Journal of Monetary Economics 163 (2026) 103987



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme





The innovation long-run risk component

Fabio Franceschini

Department of Economics, University of Bologna, Piazza Scaravilli 2, Bologna, 40126, Italy

ARTICLE INFO

Dataset link: Effective R&D (Original data)

Keywords:
R&D
Long-run risk
Asset pricing
Cointegration

ABSTRACT

This paper provides empirical evidence that shocks to aggregate Research and Development (R&D) exert persistent effects on macroeconomic dynamics and represent a significant source of risk for investors, consistent with 'long-run risk' models. The analysis focuses on 'effective R&D', the unique scaling of R&D that, across a broad class of models, maintains a linear relationship with productivity growth. This measure theoretically captures the entire contribution of R&D to productivity, flexibly accounting for knowledge spillovers and product proliferation. Importantly, its deviations from equilibrium are empirically identifiable via the error correction term in the cointegration relationship among R&D, total factor productivity, and the labor force. Using U.S. data, structural R&D shocks, based on effective R&D, are shown to affect productivity and consumption growth rates beyond business cycle horizons, and to be associated with a significant risk premium (around 12% annually) in a cross section of stock and bond portfolios, with cash-flow risk a key determinant.

1. Introduction

Investments in Research and Development (R&D) have a profound and well-documented impact on the economy. At the aggregate level, R&D has been linked to slow-moving fluctuations in macroeconomic quantities such as productivity and consumption (Comin and Gertler, 2006), and theory has shown that this nexus carries substantial implications for asset prices (Kung and Schmid, 2015). The underlying mechanism is that of the 'Long-Run Risk' (LRR) framework from Bansal and Yaron (2004), who first demonstrated how persistence in consumption growth enables macroeconomic models to account for the large equity risk premium observed in financial markets, far exceeding what consumption growth volatility alone would imply. This framework has since become a central reference in macro-finance research,[1] but has also faced sustained criticism,[2] making rigorous empirical validation of models built on it essential.

This paper documents a significant link between the dynamics of aggregate R&D and asset prices, empirically corroborating the LRR approach. Specifically, it provides robust empirical evidence that shocks to US aggregate R&D have persistent effects on productivity and consumption growth, and that such shocks constitute a risk priced by investors. The analysis pivots on a theoretical measure of aggregate R&D – 'effective R&D' – that uniquely summarizes the time-varying impact of R&D on productivity growth, emerging as the only transformation of aggregate R&D that is linearly related to it. Building on this, I propose an empirical framework to recover the structural R&D shocks embedded in effective R&D, thus scaled by their effectiveness, and illustrate the persistent effects these have on the economy. In turn, these shocks prove a significant risk factor in pricing a cross-section of financial assets.

E-mail address: franceschini-f@pm.me.
URL: https://www.fabiofranceschini.com/.

[1] Applications include exchange rate dynamics (Colacito and Croce, 2011), climate change asset pricing (Bansal et al., 2016), term structure models (Ai et al., 2018), and oil price dynamics (Ready, 2018).

[2] Notably, Constantinides and Ghosh (2011), Beeler and Campbell (2012), and Epstein et al. (2014).

https://doi.org/10.1016/j.jmoneco.2026.103987
Received 31 October 2025; Received in revised form 19 August 2026; Accepted 19 August 2026
Available online 27 August 2026
0304-3932/© 2026 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

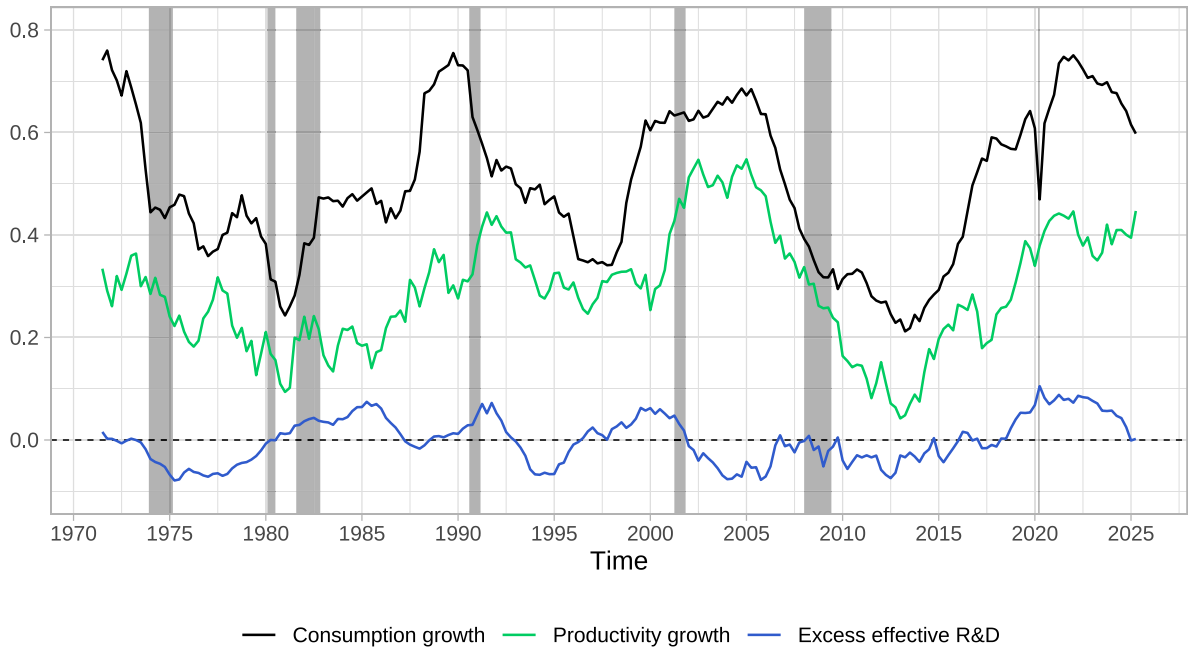


Fig. 1. The innovation component of productivity long-run risk. The figure plots excess effective R&D (its demeaned level) alongside low-frequency components of consumption and productivity growth. Effective R&D is estimated as in Section 4.2; the low-frequency components are the 6th component of the Ortu et al. (2013) decomposition (half-life 8–16 years), applied to the growth rates of U.S. real per capita consumption (BEA) and utilization-adjusted TFP (Fernald, 2012).

This constitutes a twofold novelty, introducing the first empirically significant risk factor that is (1) based on aggregate R&D, and (2) constructed from structural shocks in a theoretically identified system.

Shocks with highly persistent effects on growth have long been a focus of the macro-finance literature because they carry substantial theoretical implications while being barely detectable in the data. Specifically, they generate large fluctuations in long-run consumption forecasts even when their variance is small. This feature becomes theoretically powerful when combined with recursive preferences à la Epstein and Zin (1989), which feature aversion to forecast swings, allowing models to account for sizable risk premia while preserving consistency with standard macroeconomic unconditional moments. This feature, however, also renders empirical validation of these models elusive, precisely because persistent consumption growth shocks are inherently difficult to identify in finite samples. A growing literature has tackled this empirical challenge directly, either by developing refined econometric approaches (Ortu et al., 2013; Dew-Becker and Giglio, 2016; Schorfheide et al., 2018; Gourieroux and Jasiak, 2024; Bryzgalova et al., 2025) or by constructing new data (Liu and Matthies, 2022). A complementary strand has instead focused on deriving the key LRR conditions as reduced forms of richer structural models, thereby yielding additional testable implications (notably Kaltenbrunner and Lochstoer, 2010; Croce, 2014). This paper contributes to both strands. On the empirical side, it proposes a novel econometric approach, built on theory-implied conditions, to identify a process driving revisions in long-run expectations – effective R&D. On the theoretical side, it builds on a set of economic mechanisms highlighted by the LRR literature and provides evidence supporting its central predictions.

More specifically, the link between aggregate R&D and asset prices was first formalized by Kung and Schmid (2015), who developed a fully endogenous growth model in which R&D persistence can emerge, generating a 'productivity LRR component', i.e. low-frequency fluctuations in productivity growth transmitted to consumption growth.[3] This paper builds on their framework and provides broad empirical evidence consistent with its main predictions, while departing from it in two key respects. First, in light of the ongoing debate on scale effects in the growth literature (see Bloom et al., 2020), the analysis draws on the framework of Jones (1999), which accommodates both fully and semi-endogenous growth mechanisms, enabling a more robust empirical characterization of innovation dynamics. This specification allows effective R&D to account for spillover (net of fishing-out) and product-proliferation effects of arbitrary strength, absorbing the nonlinearities both introduce, and nesting Kung and Schmid (2015) as a special case. The resulting effective R&D measure is stationary, enabling the use of standard econometric methods without risk of spurious inference.[4] Second, R&D dynamics are modeled in a multivariate setting. This enables the isolation of unexpected

[3] Further implications of persistent R&D fluctuations for the broader macroeconomy have been studied, among others, by Benigno and Fornaro (2018), Aksoy et al. (2019), and Beqiraj et al. (2025).

[4] Kung and Schmid (2015) refer to effective R&D as 'R&D intensity', which is non-stationary in their empirical specification. Further details on the comparison appear in Appendix D.1.

variation in R&D, scaled by its effectiveness, net of contemporaneous productivity fluctuations. By preventing non-R&D productivity risk from leaking into reduced-form estimates of the risk premium associated with R&D shocks, this approach yields estimates more cleanly attributable to innovation risk and its long-run risk channel.

Effective R&D expresses aggregate R&D in units of expected productivity growth, up to scale, and is defined as a log-linear combination of aggregate R&D, the stock of ideas, and product variety. The well-established stationarity of productivity growth implies that these variables form a cointegrating relationship that can be estimated.[5] I then obtain a gross estimate of the effective R&D time series by mapping the stock of ideas and product variety into observable variables – the productivity level itself and labor input – following standard approaches in the literature. Because the cointegrating relationship is estimated with these substitutions, the resulting effective R&D measure is contaminated by non-R&D productivity drivers, departing from the theoretical object of interest. Crucially, the R&D shocks obtained from gross effective R&D are properly effectiveness-adjusted regardless of this contamination, since the structural shocks are mapped recursively to the observed variables, with the non-R&D productivity shocks absorbing the contamination. Within this setting, the effects of R&D shocks on growth build only gradually and persist over many years, possibly beyond ten, in line with the evidence in Antolin-Diaz and Surico (2025) on public R&D shocks and with the adoption-driven mechanism of Anzoategui et al. (2019). Under stricter assumptions, I further propose an estimator that recovers net effective R&D levels, used throughout the analysis. Fig. 1 plots measures of the productivity LRR component alongside the estimated effective R&D, which underlies it as the 'innovation component' of long-run risk.

On the asset pricing side, the cross-sectional risk premium associated with innovation long-run risk is estimated on a wide cross-section of assets using the procedure of Giglio and Xiu (2021), which is designed to mitigate omitted risk factor bias. The resulting risk premia are not statistically significant when the innovation risk factor is based on contemporaneous shocks alone, but highly significant for multi-year rolling means, at about 12% annually, consistent with the investor underreaction to R&D news documented in Eberhart et al. (2004). Following Bansal et al. (2005), I also find that a substantial part of innovation long-run risk originates in cash flows, a key feature of long-run risk models, augmenting standard portfolio sortings with firm characteristics previously linked to R&D investment. These results connect to both the corporate finance literature on R&D investment and the asset pricing literature on firm-specific R&D risk premia (e.g. Chan et al., 2001). Hsu (2009) and Kogan et al. (2017) likewise relate aggregate innovation measures to asset prices empirically, but neither provides explicit cross-sectional risk premium estimates. Moreover, relative to the former, the effective R&D measure is derived from theory; relative to the latter, it is complementary, since their measure captures innovation outcomes while effective R&D captures inputs (effort).

The paper proceeds as follows. Section 2 presents the theoretical framework, defining effective R&D, and relating it to macroeconomic dynamics and asset prices. Section 3 illustrates the mapping between theoretical and observable quantities. Section 4 estimates the effective R&D measures. Section 5 identifies R&D shocks and documents their macroeconomic impact. Section 6 reports the cross-sectional asset pricing results. Section 7 concludes.

2. Theoretical framework

2.1. The price of risk from persistent macroeconomic shocks

In a discrete-time, arbitrage-free economy, the expected return of any asset i between times $t-1$ and t, in excess of the risk-free rate R^f, is proportional to the covariance between its return R^i_t and a stochastic discount factor (SDF) M_t[6]:

$$\mathbb{E}\left[R^i_t\right] - R^f = -R^f \cdot \mathrm{Cov}\left[M_t, R^i_t\right] . \tag{1}$$

To discipline and better understand the dynamics of the SDF, asset pricing theory often relates it to the intertemporal marginal rate of substitution (IMRS) of a representative agent with preferences over consumption streams. Shocks to state variables that matter for investor welfare become the key drivers of asset valuations.

As illustrated by Bansal and Yaron (2004), persistency in consumption growth strongly affects the IMRS and, consequently, asset valuations, when the representative agent has recursive preferences as in Epstein and Zin (1989). These preferences separate the intertemporal elasticity of substitution (IES) from risk aversion, and imply that the agent is sensitive both to contemporaneous consumption shocks

$$\varepsilon^c_t = \Delta \ln C_t - \mathbb{E}_{t-1}\left[\Delta \ln C_t\right] , \tag{2}$$

and to shocks to long-run consumption prospects

$$\varepsilon^x_t = \left\{\mathbb{E}_t - \mathbb{E}_{t-1}\right\}\left(\sum_{j=1}^{\infty}(\kappa_x)^j \cdot \Delta \ln C_{t+j}\right) , \tag{3}$$

as reflected in the log-SDF,

$$\ln M_t = \mathbb{E}_{t-1}\left[\ln M_t\right] - b_c\, \varepsilon^c_t - b_x\, \varepsilon^x_t , \tag{4}$$

[5] Cointegration models have been widely applied to test between endogenous growth specifications (e.g., Ha and Howitt, 2007; Bottazzi and Peri, 2007; Madsen, 2008; Kruse-Andersen, 2023); this paper instead emphasizes the dynamic properties of R&D and their financial implications.

[6] A standard reference is Cochrane (2005).

where the constant $\kappa_x \in (0,1)$ is a function of the equilibrium consumption–wealth ratio, and the constants $b_c > 0$ and $b_x > 0$ are the SDF loadings on contemporaneous and long-run consumption shocks, respectively. Then, the more persistently a shock affects consumption growth, the larger the fluctuations it generates in prospects of the whole consumption path, captured by ε_t^x, and hence the stronger its impact is on the IMRS.[7]

Assuming that ε_t^c and ε_t^x are also the key stochastic determinants of asset return dynamics,[8] the following reduced-form pricing equation holds:

$$\mathbb{E}\left[R_t^i\right] - R^f = \lambda_c \beta_c^i + \lambda_x \beta_x^i, \tag{5}$$

where the β_j^i for $j \in \{c, x\}$ measure asset i's sensitivities to each shock, and the λ_j are the market-wide prices of risk, per unit of exposure to shock j, each proportional to $\text{Var}\left[\varepsilon_t^j\right]$. The economic interpretation is that the agent dislikes exposure to fluctuations in both realized and expected future consumption growth; holding assets with such exposure therefore requires compensation in the form of higher expected returns. Importantly, the more persistently a shock propagates into future consumption growth, the larger the variance of the associated long-run shock, and the greater the market price of long-run risk λ_x, with broad implications for the cross-section of asset returns.

Since aggregate consumption expectations respond to many shocks, long-run risks can originate from multiple economic sources. Productivity is a particularly important one: positive news induces agents to postpone consumption, producing fluctuations in expected consumption growth across many horizons, when productivity shocks are sufficiently persistent and the IES is sufficiently high (Croce, 2014; Ortu et al., 2013; Kaltenbrunner and Lochstoer, 2010). Building on the idea that productivity dynamics are themselves shaped by the long-lasting effects of innovation, Kung and Schmid (2015) argue that R&D is a key source of long-run risk.

2.2. The innovation component of productivity growth

The specification of the link between R&D and aggregate productivity growth is actively debated within the endogenous growth literature. To investigate how innovation dynamics drive long-run risk, I build on the framework of Jones (1999), which nests the main alternatives as special cases without committing to any one of them.[9] Final goods output Y_t is produced according to:

$$Y_t = Z_t L_t, \tag{6}$$

where L_t denotes labor and Z_t is aggregate productivity.[10] Productivity is determined by the average stock of ideas I_t over the mass of product varieties Q_t, and by a zero-mean process a_t separately capturing the contribution of factors outside the technological domain (such as misallocation):

$$Z_t = I_t^{\theta} Q_t^{\xi} e^{a_t}, \tag{7}$$

where $\theta > 0$ and $\xi > 0$ capture the returns to the stock of ideas and the love of variety, respectively. In this framework, while Q_t shapes the technological environment, only I_t is driven by aggregate R&D expenditures S_t and constitutes the innovation component of technology,[11] with the law of motion

$$I_{t+1} = (1-\phi) I_t + \chi \cdot S_t^{\eta} I_t^{\psi} Q_t^{-1}, \tag{8}$$

where $\phi \in [0,1)$ is the probability of ideas becoming obsolete, $\chi > 0$ is a scale parameter, $\eta \in (0,1]$ controls duplication in R&D efforts, and ψ governs the strength of spillovers from past ideas (net of fishing-out effects).[12] The term $S_t^{\eta} I_t^{\psi} Q_t^{-1}$ reflects the gross new average ideas, where the technological frontier determines the effectiveness of R&D in creating them: past ideas I_t raise it (in levels), while product variety Q_t erodes it by converting gross ideas into per-variety terms.

Together, (7) and (8) provide enough structure to characterize productivity dynamics, yielding the approximation[13]

$$\Delta \ln Z_{t+1} \approx \gamma_0 + \gamma_1 \left(\ln S_t - \frac{1-\psi}{\eta} \ln I_t - \frac{1}{\eta} \ln Q_t \right) + \xi \Delta \ln Q_{t+1} + \Delta a_{t+1}, \tag{9}$$

[7] For instance, if consumption growth simply followed $\Delta \ln C_t = \rho_c \Delta \ln C_{t-1} + \breve{\varepsilon}_t$ with $|\rho_c| < 1$ and $\breve{\varepsilon}_t$ i.i.d., then the long-run shock would be $\varepsilon_t^x = \frac{1}{1-\kappa_x \rho_c} \breve{\varepsilon}_t$, whose magnitude, conditional on $\breve{\varepsilon}_t$, depends crucially on the persistence parameter ρ_c.

[8] A factor structure such as $R_t^i = \bar{R}^i + \beta_c^i \varepsilon_{c,t} + \beta_x^i \varepsilon_{x,t} + e_t^i$ is often assumed for returns; the LRR framework specifically assumes this structure for assets' cash-flow growth rates.

[9] The framework spans the fully-endogenous approach (e.g., Peretto, 1998), where spillovers persist indefinitely, and the semi-endogenous approach (e.g., Jones, 1995), where they eventually die out. Steady-state growth depends on R&D intensity or population growth alone, respectively.

[10] The analysis extends to any Cobb–Douglas specification with additional rival inputs, such as capital.

[11] This is a 'lab-equipment' formulation: new ideas are produced from final goods rather than labor. Kruse-Andersen (2023) adopt a similar reformulation of Jones (1999) and Sedgley and Elmslie (2013) study the dynamics of this class of models.

[12] This specification nests Kung and Schmid (2015) as a special case and admits balanced growth paths under regime-specific additional restrictions; see Appendix A.2.

[13] Ha and Howitt (2007) and Kruse-Andersen (2023) employ similar approximations on annual data; the quarterly frequency used here likely improves its accuracy. Full derivation in Appendix A.1.

where γ_0 and γ_1 are composite parameters, with $\gamma_1 > 0$, whose expressions in terms of deeper structural parameters appear in Appendix A.1. Beyond non-idea changes Δa_{t+1} and $\Delta \ln Q_{t+1}$, productivity growth is driven by the 'effective R&D' term,

$$s_t = \ln S_t - \frac{1-\psi}{\eta} \ln I_t - \frac{1}{\eta} \ln Q_t \,. \tag{10}$$

This term reflects R&D's contribution to productivity growth, with aggregate investments scaled down by product variety and, when spillovers are less than linear, by the stock of average ideas, capturing the increasing difficulty of finding new ideas. Because the effect of R&D on expected productivity growth depends on the contemporaneous values of I_t and Q_t, any alternative transformation of R&D expenditure – such as its growth rate or its ratio to GDP, both common in the literature – fails to capture it faithfully. The scaling in s_t is the only one that does so, regardless of whether the economy is in a fully or semi-endogenous growth regime. s_t is therefore termed the *innovation component* of productivity growth – the central focus of this study.

Rearranging (10) reveals that effective R&D admits a natural interpretation as the stochastic scaler of a rule-of-thumb policy function for aggregate R&D:

$$S_t = e^{s_t} \cdot \left(I_t^{1-\psi} Q_t \right)^{1/\eta} \,. \tag{11}$$

This highlights how a constant s_t, and hence constant productivity growth, requires the absolute level of R&D investment to rise alongside the economy's stock of ideas and product variety; any lower level of investment, however high in absolute terms, is associated with a fall in s_t and, consequently, in productivity growth, as R&D loses ground against the increasing difficulty of finding new ideas and the expansion of product variety.

Assuming productivity growth is stationary, as widely supported by the data, a literal interpretation of (9) dictates that effective R&D is stationary too, unless $\Delta \ln Q_{t+1}$ or Δa_{t+1} are themselves non-stationary, a non-standard case. On the other hand, sustained productivity growth implies exponential growth in aggregate R&D and average ideas, as well as in product variety, which is typically assumed to scale with the economy, so their logarithms are predicted to be unit-root processes – also widely supported by the data. s_t is therefore stationary but also a linear combination of three unit-root variables, which implies a cointegrating relationship among them, with s_t as the associated error correction term (ECT). Intuitively, high R&D cannot systematically outpace the frontier, not least because R&D itself raises the stock of ideas; conversely, low R&D lets the stock of ideas erode, and as the frontier falls, the effectiveness of R&D rises. Crucially, the parameters of cointegrating relationships can be estimated using well-established techniques, making s_t in principle recoverable from the data.

2.3. The long-run risk from innovation

The stationarity of s_t implies the existence of an unconditional mean $\bar{s}$, which represents the infinite-horizon equilibrium level around which s_t fluctuates. Deviations from this equilibrium, $\tilde{s}_t = s_t - \bar{s}$, are termed 'excess effective R&D' and drive expected future productivity growth:

$$\mathbb{E}_t \left[\widetilde{\Delta \ln Z}_{t+1} \right] \approx \gamma_1 \cdot \tilde{s}_t + \xi \cdot \mathbb{E}_t \left[\widetilde{\Delta \ln Q}_{t+1} \right] + \mathbb{E}_t \left[\Delta a_{t+1} \right] \,, \tag{12}$$

where $\widetilde{\Delta \ln Z}_{t+1}$ and $\widetilde{\Delta \ln Q}_{t+1}$ denote the demeaned counterparts of $\Delta \ln Z_{t+1}$ and $\Delta \ln Q_{t+1}$, respectively. Note that $\tilde{s}_t$ is the only variable on the right-hand side directly contained in the time-t information set, playing the same role in forecasting future growth as the 'productivity LRR component' in Croce (2014), lending it a deeper economic interpretation.

For illustrative purposes, this section follows the long-run risk literature and assumes that excess effective R&D follows a univariate AR(1) process:

$$\tilde{s}_t = \rho_s \tilde{s}_{t-1} + \tilde{\varepsilon}_t, \qquad \tilde{\varepsilon}_t \sim \mathcal{N}(0, \sigma_s^2) \,. \tag{13}$$

This law of motion implies that, while R&D expenditures keep pace with the technological level asymptotically, their *temporary deviations* from the infinite-horizon equilibrium proportion evolve autonomously, depending only on their own past values. Moreover, as the technological frontier is predetermined at the time R&D is chosen, effective R&D shocks $\tilde{\varepsilon}_t$ can be interpreted as R&D shocks pre-scaled to reflect their effectiveness. In this setting, the effects of these shocks propagate solely through the persistence of effective R&D itself, which Kung and Schmid (2015) shows can indeed emerge endogenously. The news about long-run productivity in response to an effective R&D shock is then

$$\left\{ \mathbb{E}_t - \mathbb{E}_{t-1} \right\} \left(\sum_{j=1}^{\infty} \Delta \ln Z_{t+j} \right) = \frac{\gamma_1}{1-\rho_s} \tilde{\varepsilon}_t \,. \tag{14}$$

Under the simplifying assumption that consumption is a constant fraction of final goods produced, ε_t^x is affine in $\tilde{\varepsilon}_t$ with slope $\frac{\gamma_1 \kappa_x}{1-\rho_s \kappa_x}$, which implies $\mathrm{Var}\left[\varepsilon_t^x\right]$ being proportional to $\left(\frac{\gamma_1 \kappa_x}{1-\rho_s \kappa_x}\right)^2 \sigma_s^2$, illustrating how the persistence of R&D effects determines the impact of its shocks on the IMRS and asset prices via (4) and (5). From a financial perspective, the longer these shocks continue to affect consumption growth, the higher the returns investors demand on assets whose payoffs covary positively with them, as these signal larger shifts in productivity and consumption prospects.

Note that there is a well-known tension between the degree of persistence of a series and the difficulty of empirically distinguishing it from a non-stationary one.[14] Indeed, in finite samples, effective R&D may (i) appear stationary despite being generated by a non-stationary process, due to substantial short-term fluctuations, or (ii) appear non-stationary despite being stationary but highly persistent. This poses a challenge for this univariate setting: the stationarity of s_t is crucial for validation, to avoid spurious results in standard empirical analyses, while high persistence of s_t is essential for the long-run risk mechanism to be economically significant, with these effects strongest precisely as the process approaches the unit root. Ultimately, whether a stationary effective R&D best represents reality is an empirical question, and, in any case, the persistence of the effects of R&D shocks need not coincide with the intrinsic persistence of the innovation component, as even a transitory disturbance can generate lasting effects when propagated through interactions with the broader economic system. The key characteristic for $\tilde{s}_t$ to qualify as an Innovation Long-Run Risk Component is only that its shocks forecast growth in productivity – and in consumption – over long horizons.

3. From theory to data

The theoretical framework just laid out does not, by itself, provide reduced-form conditions mapping tightly into empirical specifications that permit testing its key predictions and recovering effective R&D dynamics. To this end, I complement it with additional conditions, drawing on standard practices in the literature. The augmented framework then allows the cointegrating relationship that yields effective R&D to be recast in terms of observables and supplies the one theoretical condition invoked in estimation to recover the R&D shocks. The deep parameters underlying these objects are neither separately identified nor systematically scrutinized, except in recovering net effective R&D levels, a supplement to the main analysis.

The full structure is articulated in two blocks: a long-run block that specifies the relation among non-stationary variables, yielding an operational gross measure of effective R&D; and a short-run block that specifies a stationary system, from which the structural shocks are identified. This split formulation is designed to align with estimation, which is carried out in two steps because of the timing of the variables entering it. Indeed, effective R&D combines R&D expenditure, a flow variable measured over a given period, and empirical counterparts of ideas and product variety, stock variables observed only at the end of each period. In standard theoretical models, however, end-of-period stock levels cannot be part of the information set underlying R&D decisions, made either at the beginning of a period or continuously throughout it, nor can they affect the effectiveness of previous expenditures. To maintain theoretical consistency, then, R&D has to be empirically aligned with stock levels at the *beginning* of the period, i.e., the end of the previous period, since intermediate stock values are unavailable and interpolation would distort the series' dynamic properties. This timing convention, in turn, is problematic for standard Vector Error Correction Models, which jointly estimate long-run relations and short-run dynamics, since these would effectively use R&D between $t-1$ and t to forecast growth in the ideas stock over the same interval, which is not the predictive relationship described by endogenous growth theory.

3.1. An empirical measure of effective R&D

Of the three theoretical quantities composing effective R&D, as defined in (10), only R&D expenditure maps cleanly to observables. The average stock of ideas and product variety are harder to measure and are thus mapped to more observable quantities. First, the stock of ideas is replaced by the productivity level, exploiting the definition of TFP in (7) and following numerous studies that have argued the most common measure of new ideas, patents, to be a poor proxy for successful innovation.[15] Moreover, even with a reliable proxy for new ideas, the stock would need to be constructed manually using the very law of motion under estimation. This would not only make the estimation procedure more convoluted, but also require a more detailed law of motion (including, most notably, foreign ideas), likely limiting the sample timespan and carrying its own misspecification risk. By contrast, TFP, measured as a Solow residual, is much easier to measure empirically: its very concept is grounded in the data, and its definition is shared across many models, enhancing the external validity of the analysis. Second, product variety is replaced by labor force L_t, under the assumption

$$\ln Q_t = \omega \cdot \ln L_t + \pi \cdot a_t\,, \tag{15}$$

for some $\omega, \pi > 0$. This reduced-form approximation follows Jones (1999) in anchoring the trend of product variety to that of the labor input, reflecting the standard result that steady-state product variety scales with labor, in line with the cointegration-based empirical literature on endogenous growth (Ha and Howitt, 2007; Madsen, 2008; Kruse-Andersen, 2023, among others). This formulation also accommodates deviations of product variety from the labor input at higher frequencies, characterizing their dynamics as stationary fluctuations associated with the non-technological process a_t, in line with its pronounced cyclicality (Broda and Weinstein, 2010).[16]

Effective R&D can then be expressed almost entirely in terms of observable variables:

$$s_t = \ln S_t - \alpha_Z \cdot \ln Z_t - \alpha_L \cdot \ln L_t + \left(\alpha_Z + \alpha_L \cdot \frac{\pi}{\omega}\right) a_t\,, \tag{16}$$

[14] See, for example, Müller (2005).

[15] See, for example, Reeb and Zhao (2020) and Herzer (2022).

[16] Including lagged effective R&D as an additional determinant of product variety *levels* would merely relabel reduced-form parameters, leaving the empirical specifications and interpretation unchanged.

where α_Z and α_L are composite parameters, functions of the deep parameters of the theoretical framework, with expressions reported in Appendix A.1. Given stationarity of productivity growth and the non-technological process, the linear combination of S_t, Z_t, and L_t inherits the theoretical prediction of stationarity, motivating an empirical cointegration model for these three variables, which also tests the stationarity prediction. What is left random from this cointegrating relation, its ECT, is interpreted by theory as a version of excess effective R&D contaminated by a non-technological component, henceforth 'gross effective R&D':

$$\hat{s}_t = \tilde{s}_t - \left(\alpha_Z + \alpha_L \cdot \frac{\pi}{\omega}\right) a_t \,. \tag{17}$$

3.2. Innovation shocks and their identification

Theory predicts that unanticipated changes to *net* effective R&D represent a substantial risk factor in financial markets. These can conveniently be identified from *gross* effective R&D in a system that spans the non-technological process, isolating and removing the contamination. The strategy requires R&D fluctuations to affect real macroeconomic quantities with a delay, which is a common assumption in the literature (e.g., Moran and Queralto (2018)) and is, in fact, the only restriction imposed on the system governing the short-run dynamics of the state variables underlying the entire theoretical framework.[17] The assumed system is indeed otherwise agnostic, accommodating standard time dependencies and feedback effects:

$$\tilde{s}_{t+1} = \rho_s \tilde{s}_t + \beta_{s,a} a_t + \beta_{s,L} \widetilde{\Delta \ln L}_t + \delta_{s,a} \varepsilon^a_{t+1} + \delta_{s,s} \varepsilon^s_{t+1} \tag{18a}$$

$$a_{t+1} = \beta_{a,s} \tilde{s}_t + \rho_a a_t + \beta_{a,L} \widetilde{\Delta \ln L}_t + \delta_{a,a} \varepsilon^a_{t+1} \tag{18b}$$

$$\widetilde{\Delta \ln L}_{t+1} = \beta_{L,s} \tilde{s}_t + \beta_{L,a} a_t + \rho_L \widetilde{\Delta \ln L}_t + \delta_{L,a} \varepsilon^a_{t+1} \,, \tag{18c}$$

where ρ_i are the autoregressive coefficients, $\beta_{i,j}$ the cross-variable feedback coefficients, $\delta_{i,j}$ the contemporaneous shock loadings, ε^i_{t+1} i.i.d. standard normal shocks, for $i, j \in \{L, a, s\}$, and $\widetilde{\Delta \ln L}_t$ denotes labor growth deviations from its unconditional mean. This system embeds a further assumption, abstracting from idiosyncratic shocks to the labor input, which is adopted on estimation and expositional grounds but is not strictly necessary. Indeed, labor force dynamics have been shown to be largely driven by near-deterministic demographic factors (Aaronson et al., 2014; Fujita, 2014), and thus provide little additional short-run information, while substantially increasing the system dimension at an appreciable cost in finite-sample estimation efficiency. Labor input is therefore assumed to follow $\Delta \ln L_{t+1} = \mu_L + \nu \cdot a_{t+1}$, where μ_L captures the demographic component and $\nu \cdot a_{t+1}$ the non-negligible cyclical fluctuations in the residual variation (Erceg and Levin, 2014). The resulting system features two primitive shocks, ε^s_t and ε^a_t, with the latter now absorbing all business-cycle variation in productivity that is orthogonal to innovation efforts. Appendix D.2 reports theoretical derivations and empirical results allowing for an independent labor shock; findings are qualitatively and quantitatively very similar, with slightly wider confidence intervals owing to the larger system size.

Note that, in this setting, structural shocks ε^s_t capture only the unforeseen variation in R&D (scaled by its contemporaneous effectiveness) that is not driven by simultaneous unanticipated changes in the other variables, effectively isolating the long-run risk channel in reduced-form estimates of the innovation risk premium. Indeed, total effective R&D shocks decompose as $\tilde{s}_t - \mathbb{E}_{t-1}\left[\tilde{s}_t\right] = \delta_{s,a}\varepsilon^a_t + \delta_{s,s}\varepsilon^s_t$, and while the component driven by ε^a_t is associated with contemporaneous changes in other real macroeconomic quantities, the one driven by ε^s_t does not leak to them contemporaneously, affecting only future growth (see (9)). In the language of Section 2.1, ε^s_t operates only through the long-run risk shock ε^x_t, in isolation from short-run shocks ε^c_t.[18] By contrast, ε^a_t drives contemporaneous changes in both effective R&D, with long-run repercussions, and consumption, so the risk premium associated with it is potentially inflated by distinct risks. Clearly, the analysis does not claim ε^s_t to be fundamental in an absolute sense, as R&D plausibly reacts to economic forces outside of this framework, most notably financial conditions. Nonetheless, the one driver it is explicitly distinguished from, productivity growth, would otherwise be the most troublesome confounder, being both a recognized driver of R&D investment decisions and the very channel through which innovation creates long-run risk.

Using (9) and (15), the system in (18) can be represented as

$$\begin{bmatrix} \widetilde{\Delta \ln Z}_{t+1} \\ \hat{s}_{t+1} \end{bmatrix} = \Theta_1 \begin{bmatrix} \widetilde{\Delta \ln Z}_t \\ \hat{s}_t \end{bmatrix} + \Theta_2 \begin{bmatrix} \widetilde{\Delta \ln Z}_{t-1} \\ \hat{s}_{t-1} \end{bmatrix} + \begin{pmatrix} \lambda^1_{Z,a} & 0 \\ \lambda^1_{s,a} & \lambda^1_{s,s} \end{pmatrix} \begin{bmatrix} \varepsilon^a_{t+1} \\ \varepsilon^s_{t+1} \end{bmatrix} + \Lambda_2 \begin{bmatrix} \varepsilon^a_t \\ \varepsilon^s_t \end{bmatrix} , \tag{19}$$

where Θ_1, Θ_2 and Λ_2 are 2×2 matrices of composite parameters whose derivation is in Appendix A.3. Importantly, R&D shocks ε^s_t can be identified using the gross effective R&D measure by virtue of the triangularity of the contemporaneous impact matrix, which holds because both the gross measure and the structural system are linear in net effective R&D and the non-technological process. Notice that the identification scheme here jointly performs two distinct tasks: it removes the contamination from effective R&D to correctly measure the dynamics of its net values, and it identifies its unexpected changes not attributable to other variables in the system.

[17] Implausibility of contemporaneous outcomes is further supported by the significant lag in the diffusion of new technologies, as stressed in Rotemberg (2003) and Anzoategui et al. (2019).

[18] In the theoretical framework of this paper, a positive shock ε^s_t actually *reduces* contemporaneous consumption growth via the budget constraint, implying a negative short-run risk that counteracts, rather than inflates, the long-run risk from ε^s_t. The risk premium associated with these shocks in reduced-form estimates would therefore represent a lower bound on its LRR component.

3.3. Recovery of net effective R&D levels

Conditions (9) and (15) also permit expressing effective R&D *levels* almost exclusively in terms of observables (see Appendix A.4):

$$\tilde{s}_t = \sum_{j=0}^{t-1} (\kappa_s)^j \left\{ \kappa_Z \cdot \widetilde{\Delta \ln Z}_{t-j} - \kappa_L \cdot \widetilde{\Delta \ln L}_{t-j} + \Delta \hat{s}_{t-j} \right\} + (\kappa_s)^t \cdot \tilde{s}_0 \,, \tag{20}$$

where $\kappa_Z = \frac{\alpha_Z + \alpha_L \frac{\pi}{\omega}}{1+\xi\pi}$, $\kappa_L = (\xi\omega)\kappa_Z$ and $\kappa_s = 1 - \gamma_1 \cdot \kappa_Z$. Though not strictly required for the main analysis, this series provides broader insight on R&D dynamics and serves as an additional robustness check for the empirical tests of the framework. Since the computation of this structural object depends on deep parameters whose identification lies outside the scope of this paper, a minimal set of auxiliary assumptions is imposed, only in recovering net effective R&D levels.

First, setting $\eta = 1$ abstracts from duplication of innovative efforts, a common assumption in the empirical literature (e.g., Bloom et al., 2020). Second, setting $\omega = 1$ implies that product variety is precisely proportional to the labor input in the long run. In this lab-equipment setup, this restriction simply specializes the framework, which remains compatible with both fully and semi-endogenous growth. Third, setting $\pi = 0$ abstracts from fluctuations in product variety not captured by the labor input, leaving residual contamination in the very short-run.[19] Under these restrictions, $\kappa_Z = \alpha_Z$ and $\kappa_L = 1 - \alpha_L$.

To identify γ_1, and thus κ_S, two further assumptions are imposed. Fourth, non-technological productivity variation is spanned by observable pervasive macroeconomic factors $\boldsymbol{f}_t$, i.e., $a_t = \boldsymbol{\gamma}_f' \boldsymbol{f}_t$ for some $\boldsymbol{\gamma}_f$; fifth, R&D's effect on real variables is negligible even one period ahead, i.e., $\beta_{L,s} = \beta_{a,s} = 0$ in (18). Under these assumptions,

$$\widetilde{\Delta \ln Z}_{t+1} = \gamma_1 \hat{s}_t + \gamma_L \widetilde{\Delta \ln L}_t + \gamma_a \boldsymbol{\gamma}_f' \boldsymbol{f}_t + \gamma_e \varepsilon^a_{t+1} \,, \tag{21}$$

with γ_1 the sole coefficient on $\hat{s}_t$, hence identifiable from data (see Appendix A.3 for γ_L, γ_a, and γ_e). The fourth assumption is common in empirical applications (most notably Ai et al., 2018). The fifth one is conceptually motivated by a_t capturing numerous aggregate factors and, consequently, R&D's one-period contribution being likely negligible, even if feedback mechanisms amplify it at longer horizons, as corroborated by results below.

The only unobservable term left in (20) is the initial value $\tilde{s}_0$, whose influence decays exponentially through κ_S^t under stationarity of $\tilde{s}_t$, as verified empirically below.

4. The empirical innovation component

4.1. The gross effective R&D

4.1.1. Empirical implementation

The cointegrating relationship in (16) is estimated as

$$\ln S_t = b_0 + b_Z \cdot \ln Z_t + b_L \cdot \ln L_t + e_t \,, \tag{22}$$

where b_Z and b_L correspond to α_Z and α_L, and the ECT e_t identifies gross effective R&D $\hat{s}_t$. The regression is estimated using the Fully Modified Ordinary Least Squares (FM-OLS) method (Phillips and Hansen, 1990), which corrects for endogeneity in cointegrating regressions by adjusting the dependent variable using the long-run covariance between the error term and regressor innovations. FM-OLS is particularly suitable here because the error correction term is expected to be highly persistent. By contrast, addressing endogeneity with Dynamic OLS (Stock and Watson, 1993), which requires including leads and lags of the differenced regressors, would consume many degrees of freedom and substantially reduce the effective sample size. For robustness, the Integrated Modified (IM) OLS of Vogelsang and Wagner (2014) is also considered, although it is expected to perform less efficiently than FM-OLS given the sample size, over 300 observations.

The baseline measure of R&D is real US private R&D expenditure, in line with related studies (Kung and Schmid, 2015; Moran and Queralto, 2018; Beqiraj et al., 2025). Indeed, the endogenous growth models underlying the framework of Section 2 rest on profit-maximizing allocation and innovation decisions, abstracting from the institutional and policy considerations central to government expenditure, and from the possibly distinct knowledge production function the latter may follow. On these grounds, public R&D effort is treated as part of the non-technological productivity component, and total R&D is left as a robustness check. Total Factor Productivity is taken from Fernald (2012). The baseline series is TFP growth adjusted for utilization and excluding R&D from capital; robustness checks use the raw TFP series. The utilization adjustment is preferred because removing cyclical utilization dynamics enhances the signal-to-noise ratio of the underlying technological component of productivity, thereby improving the precision of the estimates; R&D capital is excluded because its construction through simple cumulation and depreciation of R&D flows is inconsistent with the knowledge production function studied in this work. Differences between the series arise mainly from the utilization adjustment, with the R&D-capital exclusion contributing little. The labor force is measured by total employment, with non-farm employment used as a robustness check. Neither series includes military employment – a further mismatch with total R&D, given the substantial defense share of public R&D expenditure. The combined dataset is quarterly and spans 1948Q1 to 2025Q1; further details are provided in Appendix B.1.

[19] For comparison, assuming $\pi = \omega = 1$ yields a negative empirical impulse response of $\tilde{s}_t$ to a_t, nearly identical to that of $\hat{s}_t$ and counter to theory, suggesting that this alternative value of π is an upper bound leading to insufficient correction for the cyclical component of productivity, via $\xi\pi$ in κ_Z.

Table 1
Cointegration results. The regression is $\ln S_t = b_0 + b_Z \ln Z_t + b_L \ln L_t + e_t$, estimated on quarterly data over 1948Q1–2025Q1 via Fully Modified OLS (Phillips and Hansen, 1990). Columns show baseline estimates (Private R&D as S_t, Utilization-adjusted TFP as Z_t, Total Employment as L_t) and four robustness checks, each varying R&D (Total R&D), productivity (Raw TFP), labor (Nonfarm Employment), or the estimation method (Vogelsang and Wagner, 2014). The bottom panel reports statistics for e_t: σ is the standard deviation; *ADF* and *KPSS* are stationarity tests (detecting stochastic and deterministic trends); *AR(1)* is the fitted autoregressive coefficient; T is the number of observations.

	Baseline	S: Tot. R&D	Z: Raw TFP	L: N.F. Empl.	Est. Meth.: IM
b_Z	3.526***	4.197***	3.655***	3.349***	2.821***
	(0.439)	(0.516)	(0.490)	(0.464)	(0.552)
b_L	0.909***	−0.354	0.956***	0.953***	1.387***
	(0.336)	(0.395)	(0.356)	(0.325)	(0.398)
e_t					
σ	0.130	0.144	0.128	0.129	0.253
ADF	−2.57**	−2.45**	−2.92***	−2.45**	−9.18***
KPSS	0.09	0.09	0.09	0.10	0.29
AR(1)	0.96	0.96	0.95	0.97	0.15
T	309	309	309	309	309

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4.1.2. Results

Table 1 reports estimates of the cointegrating regression in (22). The first column shows the baseline; the rest assess robustness to alternative data and estimator.

The b_Z estimates are all positive and significantly different from zero, as expected, implying that R&D expenditure and TFP levels move together in the long run, i.e., R&D rises with the scale of the economy. The labor coefficient, b_L, is similarly positive, significant, and consistent across specifications, underscoring the dilution effect of product variety on R&D's frontier-advancing power, and product variety's relevance in shaping R&D effectiveness. The sole exception is total R&D, under which b_L becomes insignificant and negative, consistent with labor force measures excluding part of the public labor input, most notably the military, and public R&D potentially operating through a different production function.[20] The IM-OLS estimates corroborate the baseline in coefficient values but are clearly more unstable, as revealed by visual inspection of the ECT time series, which are markedly noisier and less consistent across specifications when estimated with IM-OLS (Figure D.9, Appendix D). Visual inspection of the ECTs further shows that the nonfarm-employment variant is the least consequential, leaving both ECT dynamics and coefficient estimates indistinguishable from the baseline, whereas the raw-TFP variant, while leaving the coefficients essentially unchanged, does noticeably alter the ECT dynamics.

Within the framework adopted here, these estimates cannot reject a semi-endogenous balanced growth path (BGP) but do reject a fully endogenous one at any conventional significance level, since the latter requires $b_Z = 1$ (see Appendix A.2).[21] These results, along with the coefficient estimates, are in line with Kruse-Andersen (2023). Beyond the BGP rejection, the estimates deliver two bounds on deep theoretical parameters under positive past-ideas spillovers $\psi > 0$, namely $\theta\eta < 1/\alpha_Z \approx 0.28$, bounding R&D's overall transmission to productivity, and $\theta - \xi\omega < \alpha_L/\alpha_Z \approx 0.26$, bounding the love-of-variety strength relative to ideas' contribution to productivity.

The bottom panel of Table 1 reports descriptive statistics for the ECT, which identifies the time series of $\hat{s}_t$. Across specifications, these series are highly persistent but stationary by both ADF and KPSS tests, with no deterministic trend. They are also correlated above 86% with one another (Table D.14), except for the IM-OLS variant.

4.2. The net effective R&D

4.2.1. Empirical implementation

Section 3.3 motivates a plug-in estimator[22]:

$$\hat{\tilde{s}}_t = \sum_{j=0}^{t-1} (\hat{\kappa}_s)^j \left\{ \hat{b}_Z \cdot \widetilde{\Delta \ln Z}_{t-j} - (1 - \hat{b}_L) \cdot \widetilde{\Delta \ln L}_{t-j} + \Delta \hat{e}_{t-j} \right\}, \tag{23}$$

[20] Unreported estimates are significantly negative when only public R&D is used.

[21] The rejection is not overturned by b_L being statistically indistinguishable from 1: although, under specific assumptions, $b_L = 1$ is consistent with $\omega = 1$ and a fully endogenous BGP, b_L identifies a composite parameter rather than ω itself.

[22] Note that an alternative approach to recovering $\tilde{s}$ in levels is to include the factors $\boldsymbol{f}_t$ in the cointegration regression among $\ln S_t$, $\ln Z_t$, and $\ln L_t$. However, this would require estimating a long-run covariance matrix with at least 144 elements from fewer than 280 observations, yielding imprecise and unstable cointegrating-parameter estimates.

Table 2
One-step productivity growth forecast results. The top panel reports estimates and diagnostics for the regression $\Delta \ln Z_{t+1} = b_0 + b_s \hat{s}_t + \sum_{j=0}^{P} \left(\boldsymbol{b}'_{f,j} \boldsymbol{f}_{t-j} + b_{Z,j} \Delta \ln Z_{t-j} \right) + u_{t+1}$, estimated on quarterly data from 1960Q1 to 2025Q2; HAC standard errors are reported in parentheses. Columns show the baseline specification (Private R&D as S_t, Utilization-adjusted TFP as Z_t, Total Employment as L_t, LN factors as $\boldsymbol{f}_t$) and four robustness checks, each varying one input: R&D (Total R&D), productivity (Raw TFP), labor (Nonfarm Employment), or factors (BS factors, 1952Q1–2025Q2). T is the observation count, P the lag length (see Figure B.3 for selection), and R^2 the goodness-of-fit. The intermediate row reports the estimated κ_s and its standard error, and the bottom panel reports statistics for the recovered series $\hat{\hat{s}}_t$ (both detailed in Appendix B.2): σ is its standard deviation; *ADF* and *KPSS* are stationarity tests (detecting stochastic and deterministic trends); *AR(1)* is the first-order autoregressive coefficient.

	Baseline	S: Tot. R&D	Z: Raw TFP	L: N.F. Empl.	$\boldsymbol{f}$: BS
b_s (%)	1.549***	1.223***	0.794**	1.520***	1.530***
	(0.285)	(0.254)	(0.337)	(0.289)	(0.409)
T	261	261	260	261	292
P	1	1	2	1	1
R^2 (%)	12.4	11.8	41.7	12.2	9.3
κ_s	0.945	0.949	0.971	0.950	0.946
	(0.012)	(0.012)	(0.013)	(0.012)	(0.016)
$\hat{\hat{s}}_t$					
T	218	212	134	211	216
σ	0.045	0.054	0.052	0.047	0.045
ADF	−2.61***	−1.65*	−2.08**	−2.40**	−2.49**
KPSS	0.14	0.47**	0.22	0.17	0.15
AR(1)	0.95	0.98	0.96	0.96	0.95

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

where $\hat{\kappa}_s = 1 - \hat{b}_s \cdot \hat{b}_Z$, with hats denoting parameter estimates. As discussed in Appendix B.2, this omits the unobservable initial condition, so the recovered series is trimmed at the beginning of the sample to ensure consistency, and the standard errors are derived to reflect the resulting uncertainty, beyond the sampling variability in the estimates.

The coefficient b_s, which only identifies the composite $\gamma_1 = (\theta\chi)\,\mathbb{E}[s_t]$, is estimated from the empirical counterpart of (21):

$$\Delta \ln Z_{t+1} = b_0 + b_s \hat{s}_t + \sum_{j=0}^{P} \left(\boldsymbol{b}'_{f,j} \boldsymbol{f}_{t-j} + b_{Z,j} \Delta \ln Z_{t-j} \right) + u_{t+1}\,. \tag{24}$$

As the aim is to estimate b_s rather than to characterize the full dynamics of the R&D impact, a single-equation approach is preferred to a system, since it can accommodate more controls in $\boldsymbol{f}_t$ to reduce bias at a smaller increase in estimation variance than the latter would. In line with Ai et al. (2018), the baseline factors are the LN factors of Ludvigson and Ng (2009) – nine principal components of a broad macroeconomic and financial panel – with a similarly-sized expansion of the BS factors of Bansal and Shaliastovich (2013) used for robustness; see Appendix B.1 for details. The LN factors are preferred despite their shorter sample (1960Q1–2025Q2 vs. 1952Q1–2025Q2) because they summarize a richer information set, as confirmed by model-selection comparisons. Both the factors and the dependent variable enter as regressors with P lags, chosen by standard information criteria; labor growth $\Delta \ln L_t$ is excluded as insignificant in all specifications (Wald p-values > 0.5).

4.2.2. Results

Table 2 reports the estimation results of (24).

Information criteria select $P = 1$ under utilization-adjusted TFP and $P = 2$ under raw TFP (Figure B.3, Appendix D), with LN-based specifications always achieving higher explanatory power than their BS-based counterparts. The R&D coefficient b_s is positive and statistically significant across all specifications, consistent with effective R&D being a meaningful driver of productivity growth. Its magnitude closely matches that of Kung and Schmid (2015) and implies an annualized 0.8% increase in TFP growth per one-standard-deviation rise in $\hat{s}_t$; only under raw TFP is the magnitude about half.

Table 2 also reports the implied κ_s and summary statistics for the recovered $\hat{\hat{s}}_t$ series. Across all specifications, κ_s is estimated with small standard errors, and $\hat{\kappa}_s$ lies consistently close to but below one. Its highest values arise under Raw TFP, where $\hat{b}_s$ is about half its value elsewhere, pushing $\hat{\kappa}_s$ toward one. The recovered series are stationary by the ADF and KPSS tests, with the sole exception of Total R&D, and display persistence comparable to $\hat{s}_t$. Notably, the implied half-lives range from roughly 2 to 20 years at the 95% confidence level, a span that contains the theoretical effective-R&D analogue calibrated in Kung and Schmid (2015) and runs from the empirical productivity component of Croce (2014) at the lower end to that of Ortu et al. (2013) at the upper end. The persistence of R&D effects, however, need not match the univariate persistence of $\tilde{s}_t$, since shocks propagate through the broader economy and their effects hinge on the feedback between the innovation component and external processes.

Fig. 2 plots $\tilde{s}_t$ alongside $\hat{s}_t$, both from the baseline specification (all $\tilde{s}_t$ series in Figure D.10). Because the theoretical framework ties the difference between $\tilde{s}_t$ and $\hat{s}_t$ to the external process a_t, the figure suggests a pro-cyclical behavior of innovation efforts, consistent with the predictions of Kung and Schmid (2015) and related contributions.

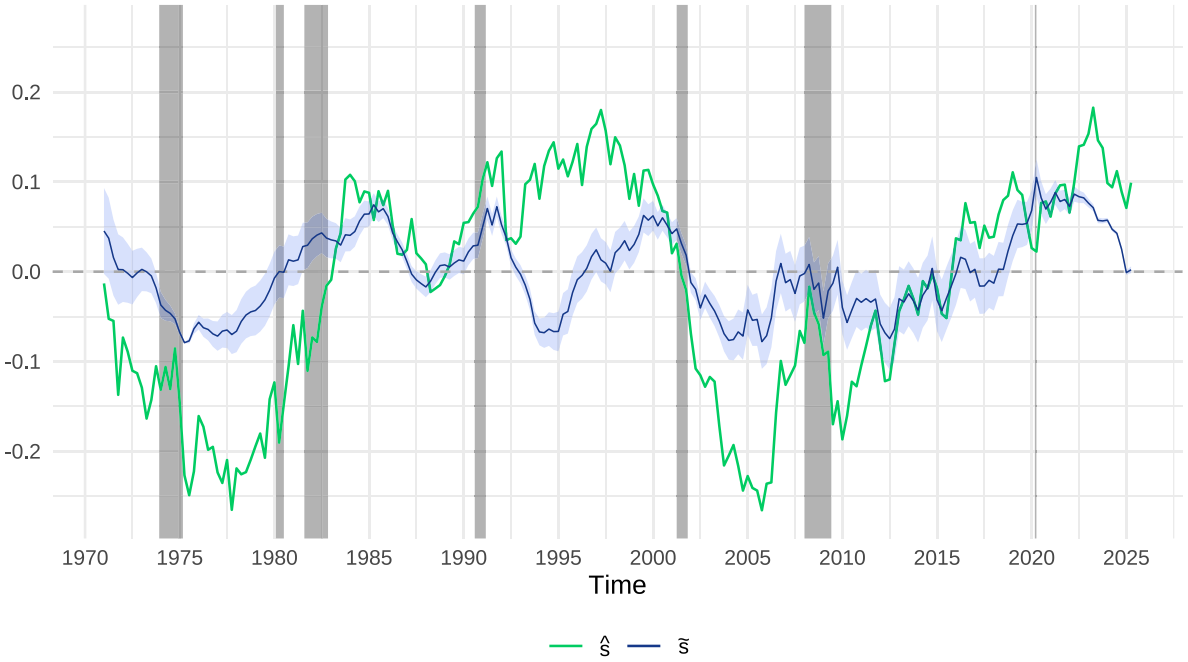


Fig. 2. Baseline effective R&D measures. In green, the gross measure from estimates in Table 1, in blue, the net measure from estimates in Table 2, with 95% confidence interval bands, obtained as illustrated in Appendix B.2. Shaded areas mark NBER recessions.

5. The long-run risk from innovation

5.1. Effective R&D shocks

5.1.1. Empirical implementation

The Vector Auto-Regressive Moving-Average (VARMA) system in (19) is estimated using a finite-order Vector Auto-Regression (VAR) that approximates its infinite-order VAR representation,

$$\boldsymbol{y}_t = \sum_{j=1}^{P} B_{x,j}\, \boldsymbol{y}_{t-j} + \boldsymbol{w}_{x,t}\,, \tag{25}$$

where the lag length P is selected by standard information criteria and $\boldsymbol{y}_t = (\Delta \ln Z_t - \hat{\mu}_Z,\, x_t)'$, with $x_t \in \{\hat{e},\, \hat{\tilde{s}}_t\}$. Indeed, shock identification rests on the recursive ordering of the variables and is in principle invariant to the choice of effective R&D measure: both the gross and net measures yield the same structural shocks $\boldsymbol{v}_t$, differing only in the loadings d^x of the lower-triangular Cholesky factorization

$$\boldsymbol{w}_{x,t} = \begin{pmatrix} d^x_{Z,a} & 0 \\ d^x_{s,a} & d^x_{s,s} \end{pmatrix} \boldsymbol{v}_t\,, \qquad \boldsymbol{v}_t = (v^a_t, v^s_t)'\,, \tag{26}$$

where v^a_t and v^s_t identify ε^a_t and ε^s_t, respectively. Accordingly, the VAR is estimated under both measures, with the gross measure retained as the baseline; results from the net measure should be read in light of its substantially higher estimation uncertainty, given the shorter sample and additional estimation step.

Identification can be compromised by truncating the infinite-order VAR at finite P, since part of the moving-average dynamics is left in the reduced-form residuals $\boldsymbol{w}_{x,t}$. This misspecification manifests itself as residual serial correlation, so the adequacy of the approximation is assessed ex post through standard diagnostic tests.

5.1.2. Results

Table 3 reports the key statistics from the estimation of (25), for different specifications of $\hat{e}$ and, in the last column, for $\hat{\tilde{s}}$. The maximum number of lags considered is 10, with the Akaike, Hannan–Quinn, and Bayesian information criteria generally agreeing on the optimal choice; the exception is the raw-TFP specification, for which the lag order indicated by the Hannan–Quinn criterion is applied. The R^2 for TFP growth is in line with the values reported in Croce (2014) for samples of comparable length. The system is stable, with the largest companion-matrix root well within the unit circle.

The middle and bottom panels report tests on the residuals and on system dynamics. The residuals of the effective R&D equation exhibit significant heteroskedasticity across all specifications, while those of the TFP-growth equation do so only under the raw-TFP

Table 3
Estimation diagnostics for the VAR in (25). Columns correspond to alternative specifications of $\hat{e}$ (gross effective R&D), with the last reporting the estimation under $\hat{s}$ (net effective R&D). T and P are the sample size and the VAR lag order. $R^2_{\Delta Z}$ is the fit of the TFP growth equation (in percent); max |roots| is the largest root of the companion matrix. *ARCH-LM* are Lagrange-Multiplier tests for autoregressive conditional heteroskedasticity at four quarterly lags (one year), applied to the residuals of the TFP growth and effective R&D equations, respectively. *AC-LM* are the wild-bootstrap, heteroskedasticity-robust autocorrelation tests of Ahlgren and Catani (2017) at the indicated lags. *GC-F(s)* is the F-test for Granger causality from effective R&D to TFP growth, based on heteroskedasticity-consistent (HC) covariance.

	Baseline	S: Tot. R&D	Z: Raw TFP	L: N.F. Empl.	s: Net
T	305	305	304	305	216
P	3	3	4	3	2
$R^2_{\Delta Z}$ (%)	5.3	4.4	5.7	5.3	2.4
max |roots|	0.92	0.90	0.94	0.92	0.94
ARCH-LM($\Delta Z, 4$)	3.2	2.3	26.8***	3.3	5.2
ARCH-LM($x, 4$)	30.2***	27.0***	25.1***	29.4***	17.9***
AC-LM(1)	5.8	6.3	1.9	5.1	5.7
AC-LM(4)	14.6	19.3	14.7	15.3	22.3
AC-LM(16)	69.6	68.6	60.4	69.6	63.8
AC-LM(40)	169.4	155.7	157.3	165.9	139.0
GC-F(s)	5.3***	6.5***	2.4**	5.5***	1.1

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

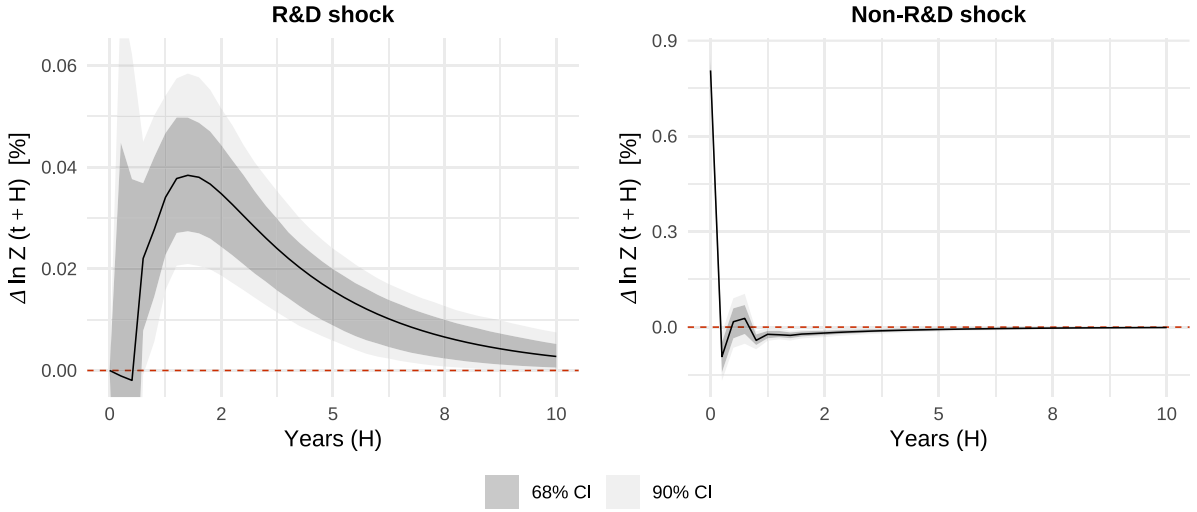


Fig. 3. Productivity impulse response functions from the VAR in (25), identified as in (26). The left panel reports the response of productivity growth ($\Delta \ln Z$) to a one-standard-deviation R&D shock (ε^s_t); the right panel reports the response to a non-R&D shock (ε^a_t). The VAR is estimated under the baseline gross effective R&D specification of Table 3. Dark and light shaded areas show 68% and 90% percentile intervals from a residual bootstrap with 1099 replications. The horizon H is in years.

specification; no significant autocorrelation is detected at horizons up to 10 years. The F-test for Granger causality from effective R&D to TFP growth (HC covariance) is highly significant under every specification of $\hat{e}$.

Fig. 3 displays the impulse response functions (IRFs) of productivity growth in the baseline specification. The effects of R&D shocks are highly persistent, with the 90% confidence interval still excluding zero at the 10-year horizon. By contrast, the effects of non-R&D TFP shocks materialize almost entirely on impact, followed by a small negative dip that dies out completely within five years. Moreover, the R&D-shock IRF is hump-shaped, remaining insignificant through the first year before peaking at two years. Since the identifying restriction fixes only the impact response, this empirical lag supports the slow-propagation assumption of Section 3.3. As a further illustration of the gross–net decomposition, Figure D.12 in the Appendix D reports the response of each effective R&D measure to non-R&D productivity shocks: the gross measure $\hat{s}_t$ responds negatively, consistent with its loading negatively on the non-R&D productivity process, while the net measure $\tilde{s}_t$, which excludes this component, shows a weakly positive response.

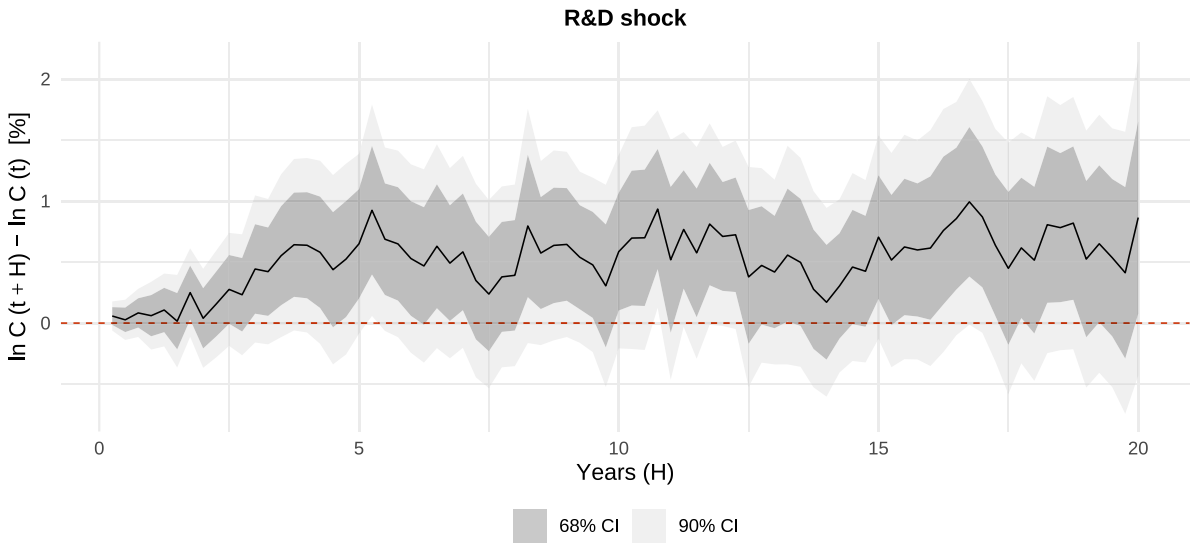


Fig. 4. Cumulative impulse response of consumption growth ($\Delta \ln C$) to the R&D shock ($\hat{v}^s$) from (27), $P = 5$, 1960 Q1–2025 Q2. C is real per-capita PCE; controls are the LN factors. Shaded areas: 68% and 90% confidence bands from Montiel Olea and Plagborg-Møller (2021). Horizon H in years.

5.2. Long-run risk from innovation

5.2.1. Empirical implementation

The theoretical framework, drawing on established results on how productivity transmits to consumption, delivers its key predictions while modeling consumption choices in a deliberately stylized way. The impact of R&D shocks on long-run consumption can nonetheless be credibly assessed with local projections (LP), which estimate impulse responses with valid inference under minimal assumptions on the underlying dynamics, at the cost of higher sampling variability (Montiel Olea et al., 2025). Accordingly, the cumulative response of consumption growth to an R&D shock is estimated by a separate regression at each horizon H, controlling for lagged consumption levels and other macroeconomic factors,

$$\sum_{h=1}^{H} \Delta \ln C_{t+h} = b_{c,H,0} + b_{c,H,s}\, \hat{v}_t^s + \sum_{j=1}^{P} \left(\boldsymbol{b}'_{c,H,f,j}\, \boldsymbol{f}_{t-j} + b_{c,H,c,j}\, \ln C_{t-j} \right) + v_{c,H,t+H}\,, \tag{27}$$

where the lag order P is selected by the AIC evaluated on the $H = 1$ regression, as recommended by Montiel Olea et al. (2025). Then, the coefficient $b_{c,H,s}$, measuring the cumulative consumption response to R&D shocks through horizon H, is informative about how strongly they pass through to long-run consumption news ε_t^x. Indeed, for horizons H where $(\kappa_x)^H \approx 1$ and the response is flat up to some $\bar{H} > H$ with $(\kappa_x)^{\bar{H}} \approx 0$, one has $b_{c,H,s}\, \varepsilon_t^s \approx \varepsilon_t^x$ (Appendix B.3). In the data, C_t is real per-capita U.S. personal consumption expenditures (PCE), 1948Q1–2025Q2, and $\boldsymbol{f}_t$ is the LN factor set discussed in Section 4.2.1, which dominates the BS alternative.

5.2.2. Results

The consumption cumulative IRFs to R&D shocks are shown in Fig. 4, at the selected lag length $P = 5$ (Figure D.13). The cumulative effect of R&D on consumption growth is positive and statistically significant, reaching the 68% level only after roughly three years. The response then increases until about six years and stabilizes around a positive level, consistent with a permanent shift in long-run consumption; estimates beyond 10 years are reported only to corroborate this pattern, albeit necessarily imprecise at this sample length. Overall, R&D shocks, scaled by their effectiveness and orthogonalized to productivity shocks, emerge as a significant driver of long-run consumption prospects.

6. The innovation risk premium

R&D shocks ε_t^s command a risk premium that identifies λ_x in (5): they carry substantial long-run consumption news ε_t^x and a negligible contemporaneous component ε_t^c, so their premium prices long-run risk alone, predicted positive and significant. Since this stylized framework leaves priced risks beyond aggregate innovation unmodeled, this premium is estimated following Giglio and Xiu (2021), whose method is designed to address the resulting omitted-variable bias.

The section then turns to the origins of equity risk, which in LRR models stems from cash-flow exposure, with return betas endogenously determined. The premium carried by the cash-flow channel is estimated following Bansal et al. (2005), whose method, though less robust, directly isolates the cash-flow source of risk from total return variation.

Table 4
Risk premia estimated following Giglio and Xiu (2021). Each entry is the estimated premium (annualized, in percent) on the innovation risk factor x_t; t-statistics in square brackets. The columns correspond to alternative definitions of the R&D shock from Table 3: the baseline, total R&D (S), raw TFP (Z), non-farm employment (L), and net effective R&D ($\tilde{s}$). The two panels differ in the construction of x_t: the contemporaneous R&D shock (horizon 1 quarter) and its four-year rolling mean (horizon 4 years). The factor counts, 6 and 14, are those selected by the criterion of Alessi et al. (2010). Test assets are 183 portfolios (153 anomaly, 17 industry, 13 bond) over 1971 Q4–2024 Q4; T is the number of quarters.

	Baseline	S: Tot. R&D	Z: Raw TFP	L: N.F. Empl.	s: Net
Horizon: 1 quarter					
6 factors	5.63	6.40	8.18	6.12	6.30
	[0.93]	[1.02]	[1.41]	[1.03]	[0.91]
14 factors	17.21	4.62	9.69	14.18	7.15
	[1.32]	[0.35]	[0.69]	[1.12]	[0.39]
T	213	213	213	213	213
Horizon: 4 years					
6 factors	2.07	1.71	2.28	2.04	2.17
	[1.33]	[1.10]	[1.29]	[1.36]	[0.85]
14 factors	12.06***	7.54**	12.41***	11.27***	16.78***
	[3.28]	[2.29]	[3.32]	[3.11]	[2.96]
T	213	213	213	213	199

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.1. A robust estimate

6.1.1. Empirical implementation

The estimator of Giglio and Xiu (2021) proceeds in three passes. It first recovers the full space of 'true' risk factors from a panel of returns by principal component analysis; then, it estimates their premia as in Fama and Macbeth (1973); finally, it recovers the premium λ_x of an observable factor x_t by projecting x_t on the true factors and combining its loadings with their premia. The resulting λ_x is the expected excess return on an asset with unit beta on x_t and zero beta on every orthogonal factor. Because the full factor space is recovered asymptotically, no priced factor is omitted, and λ_x is estimated consistently and free of omitted-variable bias. Further details are in Appendix C.2.

The estimator yields generalizable, robust estimates under two conditions: that the test assets span much of the economic state space, and that the appropriate number of true factors is selected. Expanding on the approach of Bryzgalova et al. (2025), this analysis uses 153 equity anomaly portfolios from Jensen et al. (2021), 17 industry-sorted portfolios, and 13 bond portfolios. The dataset spans 1971 Q4 to 2024 Q4; further details in Appendix C.1. The premium is estimated for the two factor counts selected by the criterion of Alessi et al. (2010); the criteria of Bai and Ng (2002) select substantially larger counts and yield premia comparable to those at the higher reported count, so are omitted.

Moreover, information about R&D is known to be incorporated into prices gradually, with Eberhart et al. (2004) documenting under-reaction lasting up to five years, so the R&D shock ε_t^s is used as the innovation long-run-risk factor both contemporaneously and as a rolling mean. Results are reported for $x_t = \hat{v}_t^s$ and $x_t = \frac{1}{16}\sum_{j=0}^{15} \hat{v}_{t-j}^s$, the latter a conservative four-year window corresponding to a typical business cycle.

6.1.2. Results

The criterion of Alessi et al. (2010) selects 6 and 14 factors in the cross-section of the 183 test assets, explaining 35.4% and 51.1% of return variance, respectively (Figure C.5).[23] In pricing, however, the first 6 and 14 components explain 47.0% and 70.7% of the variance of cross-sectional average returns (Figure C.6). As factors 7 through 14 account for only 15% of time-series variation but 24% of cross-sectional variation, the six-factor selection likely understates their relevance in approximating the 'true' factors.

Table 4 reports the estimated risk premia on the innovation factor. In the first panel, where the contemporaneous R&D shock is used, the premium is positive but not significant at conventional levels, consistent with investors' underreaction to R&D news (Eberhart et al., 2004). In the second panel, where the factor is the four-year rolling mean of the R&D shocks, the premia are sizable and significant when 14 factors are included. At 6 factors they fall short of conventional significance, but given the pricing power of factors 7 through 14, their t-statistics above one also lend mild support to the relevance of innovation risk. Overall, the results strongly support the prediction that innovation risk is positively priced.

[23] The criteria of Bai and Ng (2002) give 22 and 39 factors, and one does not converge.

Table 5
Cash-flow betas of selected portfolios from the extended pool. Each entry is the loading $b^i_{x,D}$ from the time-series regression (28), of portfolio cash-flow growth rates on the eight-quarter moving average of the innovation risk factor. Stocks are double-sorted into 2 × 3 portfolios by size (NYSE median) and one accounting characteristic: *Firm R&D* (R&D/market cap), *Profitability* (gross profits/assets), and *Tobin Q*; Lo, Mid, and Hi index the characteristic terciles, *Small* and *Big* the size halves. The risk factors are the R&D shock (ϵ^s, Table 3) and effective R&D level ($\tilde{s}$, Table 2), and the sample runs 1975 Q1–2022 Q4.

		Firm R&D			*Profitability*			*Tobin Q*		
		Lo	Mid	Hi	Lo	Mid	Hi	Lo	Mid	Hi
ϵ^s	*Small*	−0.00	−0.02	1.09	−0.18	0.42	0.71	0.27	0.20	−0.06
$\tilde{s}$		−0.58	−0.73	−3.37	−0.82	0.56	6.16	1.92	0.63	−0.37
ϵ^s	*Big*	−0.29	−0.28	−1.20	−0.22	−0.67	−0.07	−0.47	−0.52	0.03
$\tilde{s}$		−0.03	−0.07	−0.17	−0.02	−0.12	−0.05	−0.06	−0.05	−0.07

6.2. The fundamental origins of innovation risk

6.2.1. Empirical implementation

The premise of Bansal et al. (2005) is that returns can be approximated as the news about future cash-flow growth minus the news about future returns. The beta of an asset on a risk factor x_t therefore decomposes into a cash-flow beta and a discount-rate beta, $\beta^i_x \approx \beta^i_{x,D} - \beta^i_{x,R}$, with the cash-flow beta theoretically inheriting the price λ_x from (5). The cash-flow beta is measured by the loading $b^i_{x,D}$ in the univariate time-series regression

$$\Delta \ln D^i_t = b^i_{0,D} + b^i_{x,D}\left(\frac{1}{K}\sum_{l=1}^{K} x_{t-l}\right) + u^i_t, \tag{28}$$

where D^i_t denotes the cash flows of asset i and $K = 8$ to match (Bansal et al., 2005). Indeed, when x_t is serially uncorrelated – as the R&D shock ϵ^s_t is – $b^i_{x,D}$ is asymptotically proportional to the loading of the average future cash-flow growth $\frac{1}{K}\sum_{j=1}^{K}\Delta \ln D^i_{t+j}$ on x_t, i.e. to the long-run response of cash flows to the risk factor, an approximation of the cash-flow beta. The estimates $\hat{b}^i_{x,D}$ are then used in the second step of the Fama and Macbeth (1973) procedure to obtain an estimate of λ_x. The procedure abstracts from short-run consumption risks, as their prices are set to zero, $\lambda_c \approx 0$, consistent with their negligible contribution to the equity premium (Mehra and Prescott, 1985). Further details are in Appendix C.3.

The factors considered are the baseline R&D shock and the effective R&D level, with consumption and raw productivity growth as benchmarks, i.e. $x_t \in \{\hat{\upsilon}^s_t, \hat{\tilde{s}}_t, \Delta \ln C_t, \Delta \ln Z_t\}$. The main factor of interest, the R&D shock, is used bare because the estimation procedure, by design, nets out discount-rate news and cumulates the cash-flow response over time, leaving no need to treat investor under-reaction to R&D separately. The effective R&D level, which retains the predictable component of the series, is considered only to support the interpretation of the cash-flow betas, since its serial correlation distorts their magnitudes. The test assets are stock portfolios, grouped into three nested pools: the 'legacy pool' contains 15 portfolios sorted by size, book-to-market, and past-year returns, as in the original study; the 'extended pool' adds 36 portfolios sorted 2 × 3 on size and characteristics related to R&D investment; and the 'wide pool' adds a further 17 industry portfolios. The additional characteristics considered in the extended pool span innovation intensity (firm-specific R&D over market value), financing capacity (leverage, turnover, and profitability), and growth opportunities (asset growth and Tobin's Q), which are dimensions that are likely to generate heterogeneous exposure to the LRR carried by aggregate R&D.[24] Return and cash-flow growth series begin in 1967 Q1 for legacy and industry portfolios and 1975 Q1 for the others, ending 2022 Q4. Construction details and series statistics are in Appendix C.1.2.

6.2.2. Results

Table 5 reports the estimated cash-flow betas of selected portfolios from the extended pool, on the R&D shock and the effective R&D level. These are the portfolios sorted on firm-specific R&D intensity, profitability, and Tobin Q, representing innovation intensity, financing capacity, and growth opportunities, respectively. Appendix C.1.2 reports cash-flow betas for all test assets and risk factors.

Three patterns stand out. First, larger firms show little dispersion in their cash-flow betas across all sortings and all risk factors, with those on innovation mostly negative. Second, among smaller firms, the cash-flow betas on the R&D factors rise with financing capacity and fall with growth opportunities, consistent with increases in aggregate R&D being easier to act on for better-financed firms, while eroding the rents of firms already positioned for growth.[25] Third, more innovation-intensive firms carry higher cash-flow

[24] Variation in R&D intensity governs the spillover and 'fishing-out' effects a firm experiences, shown to be priced by Jiang et al. (2016), while financing capacity and growth opportunities shape how firms respond to innovation shocks (Brown et al., 2009; Maletic, 2018; Li, 2011; Zhang, 2014; Hall, 2002; Aghion et al., 2012). These sorts also yield return statistics bearing on whether firm-specific R&D intensity is priced, an open question since Chan et al. (2001) Leung et al. (2020) and Ahmed et al. (2025).

[25] R&D cash-flow betas also rise monotonically with leverage, which is harder to read since higher leverage can signal either exhausted debt capacity or better access to finance. The absence of such monotonicity in consumption betas favors the access interpretation over the distress one, though it may also reflect other correlated firm characteristics, an issue the coarse sorting aggravates.

Table 6
Cash-flow risk premia estimated following Bansal et al. (2005). Each entry is the estimated risk premium from the second pass of the Fama and Macbeth (1973) procedure, regressing portfolio returns on the cash-flow beta estimates from (28); HAC t-statistics are in square brackets, and R^2 and MAPE (mean absolute pricing error) measure the cross-sectional fit. Premia and MAPE are annualized, in percent. The factors are consumption growth (Cons.), raw TFP growth (Prod.), and the R&D shock (R&D, Table 3). The legacy pool comprises 15 test assets (224 quarters, 1967 Q1–2022 Q4); the extended pool, 51 (192 quarters, 1975 Q1–2022 Q4); and the wide pool, 68 (192 quarters, 1975 Q1–2022 Q4).

	Legacy pool			*Ext. pool*			*Wide pool*		
	Cons.	Prod.	R&D	Cons.	Prod.	R&D	Cons.	Prod.	R&D
λ_x	1.04	0.82	1.39	1.27*	1.59**	2.84**	1.13*	1.43**	2.39*
	[1.42]	[1.48]	[1.40]	[1.65]	[2.15]	[2.14]	[1.93]	[2.53]	[1.92]
R^2 (%)	54.21	59.62	55.46	3.06	23.66	20.32	2.17	18.61	17.33
MAPE	0.59	0.48	0.61	1.91	1.65	1.76	1.96	1.80	1.86

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

betas on the R&D shock but lower cash-flow sensitivity to the effective R&D level, suggesting they enjoy spillovers when aggregate R&D rises unexpectedly, yet must fund their own R&D when aggregate increases are anticipated, at the expense of payouts.

Table 6 reports the cross-sectional long-run risk premia originated in cash flows, estimated from consumption, productivity, and R&D. The premia measured directly from consumption are stable across asset pools, around 1% annually, above the estimate of Bansal et al. (2005) but well within its confidence interval, and significant only for the larger pools. In the same pools, the premia from productivity and R&D explain a significantly larger share of return variability than consumption does, with productivity achieving the best fit of all specifications across all pools. The R&D premia, in turn, are consistently the highest, significant in both larger pools and nearly double the productivity-based premium (close to three percent annually) in the extended pool, where exposure to aggregate R&D varies most. Though only a fraction of the return-based premia on R&D shocks, these R&D cash-flow premia support fundamentals as a genuine source of innovation risk, further confirming the pass-through of long-run risk via cash flows.

7. Conclusion

This paper shows that fluctuations in R&D drive persistent movements in productivity and consumption growth, and supports the theoretical prediction that this makes R&D a significant risk factor in financial markets. The analysis centers on a theoretical transformation of aggregate R&D, termed 'effective R&D', that maintains a linear relationship with future productivity growth by absorbing, under both fully and semi-endogenous growth mechanisms, the nonlinear effectiveness arising from spillovers and product proliferation. A contaminated version of effective R&D is recovered empirically from a single-equation estimate of the cointegrating relationship among R&D, productivity, and the labor force; an estimator for the net version, relying on a one-period TFP growth forecast and subject to more stringent assumptions, is also proposed. Embedding either series in a VAR with productivity growth identifies structural shocks to R&D, scaled by their effectiveness, which are shown to affect long-run productivity and consumption prospects beyond business-cycle horizons while leaving the short run largely unaffected. The risk premium on these shocks, estimated robustly to omitted factors, is positive and highly significant, underscoring the relevance of both aggregate R&D fluctuations and long-run risk for investors. The analysis further highlights the cash-flow channel in transmitting innovation long-run risk, with exposures that vary across firm characteristics related to R&D, spanning innovation intensity, financing capacity, and growth opportunities. Taken together, these findings establish aggregate innovation as a central source of macroeconomic risk with distinct asset-pricing implications, while the framework developed here provides a foundation for future work, including applications to other economies and the study of different forms of R&D.

Declaration of generative AI use

During the preparation of this work the author used ChatGPT, Claude and Gemini to improve language and readability. The author reviewed and edited the content after using these tools and takes full responsibility for the content of the published article.

Acknowledgments

This paper was a chapter of my Ph.D. thesis and was funded by the European Union – NextGenerationEU, Mission 4, Component 2, in the framework of the GRINS project (Growing Resilient, INclusive and Sustainable; PE00000018 – CUP J33C22002910001). Special thanks to my advisors Max Croce and Martín Gonzalez-Eiras for many helpful discussions. I also thank Eric Swanson (the editor) and an anonymous referee, as well as Emanuele Bacchiocchi, Frederico Belo, Jeremy Boccanfuso, Laura Bottazzi, Marco Brianti, Svetlana Bryzgalova, Giuseppe Cavaliere, Enzo Dia, Luca Fanelli, Luca Gemmi, Howard Kung, Filippo Massari, Antonio Minniti, Elisa Ossola, Oliviero Pallanch, Andrea Renzetti and Alwyn Young, in addition to seminar participants at the University of Bologna, the University of Milano-Bicocca, the Catholic University of Milan, the INSPIRE seminar series, and the 27th INFER Annual Conference for comments and discussions from which I greatly benefited. The views and opinions expressed are solely those of the author and do not necessarily reflect those of the European Union, nor can the European Union be held responsible for them.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jmoneco.2026.103987.

Data availability

Generated data is publicly available and provided among the attached files. Input data is mostly publicly available. Portfolios data in cash-flow analysis can be made available upon request.

Effective R&D (Original data) (Zenodo)

References

Aaronson, S., Cajner, T., Fallick, B., Galbis-Reig, F., Smith, C., Wascher, W., 2014. Labor force participation: Recent developments and future prospects. Brook. Pap. Econ. Act. 45 (2), 197–275.

Aghion, P., Askenazy, P., Berman, N., Cette, G., Eymard, L., 2012. Credit constraints and the cyclicality of R&D investment: Evidence from France. J. Eur. Econ. Assoc. 10 (5), 1001–1024. http://dx.doi.org/10.1111/j.1542-4774.2012.01093.x.

Ahlgren, N., Catani, P., 2017. Wild bootstrap tests for autocorrelation in vector autoregressive models. Statist. Papers 58 (4), 1189–1216. http://dx.doi.org/10.1007/s00362-016-0744-0.

Ahmed, S., Bu, Z., Ye, X., 2025. Illiquidity, R&D investment, and stock returns. J. Money Credit. Bank. 57 (4), 981–1022. http://dx.doi.org/10.1111/jmcb.13053.

Ai, H., Croce, M.M., Diercks, A.M., Li, K., 2018. News shocks and the Production-Based term structure of equity returns. Rev. Financ. Stud. 31, 2423–2467. http://dx.doi.org/10.1093/rfs/hhy015.

Aksoy, Y., Basso, H.S., Smith, R.P., Grasl, T., 2019. Demographic structure and macroeconomic trends. Am. Econ. J.: Macroecon. 11 (1), 193–222. http://dx.doi.org/10.1257/mac.20170114.

Alessi, L., Barigozzi, M., Capasso, M., 2010. Improved penalization for determining the number of factors in approximate factor models. Statist. Probab. Lett. 80 (23), 1806–1813. http://dx.doi.org/10.1016/j.spl.2010.08.005.

Antolin-Diaz, J., Surico, P., 2025. The Long-Run effects of government spending. Am. Econ. Rev. 115 (7), 2376–2413. http://dx.doi.org/10.1257/aer.20231278.

Anzoategui, D., Comin, D., Gertler, M., Martinez, J., 2019. Endogenous technology adoption and R&D as sources of business cycle persistence. Am. Econ. J.: Macroecon. 11 (3), 67–110. http://dx.doi.org/10.1257/mac.20170269.

Bai, J., Ng, S., 2002. Determining the Number of Factors in Approximate Factor Models. Econometrica 70 (1), 191–221. http://dx.doi.org/10.1111/1468-0262.00273.

Bansal, R., Dittmar, R.F., Lundblad, C.T., 2005. Consumption, dividends, and the cross section of equity returns. J. Financ. 60, 1639–1672. http://dx.doi.org/10.1111/j.1540-6261.2005.00776.x.

Bansal, R., Kiku, D., Ochoa, M., 2016. Price of Long-Run Temperature Shifts in Capital Markets. Technical Report w22529, National Bureau of Economic Research, http://dx.doi.org/10.3386/w22529.

Bansal, R., Shaliastovich, I., 2013. A Long-Run risks explanation of predictability puzzles in bond and currency markets. Rev. Financ. Stud. 26, 1–33. http://dx.doi.org/10.1093/rfs/hhs108.

Bansal, R., Yaron, A., 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. J. Financ. 59, 1481–1509. http://dx.doi.org/10.1111/j.1540-6261.2004.00670.x.

Beeler, J., Campbell, J.Y., 2012. The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment. Crit. Financ. Rev. 1 (1), 141–182. http://dx.doi.org/10.1561/104.00000004.

Benigno, G., Fornaro, L., 2018. Stagnation traps. Rev. Econ. Stud. 85 (3), 1425–1470. http://dx.doi.org/10.1093/restud/rdx063.

Beqiraj, E., Cao, Q., Minetti, R., Tarquini, G., 2025. Persistent slumps: Innovation and the credit channel of monetary policy. Eur. Econ. Rev. 172, http://dx.doi.org/10.1016/j.euroecorev.2024.104946.

Bloom, N., Jones, C.I., Van Reenen, J., Webb, M., 2020. Are ideas getting harder to find? Am. Econ. Rev. 110, 1104–1144. http://dx.doi.org/10.1257/aer.20180338.

Bottazzi, L., Peri, G., 2007. The International Dynamics of R&D and Innovation in the Long Run and in the Short Run. Econ. J. 117, 486–511. http://dx.doi.org/10.1111/j.1468-0297.2007.02027.x.

Broda, C., Weinstein, D.E., 2010. Product creation and destruction: Evidence and price implications. Am. Econ. Rev. 100 (3), 691–723. http://dx.doi.org/10.1257/aer.100.3.691.

Brown, J.R., Fazzari, S.M., Petersen, B.C., 2009. Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom. J. Financ. 64 (1), 151–185. http://dx.doi.org/10.1111/j.1540-6261.2008.01431.x.

Bryzgalova, S., Julliard, C., Huang, J., 2025. Consumption in Asset Returns. SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.3783070.

Chan, L.K.C., Lakonishok, J., Sougiannis, T., 2001. The stock market valuation of research and development expenditures. J. Financ. 56, 2431–2456. http://dx.doi.org/10.1111/0022-1082.00411.

Cochrane, J.H., 2005. Asset Pricing, Rev. ed ed. Princeton University Press.

Colacito, R., Croce, M.M., 2011. Risks for the Long Run and the Real Exchange Rate. J. Political Econ. 119 (1), 153–181. http://dx.doi.org/10.1086/659238.

Comin, D., Gertler, M., 2006. Medium-Term Business Cycles. Am. Econ. Rev. 96 (3), 523–551. http://dx.doi.org/10.1257/aer.96.3.523.

Constantinides, G.M., Ghosh, A., 2011. Asset Pricing Tests with Long-run Risks in Consumption Growth. Rev. Asset Pricing Stud. 1 (1), 96–136. http://dx.doi.org/10.1093/rapstu/rar004.

Croce, M.M., 2014. Long-run productivity risk: A new hope for production-based asset pricing? J. Monet. Econ. 66, 13–31. http://dx.doi.org/10.1016/j.jmoneco.2014.04.001.

Dew-Becker, I., Giglio, S., 2016. Asset pricing in the frequency domain: Theory and empirics. Rev. Financ. Stud. 29 (8), 2029–2068. http://dx.doi.org/10.1093/rfs/hhw027.

Eberhart, A.C., Maxwell, W.F., Siddique, A.R., 2004. An Examination of Long-Term Abnormal Stock Returns and Operating Performance Following R&D Increases. J. Financ. 59 (2), 623–650. http://dx.doi.org/10.1111/j.1540-6261.2004.00644.x.

Epstein, L.G., Farhi, E., Strzalecki, T., 2014. How much would you pay to resolve long-run risk? Am. Econ. Rev. 104 (9), 2680–2697. http://dx.doi.org/10.1257/aer.104.9.2680.

Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. Econometrica 57, 937. http://dx.doi.org/10.2307/1913778.

Erceg, C.J., Levin, A.T., 2014. Labor force participation and monetary policy in the wake of the great recession. J. Money Credit. Bank. 46 (S2), 3–49.

Fama, E.F., Macbeth, J.D., 1973. Risk, return, and equilibrium: Empirical tests. J. Political Econ. 81, 607–636.
Fernald, J.G., 2012. A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Working Paper Series, 1–28, Federal Reserve Bank of San Francisco, http://dx.doi.org/10.24148/wp2012-19.
Fujita, S., 2014. On the Causes of Declines in the Labor Force Participation Rate. Research Rap Special Report, Federal Reserve Bank of Philadelphia.
Giglio, S., Xiu, D., 2021. Asset pricing with omitted factors. J. Political Econ. 129 (7), 1947–1990. http://dx.doi.org/10.1086/714090.
Gourieroux, C., Jasiak, J., 2024. Long-run risk in stationary vector autoregressive models. J. Econometrics http://dx.doi.org/10.1016/j.jeconom.2024.105905.
Ha, J., Howitt, P., 2007. Accounting for Trends in Productivity and R&D: A Schumpeterian Critique of Semi-Endogenous Growth Theory. J. Money Credit. Bank. 39, 733–774. http://dx.doi.org/10.1111/j.1538-4616.2007.00045.x.
Hall, B.H., 2002. The financing of research and development. Oxf. Rev. Econ. Policy 18 (1), 35–51. http://dx.doi.org/10.1093/oxrep/18.1.35.
Herzer, D., 2022. Semi-endogenous Versus Schumpeterian growth models: A critical review of the literature and new evidence. Rev. Econ. 73 (1), 1–55. http://dx.doi.org/10.1515/roe-2021-0023.
Hsu, P.-H., 2009. Technological innovations and aggregate risk premiums. J. Financ. Econ. 94 (2), 264–279. http://dx.doi.org/10.1016/j.jfineco.2009.01.002.
Jensen, T.I., Kelly, B.T., Pedersen, L.H., 2021. Is there a replication crisis in finance? SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.3774514.
Jiang, Y., Qian, Y., Yao, T., 2016. R&D spillover and predictable returns. Rev. Financ. 20, 1769–1797. http://dx.doi.org/10.1093/rof/rfv050.
Jones, C.I., 1995. R&D-based models of economic growth. J. Political Econ. 103 (4), 759–784. http://dx.doi.org/10.1086/262002.
Jones, C.I., 1999. Growth: With or without scale effects? Am. Econ. Rev. 89 (2), 139–144. http://dx.doi.org/10.1257/aer.89.2.139.
Kaltenbrunner, G., Lochstoer, L.A., 2010. Long-Run Risk through Consumption Smoothing. Rev. Financ. Stud. 23, 3190–3224. http://dx.doi.org/10.1093/rfs/hhq033.
Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N., 2017. Technological innovation, resource allocation, and growth. Q. J. Econ. 132 (2), 665–712. http://dx.doi.org/10.1093/qje/qjw040.
Kruse-Andersen, P.K., 2023. Testing R&D-based endogenous growth models. Oxf. Bull. Econ. Stat. http://dx.doi.org/10.1111/obes.12552.
Kung, H., Schmid, L., 2015. Innovation, growth, and asset prices. J. Financ. 70, 1001–1037. http://dx.doi.org/10.1111/jofi.12241.
Leung, W.S., Evans, K.P., Mazouz, K., 2020. The R&D anomaly: Risk or mispricing? J. Bank. Financ. 115, http://dx.doi.org/10.1016/j.jbankfin.2020.105815.
Li, D., 2011. Financial Constraints, R&D Investment, and Stock Returns. Rev. Financ. Stud. 24 (9), 2974–3007. http://dx.doi.org/10.1093/rfs/hhr043.
Liu, Y., Matthies, B., 2022. Long-Run risk: Is it there? J. Financ. 77 (3), 1587–1633. http://dx.doi.org/10.1111/jofi.13126.
Ludvigson, S.C., Ng, S., 2009. Macro Factors in Bond Risk Premia. Rev. Financ. Stud. 22, 5027–5067. http://dx.doi.org/10.1093/rfs/hhp081.
Madsen, J.B., 2008. Semi-endogenous versus Schumpeterian growth models: testing the knowledge production function using international data. J. Econ. Growth 13 (1), 1–26. http://dx.doi.org/10.1007/s10887-007-9024-0.
Maletic, M., 2018. Profitability, R&D investments and the cross-section of stock returns. SSRN Electron. J. http://dx.doi.org/10.2139/ssrn.3178186.
Mehra, R., Prescott, E.C., 1985. The equity premium: A puzzle. J. Monet. Econ. 15 (2), 145–161. http://dx.doi.org/10.1016/0304-3932(85)90061-3.
Montiel Olea, J.L., Plagborg-Møller, M., 2021. Local projection inference is simpler and more robust than you think. Econometrica 89 (4), 1789–1823. http://dx.doi.org/10.3982/ECTA18756.
Montiel Olea, J.L., Plagborg-Møller, M., Qian, E., Wolf, C.K., 2025. Local Projections or VARs? A Primer for Macroeconomists. NBER Working Paper 33871, National Bureau of Economic Research, Cambridge, MA, http://dx.doi.org/10.3386/w33871.
Moran, P., Queralto, A., 2018. Innovation, productivity, and monetary policy. J. Monet. Econ. 93, 24–41. http://dx.doi.org/10.1016/j.jmoneco.2017.10.006.
Müller, U.K., 2005. Size and power of tests of stationarity in highly autocorrelated time series. J. Econometrics 128 (2), 195–213. http://dx.doi.org/10.1016/j.jeconom.2004.08.012.
Ortu, F., Tamoni, A., Tebaldi, C., 2013. Long-Run risk and the persistence of consumption shocks. Rev. Financ. Stud. 26, 2876–2915. http://dx.doi.org/10.1093/rfs/hht038.
Peretto, P.F., 1998. Technological change and population growth. J. Econ. Growth 3 (4), 283–311. http://dx.doi.org/10.1023/A:1009799405456.
Phillips, P.C.B., Hansen, B.E., 1990. Statistical inference in instrumental variables regression with I(1) processes. Rev. Econ. Stud. 57 (1), 99. http://dx.doi.org/10.2307/2297545.
Ready, R.C., 2018. Oil consumption, economic growth, and oil futures: The impact of long-run oil supply uncertainty on asset prices. J. Monet. Econ. 94, 1–26. http://dx.doi.org/10.1016/j.jmoneco.2017.09.010.
Reeb, D.M., Zhao, W., 2020. Patents do not measure innovation success. Crit. Financ. Rev. 9 (1), 157–199. http://dx.doi.org/10.1561/104.00000087.
Rotemberg, J.J., 2003. Stochastic technical progress, smooth trends, and nearly distinct business cycles. Am. Econ. Rev. 93 (5), 1543–1559. http://dx.doi.org/10.1257/000282803322655437.
Schorfheide, F., Song, D., Yaron, A., 2018. Identifying Long-Run risks: A Bayesian mixed-frequency approach. Econometrica 86 (2), 617–654. http://dx.doi.org/10.3982/ecta14308.
Sedgley, N., Elmslie, B., 2013. The dynamic properties of endogenous growth models. Macroecon. Dyn. 17 (5), 1118–1134. http://dx.doi.org/10.1017/S1365100512000119.
Stock, J.H., Watson, M.W., 1993. A simple estimator of cointegrating vectors in higher order integrated systems. Econometrica 61, 783. http://dx.doi.org/10.2307/2951763.
Vogelsang, T.J., Wagner, M., 2014. Integrated modified OLS estimation and fixed- inference for cointegrating regressions. J. Econometrics 178 (2), 741–760. http://dx.doi.org/10.1016/j.jeconom.2013.10.015.
Zhang, W., 2014. R&D investment and distress risk. J. Empir. Financ. 32, 94–114. http://dx.doi.org/10.1016/J.JEMPFIN.2015.03.009.