Model-CC or KRP does not imply that aggregate investment created in a model with heterogenous plants and non-convex adjustment costs can be replicated by a model with homogenous firms and quadratic adjustment costs.

6.4 Alternative Adjustment Costs

Thus far, the non-convexities at the plant-level are not smoothed. This section addresses the robustness of our findings on the aggregate implications of the plant-level non-convexities by looking at alternative adjustment costs.[24]

---

[24]We are grateful to Mike Elsby for discussions leading to the development of this section.

To do so, we study the components of the adjustment costs estimated by Cooper and Haltiwanger (2006) to determine which key factors in the lack of smoothing. It could be, for example, that in a model with only irreversible investment, the plant-level lumpiness is indeed smoothed.

Our findings are summarized in Table 12. Here we present the fit between a quadratic adjustment cost model with homogenous plants, denoted 'QAC' as above, and various models of adjustment costs. For each model of adjustment costs, we include three specifications of the SDF.[25]

There are a couple of key points. First, the fit is not unity for any of the specifications of the SDF joint with adjustment costs. This is even the case for the non-adjustment cost model. The fit is best for the case of $\lambda = 1$ so that there are only quadratic adjustment costs and irreversibility. In this case, both the HE and Model-CC specifications of the SDF generate a fit near the best fit.

Table 12: Goodness of Fit: Alternative Adjustment Costs

Model	$\Theta_B(A)$	$\Theta_B(A')$	$\Theta_B(K)$	ν	R^2
best fit	0.39	-0.41	na	0.15	0.832
CH					
HE	0.34	-0.46	na	0.15	0.51
Portfolio FF	4.60	-4.67	na	0.15	0.41
Model-CC	0.43	-0.59	0.10	0.075	0.555
NO AC					
HE	0.34	-0.46	na	0.15	0.386
Portfolio FF	4.60	-4.67	na	0.15	0.258
Model-CC	0.43	-0.59	0.10	0.075	0.636
λ only					
HE	0.34	-0.46	na	0.15	0.520
Portfolio FF	4.60	-4.67	na	0.15	0.416
Model-CC	0.43	-0.59	0.10	0.075	0.478
$\lambda = 1$					
HE	0.34	-0.46	na	0.15	0.841
Portfolio FF	4.60	-4.67	na	0.15	0.617
Model-CC	0.43	-0.59	0.10	0.075	0.847

This table shows the fit between the quadratic adjustment cost model and alternative models of adjustment costs for three representations of the stochastic discount factor. This analysis assumes $\rho_A = 0.84$.

---

[25]The table does not yet include the best-fit for each specification of adjustment costs.

7   Conclusion

This paper studies the implications of fluctuations in the SDF on plant-level and aggregate investment. In particular, we highlight how alternative specifications of the SDF influences the response of investment to variations in aggregate productivity. The smoothing effects of interest rates depend on the determination of that process.

If, as we have emphasized here, the state dependent discount factor is determined from the data, then there is little smoothing of investment due to SDF movements. In particular, the data based measures of the SDF do not smooth the non-convexities.

From this analysis, it appears that investment is very sensitive to the SDF. That is, relatively small changes in the response of the SDF to, say, variations in productivity, can have significant effects on aggregate investment.

As this work proceeds, we will turn to an analysis of monetary policy which presumably underlies the interest rate process uncovered in the data. We can use our model to see how alternative monetary policies can influence investment behavior. Our results indicate an important channel for monetary policy: influencing the amount of plant-level lumpy investment that is smoothed through interest rate movements. When the lumpiness is not smoothed, the impact of monetary policy can itself be state-dependent.

8 Appendix: TFP process

The measure of total factor productivity used in this analysis is constructed as a Solow residual following Stock and Watson (1999). The calculation includes nonfarm real GDP (source: BEA), nonfarm payroll employment (source: BLS), real nonresidential private fixed capital stock (source: BEA and authors' calculations), and a labor share of 0.65. The data sample is annual frequency from 1948 to 2008.

The profitability shock, denoted A_t, is calculated as a residual assuming a Cobb-Douglas production function using raw data on output and factor inputs. Our analysis is focused on business-cycle dynamics, so here we abstract from long-term growth in TFP by detrending A_t. Various approaches have been used in the literature for detrending, so we employ three different detrending specifications to examine the sensitivity of the results. The first approach focuses specifically on business cycle frequencies (between 3 and 8 years). For this case, we detrend using the HP filter with the λ parameter set to 7, which closely approximates a band-pass filter on annual data. The second approach is to remove a linear trend from the data, which we approximate by setting the HP filter parameter for λ to 100,000. The third approach uses an intermediate value of λ that is commonly used to filter annual data, $\lambda = 100$.

The parameters of the TFP process are estimated based on a log-normal AR(1) specification.

$$\log A_t = \rho_A \log A_{t-1} + \epsilon_{A,t}, \ \ \epsilon_A \sim N(0, \sigma^2_{\epsilon_A}) \tag{10}$$

Estimates of the shock process parameters are displayed in Table 13 for the three different detrending specifications. The estimate of serial correlation in TFP, ρ_A, is very sensitive to the detrending specification. If focusing on business cycle frequencies, $\lambda = 7$, there is little serial correlation in detrended TFP. On the other hand, detrended TFP has a much higher serial correlation when approximating the removal of a simple time trend ($\lambda = 100,000$). We consider the process parameter estimates from the two extreme detrending specifications, ($\lambda = 7; 100,000$), in our analysis to examine the role of the detrending assumption in modeling the relationship between interest rates and investment decisions.

Table 13: Parameter estimates for Solow-residual technology process

λ	ρ_A	σ_{ϵ_A}
7	0.14	0.012
100	0.45	0.015
100000	0.84	0.018

References

Bansal, Ravi and S Viswanathan, "No arbitrage and arbitrage pricing: A new approach," *The Journal of Finance*, 1993, *48* (4), 1231–1262.

Beaudry, P. and A. Guay, "What do interest rates reveal about the functioning of real business cycle models?," *Journal of Economic Dynamics and Control*, 1996, *20* (9-10), 1661–1682.

Caballero, R.J. and E.M.R.A. Engel, "Explaining investment dynamics in US manufacturing: a generalized (S, s) approach," *Econometrica*, 1999, *67* (4), 783–826.

Chatterjee, S. and R. Cooper, "Entry and exit, product variety and the business cycle," 1993. National Bureau of Economic Research Working Paper No. 4562.

Cochrane, John H, "Presidential address: Discount rates," *The Journal of Finance*, 2011, *66* (4), 1047–1108.

Cooper, R. and J. Haltiwanger, "On the Nature of the Capital Adjustment Costs," *Review of Economic Studies*, 2006, *73*, 611–34.

Cooper, Russell, John Haltiwanger, and Laura Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, 1999, *89*, 921–946.

Gallant, A.R. and H. Hong, "A Statistical Inquiry into the Plausibility of Recursive Utility," *Journal of Financial Econometrics*, 2007, *5* (4), 523–559.

Gourio, F. and A. Kashyap, "Investment spikes: New facts and a general equilibrium exploration," *Journal of Monetary Economics*, 2007, *54, supplement*, 1–22.

Khan, A. and J.K. Thomas, "Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?," *Journal of monetary economics*, 2003, *50* (2), 331–360.

− **and** − , "Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics," *Econometrica*, 2008, *76* (2), 395–436.

King, R.G., C.I. Plosser, and S.T. Rebelo, "Production, growth and business cycles:: I. The basic neoclassical model," *Journal of Monetary Economics*, 1988, *21* (2-3), 195–232.

Stock, James H. and Mark W. Watson, "Business cycle fluctuations in US macroeconomic time series," in John B. Taylor and Michael Woodford, eds., , Vol. 1, Part A of Handbook of Macroeconomics, Elsevier, 1999, pp. 3 – 64.

Thomas, Julia, "Is Lumpy Investment Relevant for the Business Cycle?," *Journal of Political Economy*, 2002, *110*, 508–34.

Zhang, Lu, "The value premium," *The Journal of Finance*, 2005, *60* (1), 67–103.