Online Appendix to Winners and Losers: Creative Destruction and the Stock Market

Leonid Kogan[*] Dimitris Papanikolaou[†]
Noah Stoffman[‡]

June 8, 2016

[*]MIT Sloan School of Management and NBER, lkogan@mit.edu
[†]Kellogg School of Management and NBER, d-papanikolaou@kellogg.northwestern.edu
[‡]Kelley School of Business, nstoffma@indiana.edu

Outline

In Section 1 we include some additional results that complement the results in the main paper. In Section 2 we discuss details of the estimation procedure.

1 Additional Results

This section is organized as follows. In Section 1.1 we analyze the cashflow risk implied by the model. In Section 1.2 we demonstrate by the market to book ratio is a good proxy for the firms' technology risk exposures. In Section 1.3 we estimate a linearized version of the SDF implied by the model and compare the results between the model and the data.

1.1 Cashflow risk

We begin by analyzing the risk properties of the model's implied cashflow process.

1.1.1 Cashflows of value and growth portfolios

We compute portfolio dividends as in Hansen, Heaton, and Li (2005). We denote by D_t^G and D_t^V the net payout of the growth and value portfolio, respectively. Following Hansen et al. (2005), we focus on the top and bottom quintile of firms sorted on book-to-market. Since payout from each portfolio can become negative, we need to modify the approach in Hansen et al. (2005). Rather than considering VARs in log dividends and the permanent component of consumption, we consider VARs in

$$d_t^i \equiv D_t^i / C_t, \qquad i \in [G, V] \tag{1}$$

and the permanent component of consumption (in our model, χ_t). We then closely follow the procedure in Hansen et al. (2005). We normalize $d_0^i = 1$. We estimate two bi-variate VARs, one for each portfolio. Each VAR has 5 lags. We estimate the VAR in levels and do not include a deterministic trend. We repeat this procedure in 1,000 samples from the model, each with length of 50 years. In Figure 1, we plot the mean response of the net payout of the value and growth portfolio (scaled by consumption) to a shock in χ_t, along with the 5% and 95% across simulations.

Examining Figure 1, we see that, in the long run, the payout series of the value portfolio responds more than consumption to a shock to χ. By contrast, the payout from the growth strategy responds less than consumption to χ. These results imply that dividends of the value strategy have more 'long run risk' than the dividends of the growth strategy, consistent with the evidence reported in Bansal, Dittmar, and Lundblad (2005) and Hansen et al. (2005). However, there is substantial dispersion of these estimates across model simulations.

Figure 1: Long run cashflow risk of value and growth portfolios. Figure plots the mean impulse response across simulations of D_t^V/C_t and D_t^G/C_t to a unit standard deviation shock to the permanent component of consumption, χ_t. The impulse responses are computed from two bi-variate VARs, one for each portfolio. Each VAR has 5 lags . We estimate the VAR in levels and do not include a deterministic trend. We repeat this procedure in 1,000 samples from the model, each with length of 50 years. In addition to the mean response, we also plot the 5% and 95% estimates across simulations.

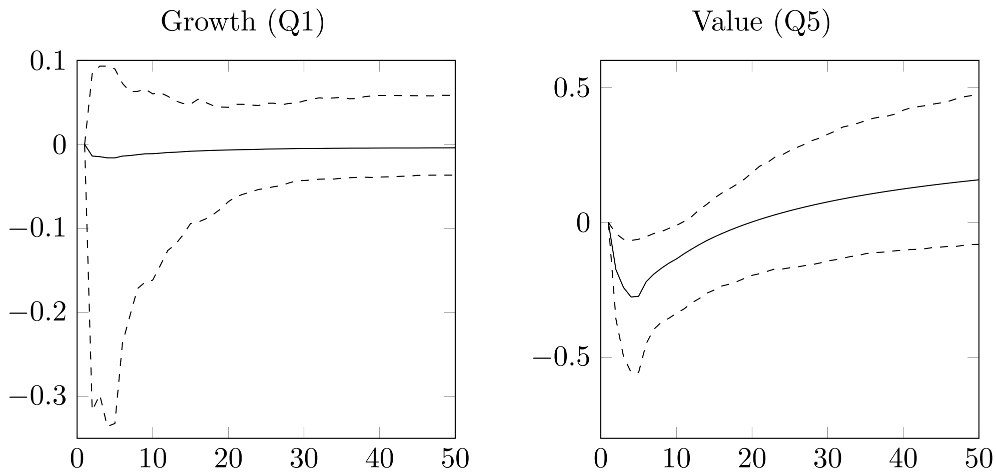

In interpreting these results, it is important to distinguish between the payout from the growth or value trading strategy and the payout from owning a growth or value firm. The two are very different due to the nature of the rebalancing strategy. Specifically, growth firms in the model have on average more investment opportunities, hence they have lower payout than value firms. Eventually, the investments that growth firms make are realized and these firms will pay higher dividends. However, by that time those firms will have transitioned out of the growth portfolio – they will become more like value firms. An investor in these growth firms will eventually capture those higher dividends; an investor in the growth strategy will not. Instead, those higher future expected dividends will manifest as capital gains for the investor in the growth strategy.

To illustrate these ideas more clearly, we plot the distribution of d_τ^V and d_τ^G across 1,000 simulations for $\tau = 50$ in Figure 2. We note that the dividends of the value strategy grow faster, on average, than the dividends of the growth strategy, consistent with the findings of Hansen et al. (2005).

Figure 2: Dividend growth of value and growth portfolios. Panel A plots the mean value of d_t^V and d_t^G across simulations for horizons up to 50 years. Panel B plots the distribution of d_T^V and d_T^G across simulations for $T = 50$. Both series at initialized to 1 at $t = 0$.

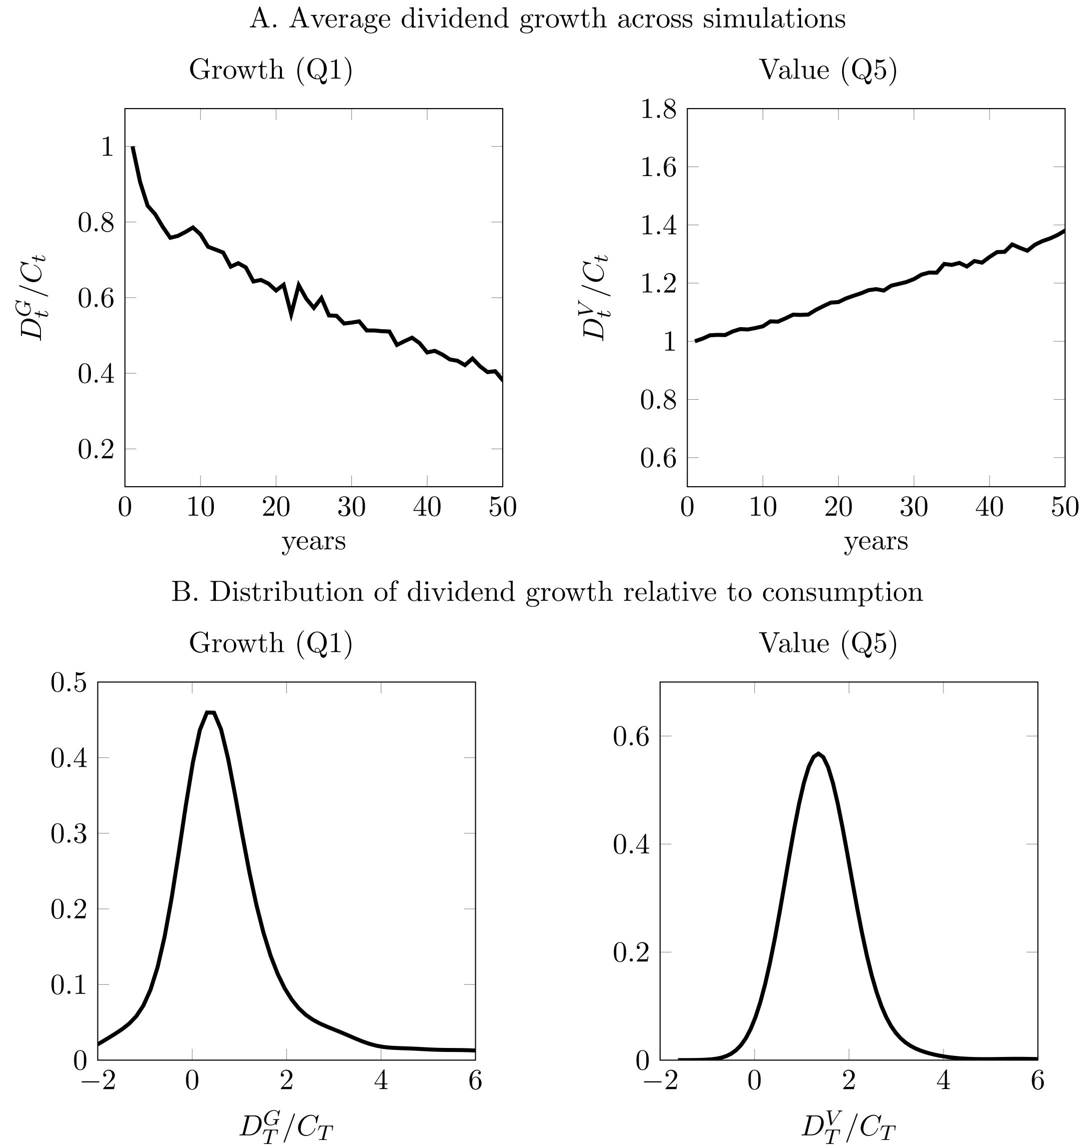

In sum, we see that our model can potentially replicate the stylized features documented by Bansal et al. (2005) and Hansen et al. (2005) regarding the long-run cashflow risk of the value and growth portfolios. However, there are two important caveats. First, there is significant variation of point estimates across simulations, so these comparisons are not a very informative test of the model. Second, and more importantly, the empirical facts concern firm dividends. In our model, the Modigliani-Miller theorem holds, and dividends are indeterminate as firms can substitute dividends for share repurchases. Further, even net firm payout is indeterminate. For instance, firms can hold cash inside the firm, or borrow or lend to each other. A more thorough empirical investigation of this aspect of the model will require additional theoretical assumptions to pin down the process of firm dividends. We think this is an interesting extension that is outside the scope of this paper.

1.1.2 Risk and risk prices across horizons

Our model generates a sizable equity premium due to joint movements in aggregate dividends and the stochastic discount factor. Here, we briefly examine how risk at different horizons contributes to asset prices. To do so, we follow Borovicka, Hansen, and Scheinkman (2014) and construct shock exposure

$$\varepsilon_q(T-t, X_t) = \eta(X_t)\frac{E_t(D_T \mathcal{M}_t \log D_T | X_t)}{E_t(D_T | X_t)} \tag{2}$$

and shock-price elasticities,

$$\varepsilon_p(T-t, X_t) = \eta(X_t)\frac{E_t(D_T \mathcal{M}_t \log D_T | X_t)}{E_t(D_T | X_t)} - \eta(X_t)\frac{E_t(\pi_T D_T(\mathcal{M}_t \log D_T + \mathcal{M}_t \log \pi_T) | X_t)}{E_t(\pi_T D_T | X_t)}. \tag{3}$$

Here, $\mathcal{M}_t \log D_T$ is a Malliavin derivative – that is, it measures the contribution of a shock dB_t to the stochastic process D at time $T > t$. Here, $\eta(X)$ indexes the direction and size of the shock. Expression (2) is very similar to a non-linear impulse response function; it examines the effect of a shock today to future values of D_t. Expression (3) represents the sensitivity of the expected log return associated with cashflow equal to D_T to a marginal increase in the exposure of that cashflow to a time-t shock. These marginal risk prices potentially vary across horizons based on how the shock D propagates.

We focus on two cashflow processes, the total payout to holders of the market portfolio,

$$D_t = \phi\, Y_t - I_t - \eta\,\lambda\,\nu_t, \tag{4}$$

and the payout accruing from assets in place at time t,

$$D^{ap}_s = p_{Z,s}\, e^{-\delta(s-t)} K_t. \tag{5}$$

Even though aggregate dividends D can potentially become negative, this does not happen near the mean of the stationary distribution of ω. It only happens at extreme ranges of the state space that were not reached across 1,000 model simulations. We compute shock-exposure ε_q and shock-price ε_p elasticities for the two processes D and D^{ap} at the mean of the stationary distribution of ω using Monte Carlo simulations. We plot the estimated elasticities in Figure 3.

The first two columns show the impulse response ε_q of the two cashflow processes to a technology shock. A positive disembodied shock (x) increases dividends, both for the overall market and for assets in place. By contrast, improvements in technology that are embodied in new vintages (ξ) lead to a decline in aggregate dividends in the short term – as firms fund new investments – and an increase in the long run. By contrast, the cash flows accruing from installed capital falls due to competition – the equilibrium price p_Z falls as the economy acquires more and better new capital. The last two columns show the shock-price elasticities ε_p. We see that the marginal risk prices are essentially flat across horizons.

Figure 3: Shock-exposure and shock-price elasticities

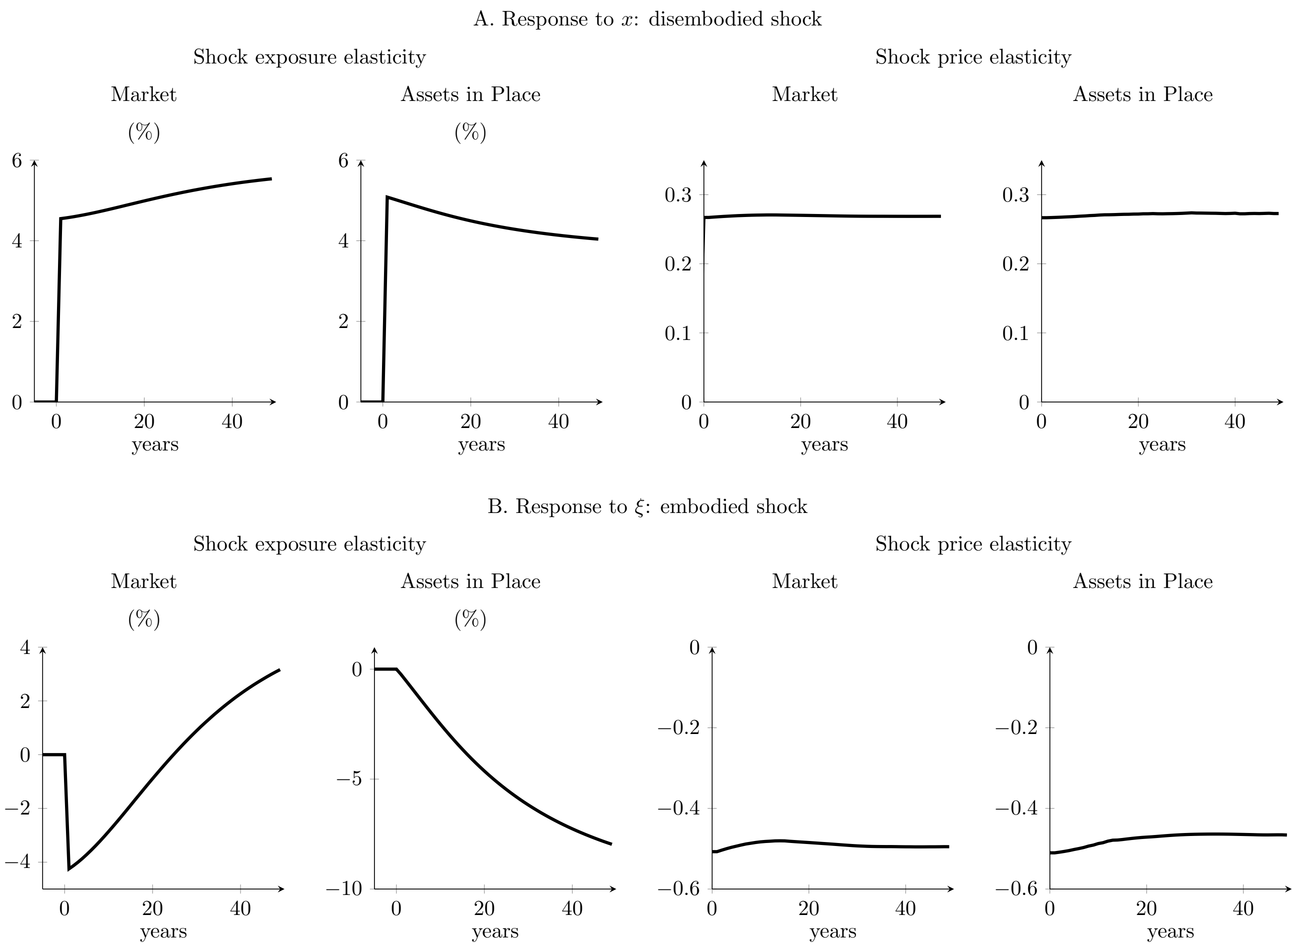

Figure plots shock-exposure and shock-price elasticities of the aggregate dividend process D, and the dividends from asset in place D^{ap} to the two technology shocks in the model. We construct the shock exposures taking into account the nonlinear nature of equilibrium dynamics: we introduce an additional shock of magnitude $\sigma\sqrt{dt}$ at time $0+dt$ without altering the realizations of all future shocks. We then scale the resulting impulse responses by $1/\sqrt{dt}$. We compute these elasticities at the mean of the stationary distribution of ω.

We conclude that the model's implications about the term structure of risk premia stem mostly from the dynamics of cash flows. We can thus see that the contribution of the dividend dynamics induced by x to the equity premium rises modestly with the horizon. Panel B implies the opposite pattern; the contribution of the dividend dynamics induced by ξ to the equity premium is concentrated in the short and medium run, and the rise in long-run dividends contributes negatively to the equity premium. Thus, the equity premium in the model is concentrated at shorter maturities. In terms of assets in place, the risk due to x is somewhat higher in the short run, while the exposure to ξ-shocks is negative and increases in magnitude with maturity. These patterns imply that, in the model, the value premium is most pronounced at longer maturities.

1.2 Firm risk exposures and the market-to-book ratio

We next examine the extent to which Q is a useful summary statistic for firm risk and risk premia in our model. Recall that a firm's log market-to-book ratio can be written as

$$\log Q_{ft} - \log Q_t = \log \left[\frac{V_t}{V_t + G_t} \, (1 + \widetilde{p}(\omega_t) \, (\bar{u}_{ft} - 1)) + \frac{G_t}{P_t + G_t} \frac{\bar{u}_{ft}}{z_{ft}} \left(1 + \widetilde{g}(\omega_t) \, \left(\frac{\lambda_{ft}}{\lambda} - 1 \right) \right) \right], \tag{6}$$

Examining (6), we note that a firm's market to book ratio is increasing in the likelihood of future growth λ_f, decreasing in the firm's relative size z_f, and increasing in the firm's current productivity $\bar{u}_f$. This latter effect prevents Tobin's Q from being an ideal measure of growth opportunities, since it is contaminated with the profitability of existing assets.

To examine the extent to which Q is correlated with technology risk exposures, we examine how changes in the firm's current state $(\lambda_{ft}, k_{ft}, \bar{u}_{ft})$ jointly affect both firm Q and risk exposures. We plot the results in Figure 4. In each of the three columns, we vary one of the elements of the firm's current state (λ_{ft}, z_{ft}, or $\bar{u}_{ft}$) and keep the other two constant at their steady-state mean. On the horizontal axis, we plot the change in the firm's Q (relative to the market). On the vertical axis we plot the firm's return exposure to x and ξ (panels A and B, respectively) and the firm's risk premium (panel C). We scale the horizontal axis so that it covers the 0.5% and 99.5% of the steady-state distribution of the each of the firm's state variables. The resulting cross-sectional distribution of Tobin's Q is highly skewed.

Examining Figure 4, we see that regardless of the source of the cross-sectional dispersion in Q, the relation between Q and risk exposures is positive. The pattern in the first two columns is consistent with the impulse responses in the previous section – small firms with high probability of acquiring future projects have higher technology risk exposures than large firms with low growth potential. The last column shows that increasing the firm's current productivity $\bar{u}_f$ – holding λ_f and k_f constant increases both Q and risk exposures. This pattern might seem puzzling initially, since increasing productivity $\bar{u}_f$ while holding size k_f and investment opportunities λ_f constant will lower λ_f/z_f. However, altering $\bar{u}_f$ also has a cashflow duration effect: due to mean reversion in productivity, profitable firms have lower cashflow duration – their cashflows are expected to mean-revert to a lower level. This lower duration of high $\bar{u}_f$ firms implies that their valuations are less sensitive to the rise in discount rates following a positive technology shock– see the response of the interest rate in the paper. In our calibration, this duration effect overcomes the effect due to λ_f/z_f, implying a somewhat more positive stock price response for high-$\bar{u}_f$ firms. However, the magnitude of this effect is quantitatively minor.

The last row of Figure 4 shows how the firms' risk premium (unlevered) is related to cross-sectional differences in Q. Recall that the two technology shocks carry risk premia of the opposite sign. The disembodied shock carries a positive risk premium; in the absence of other technology shocks, this would imply that firms' risk premia rise with their market-to-book ratio. However, the fact that the embodied shock carries a negative risk premium – coupled with its higher volatility – implies a lower risk premium for growth firms relative to value firms. Households are willing to accept lower average returns for investing in growth firms because doing so allows them to partially hedge the displacement arising from the embodied shock ξ – the decline in their continuation utility.

Figure 4: Firm risk exposures and risk premia

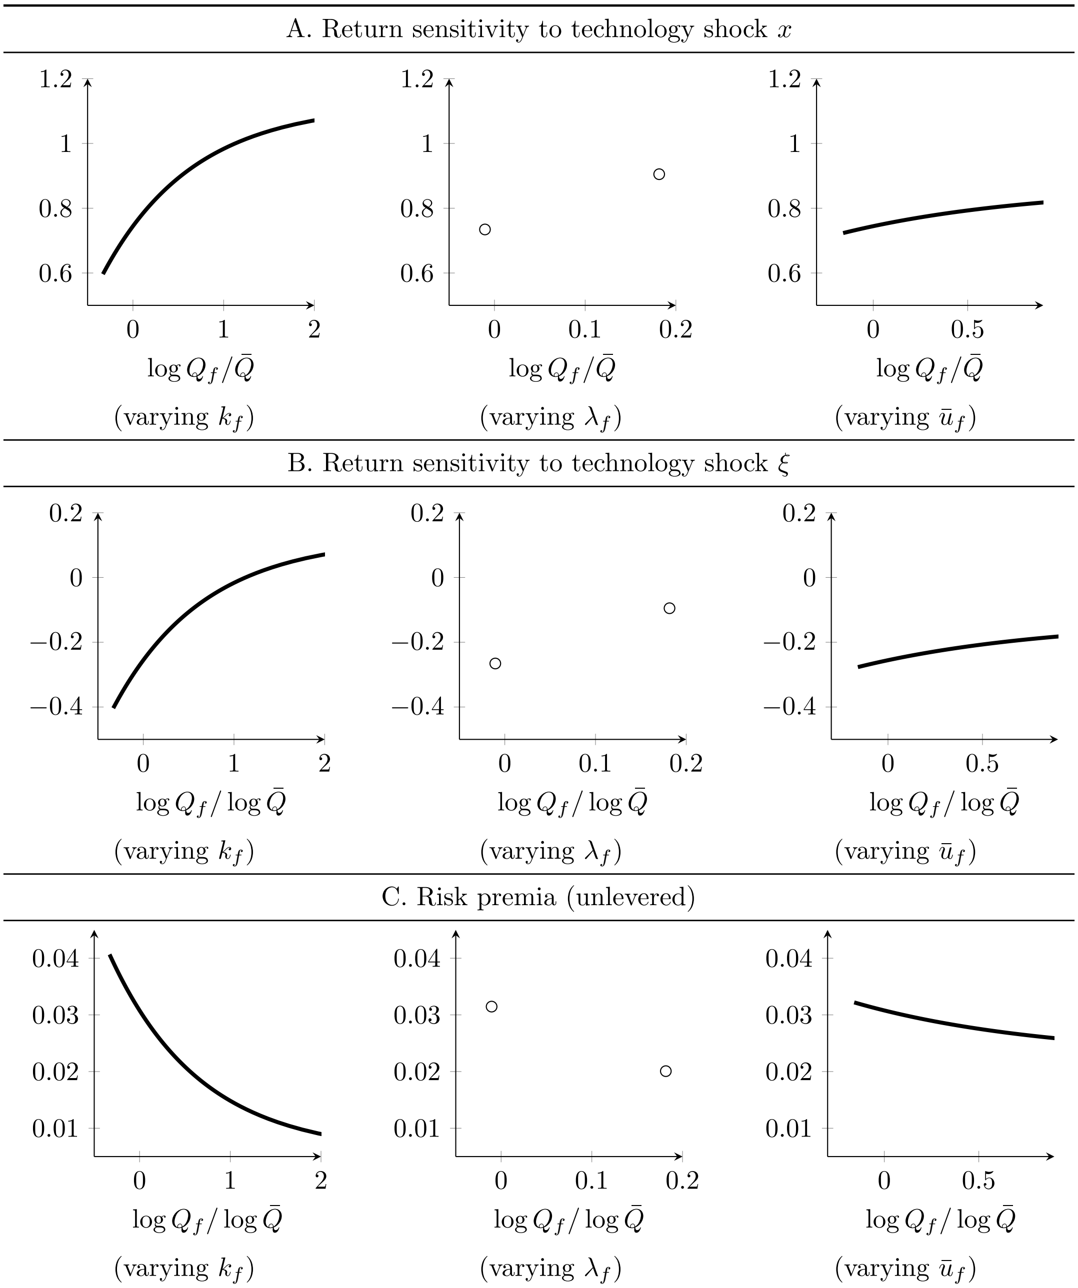

Figure illustrates how technology risk exposures (Panels A and B) and risk premia (Panel C) vary with the firm's market-to-book ratio (Q). A firm's market-to-book ratio is a function of the firm's relative size k_f, likelihood of future growth λ_f, and its current productivity $\bar{u}_f$. In each of the three columns, we examine how variation in Q due to each of these three state variables translates into variation in risk premia – while holding the other two at their average values, i.e. $\lambda_f = \lambda$, $k_f = 1$ and $\bar{u}_f = 1$. The range in the x-axis corresponds to the 0.5% and 99.5% of the range of each these three variables in simulated data. The value of the state variable ω is set to its unconditional mean, $\omega = \mathrm{E}[\omega_t]$.

1.3 Stochastic discount factor

Tests using linearized pricing models are common in the empirical asset pricing literature. Here we use the estimates of the linearized stochastic discount factor (SDF) as a reduced-form summary of the model's implications for asset prices. The two technology shocks x and ξ are unobservable in the data. In the model, the two shocks are spanned by shocks to $\hat{\omega}$, and either aggregate consumption C, aggregate output Y or labor productivity x. To compare the model with the data, we estimate the following linearized SDF which is a function of observables,

$$\hat{m}_t = \Delta \log \hat{\pi}_t = a - b\, Z_t. \tag{7}$$

where Z_t is a vector of shocks. The first element of Z is our proxy for the relative value of new blueprints, $\hat{\omega}$. The second element, is either, the annual growth rate in aggregate consumption, aggregate output or total-factor productivity (adjusted for utilization from Basu, Fernald, and Kimball (2006)). In the model, the first element is $\hat{\omega}$, while the second is either the growth rate in aggregate consumption, aggregate output or the disembodied shock x.

We follow standard practice and estimate the parameters b using the moment conditions $E[\hat{m}_t\, R^e_t] = 0$, where R^e are the excess returns on 10 value-weighted portfolios of firms sorted by book-to-market. We normalize a so that $E[\hat{m}_t] = 1$, which reduces the moment conditions to $E[R^e_t] = -cov(\hat{m}_t, R^e_t)$ In addition to the estimated risk prices b, we also report the cross-sectional R^2, the mean absolute pricing error (MAPE), Hansen's J-test and its associated p-value. We construct empirical confidence intervals for the latter three statistics using a jackknife estimator.

Table 1 compares estimates of (7) in the data and in the model. To facilitate this comparison, we normalize $\Delta\hat{\chi}$ and $\Delta\hat{\omega}$ to unit standard deviation. Examining Table 1, we see that specifications of the SDF that include only consumption, output or TFP growth – those typically implied by the existing general equilibrium models – are rejected by the data. The same is true in our model. By contrast, the two-factor specification (7) captures return differences on the test portfolios relatively well. Moreover, the magnitude of the estimated coefficients b_2 is quite similar between the model and the data.

Table 1: Estimates of a linearized Stochastic Discount Factor (Data vs Model)

	A. Data						B. Model					
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Δx	1.89	0.41					0.50	0.28				
	[0.85]	[0.76]					[0.18]	[0.21]				
$\Delta \log C$			1.79	0.01					0.80	0.46		
			[0.95]	[0.35]					[0.31]	[0.32]		
$\Delta \log Y$					2.06	0.37					0.83	0.45
					[1.2]	[0.35]					[0.33]	[0.32]
$\Delta\hat{\omega}$		-0.80		-0.83		-0.69		-0.69		-0.63		-0.69
		[0.38]		[0.27]		[0.24]		[0.38]		[0.39]		[0.39]
R2	0.35	0.86	-0.20	0.84	-1.38	0.80	-0.29	0.50	-0.20	0.50	-0.48	0.50
	[3.12]	[0.22]	[3.63]	[0.1]	[3.53]	[0.11]	[0.44]	[0.4]	[0.57]	[0.38]	[0.74]	[0.38]
MAPE (%)	1.28	0.66	1.93	0.72	2.89	0.86	1.68	0.93	1.60	0.92	1.76	0.92
	[1.13]	[0.32]	[1.69]	[0.28]	[1.67]	[0.25]	[0.53]	[0.31]	[0.53]	[0.31]	[0.58]	[0.31]
J-stat	21.71	9.05	39.55	9.99	41.08	9.99	39.75	26.23	40.45	26.54	42.30	26.54
	[21.35]	[14.3]	[13.74]	[8.79]	[14.06]	[9.03]	[24.23]	[20.71]	[26.04]	[21.11]	[26.15]	[21.1]
(p-val)	0.01	0.43	0.00	0.35	0.00	0.35	0.02	0.13	0.02	0.13	0.02	0.13

Table shows estimates of a linearized SDF $m = a - bF$ using the cross-section of 10 portfolios sorted on book-to-market. We normalize a so that $E[m] = 1$, which yields the moment conditions $E[R^e] = -cov(R^e, m)$. We consider various combinations of F, including utilization-adjusted TFP (proxied by Δx in the model), aggregate consumption growth, aggregate output (GDP) growth and changes in $\hat{\omega}$. Standard errors in the data are HAC (3-lags) and include a Shanken adjustment. We report mean parameter estimates across simulations and 90% confidence intervals. In addition, we report the cross-sectional R^2, the mean absolute pricing error (MAPE), and the Hansen J statistic. Standard errors for the last three statistics are computed using a delete-d Jackknife estimator, where $d = 2\sqrt{T}$ rounded to the nearest integer.

2 Model Estimation

We first describe the estimation procedure. We then discuss the sensitivity of the moments to parameters.

2.1 Estimation Details

Here, we discuss some details of the estimation procedure. First, it is convenient to transform some of the parameters of the model. Specifically, we replace the volatility of the idiosyncratic shock with the variance of its ergodic distribution,

$$v \equiv \frac{\sigma_u^2}{2\,\kappa_u - \sigma_u^2}. \tag{8}$$

Further, in place of the project arrival rate parameters $[\lambda_L, \lambda_H]$, we formulate the model in terms of the mean arrival rate,

$$\lambda \equiv \frac{\mu_L}{\mu_L + \mu_H}\lambda_L + \frac{\mu_H}{\mu_L + \mu_H}\lambda_H \tag{9}$$

and the relative difference,

$$\lambda^R \equiv \frac{\lambda_H - \lambda_L}{\lambda_L}. \tag{10}$$

The RBF algorithm searches for an approximate minimum over a rectangular set. Table 2 reports the bounds of this set. We confirm that the estimated parameters lie in the interior of the set.

Table 2: Parameter Set

Parameter	Lower Bound	Upper Bound	Estimate
γ	5	165	105.856
θ	0.1	3	2.207
h	0	1	0.947
α	0.1	0.7	0.362
η	0	1	0.785
σ_x	0.02	0.2	0.077
σ_ξ	0.02	0.2	0.137
μ_L	0.005	0.5	0.283
μ_H	0.005	0.5	0.015
λ^R	0	100	69.237
λ	0.1	6	0.543
δ	0.02	0.1	0.033
μ_x	0	0.05	0.016
μ_ξ	0	0.05	0.004
υ	0.1	2.5	2.000
ρ	0.025	0.07	0.040
κ_u	0.1	0.5	0.303

2.2 Elasticity of moments to parameters

We examine the sensitivity of model-implied moments $\mathcal{X}(p)$ to small changes in parameters

$$\frac{\partial \mathcal{X}_i(p)}{\partial p_j} \frac{p_j}{\mathcal{X}_i(p)} \tag{11}$$

We computed elasticities with respect to the original model parameters. We emphasize that these are local sensitivities, that is, they are computed at the estimated parameter values. With that caveat in mind, we plot the elasticities in Figures 5 through 7.

Examining Figure 5, we see that the moments of aggregate consumption growth are particularly sensitive to the parameters governing the disembodied shock – μ_x and σ_x. The long run volatility of consumption is highly sensitive to σ_x and to some degree σ_ξ. The volatility of the embodied shock σ_ξ is particularly important for the volatility of investment growth – and to some extent so is the adjustment cost parameter α. The correlation between investment and consumption is thus determined by σ_x, σ_ξ and to some extend α. Last, the volatility of the embodied shock ξ, the adjustment cost parameter α and the preference share of relative consumption h are important in determining the moments of net payout. The preference for relative consumption affects households' effective elasticity of inter-temporal substitution much more than θ at the optimum because h is close to one.

In Figure 6 we see that utility curvature γ, the preference share of relative consumption h and the degree of market incompleteness η matter for the equity premium. The volatility of the market portfolio is sensitive to the volatilities of the two technology shocks and η and the surplus that accrues to innovators. The mean of the risk-free rate depends on ρ, γ and h, while its volatility depends on σ_ξ, the curvature parameter α, the rate of depreciation, and preference parameters. The magnitude of the value premium is a function of preference parameters (h), the volatility of the embodied shock σ_ξ the moments governing the evolution of $\lambda_{f,t}$, and the degree of market incompleteness. The volatility of the value factor depends on σ_ξ and the dynamics of u. The last panel in Figure 6 shows that the CAPM alpha is locally most sensitive to h.

Last, Figure 7 shows how the cross-sectional firm moments depend on parameters. The dispersion in firm investment rates depends on the volatility of the embodied shock ξ, λ_L and λ_H; the persistence of investment depends on the moments of $\lambda_{f,t}$ and the rate of depreciation. The dispersion and persistence of Tobin's Q depends on the properties of $\lambda_{f,t}$ and preference parameters. The sensitivity of firm investment to Tobin's Q depends largely on λ_L. Last, the mean and persistence of firm profitability largely depend on the parameters governing firm productivity – σ_u and κ_u.

References

Bansal, R., R. F. Dittmar, and C. T. Lundblad (2005, 08). Consumption, Dividends, and the Cross Section of Equity Returns. *Journal of Finance 60*(4), 1639–1672.

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Borovicka, J., L. P. Hansen, and J. A. Scheinkman (2014, May). Shock Elasticities and Impulse Responses. NBER Working Papers 20104, National Bureau of Economic Research, Inc.

Hansen, L. P., J. C. Heaton, and N. Li (2005). Intangible risk. In *Measuring Capital in the New Economy*, NBER Chapters, pp. 111–152. National Bureau of Economic Research, Inc.

Figure 5: Sensitivity of moments to parameters

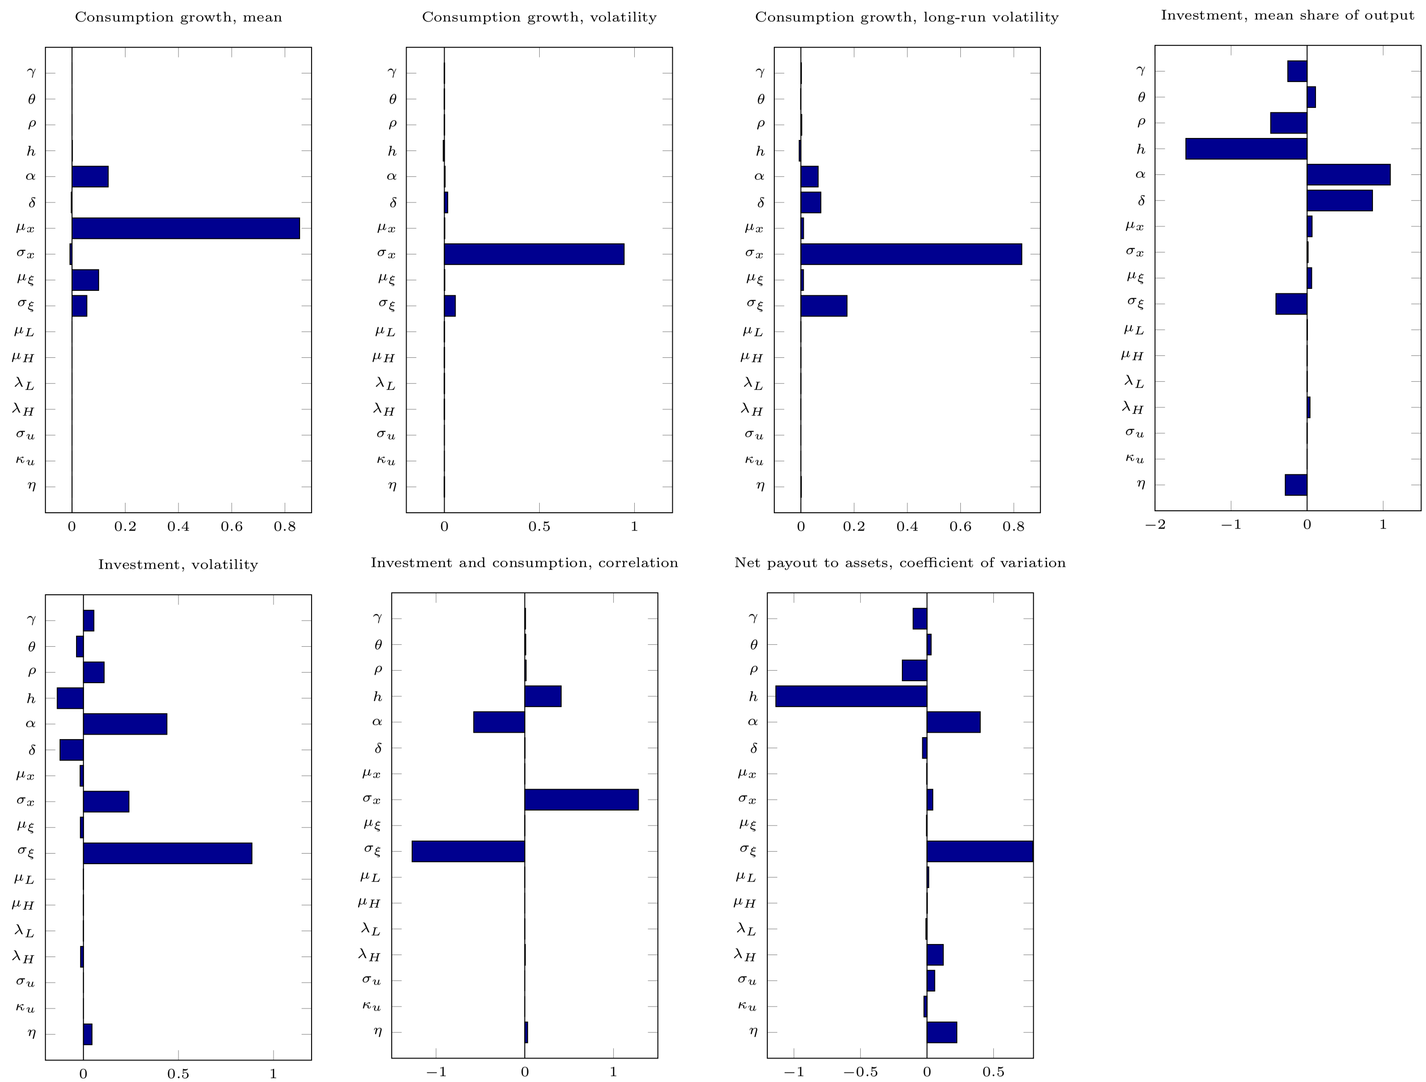

We report the elasticity of moment estimates $\mathcal{X}(p)$ to parameters p.

Figure 6: Sensitivity of moments to parameters (cont)

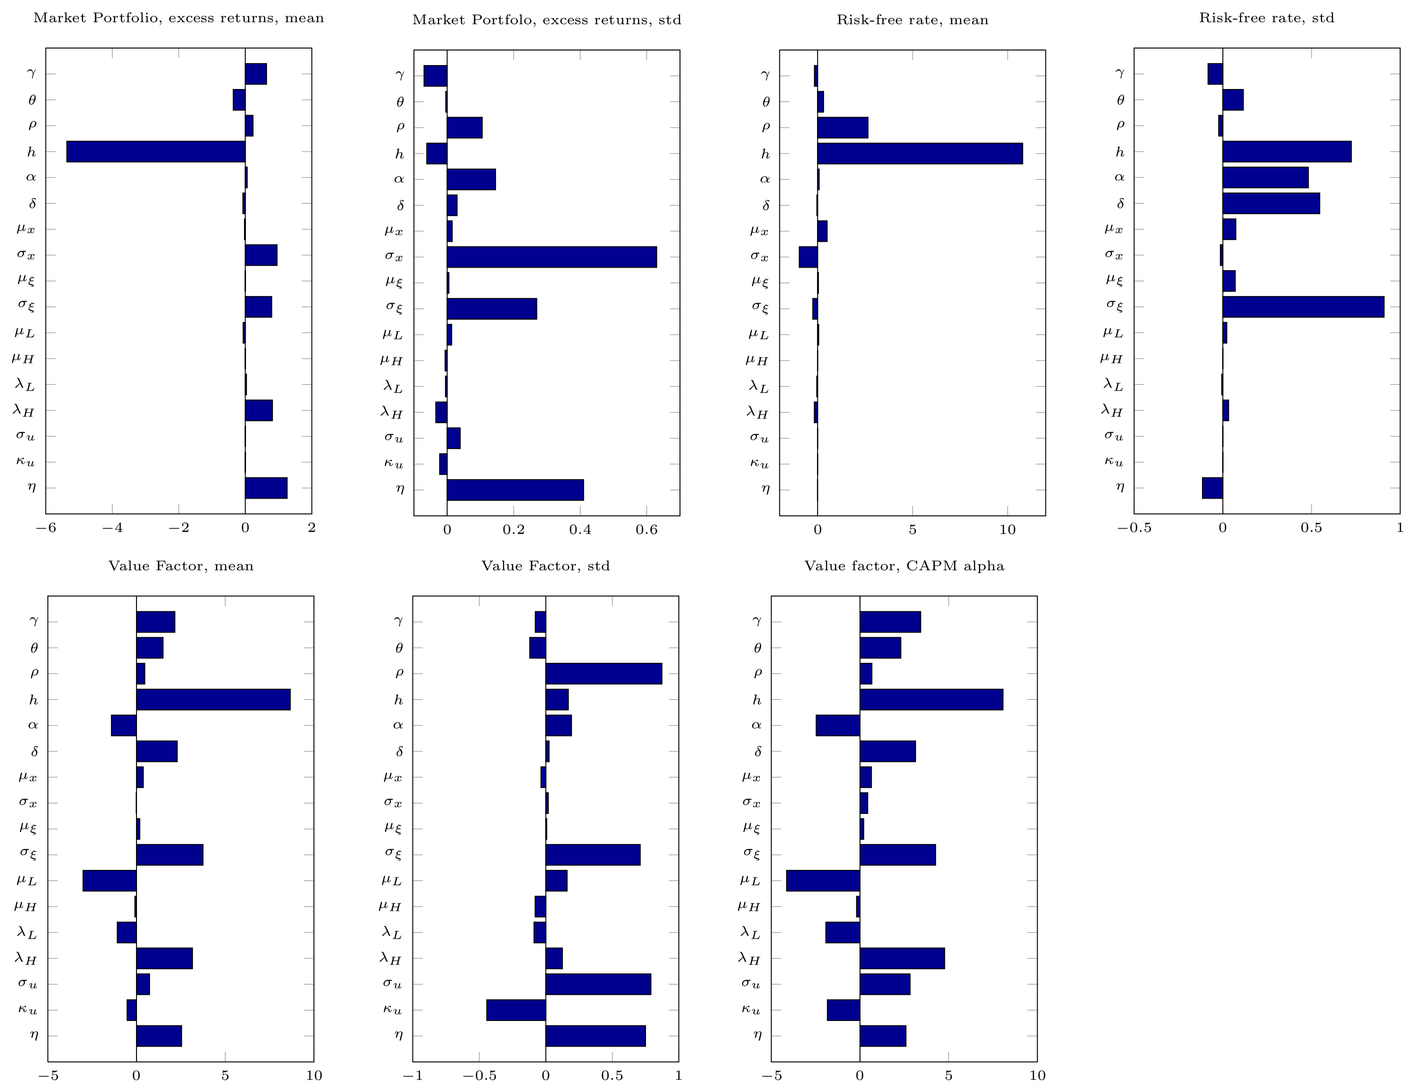

We report the elasticity of moment estimates $\mathcal{X}(p)$ to parameters p.

Figure 7: Sensitivity of moments to parameters (cont)

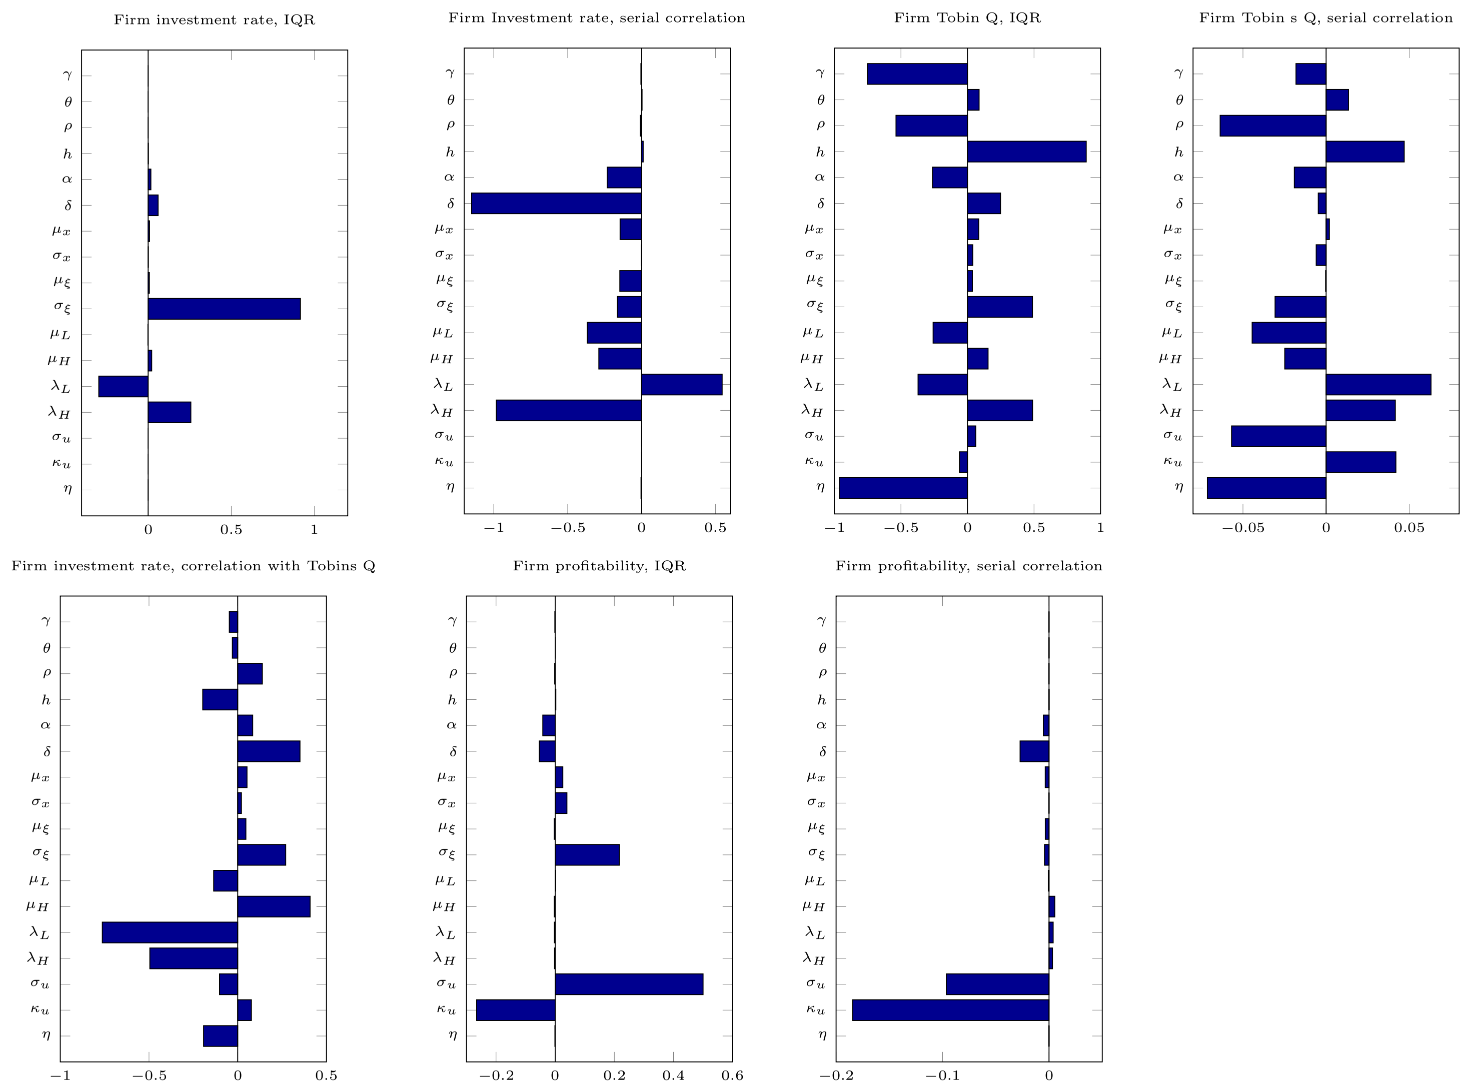

We report the elasticity of moment estimates $\mathcal{X}(p)$ to parameters p.