Search Costs, Working Capital, and Asset Pricing

Matthew L. Kelly

Campbell University

November, 21 2025

Abstract

How might the cost of searching for information influence asset prices? To answer that question, this paper develops and tests a search-theoretic asset pricing model with imperfect information. In the model, firms can wait to buy inputs and sell inventories while searching for low cost suppliers and high value customers, respectively. Simultaneously, investors can wait to buy assets while searching for bargain prices and low borrowing costs. The model predicts that when the cost of searching for information rises, equity returns fall and the marginal utility of consumption spikes, leading investors to discount firms whose cash flows are especially sensitive to search cost ex ante. The model also implies that observable firm characteristics related to working capital management are indicators of search cost risk, and empirical evidence is consistent with the hypothesis that search cost risk is priced. Results suggest that excess inventory can be a hedge against search-related risks accompanying supply chain fragility or information suppression.

1 Introduction

Uncertainty pervades economic decision-making. Because knowledge is diffuse and costly to acquire, agents spend time searching for information about prices, opportunities, and counterparties (Stigler, 1961; Hayek, 1948). Information is unevenly distributed (Akerlof, 1970), and search is mutual and ubiquitous (McCall, 1970): firms seek customers and suppliers, workers and employers must find one another, consumers and investors hunt for firms offering bargain sales or investment opportunities. In the limiting benchmark of perfect informational efficiency, prices reflect all available information (Fama, 1970), effectively eliminating the need to locate counterparties. In that frictionless world, the institutional nexus of the firm is not even necessary (Coase, 1937). In reality, search frictions are persistent and heterogeneous, shaping industrial

organization (Ellison and Ellison, 2009; Wilson, 2012). This paper argues that search frictions also shape asset prices.

I develop a search-theoretic asset pricing model in which both firms and investors face costly information acquisition. Firms can delay purchases while searching for lower input costs and delay sales while searching for higher-paying customers. Faster search improves margins in ways reminiscent of market power. However, unlike monopolies, firms with search advantages raise output rather than restrict it. Investors also search: cross-sectionally for underpriced securities and intertemporally by waiting to buy when consumption growth (and thus borrowing costs) are low. Search costs therefore depress firm profits while increasing the marginal utility of consumption, generating a systematic risk channel. The model predicts that firms with greater exposure to search-cost shocks have a higher cost of capital, and that this exposure is mediated by working capital composition.

Working capital is integral to search. Extending search over suppliers or customers requires short-term resources like inventory, cash, or credit lines that provide ready liquidity. The model delivers endogenous turnover policies for cash and inventories and implies a tight mapping from search-risk exposure to balance-sheet composition. A key theoretical result is that sensitivity of returns to search costs is closely related to the inventory-to-current-assets ratio: firms with lower inventories relative to other current assets are more exposed to search-cost risk.

Empirically, I focus on two firm characteristics derived in Section 3: the ratio of inventories to current assets (INVT/ACT) and the ratio of inventories to market equity (INVT/ME). I include INVT/ME in the analysis as a sanity check, because it is similar to INVT/ACT but has no clear theoretically justified reason to earn excess returns. Long–short portfolios sorted on INVT/ACT (ILMH: inventory lean minus heavy) and INVT/ME (IGMV: inventory growth minus value), rebalanced monthly using data from the previous quarter, both earn high returns. Despite a low Sharpe ratio of 0.28 (0.22), ILMH (IVMG) earns an annualized alpha of about 10% (13%) relative to the Fama and French (2014) five-factor model, and about 7% (6%) relative to the Huo, Mo, Zue, and Zhang (2021) q-factor model. Results are robust across numerous time-series specifications. Fama–MacBeth regressions largely corroborate this evidence in the cross-section: controlling for past returns, time-varying CAPM betas, size, book-to-market, profitability, investment growth, leverage, time-varying liquidity betas, and monopoly power, a one-standard-deviation decline in INVT/ACT (INVT/ME) is associated with a 1.2% (0.2%) increase in returns one quarter ahead. The theory-driven, hypothesis-based INVT/ACT premium is statistically significant and very robust. The more speculatively postulated INVT/ME premium is less robust, less economically significant, and partly explained by industry effects and outliers.

These ratios vary substantially across industries and over time, and are not subsumed by existing

inventory-growth factors. The inventory growth factor of Belo and Lin (2012) and the inventory-as-a-share-of-total-assets growth factor of Thomas and Zhang (2013) load with opposite signs on both long–short portfolios, indicating that INVT/ACT and INVT/ME capture something orthogonal to inventory investment growth. Guided by the model, I interpret these ratios as measures of exposure to systematic search costs.

Surely many other factors unrelated to search influence inventory investment and other components of working capital. These factors are discussed in the next section, but I bracket them theoretically and control for them empirically as much as possible in order to focus on search costs.

Section 2 reviews related literature. Section 3 presents the model. Section 4 describes the data, empirical design, and results. Section 5 concludes.

2 Literature Review

As Blinder (1981) memorably notes, "the business cycle, to a surprisingly large degree, is an inventory cycle." Net inventory investment is small in level yet highly volatile and procyclical, accounting for a substantial share of output declines in recessions; recoveries are often marked by liquidations of excess stocks (Bosworth, 1970; Ramey and West, 1999). A large macro tradition models inventories using progressively richer micro-foundations. Early Keynesian "stock-adjustment" models (Metzler, 1941; Lovell, 1960) generated cyclical inventory dynamics but faced well-known tensions, including implausibly slow estimated adjustment speeds and production volatility exceeding sales volatility. Later neoclassical explanations emphasize either production smoothing (inventories economize on costs of adjusting output) under rational expectations (Kahn, 1992; Bils and Kahn, 2000; Khan and Thomas, 2007), or (s,S) policies (inventories economize on fixed delivery costs by inducing lumpy replenishment) meant at avoid stockouts (Arrow, Harris, and Marschak, 1951; Holt et al., 1960). While both mechanisms are plausible and empirically important, they largely ignore a distinct motive central to this paper: inventories (and liquid working capital more broadly) may be held to support costly search for customers, suppliers, and contracting terms. Firms may delay sale until high-paying buyers are found, or even maintain stocks to attract bids and learn about demand when price information is noisy, incomplete, or slow to emerge.

A related macro literature does connect inventory behavior to structural change in the information and distribution environment. The "Great Moderation" evidence documents a sharp decline in U.S. output volatility and inflation beginning in the mid-1980s, partly attributed to inventory-intensive (especially durable-goods) sectors. McConnell and Perez-Quiros (2000) tie the break to inventory dynamics, and subsequent work argues that improved information technology and supply-chain management altered the timing and amplitude of inventory adjustment. Chen, Frank, and Wu (2005) document large secular declines in

inventory holding periods. Davis and Kahn (2008) provide complementary evidence consistent with better inventory control contributing to lower volatility. Industry studies reinforce this view: Ramey and Vine (2006) document changing interactions among sales, inventories, and production in the auto industry. This body of work strongly suggests that information/coordination frictions lessened over time to allow leaner inventory management, but it rarely models underlying information search frictions explicitly.

Inventories are remarkably elastic. Carpenter, Fazzari, and Petersen (1994) emphasize that inventories can adjust more quickly than fixed capital and therefore absorb cash-flow shocks. Kashyap, Lamont, and Stein (1994) show that inventory investment is highly sensitive to credit conditions, especially for bank-dependent firms, highlighting inventories as a margin through which monetary tightening transmits to real activity. Gertler and Gilchrist (1994) similarly find that small firms drive inventory drawdowns after monetary contractions. Small firms might be more responsive because they are more financially constrained. The model developed by Dasgupta, Li, and Yan (2019) suggests that firms with borrowing constraints exhibit more volatile inventory investment, with slower accumulation and quicker depletion in response to demand shocks.

Finance and accounting studies find that inventories forecast cash flows and returns, but prevailing interpretations typically run through investment frictions or mispricing rather than systematic search risk. Thomas and Zhang (2002) show that the accrual anomaly documented by Sloan (1996) is largely attributable to inventory changes, underscoring that inventory dynamics embed information the market processes imperfectly. Excessive accumulation of inventories, relative to sales, is usually perceived as a negative signal of firm performance, affecting stock returns (Abarbanell and Bushee, 1998). Production-based asset pricing links firm policies and cost structures to exposure to aggregate risk (Cochrane, 1991). In this framework, inventory investment often features as a production input with special adjustment costs. Chen (2016) shows that introducing inventories into a time-to-build/time-to-produce environment can shape equilibrium asset prices and term premia. Inventory growth predicts returns (Belo and Lin, 2012), with related evidence in Jones and Tuzel (2013). Separately, a growing operations-finance interface documents that inventory efficiency and turnover predict returns: Chen, Frank, and Wu (2005) show that abnormally high inventories predict poor long-run stock performance, while Alan, Gao, and Gaur (2014) find strong return predictability from inventory productivity in retail. Beyond inventories, complementary production-based mechanisms emphasize how the composition of costs and the durability of output affect risk premia (Gomes, Kogan, and Yogo, 2009). The working-capital and liquidity literature also documents return predictability tied directly to short-term balance-sheet dynamics. Pastor and Stambaugh (2003) establish that market-wide liquidity risk is priced in expected returns, and Wang (2019) shows that the cash conversion cycle predicts stock returns in the cross-section. These papers establish that inventory policy, and working capital policy more broadly,

is tied to risk premia, but they generally do not interpret inventory levels and working-capital composition as exposures to an economy-wide information/search-cost state variable. In my setting, higher search costs simultaneously depress firm profits and raise marginal utility by limiting the flow of economically meaningful information, and firms' balance-sheet liquidity determines the strength of this covariance channel.

On the supply side, search for productive inputs necessarily occurs in complex supply chain networks. The financing of inventories and procurement often occurs through trade credit and in-kind business-to-business arrangements. Burkart and Ellingsen (2004) provide a theory of in-kind finance that helps rationalize why short-term operating assets can be financed in ways that differ from cash, and why this matters for firm behavior under frictions. Cohen and Frazzini (2008) show return predictability across economically linked firms, consistent with information frictions in networks. Barrot and Sauvagnat (2016) identify supplier shocks using natural disasters and show they propagate to customers' output and market values, providing direct evidence that production links transmit real and financial shocks. Herskovic (2018) derives priced factors from input-output networks, emphasizing that network structure itself is a source of systematic risk. Related empirical work measures supply-chain risk and connects it to stock market outcomes. This literature formalizes linkages and propagation, but it typically takes the mechanism of contracting counterparties as primitive. In contrast, my approach emphasizes the search process - how quickly firms and investors can locate and evaluate trading partners - as a distinct, priced channel that becomes especially important when networks are strained.

My empirical measures are designed to capture the distinct search channel implied by the model. Whereas the inventory-growth literature centers on changes in inventory investment (e.g., Belo and Lin, 2012; Jones and Tuzel, 2013), I study inventory composition: inventories scaled by current assets (INVT/ACT) and by market equity (INVT/ME). These ratios isolate how much of short-term balance-sheet capacity is committed to physical stock rather than other liquid resources (cash, receivables, etc.), and therefore proxy for a firm's ability to extend customer and supplier search, buffer delays, and transact quickly when information is scarce. The ratios also align with a broader working-capital perspective in which the timing of cash flows—and hence effective liquidity—matters for risk (Wang, 2019). In short, the contribution is not just "inventories matter," but that the level and composition of working capital carries distinct information about exposure to systematic search-cost shocks.

The economic theory of search has deep roots. Stigler (1961, 1962) establishes the canonical point that information is costly and price dispersion is a natural consequence of search. McCall (1965) formalizes optimal search stopping times in a dynamic framework, with the most developed applications in labor markets (Mortensen, 1986; Mortensen and Pissarides, 1994). In a product-market context, Alchian (1969) argues that inventories can facilitate bid solicitation and information discovery when reliable price data

are scarce, implying that unsold goods may serve an information-gathering function. Lachmann's work, especially "Commodity Stocks and Equilibrium" (1936), "Uncertainty and Liquidity Preference" (1937), and "Commodity Stocks in the Trade Cycle" (Lachmann and Snapper, 1938), treats inventories as instruments of flexibility that support planning and adaptation when price signals are incomplete. Like Lachmann, Hutt's *The Theory of Idle Resources* (1939) frames idle stocks and capacity as potentially equilibrium responses when adjustment requires time, information, and re-coordination. The gap is that these information-centered perspectives are rarely embedded in modern asset-pricing models as a systematic state variable that jointly affects firm profit opportunities and investors' intertemporal marginal rates of substitution. This paper takes a step toward formalizing that idea by modeling investor, customer and supplier search as a technology that maps liquid working capital into the option value of waiting and substitution.

Surely, other factors unrelated to search also play a role in determining inventories and other current assets. Firms hold cash to economize on fixed transactions costs of withdrawal (Baumol, 1952), as a hedge against capital supply risk (Hugonnier, Malamud and Morellec, 2015) and as a slush fund for empire building where agency frictions are present (Jensen, 1986). Without denying the importance of these confounding factors, I bracket them theoretically and empirically control for them as much as possible in order to isolate the effects of search costs. I do not claim that search costs entirely subsume other factors influencing working capital. I merely submit that search costs play a role worth considering.

Empirical identification faces measurement error and omitted variables—especially industry competitiveness and the extent/completeness of input and product markets. Inventory measurement quality varies within and across industries.[1] Imperfect accounting measurement can mute statistical significance, and if mismeasurement correlates with inventory levels or returns it can bias estimates. Even classification into finished goods, work-in-process, and raw materials is imperfect and may not match an item's economic role in production. Market power further confounds identification: monopolies/monopsonies may optimally hold atypical inventories (and induce adjustments along the supply chain), with return effects that could vary with cost structure and vertical integration. Finally, forward and derivatives markets can buffer suppliers and customers from shocks that would otherwise lengthen cash cycles, while also lowering search costs by making prices common knowledge. Industries without access to such derivatives could therefore manage working capital differently.

Finally, this paper relates to research on operating leverage and operating "hedges." Profits of high operating-leverage firms can be more sensitive to changes in revenue, raising cash-flow risk (Lev, 1974;

[1]The medical device industry is a good example of heterogeneity in this measurement error. Through conversations with medical device salespeople, it was learned that inventories are frequently used up during sales demonstrations but not accounted for. Part of what makes this tolerable in that industry is the high markups associated with each unit, which easily offset the costs of loss and accounting inaccuracy. Measurement problems also probably manifest themselves differently in industries with a dealership sales model verses those where manufacturers sell directly.

Carlson et al., 2004), and operating leverage has been used to rationalize cross-sectional return premia (Novy-Marx, 2010). More recent theory shows that in some settings higher variable costs can instead create an operating hedge that attenuates profit sensitivity to shocks (Kogan, Li, and Zhang, 2022; Kogan, Li, Zhang, and Zhu, 2023). The mechanism laid out in this paper is analogous, but it operates through the time structure of production and exchange embedded in working capital. When search costs rise, firms whose operations require rapid turnovers and lean inventories, i.e., those that must continually refinance purchases and restock quickly, experience more disruptions in marketing and procurement. By contrast, holding more inventory (and for a longer time) grants the option to continue operating (and learning about demand) while search is impaired, and it relaxes the need to transact with suppliers immediately at unfavorable terms. In the model, this type of operating hedge arises not so much by lowering average profitability, but rather by lowering the covariance of profits with the search-cost state variable.

3 Model

There are two agents in the model: a firm and an investor, both of whom must search for uncertain information relevant to their decisions. I begin with the firm, and focus on the special properties of its input supply and output demand. After solving the firm's profit maximization problem, I theoretically derive inventory and cash turnovers, working capital composition ratios, and calculate elasticities of profit with respect to search costs. Turning to the investor's utility maximization problem, I derive an asset pricing equation and arrive at a single beta representation of expected excess returns.

3.1 The Firm's profit maximization problem

The firm exists in a market characterized by a horizontal but stochastic demand curve. At every quantity Y the firm can produce, consumers' willingness to pay D_Y is risky, normally distributed with mean μ_Y and variance σ_Y^2. Summarizing,

$$D_Y \sim N(\mu_Y, \sigma_Y^2)$$

The firm has a production technology, represented by a Cobb-Douglass function $Y(K, L)$.

$$Y(K, L) = AK^{1-\theta}L^{\theta}$$

where A is total factor productivity, K is a fixed input, L is a variable input, and $\theta \in \{0, 1\}$ is a parameter controlling the marginal rate of technical substitution. In the long run, the firm can adjust both K and L,

but in the short run it can only adjust L. Fixing $K = \bar{K}$ in the short run, the firm decides how much L to invest. S_L, the supply of input L, is also assumed to be horizontal but stochastic. That is, for any amount of L the firm wishes to use, input suppliers' marginal costs are risky, normally distributed with mean μ_L and variance σ_L^2. Summarizing,

$$S_L \sim N(\mu_L, \sigma_L^2)$$

Assume production is instantaneous, but that the firm can wait to buy inputs and wait to sell outputs. The firm chooses to wait because it can spend that time searching for better deals. Another interpretation is equally valid: at prevailing sales prices (costs), it takes the firm some time to find customers (suppliers) willing to pay (sell), and lower search costs shorten that length of time. The sales price $P_Y = \mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y t_Y)}$ is set knowing that (on average) the firm will have to wait until time $1 + t_Y$ to find a customer willing to pay price P_Y. This is because P_Y is the expected maximum value after the $1+t_Y$'th draw from a normal distribution with mean μ_Y and standard deviation σ_Y, given a sampling rate of ζ_Y draws per unit of time.[2] During this holding time between production and sale, output is labeled "inventory" on the firm's balance sheet.[3] The parameter ζ_Y should be interpreted as the speed of potential customer arrival (or in some contexts, a rate of customer bid submission). All else equal, items with higher sales prices take a longer time to sell. Conversely, the firm pays $P_L = \mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L t_L)}$ per variable input. Waiting some amount of time before making purchases can yield savings. To remain prepared to purchase inputs when a good deal is found, the firm must hold an amount $P_L L$ in liquid capital until time t_L when it makes purchases. Firms managers who want to lower input costs (without reducing output) by spending time searching for low cost suppliers, holding non-yielding liquid assets until a satisfactory deal is found.

A note on interpretation: the firm is not perfectly price discriminating or continuously attaining the lowest feasible unit costs. It posts a single sales price and pays a single input price, chosen to maximize profits in expectation given the distributions of customer valuations and supplier costs and the sampling rate ζ; prices and costs need not vary over time unless those primitives change, even though realized quantities and the time to clear inventory fluctuate. A useful example is a retail store: setting a higher posted price means shelves empty more slowly on average, since some days few "high-valuation" customers arrive, other days many do, so the realized time to sell a batch is random, but its expected value is pinned down by the

[2]This fact, used in Alchian (1969), follows from the Fisher-Tippet-Gnedenko theorem. For more information, see Mood, A.M. (1950). "Order Statistics". Introduction to the theory of statistics. New York, NY, US: McGraw-Hill. pp. 251–270. Interested readers may want to study an online tutorial "Calculating the Gaussian Expected Maximum" which includes R code (available at https://gwern.net/order-statistic).

[3]Treating inventory as an output rather than an input avoids assumptions about the elasticity of substitution between inventories and capital or labor. There is little consensus on this point. Chen (2016) assumes it is unitary, as do Gomes, Kogan, and Yogo (2009) but he notes that Kydland and Prescott (1982) set it to 0.2, Belo and Lin (2012) set it to 2/3 and Jones and Tuzel set it between 1/3 and 1.

price rule implied by the demand distribution.[4] The sampling rate ζ summarizes search costs and visibility (foot traffic, advertising, location, etc.), treated here as a firm characteristic with sunk investment.[5] An analogous logic applies on the procurement side: the firm sets a reserve input cost, meets potential suppliers at frequency ζ_L, and purchases only when a sufficiently low-cost supplier is found, so targeting a lower input cost (or facing greater cost dispersion σ_L) requires longer expected search.

To illustrate these ideas, panels A and B of Figure 1 plots P_Y and P_L as functions of inventory and cash holding times, respectively.

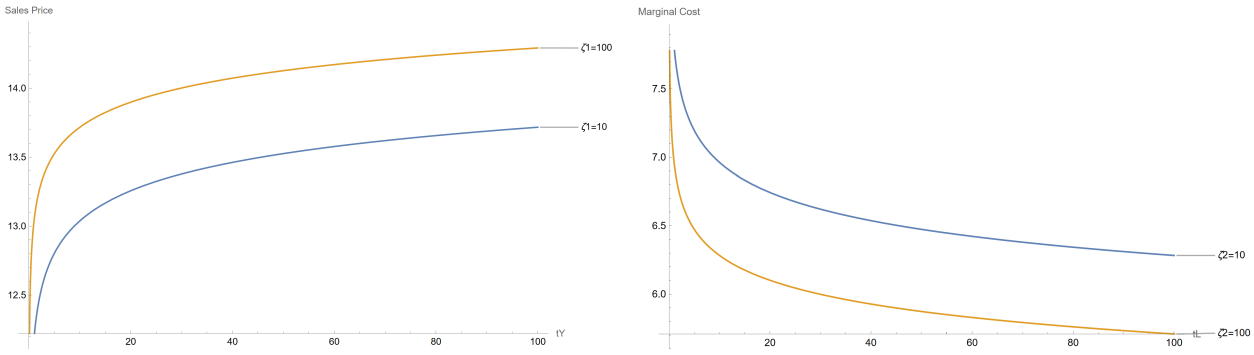


Figure 1: Sales Prices and Marginal Costs as functions of search costs. Panel A: Sales Price P_Y as a function of inventory holding time t_Y for different values of ζ_Y. Panel B: Marginal cost P_L as a function of cash holding time t_L for different values of ζ_L.

Notice that if demand or supply were completely certain ($\sigma_Y = 0$ or $\sigma_L = 0$) then there would be no benefit to searching. In this case, the firm would set its sales price (target cost) to μ_Y (μ_L) and both waiting times t_Y and t_L would be zero. On the other hand, if ζ_Y and ζ_L approach infinity, then search is effectively costless. In that case, search times would also be zero because all information is already in hand, but profits would become unbounded in the limit. It is because the firm meets only a finite amount of potential customers and suppliers each time period, rather than being constantly inundated by an infinite mass, that it takes time to search for a deal.

Let $P_K\bar{K}$ equal total fixed costs. Let the total cash cycle be $t = t_Y + t_L$, and further reparameterize $t_Y = \alpha t$ and $t_L = (1-\alpha)t$ where $\alpha \in [0,1]$. α represents the proportion of the firm's cash cycle spent searching for customers. In the time it takes the firm to purchase inputs and sell outputs (its cash cycle, t), it earns a gross profit. Multiply that per-cycle profit by the number of times the cycle is completed in a year to calculate annual gross profit. Let the amount of time in the year be normalized to $T = 1$. Subtract the firm's fixed costs from annual gross profit to arrive at the firm's annual operating profit.

[4]Non-normal distributions imply different order-statistic expressions than P_Y as written above, but the basic principle that more samples raise the maximum and lower the minimum holds broadly for distributions with finite variance.

[5]I implicitly assume that the investment in such assets is fixed, nondepreciating, and cannot be repurposed; extensions could relax this.

$$\Pi = \frac{Y(\bar{K}, L)(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)}) - L(\mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)})}{t} - P_k\bar{K}$$

First-Order Conditions.

To determine optimal decisions, differentiate the profit function with respect to arguments K, L, t and α. In the long run, optimal fixed capital is

$$K^* = L\left(\frac{A(1-\theta)(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)})}{P_K t}\right)^{1/\theta} \tag{1}$$

Fixing $K = \bar{K}$ in the short run, the firm chooses L, t and α to maximize profits. First order conditions associated with these remaining choice variables are

$$A\bar{K}\theta(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)}) = \left(\frac{\bar{K}}{L}\right)^{\theta} L(\mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)}) \tag{2}$$

$$\begin{aligned} A\bar{K}^{1-\theta}L^{\theta}\mu_Y + L\mu_L = {} & \frac{A\bar{K}^{1-\theta}L^{\theta}\sigma_Y(\alpha t\zeta_Y - 2(1+\zeta_Y\alpha t))\ln(1+\zeta_Y\alpha t)}{(1+\zeta_Y\alpha t)\sqrt{2\ln(1+\zeta_Y\alpha t)}} \\ & + \frac{L\sigma_L(\frac{t(1-\alpha)\zeta_L}{1-t(1-\alpha)\zeta_L} - 2\ln(1+\zeta_L(1-\alpha)t))}{\sqrt{2\ln(1+\zeta_L(1-\alpha)t)}} \end{aligned} \tag{3}$$

And

$$L(1+\zeta_Y\alpha t)\zeta_L\sigma_L\sqrt{\ln(1+\zeta_Y\alpha t)} = A\bar{K}^{1-\theta}L^{\theta}\zeta_Y(1+\zeta_L(1-\alpha)t)\sigma_Y\sqrt{\ln(1+\zeta_L(1-\alpha)t)} \tag{4}$$

Rearranging (2), we can write

$$L^* = \left(\frac{A\bar{K}^{1-\theta}\theta(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)})}{(\mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)})}\right)^{\frac{1}{1-\theta}}$$

Because of the combination of logarithmic and exponential terms inside and outside of radical operators in (3) and (4), no analytic expression can be found for the optimal t or α (where these variables would appears on only one side of the equation). Nonetheless, numerical solutions for both choice variables exist and are easy to find. Furthermore, the following theorem shows that a unique solution for optimal α^* exists under very general conditions.

Theorem 1. If $A > 0$, $L > 0$, $K > 0$, $t > 0$, $\sigma_Y > 0$, $\sigma_L > 0$, $\zeta_Y > 0$, $\zeta_L > 0$, then for any $t > 0$ there exists a unique optimal α.

Proof. Use a fixed point theorem. α is obviously increasing in α. Rearrange equation (1.4) to define a function f, where

$$f(\alpha) = \frac{(1+\zeta_L(1-\alpha)t)\sqrt{\ln(1+\zeta_L(1-\alpha)t)}AK^{1-\theta}L^{\theta}\sigma_Y}{L\zeta_L\sigma_L\sqrt{\ln(1+\zeta_Y\alpha t)}} - \frac{1}{\zeta_Y t}$$

which is decreasing in α. Recall that $\alpha \in \{0,1\}$. As $\alpha \to 0$ we have $f(\alpha) \to \infty$ and as $\alpha \to 1$, we have $f(\alpha) \to -\frac{1}{\zeta_Y t} < 0$. Therefore, there exists a unique value $\alpha = \alpha^*$ such that $f(\alpha^*) = \alpha^*$, and α^* is the firm's optimal proportion of time spent searching for customers by construction of f. Q.E.D.

For a given α (and other parameter settings), numerical solutions for t^* are easy to find, but a fixed-point argument for existence and uniqueness as used in theorem 1 will not apply as generally for t^*. This is because a function of t will not be nonotonic in general. Characterization of numerical solutions is also complicated by the number of parameters and functional forms involved, but the following simplified special case does clarify some of the economically important effects of key variables on a firm's optimal cash cycle.

Suppose that the firm is a pure middleman, a dealer buying and selling without any transformation happening via production. Assume also that some spoilage occurs in between purchase and resale, and that this spoilage scales with volume of business but does not scale with time.[6] In this case, the firm can only profit by exploiting its search-related advantages: finding low cost suppliers and high value customers. To capture these features of the example, set $\mu_Y = \mu_L = \mu$, $\sigma_Y = \sigma_L = \sigma$, $\theta = 1$, $A < 1$, and fix $K = 1$. For additional simplicity, let $\zeta_Y = \zeta_L = \zeta$, so the firm can search for supply equally as well as it can search for demand. Under those circumstances, for a given α, higher benefits of search (higher σ), lower costs of search (higher ζ), lower average market prices of the good (lower μ), or lower spoilage (higher A) will all serve to shorten the firm's optimal cash cycle t^*, increasing annual profits. In this special case, where the cost of supply search is identical to the cost of demand search, the length of the cash cycle is unaffected by the volume of sales. However, if demand search cost are higher than supply search costs, then, because more of the firm's cash cycle will be spent searching for demand than is spent searching for supply, an increase in sales volume will tend to shorten cash cycles.

Figure 2 visualizes the stochastic demand curve to aid the reader. Panel A plots level curves of the firm's

[6] A common and realistic example of this type of "spoilage" is present for newspaper boys, or dealers of consumer electronics products or cars. Such goods are transported from manufacturers to dealers without any "natural" spoilage as occurs for, say, milk (nor any improvement in quality as occurs for, say, wine), but nonetheless some percentage of the purchased goods are regularly lost or broken during delivery. Now, in some industries, dealers are required to pay only for the goods that arrive without defect, but the example currently under discussion would apply to a case where dealers, not producers, must bear the cost.

The example of a market maker also fits this special case in many respects. Market makers broadcast prices but don't alter the financial products they broker, profiting as middlemen only by lowering the cost of search between counterparies. The market maker's profit per match, its bid-ask spread, widens and narrows in response to the uncertainty of information and the costs of search. As sampling rates, search cost parameters ζ_Y and ζ_L could be likened to the foot traffic on a trading floor of buyers and sellers, respectively. Payment for order flow can motivate the "spoilage" cost captured by $A < 1$ in the model.

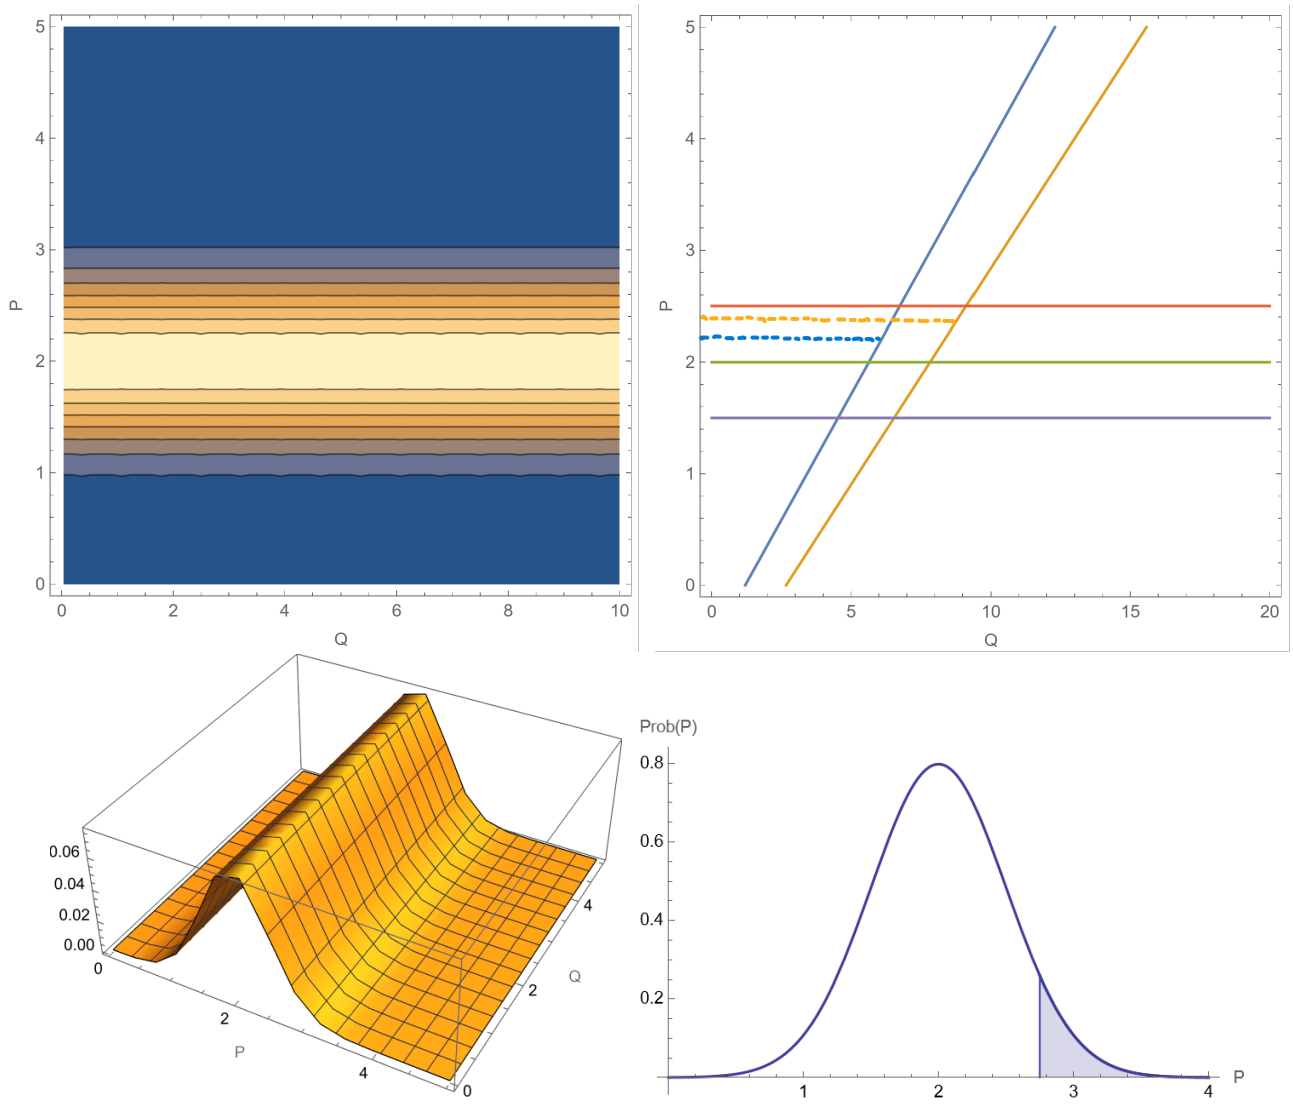


Figure 2: Demand is assumed to be perfectly elastic but stochastic. Panel A: 2-dimensional level curves of the firm's stochastic marginal revenue is distributed $N(\mu, \sigma^2)$. Parameter settings: $\theta = 0.5$, $\mu = 2$, $\sigma = 0.5$, $x = 1$, all cost parameters set to 1. Panel B: Intersection of demand and supply. Average demand μ is plotted horizontally, with $\mu + \sigma$ and $\mu - \sigma$ plotted above and below, respectively. S_1 is the firm's supply curve when $\zeta = 10$, S_2 is the firm's supply curve when $\zeta = 100$. Panel C: A 3-dimensional rendering of Panel A. Panel D: Customers with a willingness to pay below the firm's optimal price are rejected (no shading), while those customers with a willingness to pay above the firm's optimal price (blue shading) are accepted, much like a confidence interval.

marginal revenue distribution function, which is uniform in Q but normal in P. The green line in Panel B is the mean marginal revenue, and the red and purple lines above and below represent one standard deviation above and below that mean. Also plotted are two upward-sloping supply curves, corresponding to the cases where $\zeta_Y = 10$ and $\zeta_Y = 100$. Increasing ζ_Y - indicating a lower cost of customer search - shifts the firm's supply curve to the right while also allowing it to charge a higher price. The optimal sales prices associated with these two curves are marked by dotted lines. Note that price theory diagrams always involve an implicit

time period, but in this model the frequencies associated with the two supply curves depicted in panel B of Figure 2 are different because the cash cycles associated with the profit maximizing sales price are faster if search costs are lower. The blue supply curve to the left is for a longer cash cycle than for the orange supply curve to the right. With a higher cost of search it takes longer for the firm to offload its inventories, leading to lower efficiency turnovers, ceteris paribus. Panel C is a 3-D version of Panel A. Panel D illustrates how the firm's price and implied cash cycle decisions creates a cut-off point, much like a confidence interval on the demand distribution. The firm rejects consumers with a willingness to pay below this cutoff. Similar supply-side diagrams are omitted only to save space.

Optimal search times endogenously determine efficiency turnovers.

Inventory Turnover: If the total amount of time in the year is $T = 1$, inventory holding time $t_Y = \alpha t$ determines the firm's inventory turnover ratio: $\frac{1}{\alpha t}$.

Cash Turnover: If the total amount of time in the year is $T = 1$, cash holding time $t_L = (1 - \alpha)t$ determines the firm's cash turnover ratio: $\frac{1}{(1-\alpha)t}$.

Panel A of Figure 3 visualizes the fixed-point argument for α^*. Panel B of Figure 3 depicts a similar fixed-point solution for t^* as well, for the special case of the pure dealing middleman discussed in the previous paragraph. Panels C, plots efficiency turnovers for different values of α^*.

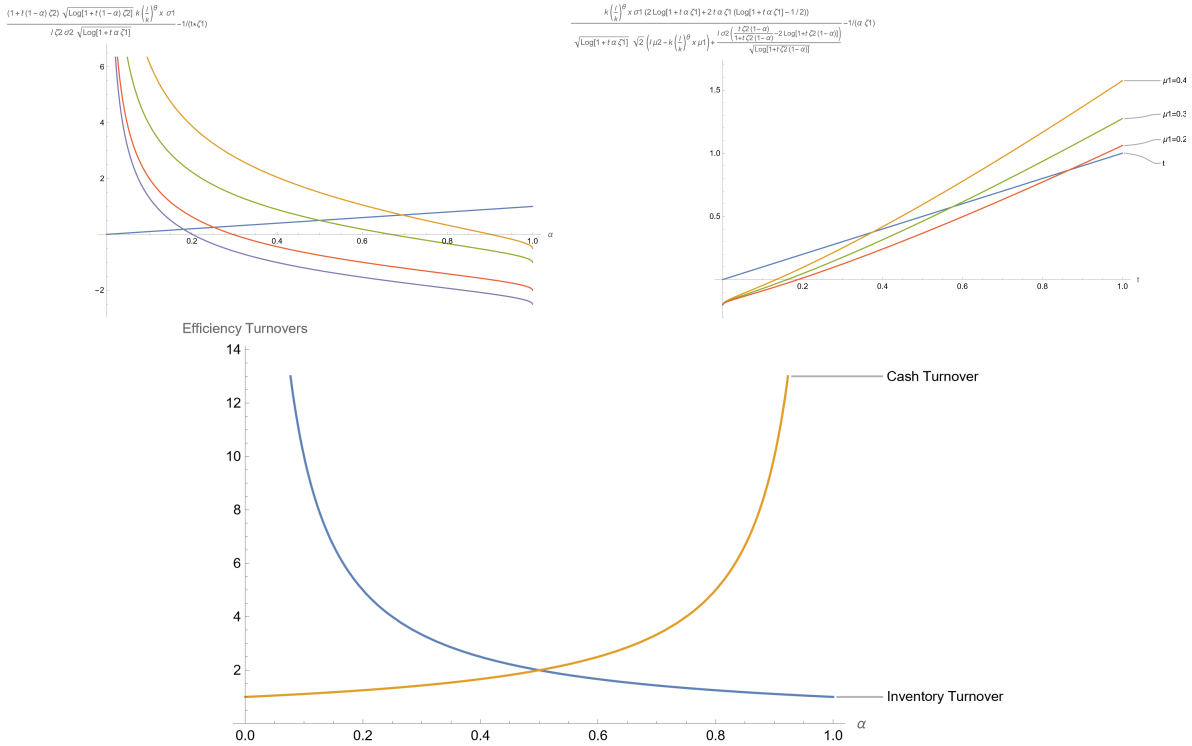


Figure 3: Illustration of fixed point theorems to show that optimal α exists (panel A) and optimal t exist (panel B) under assumptions for pure dealing middlemen. Parameter settings: $\theta = 0.5$, $\mu = 1$, $\sigma = 1$, $x = 1$, all cost parameters set to 1. Panel C shows efficiency turnovers for different α values.

3.2 The Inventory-to-Current Asset Ratio

This section shows that if the firm does not lend or borrow in-kind, then the ratio of inventories divided by current assets is simply α. Current assets are liquid, near-term assets, comprising cash and equivalents, inventories, and accounts receivable, excluding long-lived, illiquid fixed assets such as PPE. Receivables (and payables) arise from within–supply-chain credit, which is undoubtedly important, but modeling trade credit as in-kind intermediation would add substantial complexity; I therefore assume firms neither borrow from nor lend to each other, so receivables and payables are zero. This abstraction is imperfect: Figure 4 shows receivables are a meaningful share of current assets, declining more slowly than inventories as cash rose. Receivables are also plausibly tied to customer search costs, since offering trade credit can attract buyers, speed sales, and dampen sales' sensitivity to income. Extending the model to endogenize receivables—e.g., predicting that lower search costs raise their share—seems natural, but is beyond the scope of this paper.

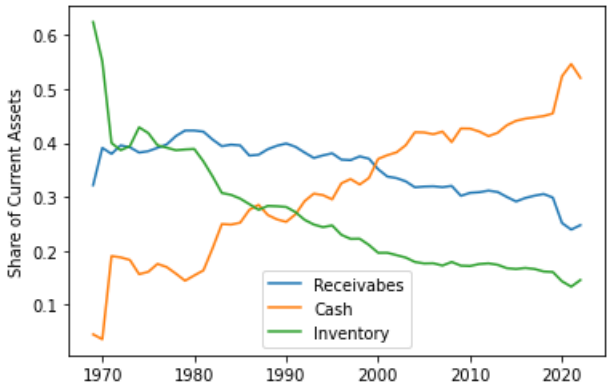


Figure 4: Receivables, Cash, and Inventory as shares of current assets (medians).

If inventory and cash (and cash equivalents) comprise total current assets, the inventory-to-current asset ratio can be calculated using the (observable) efficiency turnovers already derived.

$$\frac{\texttt{INVT}}{\texttt{ACT}} = \frac{t_Y}{t_Y + t_L} = \frac{\alpha t}{\alpha t + (1-\alpha)t} = \alpha$$

So the inventory-to-current asset ratio is equal to the proportion of the cash cycle spent searching for customers.

Figure 5 plots the optimal inventory-to-current asset ratio as a 3-dimensional surface, a function of search cost parameters ζ_Y and ζ_L. In panel A, the orange surface corresponds to a low value of t, the blue surface with a middle value, and green with the highest. In panel B, the orange surface corresponding to the case where $\sigma_Y = \sigma_L$, the blue surface corresponding to the case where $\sigma_Y < \sigma_L$ and the green surface corresponding to the case where $\sigma_Y > \sigma_L$. Shortening t tends to lower α. The ratio is lowest when supply is

riskier than demand, and when searching for demand is costlier than searching for supply. Low alpha firms thus have more difficulty and less to gain from searching for demand than from searching for supply.

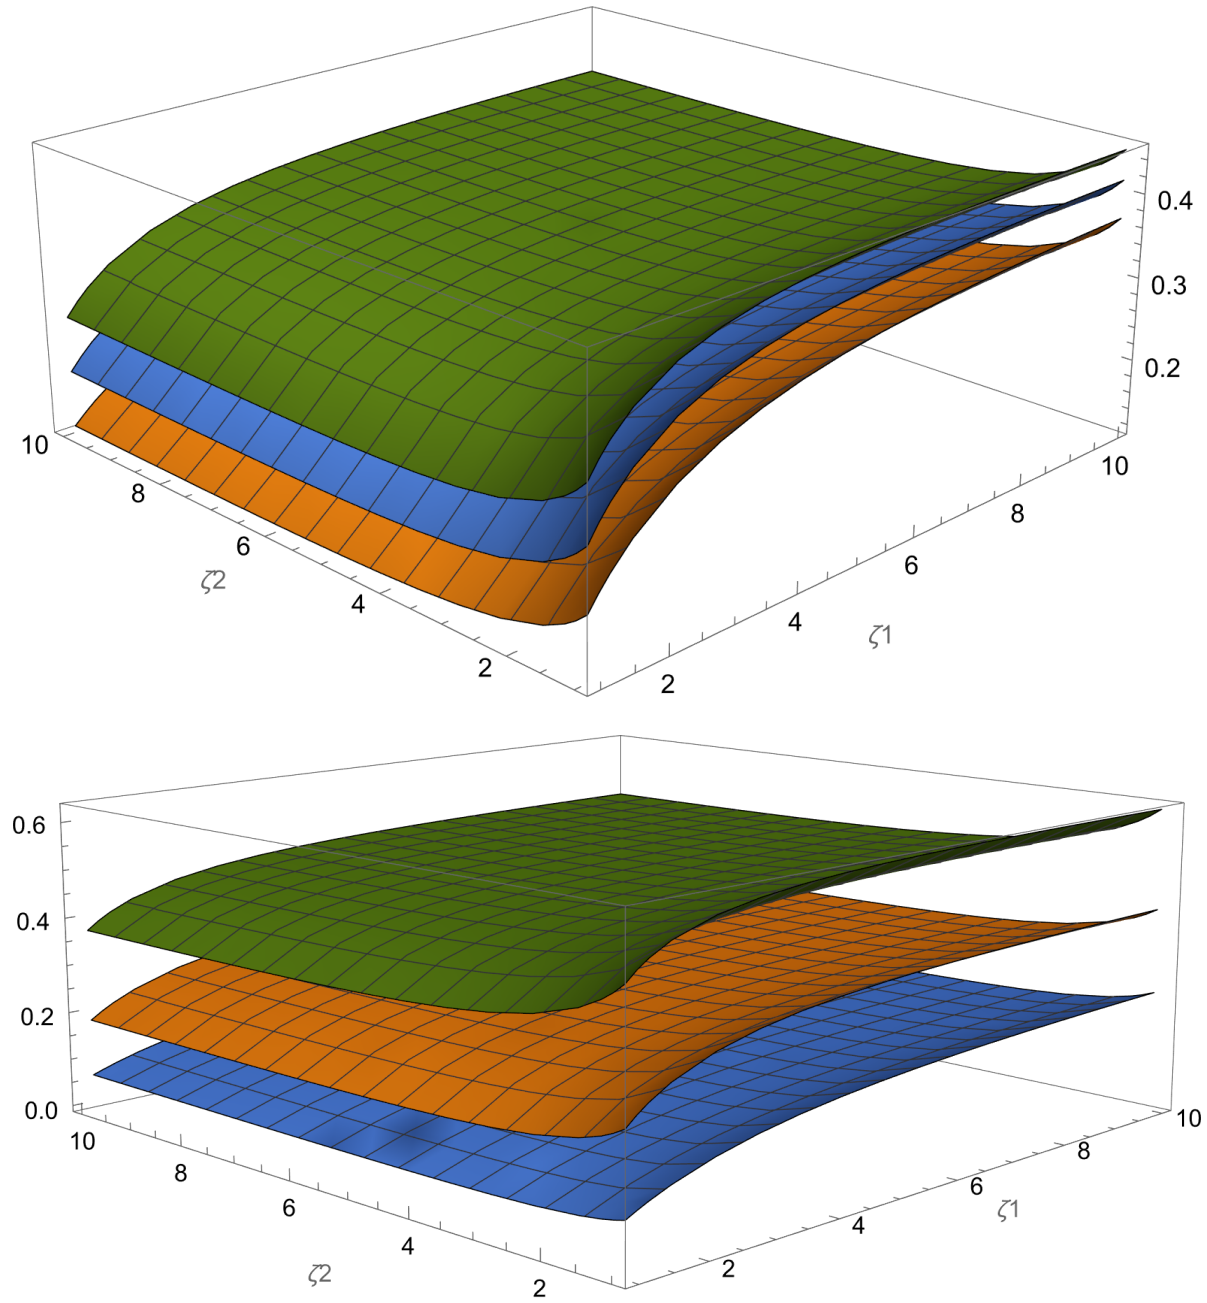


Figure 5: The inventory-to-current asset ratio as a function of ζ_Y and ζ_L. In panel A, the ratio is plotted for different values of t, $t = 0.4$ (orange), $t = 0.5$ (blue), $t = 0.6$ (green). In panel B, the ratio is plotted for different values of σ_Y and σ_L, $\sigma_Y = \sigma_L = 0.25$ (orange), $0.25 = \sigma_Y < \sigma_L = 0.5$ (blue), $0.5 = \sigma_Y > \sigma_L = 0.25$ (green).

It remains to show exactly why this firm characteristic is informative of exposure to systematic risk. The next subsection outlines more precisely the relationship between search costs and profits.

3.3 Sensitivities of Profit to Search Costs

This section calculates theoretical sensitivities of firm profit to changes in search cost parameters. These elasticities are calculated as the log-log derivative of firm profits with respect to the the shock, in this case to the parameters ζ_Y and ζ_L. This ensures that the β's will be unitless measures of sensitivity, the percent change in profits in response to a percent change in the shocked variable. The exact form that these elasticities take depend on the assumptions made about how the firm can respond to shocks.

Generally,

$$\beta_{\zeta_Y} = \frac{\partial \ln \Pi}{\partial \ln \zeta_Y} = \frac{\ln(\zeta_Y)}{-\bar{K}P_K + \frac{-L\mu_L + A\bar{K}^{1-\theta}L^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+e^{\zeta_Y}\alpha t)} + L\sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)})}{t}}$$

$$\beta_{\zeta_L} = \frac{\partial \ln \Pi}{\partial \ln \zeta_L} = \frac{\ln(\zeta_L)}{-\bar{K}P_K + \frac{-L\mu_L + A\bar{K}^{1-\theta}L^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)} + L\sigma_L\sqrt{2\ln(1+e^{\zeta_L}(1-\alpha)t)})}{t}}$$

However, this general form is an upper bound on the β's. If the firm is powerless to adjust any decisions in response to an increase in search costs, it will passively use the same amount of variable inputs at a slower pace, lowering the annualized rate of profit. In that case, either β_ζ will be every bit as large as implied above. However, if the firm can adjust its choice of variable inputs quickly in response to a change in search costs, the effect on profits will necessarily be dampened. Inserting optimal $L^*(\zeta)$ in for L in the beta equations above, sensitivities of profit to search costs will take the forms

$$\beta_{\zeta_Y} = \frac{\ln(\zeta_Y)}{\ln\left(-P_K\bar{K} + \frac{1}{t}\left(A\bar{K}^{1-\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+e^{\zeta_Y}\alpha t)})\left(\left(\frac{A\bar{K}^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+e^{\zeta_Y}\alpha t)})}{\mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)}}\right)^{1/(1-\theta)}\right)^{\theta} + \left(\frac{A\bar{K}^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+e^{\zeta_Y}\alpha t)})}{\mu_L - \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)}}\right)^{1/(1-\theta)}(-\mu_L + \sigma_L\sqrt{2\ln(1+\zeta_L(1-\alpha)t)})\right)\right)} \tag{5}$$

$$\beta_{\zeta_L} = \frac{\ln(\zeta_L)}{\ln\left(-P_K\bar{K} + \frac{1}{t}\left(A\bar{K}^{1-\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)})\left(\left(\frac{A\bar{K}^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)})}{\mu_L - \sigma_L\sqrt{2\ln(1+e^{\zeta_L}(1-\alpha)t)}}\right)^{1/(1-\theta)}\right)^{\theta} + \left(\frac{A\bar{K}^{\theta}(\mu_Y + \sigma_Y\sqrt{2\ln(1+\zeta_Y\alpha t)})}{\mu_L - \sigma_L\sqrt{2\ln(1+e^{\zeta_L}(1-\alpha)t)}}\right)^{1/(1-\theta)}(-\mu_L + \sigma_L\sqrt{2\ln(1+e^{\zeta_L}(1-\alpha)t)})\right)\right)} \tag{6}$$

Equations (5) and (6) implicitly assume, however, that prices cannot be adjusted in response to a change in search costs. One could imagine inserting an expression for optimal prices, reserve costs, cash cycles or inventory-to-current asset ratios as functions of search costs, in order to see how the ability to adjust on these margins would further dampen β_ζ's. However, such an expression cannot be written for the same reasons that analytic expressions for optimal α^* and t^* could be written. One can only deduce qualitatively that if the firm could adjust on other margins in response to a change in search costs, then the β_ζ's would be lower still.

Consider the relationship between α and the β_ζ's. As the following corollaries shows, the inventory-to-current asset ratio is close to a sufficient statistic for *relative* costs and benefits of search. Low α firms, having higher inventory turnovers than cash turnovers, are more sensitive to a higher supply search costs that would hinder the production needed to satisfy demand. High α firms, having higher cash turnovers than inventory turnovers, are most sensitive to increases in demand search cost that would slow sales and result in an over-accumulation of inventories. Whether shocks to demand or supply search costs will have a greater effect on profit depends, in part, on the ex ante level of search costs, because the very concave (convex) form of $P_Y(\zeta_Y)$ ($P_L(\zeta_L)$) means that there are diminishing returns to search. When the benefits of demand-side search dominate the benefits of supply-side search, β_ζ is decreasing in α.

I state the corollary in two ways. First, a sharp "knife edge" case where supply costs are utterly certain and thus there are no gains to searching for suppliers, and second, a more general statement.

Knife-edge Corollary: Fix (L, K, t) with $t > 0$. Assume the firm's input cost does not load on ζ at all (e.g., supplier search is ζ-insensitive), so that the only ζ-dependence of profits comes through the demand-side term. Then, up to ζ and a-independent constants,

$$\Pi(\zeta, a) = Y\sigma_Y\sqrt{2\ln(1 + a\zeta t)} \;>\; 0.$$

Define the profit elasticity with respect to ζ as

$$\varepsilon_{\Pi,\zeta}(\zeta, a) \equiv \frac{\partial \ln \Pi(\zeta, a)}{\partial \ln \zeta}.$$

Then for all $\zeta > 0$ and $a \in (0, 1)$,

$$\frac{\partial^2 \ln \Pi(\zeta, a)}{\partial \ln \zeta \, \partial a} < 0.$$

Proof. Let $x \equiv a\zeta t > 0$. Since $\Pi(\zeta, a) = Y\sigma_Y\sqrt{2\ln(1 + x)}$,

$$\ln \Pi = \text{const} + \frac{1}{2}\ln\big(2\ln(1 + x)\big).$$

Differentiating with respect to $\ln \zeta$ (noting $\partial x/\partial \ln \zeta = x$),

$$\varepsilon_{\Pi,\zeta}(\zeta, a) = \frac{\partial \ln \Pi}{\partial \ln \zeta} = \frac{x}{2(1 + x)\ln(1 + x)}.$$

Differentiating with respect to a (noting $\partial x/\partial a = \zeta t$),

$$\frac{\partial^2 \ln \Pi}{\partial \ln \zeta \, \partial a} = \zeta t \cdot \frac{d}{dx}\left[\frac{x}{2(1+x)\ln(1+x)}\right] = \zeta t \cdot \frac{\ln(1+x) - x}{2(1+x)^2\big(\ln(1+x)\big)^2}.$$

For all $x > 0$, $\ln(1+x) < x$, hence the numerator is negative and the denominator is positive, implying the stated strict inequality. Q.E.D.

more general corollary Fix (L, K, t) with $t > 0$ and let $Y = L^\theta K^{1-\theta}$. Assume the firm's demand-side and supply-side execution prices follow

$$P_Y(t, a; \zeta) = \mu_Y + \sigma_Y \sqrt{2\ln(1 + a\zeta t)}, \qquad \sigma_Y > 0,$$

$$P_L(t, a; \zeta) = \mu_L - \sigma_L \sqrt{2\ln(1 + (1-a)\zeta t)}, \qquad \sigma_L \geq 0,$$

so that, up to ζ- and a-independent constants,

$$\Pi(\zeta, a) = Y\sigma_Y \sqrt{2\ln(1 + a\zeta t)} + L\sigma_L \sqrt{2\ln(1 + (1-a)\zeta t)} \; > \; 0.$$

Let $x \equiv a\zeta t$ and $y \equiv (1-a)\zeta t$. Define the ζ-elasticity

$$\varepsilon_{\Pi,\zeta}(\zeta, a) \equiv \frac{\partial \ln \Pi(\zeta, a)}{\partial \ln \zeta}.$$

If the following dominance condition holds for all relevant (ζ, a),

$$Y\sigma_Y \, \Gamma(x, y) \; > \; L\sigma_L \, \Delta(x, y) \, , \tag{$\mathrm{D}_{\ln}$}$$

where

$$\Gamma(x, y) \equiv \frac{\sqrt{2\ln(1+y)}}{\Pi(\zeta, a)} \cdot \frac{x - \ln(1+x)}{2(1+x)^2\big(\ln(1+x)\big)^{3/2}},$$

$$\Delta(x, y) \equiv \frac{\sqrt{2\ln(1+x)}}{\Pi(\zeta, a)} \cdot \frac{y - \ln(1+y)}{2(1+y)^2\big(\ln(1+y)\big)^{3/2}},$$

then

$$\frac{\partial^2 \ln \Pi(\zeta, a)}{\partial \ln \zeta \, \partial a} < 0.$$

Proof. Write $\Pi(\zeta, a) = A(x) + B(y)$ where

$$A(x) \equiv Y\sigma_Y\sqrt{2\ln(1+x)}, \qquad B(y) \equiv L\sigma_L\sqrt{2\ln(1+y)}.$$

By the chain rule,

$$\frac{\partial \ln \Pi}{\partial \ln \zeta} = \frac{\zeta}{\Pi}\frac{\partial \Pi}{\partial \zeta} = \frac{1}{\Pi}\Big(A'(x)\,x + B'(y)\,y\Big),$$

and differentiating with respect to a using $\partial x/\partial a = \zeta t$ and $\partial y/\partial a = -\zeta t$ gives

$$\frac{\partial^2 \ln \Pi}{\partial \ln \zeta\, \partial a} = \zeta t \cdot \frac{d}{da}\left[\frac{A'(x)\,x + B'(y)\,y}{\Pi}\right]$$

$$= \zeta t \cdot \frac{\Pi\Big(A''(x)\,x + A'(x)\Big) - \Big(A'(x)\,x + B'(y)\,y\Big)A'(x)}{\Pi^2} - \zeta t \cdot \frac{\Pi\Big(B''(y)\,y + B'(y)\Big) - \Big(A'(x)\,x + B'(y)\,y\Big)B'(y)}{\Pi^2}.$$

A direct substitution using

$$A'(x) = Y\sigma_Y\frac{1}{(1+x)\sqrt{2\ln(1+x)}}, \qquad A''(x) = -Y\sigma_Y\left[\frac{1}{(1+x)^2\sqrt{2\ln(1+x)}} + \frac{1}{(1+x)^2(2\ln(1+x))^{3/2}}\right],$$

(and analogously for $B'(y), B''(y)$) yields

$$\frac{\partial^2 \ln \Pi}{\partial \ln \zeta\, \partial a} = \zeta t\left[-Y\sigma_Y\,\Gamma(x,y) + L\sigma_L\,\Delta(x,y)\right].$$

Since $x - \ln(1+x) > 0$ and $y - \ln(1+y) > 0$ for all $x, y > 0$, both Γ and Δ are strictly positive. Condition $(\mathrm{D}_{\ln})$ therefore implies the bracketed term is negative, establishing the claim. Q.E.D.

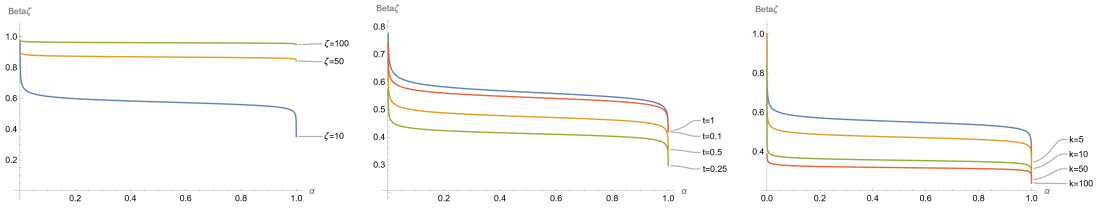


Figure 6: Sensitivity of Firm Profit to Search cost, as a function of inventory-to-current asset ratios. Panel A plots β_ζ for different levels of ζ, Panel B does the same for different levels of t, Panel C does so for different levels of K.

Figure 6 plots the relationship for the special case of the pure dealing middleman. Parameter settings are again specified in the figure caption. Panel A shows this relationship for different starting values of ζ, panel B shows the relationship for different cash cycle lengths, and panel C shows it for different levels of fixed capital. In all cases, the relationship is downward sloping, with these instruments affecting position

and slope of descent.

3.4 The Investor's Utility Maximization Problem

Shift attention from the firm to the investor. The investors choose how much of their wealth to consume today and how much to invest in a risky asset issued by the firm. Importantly, the investor can search for information about future consumption growth in order to time the market and minimize borrowing costs, and can also search for trading partners willing to offer advantageous deals. The imperfection of information in the model means that (within limits) there is good reason to try and time the market, and also good reason to search for value bargains that can be bought low and sold high.

Let consumption growth between times t and $T > t$ be $g(\zeta, b; \tau) = \mu_g - \sigma_g\sqrt{2\ln(1 + b\zeta\tau)}$, where ζ is the now familiar cost of search parameter. Let the price of the firm's equity at time of purchase be $P_t(\tau, b; \zeta) = \mu_P - \sigma_P\sqrt{2\ln(1 + (1 - b)\zeta\tau)}$.

Let preferences be time separable, represented by a utility function exhibiting constant relative risk aversion. The risk aversion parameter is γ. The investor chooses how much to consume (c_t) and invest (x_t) as if to maximize

$$U(c_t, c_T) = \delta^t \frac{c_t^{1-\gamma}}{1-\gamma} + E\left[\delta^T \frac{c_T^{1-\gamma}}{1-\gamma}\right]$$

Subject to a constraint, which contains wealth endowment w,

$$c_t = w - x_t P_t(\tau, b; \zeta)$$

$$c_T = x_t \Pi_T(t, \alpha; \zeta)$$

$\Pi_T(t, \alpha; \zeta)$ being the firm's profit, distributed entirely to shareholders as a dividend at time T. Assume the firm is financed entirely by equity. Assume subjective time preference factor $\delta < 1$. Making clear that profit and consumption growth are both functions of search cost, the usual first order conditions imply

$$P_t(\tau, b; \zeta) = E\left[\delta^{T-t} g(\tau, b; \zeta)^{-\gamma} \Pi_T(t, a; \zeta)\right]$$

Where the stochastic discount factor (SDF) for cash flows arriving at time T, $M_T(\tau, b; \zeta) = \delta^{T-t} g(\tau, b; \zeta)^{-\gamma}$ and t and b solving the equation to attain equality.[7] If the investor is risk averse, then the SDF is convex and decreasing in consumption growth. Since expected consumption growth is a concave, slowly increasing function of ζ, the SDF is lower (and thus, discount rates higher) in states of the world in which search cost is low

[7] For this to work, a condition must hold: $P_t(0, b; \zeta) > M(0, b; \zeta)$ for all $b \in [0, 1]$, so that there exists a $\tau > 0$ that makes the pricing equation hold.

(ζ is high). $M(\zeta)$ is thus decreasing in ζ. By Jensen's inequality, convex M implies that $E[M(\zeta)] \leq M(E[\zeta])$ if ζ is taken to be a random variable. Because investors have a higher marginal utility for consumption in states of the world where search cost is high, investors demand a risk premium for exposure to search cost risk. They value payoffs in bad (high search cost) states of the world more than in good (low search cost) states, making them willing to pay more for assets that provide returns in those bad states. The price of an asset depends not only on the average SDF but also on the covariance of asset returns with the SDF. Assets whose returns are especially sensitive to search cost risk ought to earn higher risk premia.

Abbreviating somewhat, go from price space to expected return space

$$E[R_T(\zeta)] = \frac{E[\Pi_T(\zeta)]}{P_t(\zeta)}$$

.

Since profits are increasing in ζ and asset price P_t is decreasing in ζ, returns are unambiguously increasing in ζ. We have already established that M is decreasing in ζ. Search costs thus induce a negative covariance between profits and the SDF. Investors discount (demand higher expected returns of) firms whose profits are more sensitive to search cost. Writing the pricing equation in terms of returns.

$$1 = E[M_T(\zeta)R_T(\zeta)] = E[M_T(\zeta)]E[R_T(\zeta)] + COV(M_T(\zeta), R_T(\zeta))$$

Divide both sides by $E[M_T(\zeta)]$

$$\frac{1}{E[M_T(\zeta)]} = E[R_T(\zeta)] + \frac{1}{E[M_T(\zeta)]}COV(M_T(\zeta), R_T(\zeta))$$

Since the gross risk-free rate is $R^f(\zeta) = \frac{1}{E[M_T(\zeta)]}$, rearrange

$$E[R_T(\zeta)] = R^f(\zeta) - R^f(\zeta)COV(M_T(\zeta), R_T(\zeta))$$

Risk-free rates are increasing in ζ, because $M'(\zeta) < 0$. Investors are less patient when they can easily search for investment opportunities. Dropping time subscripts, rewrite expected excess returns in the single beta representation.

$$E[R(\zeta)] - R^f(\zeta) = -\beta_{R,M}(\zeta)\lambda_M(\zeta)$$

where $\lambda_M(\zeta) = -\frac{\sigma^2(M(\zeta))}{E[M(\zeta)]}$ and $\beta_{R,M}(\zeta) = \frac{COV(M(\zeta),R(\zeta))}{\sigma^2(M(\zeta))}$.

And finally, note that because $\frac{\partial \ln(y)}{\partial \ln(x)} = \frac{x}{y}\frac{\partial y}{\partial x}$ and other properties of logarithms, we have

$$\beta_{M,\zeta} \equiv \frac{d \ln M}{d \ln \zeta}, \qquad \beta_{R,m} \equiv \frac{d \ln R}{d \ln M}.$$

$$\ln R = \ln \Pi - \ln P \;\; \Rightarrow \;\; \beta_{R,\zeta} = \beta_{\Pi,\zeta} - \beta_{P,\zeta}.$$

$$\beta_{R,M} = \frac{d \ln R}{d \ln M} = \frac{d \ln R / d \ln \zeta}{d \ln m / d \ln \zeta} = \frac{\beta_{\Pi,\zeta} - \beta_{P,\zeta}}{\beta_{m,\zeta}}.$$

And the corollaries above imply the main hypothesis: $\beta_{R,M}$ is decreasing in α; firms with low α have more systematic risk.

Effect of search costs on the price of risk: How does ζ affect the overall price of risk λ_M? Higher ζ increases $E[M]$ in the denominator of λ_M, leading to a lower λ_M overall. Higher ζ also decreases $\sigma^2(M)$ in the numerator. This is because investors are pricing assets as if to reject transactions in circumstances when consumption growth is too low to justify the price (recall panel D of Figure 3, which depicts such a truncation). Selection through search effectively removes low consumption growth observations from the distribution, mechanically lowering the observed variance of M. On the other hand, if search costs rise, both the mean and variance of the stochastic discount factor increase.

Figure 7 plots the stochastic discount factor (panels A and B) and asset prices (panels C and D) over values of search cost parameter ζ. The SDF falls faster when risk aversion is lower, and when σ_g is higher. Parameter settings were chosen to show that firms facing high enough search costs can be nonviable (with negative asset prices).

The next section provides some empirical evidence regarding the relationships between the inventory ratios (INVT/ACT and INVT/ME) and expected returns.

4 Empirics

This section details empirical methodology and results, which comprise a reduced form test of the model presented in chapter 3.

Hypotheses:

1. Firms with low inventory-to-current asset ratios ($\frac{INVT}{ACT}$) have higher returns than firms with high $\frac{INVT}{ACT}$.

2. A firm's inventory-to-market equity ratios ($\frac{INVT}{ME}$) have no statistical relationship with returns.

The first hypothesis is clearly predicted in chapter 3. Hypothesis 2, however, is somewhat more am-

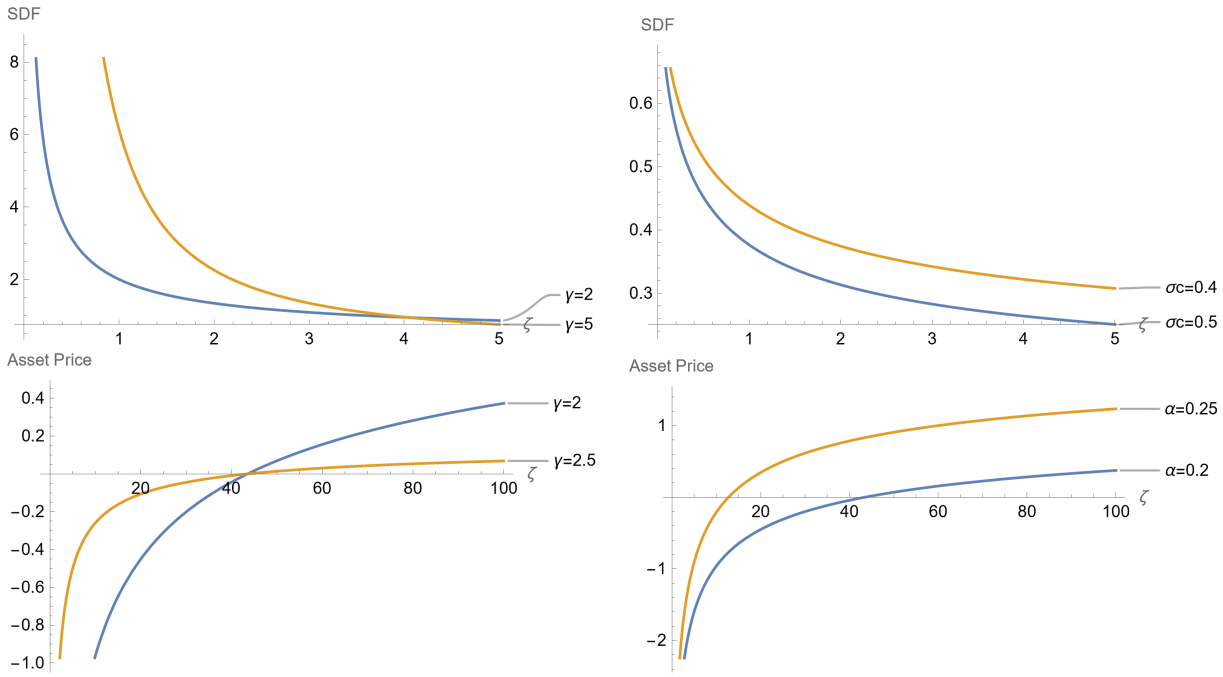


Figure 7: Stochastic discount factor (SDF) and asset price as functions of search cost.

biguous because of the potential for countervailing relationships between search costs and the numerator and denominator of the ratio. Nonetheless, I state hypothesis 2 as it is because this allows an empirical exploration of the relationship to begin from a clear starting point.

Section 4.1 describes data sources. Section 4.2 details some accounting identities relevant for analyzing the inventory ratios. Section 4.3 provides descriptive statistics of inventory-related variables. Section 4.4 reports moments of firm characteristics and portfolio statistics for decile portfolios sorted on the inventory ratios. Section 4.5 presents time-series regression results involving the long-short portfolios ILMH and IVMG. Section 4.6 presents Fama-MacBeth panel regression results.

4.1 Data Sources

Monthly stock return data are from the Center for Research in Security Prices (CRSP). Annual and quarterly accounting data are from COMPUSTAT. Factor returns for the Fama-French factors are taken from Kenneth French's website, as are the 49 industry portfolio returns. Factor returns for the q-factor investment CAPM model of Hou, Xue and Zhang (2015) are taken from their data website (global-q.org). Factor returns associated with 153 other firm characteristics were collected from the website accompanying a forthcoming *Journal of Finance* article by Jensen, Kelly, and Pedersen (jkpfactors.com). The Global Supply Chain Pressure Index, which relies mostly on shipping wait times from major ports, is taken from the New York Fed. Other data (some of which was already presented in Figure 1) are taken from the Census Bureau. The

Census Bureau's inventory data comes from its Monthly Retail Trade Survey, Monthly Wholesale Trade Survey, and the Manufacturers' Shipments, Inventories, and Orders Survey (also called the M3 survey). The M3 is based on a sample of large firms in 83 industry segments as defined by the Census Bureau. These industry segments do not necessarily correspond to Fama and French's 49 industry categories. The Census covers a larger number of firms, not all of which are publicly traded, and are generally more reliable for disaggregation of inventories by raw materials, works-in-process, and finished goods. In COMPUSTAT, that breakdown is less reliable, because many firms do not report the type of inventory they posses (although most do provide a written description in the footnotes of their 10-K forms). Valuations of WIP inventories also vary. Retail businesses covered in the Monthly Retail Trade Survey have more than one establishment, while the Monthly Wholesale Trade Survey covers merchants who take ownership of the goods from manufacturers before selling to retailers. So the wholesaler data does not include manufacturers or retailers who do their own distribution, or wholesalers for whom this trade is not their primary business. All other indicators related to inventories, labor unemployment, nondurable consumption spending, term and credit spreads are taken from the BLS and Census Bureau via the St. Louis Federal Reserve's FRED database.

4.2 Key Accounting Ratios

Financial accounting ratios are constructed from sales (COGS), market equity (ME) and components of current assets (ACT) and current liabilities (LCT). $\texttt{ACT} = \texttt{CHE} + \texttt{INVT} + \texttt{RECVT}$, where CHE is cash and cash equivalents, INVT is inventory and AR is accounts receivable. Current liabilities include accounts payable (AP) and other liabilities that must be paid within a year. The inventory-to-current asset ratio is

$$\frac{\texttt{INVT}}{\texttt{ACT}}$$

.

ME is calculated by multiplying stock prices (PRCC) by the number of shares outstanding (CSHO) with adjustments for splits and reverse splits. The inventory-to-market ratio is

$$\frac{\texttt{INVT}}{\texttt{ME}}$$

. Naturally, firms with low $\frac{\texttt{INVT}}{\texttt{ACT}}$ should have low $\frac{\texttt{INVT}}{\texttt{ME}}$. Some motivation for these ratios is warranted. Notice that the price-to-sales ratio, a commonly cited market multiple, can be expanded

$$\frac{\texttt{ME}}{\texttt{COGS}} = \frac{\texttt{ME}}{\texttt{INVT}} \Big/ \frac{\texttt{COGS}}{\texttt{INVT}}$$

Such an expansion could be performed with any asset on the firm's balance sheet. The point is to put market value in relation to some asset via that asset's turnover ratio. The same could be done for total current assets, and furthermore, could be done for current liabilities. Current liabilities turnover is calculated as follows

$$\frac{\texttt{COGS}}{\texttt{LCT}} = \frac{\texttt{COGS}}{\texttt{INVT}} \times \frac{\texttt{INVT}}{\texttt{ACT}} \times \frac{\texttt{ACT}}{\texttt{LTC}}$$

This last fraction on the end of the expression above is the current ratio, a commonly used measure of firm liquidity. $\frac{\texttt{INVT}}{\texttt{ACT}}$ can also be decomposed to make its relationship with other liquidity ratios clear.

$$\frac{\texttt{INVT}}{\texttt{ACT}} = \frac{\texttt{CurrentRatio} - \texttt{QuickRatio}}{\texttt{CurrentRatio}}$$

Where the quick ratio is $\frac{\texttt{ACT}-\texttt{INVT}}{\texttt{LCT}}$ and the current ratio is $\frac{\texttt{ACT}}{\texttt{LCT}}$, a more restrictive measure of liquidity.

The following decomposition of return on equity expands the standard DuPont analysis in order to show that, all else constant, higher $\frac{\texttt{INVT}}{\texttt{ACT}}$ is associated with higher ROE.

$$\texttt{ROE} = \texttt{ProfitMargin} \times \frac{\texttt{COGS}}{\texttt{INVT}} \times \frac{\texttt{INVT}}{\texttt{ACT}} \times \frac{\texttt{ACT}}{\texttt{AT}} \times \texttt{EquityMultiplier}$$

The model presented above, as well as empirical results below, suggest that this positive relationship between ROE and $\frac{\texttt{INVT}}{\texttt{ACT}}$ is more than just an accounting tautology.

Another useful variable is the cash conversion cycle (CCC), which measures the length of time (in days) it takes a firm to convert inventory investment into cash.

$$\texttt{CCC} = \texttt{DIO} + \texttt{DSO} - \texttt{DPO}$$

where DIO is days inventory outstanding (equal to inventory-sales ratio times 365 days), DSO is days sales outstanding and DPO is days payables outstanding. A low CCC means that the firm is quick to collect cash from sales of its inventory. Rewriting $CCC = DIO + d$, where d is the difference between days receivables outstanding and days payables outstanding, notice that

$$\frac{1}{365}\texttt{CCC} = (\frac{\texttt{INVT}}{\texttt{ME}} / \frac{\texttt{ME}}{\texttt{COGS}} + d\frac{1}{365}) = ((\frac{\texttt{INVT}}{\texttt{ACT}} \times \frac{\texttt{ACT}}{\texttt{LCT}}) / \frac{\texttt{COGS}}{\texttt{LCT}} + d\frac{1}{365})$$

.

4.3 Descriptive statistics

Figure 8 plots some time-series data on inventory in the United States. Since 1975, inventories grew at a 7.4% average annual rate. Inventories are very seasonal and volatile in the short term, with a quarterly standard deviation (44%), almost three times the annual figure (15%). In 2019 (a typical year) the median firm's inventory holdings swung from 20 to 45 percent of current assets. About 24 percent of publicly traded firms record no inventory at all, a proportion that has grown in tandem with the rise of the service sector of the US economy. Panel A plots aggregate inventory-sales ratios since the 1950s. Innovation in inventory management has gradually pulled this ratio downward.[8] Does it matter which node in the supply chain network holds the inventory? Panel B disaggregates inventory-sales ratios by manufacturers, wholesalers, and retailers, showing a sudden upstream shift during the Covid-19 pandemic. Between 1992-2020, retailers had the slowest turnovers, but after 2020 manufacturers absorbed the bulk of unsold goods as the pace of sales ground to a near-halt. These data still seem to be in a protracted process of convergence back to their historically typical order. Panel C disaggregates inventory-sales ratios by quantile to show increasing skewness associated with specialization; the push toward just-in-time coincided with fewer firms holding a larger share of total inventories. Panel D plots a similar quantile disaggregation of inventory-to-current assets ratios, and Panel E plots another for the inventory-to-market equity ratio. Panel F plots cumulative changes in ratios of nondurable-to-durable inventories for finished goods and works-in-process between 1996-2022. The nondurables-to-durables ratio is meant to capture the balance of production between output meant for near-term consumption and output meant to enhance future consumption, e.g., the stock of food in refrigerators as opposed to the refrigerators themselves.[9] The level and composition of inventories are clearly quite variable along a number of dimensions.

Table 1 reports summary statistics for the inventory ratios INVT/ACT and INVT/ME, CCC, inventory levels, quarterly turnovers, and percentage growth, as well as cash and cash equivalents levels, quarterly turnovers, and percentage growth. Since there is a trend in some of these variables, Table 1 reports cross-sectional moments for a recent, typical year, 2019. Industry adjustments are made by de-meaning using the Fama and French 49 industry categories. The statistics in Table 1 are conditional, covering only those firms who record positive inventories on their balance sheets in some quarter of the year. A natural logarithm transformation is included for INVT/ME in Table 1, but no transformation is required for INVT/ACT because the latter variable has few outliers and only takes on values between 0 and 1.

[8] Lee, Zhou and Hsu (2015) document a positive correlation between firms' innovation performance and inventory turnover.

[9] The model presented in this paper is most relevant for finished goods inventories, but could apply to durable or nondurable goods just as well. Yogo (2006) notes that expenditures on durables are more volatile and procyclical than non-durables. Jones and Tuzel (2013) find the inventory growth spread documented by Belo and Lin (2012) is stronger for durable inventories than for nondurables.

Empirical distributions of these variables have changed over time. To illustrate these changes, Figure 9 plots cross-sectional distribution moments of INVT/ACT (panels A-D), INVT/ME (panels E-H), growth in inventory (panels I-L), growth in cash (panels M-P), and cash conversion cycles (panels Q-T) within each year from 1972-2021. Where indicated, data are demeaned by industry using the Fama and French 49 industry categories. Mean INVT/ACT ratios have fallen steadily, but the cross-sectional distribution has become more positively skewed over that time. The cross-sectional standard deviation (across firms) has oscillated in a slow, persistent manner, peaking in about 1990 and rising again since about 2020. Mean logarithms of the INVT/ME ratio have also steadily declined. Standard deviations appear to rise and fall randomly, if somewhat slowly, exhibiting many peaks and valleys during the sample. Skewness also appears to vary randomly, though a slight (statistically insignificant) upward trend is visible. Mean inventory and cash growth exhibit no real trend. No trend in standard deviations of these variables emerge either. Skewness in inventory growth is large and positive, varying year to year but without a discernible trend. Skewness in cash growth is usually large and positive, but prone to rare, sharp declines into negative territory in some years. Although mean cash growth rose in 2020, panel N shows that firms in the lowest quantiles of cash growth registered the deepest decline on record. Like INVT/ACT and INVT/ME, CCC declines steadily over the period, though it remained more elevated during the 1980s. Standard deviations are high, but vary substantially year to year. Skewness in CCC swings from positive to negative quite frequently. These plots indicate that the cross-section of variables related to working capital management vary in distribution over time, punctuated by what look to be sudden reshufflings of liquid assets between firms. Generally, firms have made greater use of fewer short-term assets, becoming more liquid, and quickening their cash cycles while trending toward greater centralization of inventory holdings into the hands of fewer firms.

4.4 Decile Portfolios

To study how returns vary with inventory ratios, I sort firms into deciles each period by INVT/ACT and INVT/ME, using values lagged three months. Each decile contains roughly 50,000 firm-years, and membership changes over time as firms move across rankings. Table 2 reports, by INVT/ACT decile, average annualized value- and equal-weighted returns, CAPM alphas and betas, and medians for book-to-market, gross profitability (gross profits/assets), size (total assets, sales, market equity), fixed costs (SGA/COGS), CCC, liquidity (quick ratio), current assets/total assets, efficiency ratios (inventory, cash, and receivables turnover), trade credit (receivables/payables), leverage (debt-equity), and liquidity betas using the Pastor and Stambaugh (2003) factor; Table 3 reports the same for INVT/ME deciles.

Returns are not perfectly monotone, but generally decline with both ratios (Figure 4.3). The high-minus-

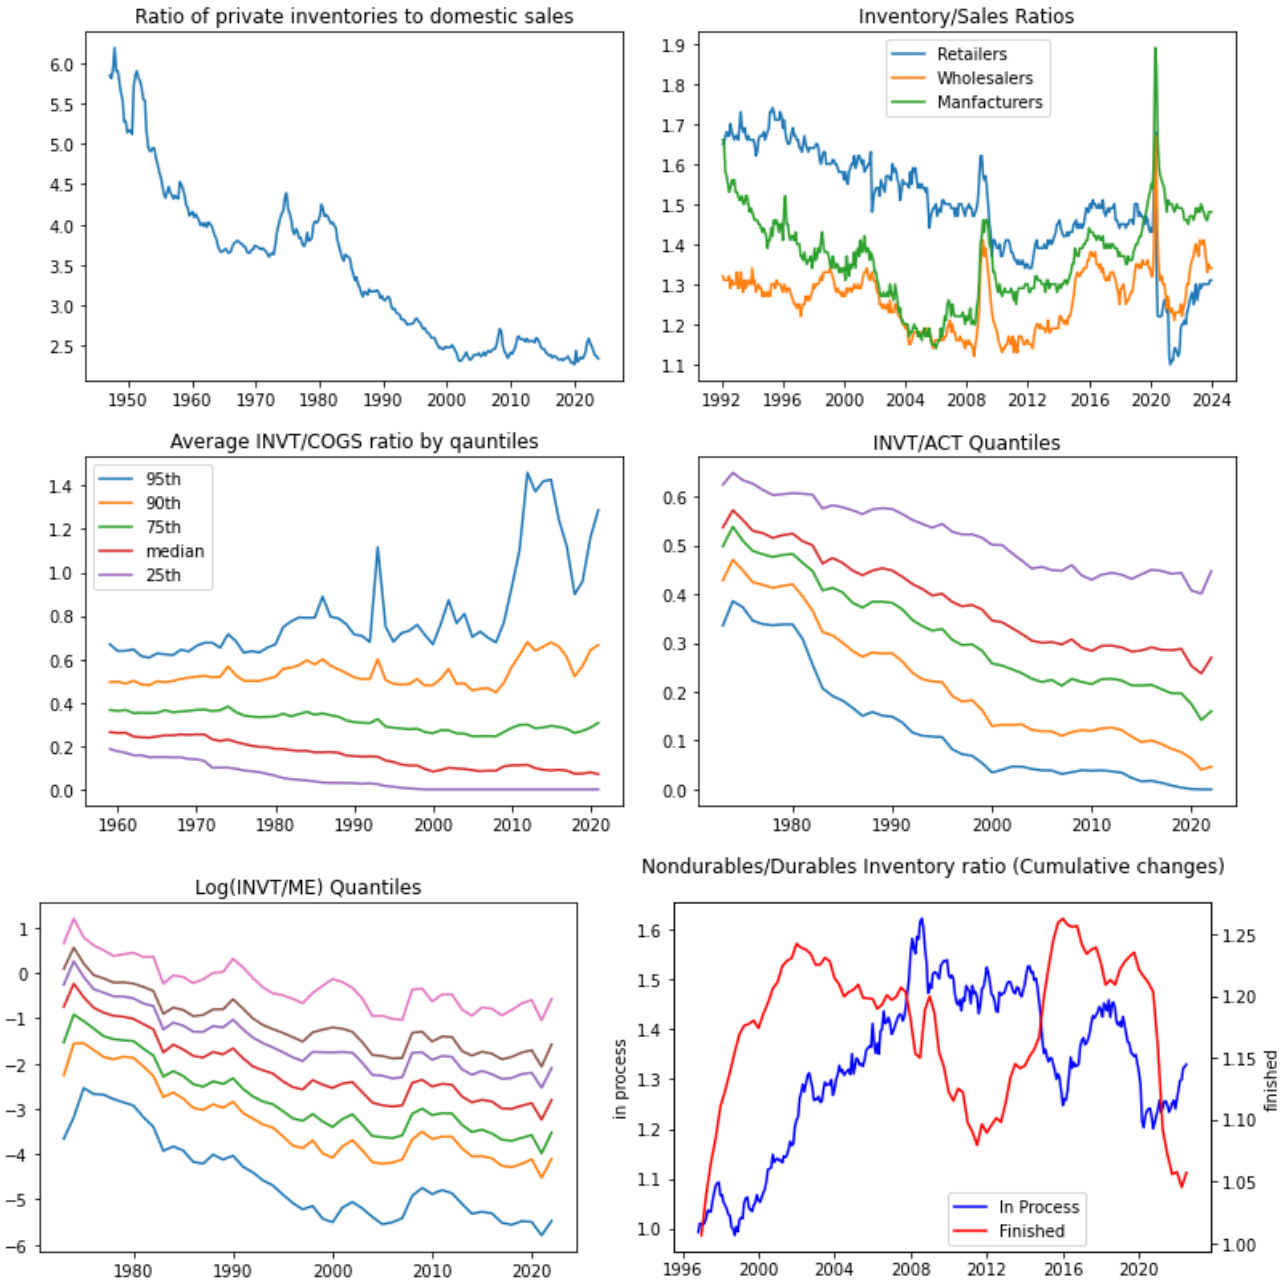


Figure 8: Inventory ratios data. Panel A: Aggregate inventory-Sales ratio 1945-2022. Panel B: Inventory-sales ratios by supply chain category, 1992-2024. Panel C: Inventory-Sales ratios by quantile, 1958-2022. Panel D: Inventory-to-Current Assets ratio by quantile, 1958-2022. Panel E: Logarithm of Inventory-to-market equity ratio by quantiles, 1972-2022. Panel F: Nondurables-to-durabales inventory ratios for finished goods (yellow) and works in process (green) alongside inventory factors ILMH (red) and IVMG (blue) portfolios returns, 1997-2022.

Table 1: Summary Statistics of Inventory Variables. To avoid problems related to nonstationarity, columns 1-6, which report cross-sectional summary statistics for INVT/ACT, INVT/ME, COGS/INVT, CCC, and INVT are all conditional, with the sample restricted to data within in the year 2019. Columns 7-9 report time-series means, standard deviations, and skewness for percent changes in total business inventories, nondurable goods inventories, and durable goods inventories for the entire sample.

	Mean	St.Dev	Skew	Min	25th pct	Median	75th pct	Max
INVT/ACT	0.242	0.190	0.787	0.000	0.078	0.213	0.362	0.948
INVT/ACT(IndAadj)	-0.000	0.149	1.197	-0.478	-0.078	-0.048	0.059	0.826
INVT/ME	31.664	1.549	31.664	0.000	0.015	0.056	0.182	90.759
INVT/ME(IndAadj)	40.236	1.221	40.236	-1.340	-0.154	-0.078	-0.024	90.452
Log(INVT/ME)(IndAadj)	-0.122	1.065	-0.122	-4.060	-0.612	0.031	0.705	3.424
CCC	615.712	1211.840	6.083	0.020	152.738	315.348	599.878	14152.034
CCC(IndAdj)	-8.963	2379.895	-24.433	-142951.917	-206.760	-42.360	130.158	51939.070
COGS/INVT	17.104	237.311	43.272	0.002	0.696	1.463	4.769	11523.500
INVT	313.100	670.581	3.363	0.004	6.679	44.384	257.531	4184.000
INVT growth	0.116	0.423	1.801	-1.636	-0.081	0.007	0.193	2.816
COGS/CHE	64.498	1787.023	40.032	0.000	0.158	0.639	2.678	82428.388
CHE	480.778	1177.911	4.548	0.012	23.578	98.630	331.822	9620.000
CHE growth	0.576	1.178	1.965	-2.445	-0.055	0.199	0.871	8.399

low spread is about 4.6% for INVT/ACT and about 4.0% for INVT/ME. Panels B–E of Figure 4.3 plot value- and equal-weighted returns, 6-factor alphas, and Sharpe ratios. Figure 4.4 plots characteristics by decile: low INVT/ACT and low INVT/ME firms are liquid, growth firms with high fixed costs, short CCCs, high inventory investment growth but low cash growth, high inventory turnover but low cash turnover. Size is hump-shaped in INVT/ACT (extremes small; large firms in the middle) and similarly for INVT/ME, though low INVT/ME firms mechanically have higher market equity than high INVT/ME firms. Low-ratio firms have low leverage (by debt-to-equity, whether book or market), and trade credit activity is hump-shaped in both ratios, with low-ratio firms lending in-kind more than they borrow (panel O). Profitability is low for low-ratio firms (with a more hump-shaped pattern for INVT/ME), and low-ratio firms also exhibit low market power, measured by the median firm's revenue share within its SIC industry (panel P), with similar results under alternative market-power measures

Advertising may help explain why low inventory-ratio firms have high fixed costs: marketing outlays (in SGA) expand a firm's reach and thus its customer-search intensity. Such firms operate with high fixed relative to variable costs and target rapid inventory turnover, potentially by extending trade credit (lower receivables turnover) to move goods quickly and attract new customers. Extensions of this research could incorporate this channel by making foot traffic (the sampling rate ζ) a choice variable purchased through advertising, however, lower advertising costs or higher advertising productivity could raise ζ even without an increasing to the marketing budget.

Low inventory-ratio firms also tend to be small relative to their industry, whereas high-ratio firms are

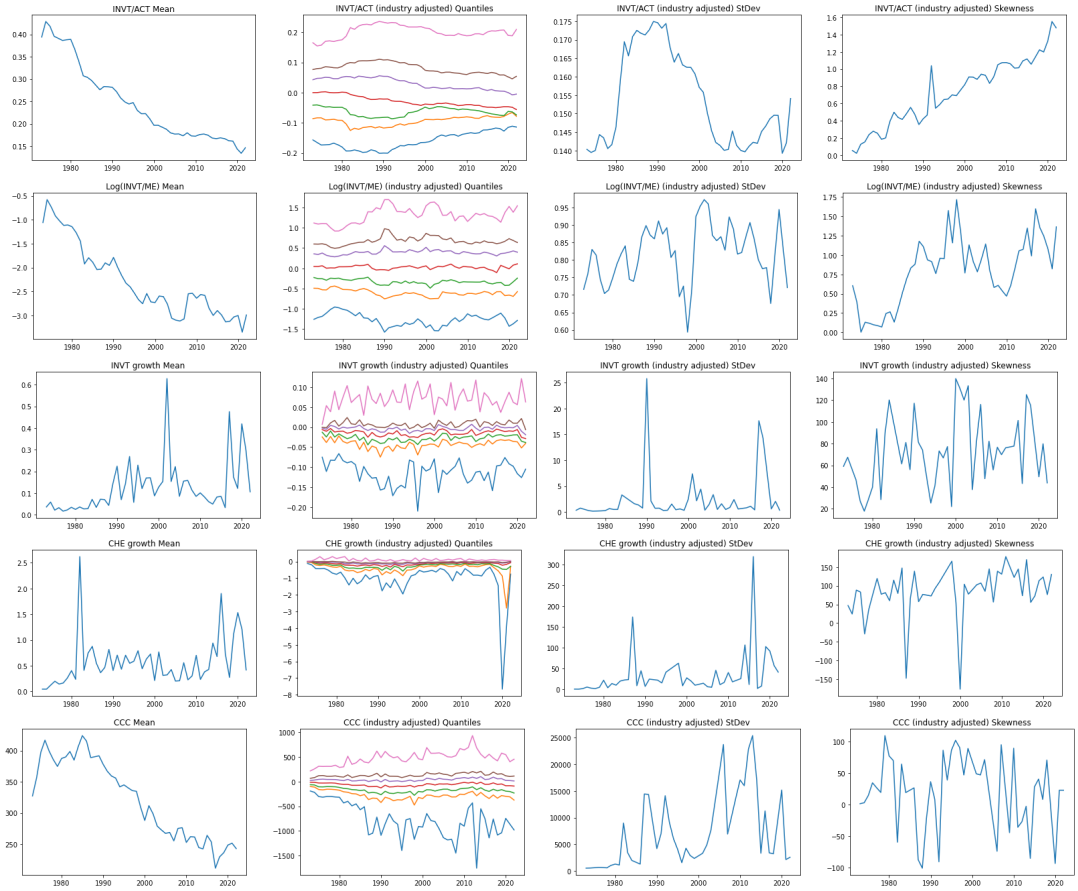


Figure 9: Cross-sectional sample moments of key variables over time. Sample period is 1972-2022. First row (panels A-D) plot means, industry adjusted quantiles (10th, 25th, 35th, 50th, 65th, 75th, and 90th percentiles), industry adjusted standard deviations, and industry adjusted skewness of $\frac{INVT}{ACT}$. Second row (panels E-H) do the same for $\frac{INVT}{ME}$. Third row (panels I-L) do the same for inventory investment growth in percent terms. Fourth row (panels M-P) does so for growth in cash and cash equivalents in percent terms. Fifth row (panels Q-T) does so for cash conversion cycles (CCC). Industry adjustments are made using the Fama and French 49 industry categories. Data are winsorized at the first and ninety-ninth percentile cutoffs, and CCC is truncated below a value of 6000 to avoid spurious outliers.

more likely to exhibit market power, reflected in profitability and revenue shares within SIC categories. Firms with market power optimally carry higher inventories because doing so exerts less competitive pressure on prices than it would in more competitive settings. These patterns suggest that cross-sectional tests of the inventory-ratio premia should control for market power, size, profitability, fixed costs, liquidity, and related firm characteristics.

4.5 Time-series results

I form long–short portfolios by subtracting the highest-decile return from the lowest-decile return for each sorting variable. The INVT/ACT spread portfolio is ILMH (inventory lean minus heavy) and the INVT/ME

Table 2: INVT/ACT Decile Portfolio Statistics and median Firm Characteristics.

	Low	2	3	4	5	6	7	8	9	High	Low-High
E[r](VW)	0.122	0.115	0.106	0.116	0.109	0.101	0.091	0.083	0.087	0.072	0.046(1.97)
E[r](EW)	0.160	0.137	0.150	0.158	0.152	0.162	0.141	0.137	0.118	0.094	
FF6 Alpha	0.069	0.034	0.043	0.056	0.038	0.041	0.014	0.009	-0.013	-0.038	
FF6 Mktrf Beta	1.002	0.990	1.000	0.980	0.980	1.008	1.004	1.006	1.000	0.956	
BE/ME	0.459	0.508	0.525	0.540	0.556	0.562	0.583	0.618	0.653	0.711	
Profitability	0.049	0.057	0.065	0.073	0.079	0.084	0.086	0.085	0.084	0.085	
AT	124.987	215.164	225.174	259.076	294.452	310.653	307.458	261.614	203.379	146.112	
COGS	11.208	20.769	24.789	34.707	45.764	53.522	57.972	55.525	49.949	39.452	
ME	146.615	207.446	234.697	276.414	290.271	280.024	243.126	190.710	135.753	81.887	
SGA/COGS	0.613	0.505	0.491	0.444	0.392	0.361	0.327	0.294	0.267	0.275	
CCC	92.153	133.335	202.021	270.410	317.044	355.852	384.030	418.228	455.988	517.125	
Quick Ratio	2.015	1.612	1.638	1.586	1.473	1.364	1.252	1.175	1.059	0.729	
ACT/AT	0.468	0.390	0.453	0.493	0.506	0.510	0.514	0.530	0.560	0.600	
INVT/COGS	11.522	4.398	2.414	1.727	1.413	1.229	1.113	0.999	0.880	0.686	
CHE/COGS	0.436	0.832	0.936	1.095	1.451	2.026	2.997	4.434	7.225	12.241	
RECT/COGS	0.976	0.984	0.976	0.988	1.010	1.029	1.063	1.119	1.236	1.694	
RECT/AP	2.029	1.916	2.021	2.071	2.078	2.041	1.988	1.944	1.801	1.297	
Debt/Equity	0.222	0.380	0.331	0.334	0.361	0.403	0.450	0.513	0.547	0.588	
FF6 Liq Beta	0.046	0.033	0.049	0.048	0.050	0.056	0.030	0.030	0.059	0.059	

Table 3: INVT/ME Decile Portfolio Statistics and median Firm Characteristics.

	Low	2	3	4	5	6	7	8	9	High	Low-High
E[r](VW)	0.122	0.115	0.124	0.106	0.117	0.116	0.100	0.106	0.099	0.092	0.039(1.22)
E[r](EW)	0.162	0.136	0.135	0.143	0.149	0.137	0.152	0.140	0.131	0.149	
FF6 Alpha	0.022	0.003	0.005	0.014	0.015	0.003	0.018	0.005	-0.004	0.011	
FF6 Mktrf Beta	1.204	1.128	1.080	1.081	1.135	1.142	1.143	1.149	1.161 *	1.187	
BE/ME	0.315	0.393	0.426	0.458	0.491	0.536	0.610	0.708	0.860	1.201	
Profitability	0.052	0.059	0.068	0.074	0.083	0.085	0.085	0.085	0.082	0.077	
AT	155.355	241.600	311.205	346.914	309.426	294.154	270.912	211.125	156.689	127.941	
COGS	11.286	21.574	35.311	46.175	46.589	50.621	51.428	47.056	38.557	37.758	
ME	271.058	325.273	384.974	396.373	330.857	273.228	213.709	141.830	77.486	35.135	
SGA/COGS	0.666	0.577	0.567	0.492	0.424	0.350	0.317	0.289	0.266	0.233	
CCC	77.463	138.438	203.687	268.954	333.411	367.837	401.345	424.267	437.618	457.796	
Quick Ratio	1.824	1.536	1.452	1.428	1.436	1.386	1.324	1.247	1.132	0.906	
ACT/AT	0.369	0.346	0.393	0.424	0.477	0.504	0.529	0.562	0.601	0.663	
INVT/COGS	8.363	3.377	2.077	1.563	1.280	1.178	1.082	1.012	0.958	0.884	
CHE/COGS	0.520	0.885	1.128	1.416	1.684	2.113	2.815	3.857	5.784	10.333	
RECT/COGS	0.977	0.998	1.025	1.043	1.007	1.045	1.066	1.126	1.199	1.314	
RECT/AP	1.759	1.823	1.871	1.969	2.068	2.069	2.055	1.985	1.845	1.606	
Debt/Equity	0.211	0.350	0.367	0.381	0.353	0.380	0.405	0.447	0.534	0.845	
FF6 Liq Beta	-0.011	0.026	0.026	0.038	0.033	0.058	0.054	0.029	0.069	0.166	

spread portfolio is IGMV (inventory growth minus value). Average annual returns are 4.6% for ILMH, and 3.9% for IGMV (see Table 3). Figure 4.5 plots cumulative log returns. ILMH and IGMV are highly correlated (Pearson correlation =0.53) and substantially more volatile than existing inventory growth factors, yielding modest Sharpe ratios (0.28 for ILMH; 0.22 for IGMV). Panel B of Figure 4.5 shows this in kernel density plots comparing ILMH, IGMV, and the inventory growth factors. Volatility is not priced uniformly, however: despite low Sharpe ratios, annualized alphas reported in Tables 4 and 5 are large, consistent with these portfolios covarying with aggregate risks investors value.

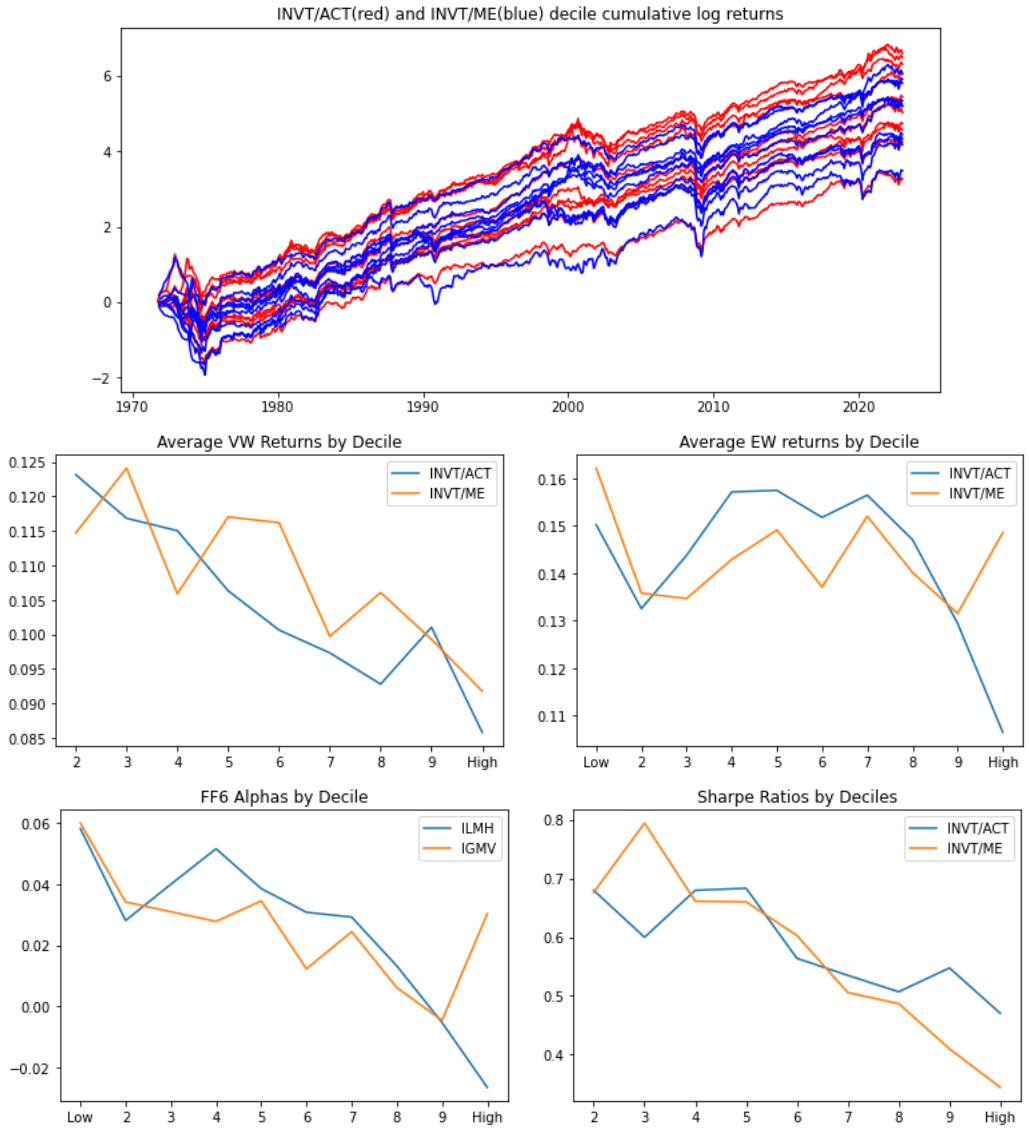


Figure 10: INVT/ACT and INVT/ME decile portfolio statistics. Panel A shows cumulative returns for each decile, with INVT/ACT deciles colored in red and INVT/ME deciles colored in blue. Panel B shows average value weighted returns. Panel C shows average equal weighted returns. Panel D shows alphas estimated using the Fama-French 5 factor model plus momentum portfolio (UMD) returns. Panel E shows Sharpe Ratios.

Table 4 reports time series regression results for the ILMH portfolio. The first row reports portfolio alphas on an annualized basis, while all other numbers in the table are unaltered regression coefficients. ILMH earns an annualized alpha of about 3.6 percent relative to the CAPM (column (1)), 6 percent relative to the Fama and French (1993) three factor model (column (2)), about 10 percent relative to the Fama and French (2014) five factor models with or without momentum (columns (3) and (4)), 8.4 percent relative to the AQR six factor model (column (5)), and about 9.6 and 7.2 percent relative to the Huo, Mo, Zue, and Zhang (2021) four and five factor q-factor models (columns (8) and (9)). Including the two inventory

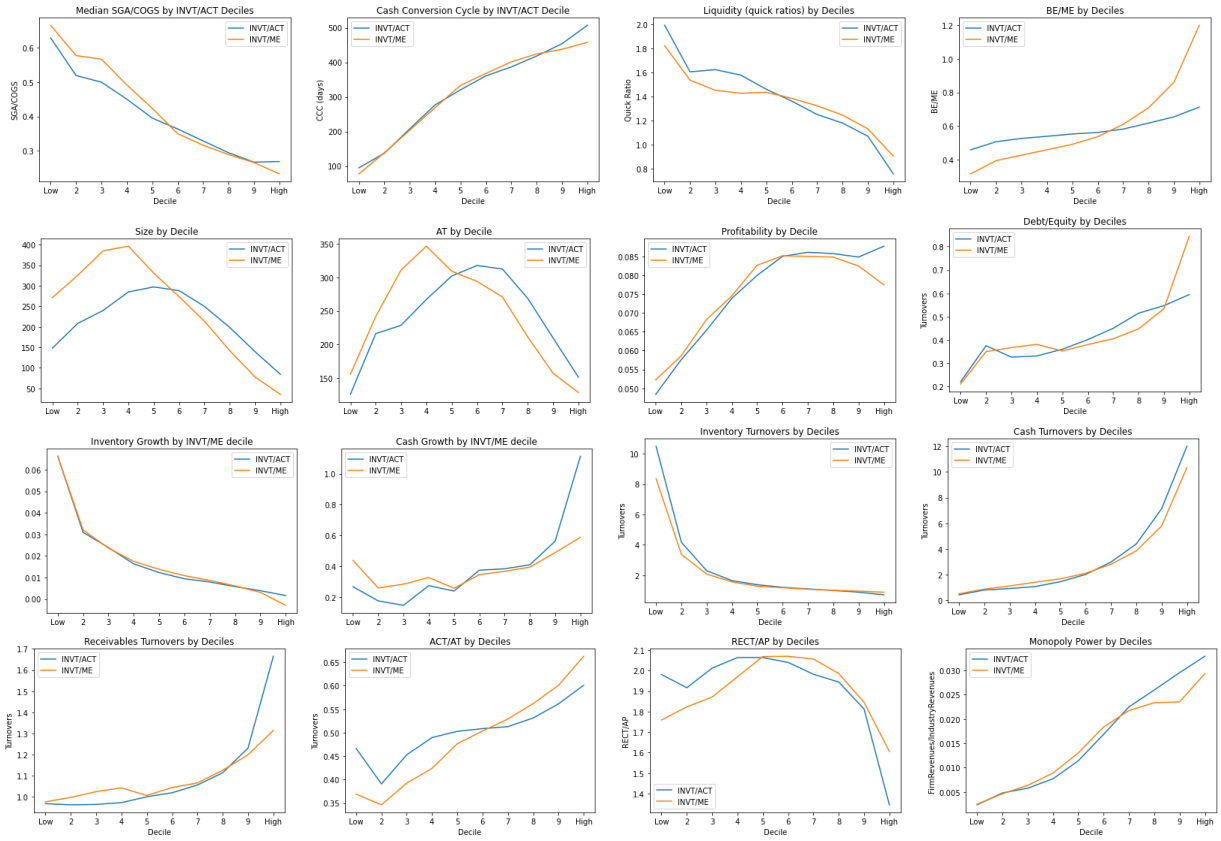


Figure 11: INVT/ACT and INVT/ME Decile Portfolio Firm Characteristics. This figure plots the statistics reported in Tables 2 and 3. Panel A plots measures of size. Panel B plots fixed costs as measured by sales, general and administrative expenses as a share of costs of goods sold. Panel C plots liquidity, as measured by the quick ratio. Panel D plots book-to-market ratios. Panel E plots profitability as measured by gross profits as a share of total assets. Panel F plots percentage growth in inventory investment. Panel G plots inventory turnover and cash turnover. Panel H plots measures of size. Panel I plots fixed costs as measured by sales, general and administrative expenses as a share of costs of goods sold. Panel J plots liquidity, as measured by the quick ratio. Panel K plots book-to-market ratios. Panel L plots profitability as measured by gross profits as a share of total assets. Panel M plots percentage growth in inventory investment. Panel N plots inventory turnover and cash turnover.

growth factors (IG1 and IG2) in the Fama-French 6-factor model (columns (6) and (7)) does little to erode the statistical significance or magnitude of the ILMH alpha, but it is interesting that these two factors load with opposite signs. The equity duration factor of Dechow, Sloan, Soliman (2004) is also included in column (7). ILMH's returns are clearly not explained by existing risk factors alone. Beyond the robust alpha result, factor loadings tell a story. ILMH returns are mostly orthogonal to the market risk premium. Recalling the hump-shaped relationship between INVT/ACT and with size, the somewhat inconsistent loading on SMB and QME factors makes sense. Value factor HML loads negatively, though this seems to be explained by duration (see column (7)). ILMH has a negative relationship with profitability factors RMW and QROE.

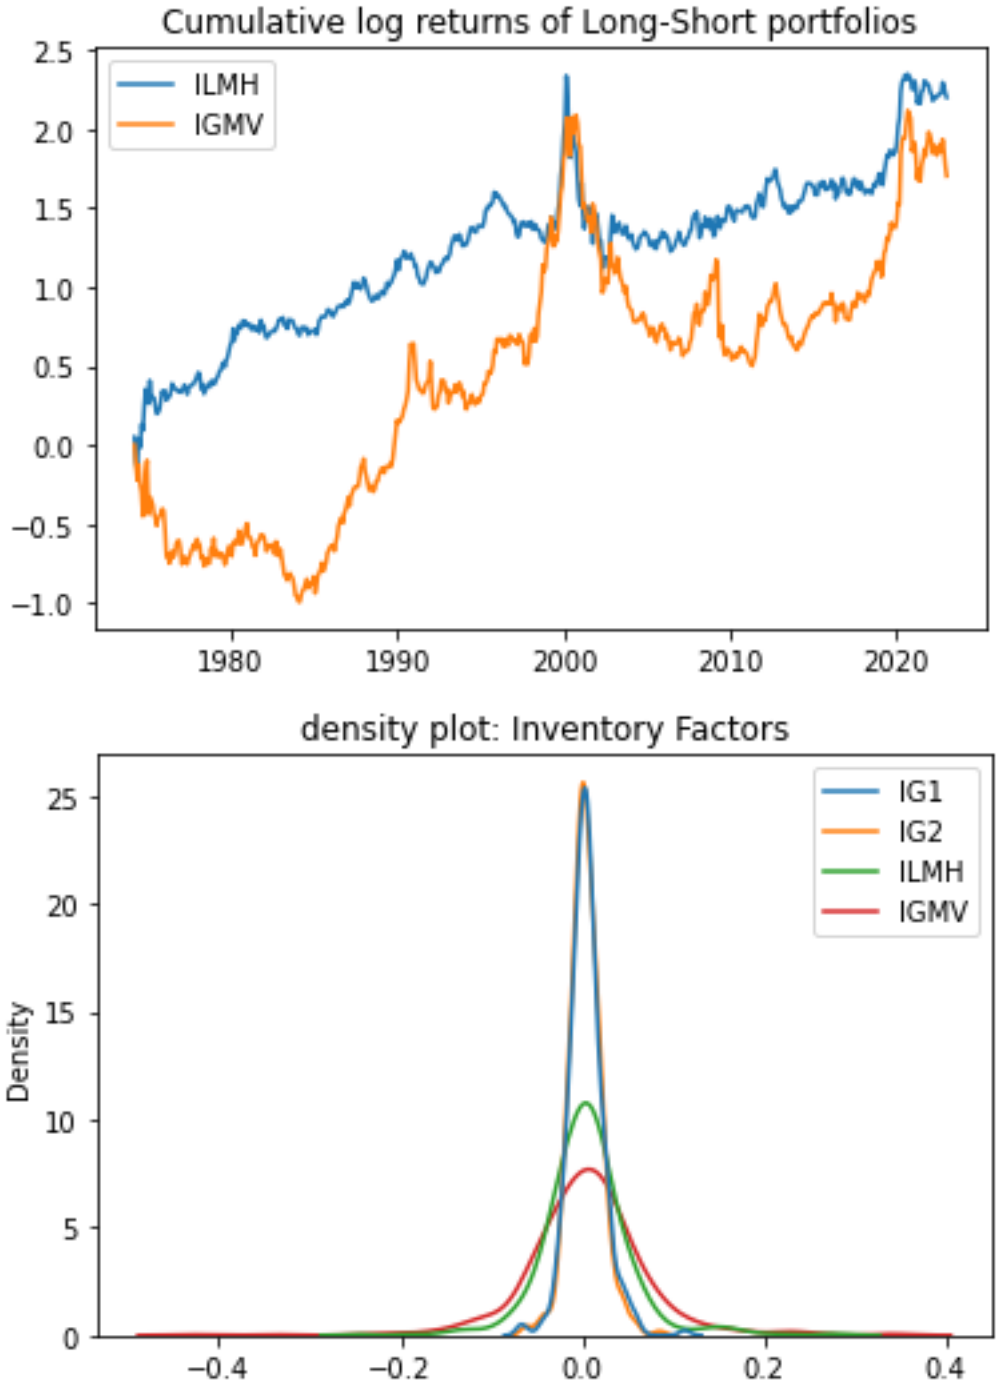


Figure 12: ILMH and IGMV factor portfolio returns. Panel A shows cumulative returns of the inventory factors (log-scale). Panel B shows empirical probability density plots of the inventory factors. Sample period is 1971-2022. ILMH is sorted on INVT/ACT, IVMG is sorted on INVT/ME, IG1 is sorted on INVT growth, and IG2 is sorted on INVT/AT.

Table 4: Time Series regressions for the ILMH long-short portfolio sorted on INVT/ACT. Dependent variable in all regressions is ILMH long-short portfolio returns. All t-statistics are robust to heteroscedasticity and autocorrelation at a one month lag. Annualized alphas are reported, all other coefficients are unaltered.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Alpha	0.036*	0.06***	0.108***	0.096***	0.084***	0.096***	0.108***	0.096***	0.072***
	1.790	3.562	5.961	5.218	4.398	4.787	5.633	4.266	3.121
MKTRF	0.153**	0.054	-0.010	0.016	-0.191***	0.027	-0.021		
	2.642	1.156	-0.262	0.439	-3.231	0.709	-0.550		
SMB		-0.003	-0.213**	-0.215***	-0.223*	-0.149*	-0.165**		
		-0.024	-2.659	-2.889	-2.279	-2.035	-2.458		
HML		-0.655***	-0.491***	-0.402***	-0.558***	-0.477***	-0.013		
		-5.977	-6.075	-5.009	-6.044	-5.109	-0.104		
RMW			-0.797***	-0.822***		-0.758***	-0.450***		
			-6.515	-7.086		-5.925	-4.205		
CMA			-0.231*	-0.297**		-0.328**	-0.296**		
			-1.801	-2.149		-2.329	-2.135		
UMD				0.159**	0.242***	0.147**	-0.020		
				2.564	3.192	2.343	-0.331		
BAB					-0.391***				
					-4.425				
QMJ					-0.289**				
					-2.303				
IG1						-0.320	-0.340		
						-1.497	-1.537		
IG2						0.602**	0.652**		
						2.397	2.575		
EQDR							-0.598***		
							-5.183		
QMKT								0.032	0.073
								0.559	1.351
QME								-0.105	-0.070
								-0.667	-0.446
QROE								-0.725***	-0.741***
								-4.756	-4.711
QIA								-0.249*	-0.373***
								-2.056	-2.875
QEG									0.358*
									2.150
R^2	0.022	0.248	0.372	0.383	0.32	0.397	0.418	0.128	0.141

This should not be surprising given what has already been established above in Tables 2 and 3 and seen in Figures 3 and 4. ILMH is a good hedge for duration, as well as profitability. Investment factors CMA and QIA both load negatively too, indicating that low INVT/ACT firms are fairly aggressive investors. The negative sign on EQDR is also consistent with the idea that inventory-lean firms should have greater sensitivity to discount rate shocks. In most specifications, the momentum factor UMD is positively related to ILMH, indicating that ILMH returns may cascade as new information about the firm's value arrives. The BAB (betting against beta) and QMJ (quality minus junk) factors are included, both of which load negatively. To the degree that quality is associated with profitability, QMJ's coefficient should come as no

surprise. The simultaneously negative loadings on MKTRF an BAB is consistent with the model, since low INVT/ACT firms should be high market beta firms, and the BAB factor is long low beta stocks and short high beta stocks. This is consistent with the notion that ILMH returns come about because of time varying sensitivity to systematic risk, caused in turn by variation in search costs and sensitivity to search cost risk. The QEG (q-factor expected growth factor) has a positive correlation with ILMH, which diminishes but does not erase or render insignificant the alpha.

Table 5: Time Series regressions for the IGMV long-short portfolio sorted on INVT/ME. Dependent variable in all regressions is ILMH long-short portfolio returns. All t-statistics are robust to heteroscedasticity and autocorrelation at a one month lag. Annualized alphas are reported, all other coefficients are unaltered.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Alpha	0.048	0.108***	0.132***	0.084***	0.06***	0.06***	0.084***	0.084**	0.06**
	1.629	4.554	4.908	3.787	2.527	3.065	3.770	2.357	2.088
MKTRF	-0.193*	-0.233***	-0.255***	-0.159	-0.143**	-0.137***	-0.184		
	-2.065	-3.742	-3.821	-2.890	-1.737	-2.656	-3.925		
SMB		-0.738***	-0.847***	-0.915***	-0.812***	-0.796***	-0.818***		
		-4.721	-7.552	-10.842	-7.687	-9.396	-10.370		
HML		-1.012***	-0.965***	-0.574***	-0.815***	-0.715***	-0.281*		
		-7.863	-7.040	-5.326	-7.958	-6.865	-1.769		
RMW			-0.381*	-0.528***	0.636***	-0.415**	-0.120		
			-2.112	-3.913	7.503	-2.793	-0.777		
CMA			-0.010	-0.313*		-0.384*	-0.378*		
			-0.051	-1.995		-2.114	-2.177		
UMD				0.650***		0.632***	0.479***		
				10.121		9.246	5.949		
BAB					-0.020				
					-0.194				
QMJ					-0.175				
					-1.136				
IG1						-0.554**	-0.563**		
						-2.423	-2.525		
IG2						1.096***	1.155***		
						4.361	4.820		
EQDR							-0.565***		
							-3.751		
QMKT								-0.185**	-0.164**
								-2.472	-2.234
QME								-0.578***	-0.561***
								-2.779	-2.653
QROE								-1.027***	-1.031***
								-5.240	-5.244
QIA								0.558	0.492**
								3.357	2.413
QEG									0.184
									0.832
R^2	0.019	0.377	0.393	0.561	0.533	0.583	0.6	0.27	0.272

Table 5 reports time-series results from the same set of factor models, with the dependent variable being IGMV returns. Unlike the ILMH portfolio, IGMV's returns are not significant when using only the CAPM

model. However, average returns are positive and jointly significant when other factors are included. These risk-adjusted returns are of a similar size as for ILMH. Market risk premium factors (MKTRF and QMKT) both load negatively in all specifications. Size factors (SMB and QME) load negatively. The negative loading on HML is unsurprising, since HML goes long value firms while IGMV goes long "inventory growth" firms. Interestingly, the CMA investment factor has a negative sign, while the QIA factor loads positively. Profitability factors both load negatively, as does equity duration. The inventory growth factors both behave similarly as in Table 4, but neither IG1 nor IG2 explain the IGMV returns entirely.

How should we interpret the opposite signs of coefficients on the two inventory growth factors? This asymmetry may provide some important clues. The raw inventory investment growth factor (IG1), which goes long firms with low inventory investment growth and shorts firms with high inventory investment growth, is negatively correlated with ILMH and IGMV. Panel F of Figure 9 shows a similar pattern with the decile portfolios. If inventory-lean firms have low INVT/ACT ratios because they are selling products quicker as they grow, then their inventory investment growth should be high in absolute terms. Should it be surprising then that IG2, a portfolio that goes long (short) firms with low (high) growth in inventory as a share of total assets, has instead a positive correlation with ILMH? Not necessarily. Growing, inventory-lean firms could be accumulating lots of new inventories while simultaneously investing even more aggressively in long term assets meant to improve future growth.

The positive and significant alphas shown in Tables 4 and 5 constitute a failure to reject hypotheses 1 and 2. However, the lack of significance of the alpha in column (1) using just the CAPM does hint at the importance of controlling for other factors and of the relatively less robust relationship between returns and the INVT/ME ratio. The next section leads to somewhat similar conclusions. Hypothesis 1 is mostly not rejected, but hypothesis 2 only passes the test with qualification. The effect of both inventory ratios on returns is more (jointly) pronounced and significant when other factors are included in regressions.

4.6 Cross-Sectional Results

Tables 6 and 7 report panel regression results, employing the Fama and MacBeth (1973) procedure. For each time period, a cross-sectional regression was run to estimate the relationship between stock returns and lagged firm characteristics. Controls include returns at a one month lag, a rolling sum of returns from 12 months previously up to one month previously, time-varying CAPM betas using data from the previous 12 months, the natural logarithm of size, the natural logarithm of book-to-market ratios, gross profits as a share of total assets, growth in total assets, debt-equity ratios, time-varying liquidity betas using data from the previous 12 months, a measure of market power. Liquidity betas are estimated using the value-weighted

liquidity factor of Pastor and Stambaugh (2003). The monopoly power variable is a firm's revenues divided by the sum of revenues of all firms in the same SIC code industry category. Table 6 is for raw data, while Table 7 is with all variables demeaned by industry using the Fama-French 49 industry categories. The table reports time-series averages of coefficients across the sample period. Since the two inventory ratios INVT/ACT and INVT/ME are highly correlated with each other, regressions were run separately for each variable. Columns (1)-(3) of Tables 6 and 7 feature the inventory-to-current asset ratio, while columns (4)-(6) feature the inventory-to-market equity ratio. Newey and West (1987) t-statistics, which are robust to heteroskedasticity and auto-correlation up to a maximum lag of 2 months, are reported beneath each coefficient. All variables are transformed to z-scores, except lagged return variables and the monopoly power variable, so coefficients can be interpreted as the effect of a one standard deviation increase in the variable on returns in decimals. All variables were cross-sectionally winsorized within each time period by removing the 1st and 99th percentiles.

Although neither variable is significant in univariate regressions without industry effects, both variables' estimated coefficients are significant and have the predicted negative sign when other standard controls are included. In Table 7, which does account for industry effects, the INVT/ACT ratio coefficient is significant and negative but the INVT/ME ratio is not. Table 6 shows that, without industry effects, and controlling for all other confounders, a one standard deviation increase in the INVT/ACT (INVT/ME) ratio reduces stock returns 3 months later by 1.2 (0.2) percent. Table 7 represents a failure to reject hypothesis 1, but a rejection of hypothesis 2. The relationship between INVT/ACT and stock returns is robust in the panel, but the relationship between INVT/ME and stock returns is explained by industry effects.

5 Discussion and Conclusion

Market efficiency implies that prices reflect available information, yet acquiring information is costly. If information search costs are systematic, they should nonetheless be reflected in asset prices. This paper develops a production-based asset pricing model under imperfect information in which firms search for customers and suppliers while investors try to time the market and hunt for bargain prices. When search is more costly, firm profits fall and the marginal utility of consumption rises, increasing required returns for firms that are more exposed to search-cost risk.

The model highlights an observable proxy for this exposure: the inventory-to-current-assets ratio (INVT/ACT). It predicts that firms with lower ratios are more sensitive to search-cost shocks and therefore must offer higher expected returns. I test these implications in reduced form using decile portfolio sorts, time-series factor regressions, and Fama–MacBeth panel regressions on accounting and stock-return data for publicly traded

Table 6: Fama-MacBeth Regression Results, no industry adjustment. Reports time series averages of coefficients estimated from T cross-sectional regressions of firm-level excess returns against firm characteristics at each time period. Data are winsorized within each time period at the 1st and 99th percentile cutoffs. All independent variables are lagged 3 months. Newey and West (1987) t-statistics are reported in parentheses, robust to heteroscedasticity and autocorrelation at a 3 month lag. Return (-1) is the previous month's returns. Return (-12-1) is the sum of the previous 12 month's returns up to one month previous. Profitability is gross profits divided by total assets. Investment is percentage growth in total assets. Debt/Equity is total liabilities divided by book equity. Liquidity Beta is estimated from univariate regressions of excess returns on the value-weighted liquidity factor constructed by Pastor and Stambaugh (2003). Monopoly power is firm revenues divided by total revenues of all firms in the same SIC code industry category. Time period is 1974-2022.

	(1)	(2)	(3)	(4)	(5)	(6)
INVT/ACT	-0.000	-0.016***	-0.012***			
	(-0.663)	(-6.826)	(-6.074)			
INVT/ME				-0.001	-0.005***	-0.002**
				(-1.034)	(-4.794)	(-2.209)
Return(-1)		-0.055***	-0.059***		-0.055***	-0.057***
		(-12.112)	(-12.860)		(-11.595)	(-12.042)
Return(-12-1)		0.002	0.002		0.003*	0.002
		(1.683)	(1.172)		(2.164)	(1.700)
CAPM Beta		-0.000	-0.000		0.001	-0.000
		(-0.162)	(-0.583)		(0.652)	(-0.035)
Log(Size)		0.008***	0.004**		0.009***	0.004**
		(3.001)	(2.710)		(3.428)	(2.628)
Log(B/M)		-0.006***	-0.005***		-0.005***	-0.005***
		(-3.784)	(-3.573)		(-3.114)	(-3.365)
GP/AT		0.070***	0.066***		0.094***	0.09***
		(13.656)	(14.589)		(12.578)	(13.630)
Investment Growth		-0.056***	-0.046***		-0.094***	-0.076***
		(-9.678)	(-6.853)		(-12.529)	(-9.755)
Debt/Equity			-0.001			-0.001
			(-1.576)			(-1.619)
Liquidity Beta			-0.000			0.000
			(-0.292)			(0.025)
Monopoly Power			-0.005*			-0.010***
			(-1.700)			(-2.994)

Table 7: Fama-MacBeth Regression Results, with industry adjustment. Reports time series averages of coefficients estimated from T cross-sectional regressions of firm-level excess returns against firm characteristics at each time period. Data is winsorized within each time period at the 1st and 99th percentile cutoffs. All independent variables are lagged 3 months. Newey and West (1987) t-statistics are reported in parentheses, robust to heteroscedasticity and autocorrelation at a 3 month lag. Return (-1) is the previous month's returns. Return (-12-1) is the sum of the previous 12 month's returns up to one month previous. Profitability is gross profits divided by total assets. Investment is percentage growth in total assets. Debt/Equity is total liabilities divided by book equity. Liquidity Beta is estimated from univariate regressions of excess returns on the value-weighted liquidity factor constructed by Pastor and Stambaugh (2003). Monopoly power is firm revenues divided by total revenues of all firms in the same SIC code industry category. Time period is 1974-2022.

	(1)	(2)	(3)	(4)	(5)	(6)
INVT/ACT	-0.001**	-0.006***	-0.007***			
	(-2.861)	(-6.498)	(-7.634)			
INVT/ME				-0.000	-0.001	-0.001
				(-0.642)	(-1.241)	(-1.387)
Return(-1)		-0.054***	-0.055***		-0.062***	-0.063***
		(-11.438)	(-11.801)		(-11.252)	(-11.379)
Return(-12-1)		0.002	0.002		0.003	0.002
		(1.554)	(1.197)		(1.524)	(1.287)
CAPM Beta		0.000	0.000		0.000	0.000
		(0.018)	(0.316)		(0.914)	(0.460)
Log(Size)		0.001	0.001		0.000	0.001
		(1.241)	(1.396)		(1.168)	(1.448)
Log(B/M)		-0.002***	-0.002***		0.001	0.001
		(-4.209)	(-4.197)		(0.776)	(1.260)
GP/AT		0.029***	0.038***		0.051***	0.071***
		(14.221)	(15.217)		(10.650)	(11.793)
Investment Growth		-0.022***	-0.028***		-0.052***	-0.069***
		(-9.203)	(-7.593)		(-10.140)	(-10.783)
Debt/Equity			-0.001***			-0.000
			(-3.363)			(-0.846)
Liquidity Beta			0.000			-0.000
			(0.141)			(-0.772)
Monopoly Power			-0.001			-0.001
			(-0.327)			(-0.846)

U.S. firms over roughly fifty years. The evidence is broadly consistent with the model. INVT/ACT is robustly and negatively related to future returns. INVT/ME is also negatively related to returns, but less robust and largely explained by industry effects. As always, the evidence cannot prove the theory; it can only reject or fail to reject it. While several alternative explanations are partly ruled out, the results should be interpreted as consistent with the model-driven hypothesis that investors price search-related risks.

The analysis also has a policy implication. The COVID-19 episode renewed attention to inventory cycles, supply-chain fragility, and proposals to "reshore" production through trade restrictions and domestic subsidies. Much of this debate treats resiliency as a justification for industrial policy distinct from traditional infant-industry or national-security arguments. Yet the return premia documented here can be interpreted as an ex ante market-based penalty on supply-chain over-extension: firms operating at the frontier of just-in-time networks - those most reliant on rapid information acquisition and rapid replenishment - face a higher cost of capital because they are the marginal searchers. The market already levies the "tax". To the extent that this pricing already disciplines fragility, additional policy aimed at discouraging supply-chain specialization may be redundant, and could even impede the very information flows that lower search costs. Whether protectionism and subsidies reduce exposure to shocks, or instead make economically meaningful information scarcer, remains an open question—and one that should be evaluated against the well-documented benefits of trade.

References

[Alchian, ArmenAlchian, Armen1969] Alchian, Armen (1969). "Information Costs, Pricing and Resource Unemployment" *Economic Enquiry* Volume 7, pages 109-128. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1465-7295.1969.tb01469.x.

[Alan, Gao, and GaurAlan, Gao, and Gaur2013] Alan, Yueh-feng, Jie Gao, and Vishal Gaur (2013). "Does Inventory Productivity Predict Future Stock Returns? Evidence from the Retail Industry." *Working Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1971774.

[Arrow K., T. Harris and J. Marschak.Arrow K., T. Harris and J. Marschak.1951] Arrow K., T. Harris and J. Marschak. (1951). "Optimal Inventory Policy." *Econometrica* Volume 19, No. 3, pages 250-272. https://www.jstor.org/stable/1906813.

[Barrot and SauvagnatBarrot and Sauvagnat2016] Barrot, Jean-Noël and Julien Sauvagnat (2016). "Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks." *The Quarterly Journal of Economics* Volume 131, No. 3, pages 1543–1592. https://doi.org/10.1093/qje/qjw018.

[Baumol, William J.Baumol, William J.1952] Baumol, William J. (1952). "The Transactions Demand for Cash: An Inventory Theoretic Approach" *The Quarterly Journal of Economics* Volume 66, No. 4., pages 545–556. https://www.jstor.org/stable/1882104.

[Belo, Federico and Xiaoji LinBelo, Federico and Xiaoji Lin2012] Belo, Federico and Xiaoji Lin (2012). "The Inventory Growth Spread". *The Review of Financial Studies* Volume 25, Issue 1, January 2012, pages 278–313. https://academic.oup.com/rfs/article-abstract/25/1/278/1571868.

[Bensousson, Alain.Bensousson, Alain.2011] Bensousson, Alain. (2011). Dynamic Programming and Inventory Control (IOS Press September 2011) Vol. 3 Studies in Probability, Optimization and Statistics. https://www.iospress.com/catalog/books/dynamic-programming-and-inventory-control.

[Blinder, Alan S.Blinder, Alan S.1981] Blinder, Alan S. (1981). "Retail Inventory Behavior and Business Fluctuations". *Brookings Papers on Economic Activity* Volume 2, pages 443–505. https://www.brookings.edu/wp-content/uploads/1981/06/1981b_bpea_blinder_lovell_summers.pdf.

[Blinder, Alan S. and Louis J. MacciniBlinder, Alan S. and Louis J. Maccini1981] Blinder, Alan S. and Louis J. Maccini (1991). "Taking Stock: A Critical Assessment of Recent Research on Inventories". *Journal of Economic Perspectives* volume 5, pages 73–96. https://www.aeaweb.org/articles?id=10.1257/jep.5.1.73.

[Bosworth, Barry P.Bosworth, Barry P.1970] Bosworth, Barry P. (1970). "Analyzing Inventory Investment" *Brookings Papers on Economic Activity* No. 2. https://www.brookings.edu/wp-content/uploads/1970/06/1970b_bpea_bosworth_duesenberry.pdf.

[Brandes, UlrikBrandes, Ulrik2022] Brandes, Ulrik (2001). "A Faster Algorithm for Betweenness Centrality." *Journal of Mathematical Sociology* Volume 25, Issue 2, 2001, pages 163-177. https://doi.org/10.1080/0022250X.2001.9990249.

[Burkart, Mike and Tore EllingsenBurkart, Mike and Tore Ellingsen2004] Burkart, Mike and Tore Ellingsen (2004). "In-Kind Finance: A Theory of Trade Credit". *American Economic Review* Volume 94, No. 3, pages 569-590. https://www.aeaweb.org/articles?id=10.1257/0002828041464579.

[Cai, YongCai, Yong2022] Cai, Yong (2022). "Linear Regression with Centrality Measures" *Working Paper* . https://arxiv.org/abs/2210.10024.

[Caplan, Andrew S.Caplan, Andrew S.2022] Caplan, Andrew S. (2022). "The Variability of Aggregate Demand with (S, s) Inventory Policies" *Econometrica* Volume 53, 1985, pages 1395-1410. https://arxiv.org/abs/2210.10024.

[Carlson, Murray, Adlai Fisher and Ron GiammarinoCarlson, Murray, Adlai Fisher and Ron Giammarino2006] Carlson, Murray, Adlai Fisher and Ron Giammarino (2006). "Corporate Investment and Asset Price Dynamics: Implications for SEO Event Studies and Long-Run Performance" *Journal of Finance* Volume 61, No. 3., pages 1009-1034. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2006.00865.x.

[Chakraborti, Rik and Gavin RobertsChakraborti, Rik and Gavin Roberts2023] Chakraborti, Rik and Gavin Roberts (2023). "How price-gouging regulation undermined COVID-19 mitigation: county-level evidence of unintended consequences". *Public Choice* 2023 Apr 18, pages 1–33. https://pubmed.ncbi.nlm.nih.gov/37360988/.

[Chen, Frank, and WuChen, Frank, and Wu2005] Chen, Hong, Murray Z. Frank, and Owen Q. Wu (2005). "What Actually Happened to the Inventories of American Retailers?" *Management Science* Volume 51, No. 7, pages 1015–1031. https://doi.org/10.1287/mnsc.1050.0385.

[Chen, ZhanhuiChen, Zhanhui2016] Chen, Zhanhui (2016). "Time-to-produce, inventory, and asset prices". *Journal of Financial Economics* Volume 120(2), pages 330-345. https://www.sciencedirect.com/science/article/abs/pii/S0304405X16000088.

[Cochrane, John H.Cochrane, John H.2008] Cochrane, John H. (2008). "The Dog That Did Not Bark: A Defense of Return Predictability." *The Review of Financial Studies* Volume 21, No. 4, pages 1533–1575. https://static1.squarespace.com/static/5e6033a4ea02d801f37e15bb/t/5edacaeb48decf49ea19692a/1591397101885/cochrane+dog+that+did+not+bark.pdf.

[CochraneCochrane1991] Cochrane, John H. (1991). "Production-Based Asset Pricing and the Link between Stock Returns and Economic Fluctuations." *The Journal of Finance* Volume 46, No. 1, pages 209–237. https://doi.org/10.1111/j.1540-6261.1991.tb03750.x.

[Cohen and FrazziniCohen and Frazzini2008] Cohen, Lauren and Andrea Frazzini (2008). "Economic Links and Predictable Returns." *The Journal of Finance* Volume 63, No. 4, pages 1977–2011. https://doi.org/10.1111/j.1540-6261.2008.01379.x.

[Dasgupta, Sudipto, Erica X N Li, Dong YanSudipto Dasgupta, Erica X N Li, Dong Yan2019] Dasgupta, Sudipto, Erica X N Li, Dong Yan. (2019). "Inventory Behavior and Financial Constraints: Theory and Evidence." *The Review of Financial Studies* Volume 32, No. 3, pages 1188–1233. https://ideas.repec.org/a/oup/rfinst/v32y2019i3p1188-1233..html.

[Davis and KahnDavis and Kahn2008] Davis, Steven J. and James A. Kahn (2008). "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels." *Journal of Economic Perspectives* Volume 22, No. 4, pages 155–180. https://www.aeaweb.org/articles?id=10.1257/jep.22.4.155.

[Eppen, G.D.Eppen, G.D.1979] Eppen, G.D. (1979). "Effects of centralization on expected costs in a multi-location newsboy problem." *Management Science* Volume 25, No. 5, pages 498-501. .

[Fama, Eugene F. and Kenneth R. FrenchFama, Eugene F. and Kenneth R. French2014] Fama, Eugene F. and Kenneth R. French (2014). "A five-factor asset pricing model". *Journal of Financial Economics* vol. 116, pages 1–22. https://www.sciencedirect.com/science/article/abs/pii/S0304405X14002323.

[Freeman, Linton C.Freeman, Linton C.1977] Freeman, Linton C. (1977). "A Faster Algorithm for Betweenness Centrality." *Sociometry* Volume 40, 1977, pages 34-41. https://doi.org/10.2307/3033543.

[Gertler and GilchristGertler and Gilchrist1994] Gertler, Mark and Simon Gilchrist (1994). "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms." *The Quarterly Journal of Economics* Volume 109, No. 2, pages 309–340. https://doi.org/10.2307/2118465.

[Gomes, João F., Leonid Kogan, and Motohiro YogoGomes, João F., Leonid Kogan, and Motohiro Yogo2009] Gomes, João F., Leonid Kogan, and Motohiro Yogo (2009). "Durability of Output and Expected Stock Returns". *Journal of Political Economy* Volume 117, Number 5. October 2009. pages 941-986. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=930242.

[HerskovicHerskovic2018] Herskovic, Boris (2018). "Networks in Production: Asset Pricing Implications." *The Journal of Finance* Volume 73, No. 4, pages 1785–1818. https://doi.org/10.1111/jofi.12684.

[Holt, Charles C., Franco Modigliani, John F. Muth and Herbert A. Simon.Holt, Charles C., Franco Modigliani, John F. Mu
Holt, Charles C., Franco Modigliani, John F. Muth and Herbert A. Simon (1960). Planning Production, Inventories, and Work Force, Englewood Cliffs, NJ, Prentice Hall Inc. https://econpapers.repec.org/article/oupajagec/v_3a44_3ay_3a1962_3ai_3a2_3ap_3a650-652..htm.

[Hornstein, AndreasHornstein, Andreas1998] Hornstein, Andreas (1998). "Inventory Investment and the Business Cycle" *Federal Reserve Bank of Richmond Economic Quarterly* Volume 84, No. 2, pages 73–96. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2126280.

[Hugonnier, Julien, Semyon Malamud and Erwan MorellecHugonnier, Julien, Semyon Malamud and Erwan Morellec2015] Hugonnier, Julien, Semyon Malamud and Erwan Morellec (2015). "Capital Supply Uncertainty, Cash

Holdings, and Investment" *Review of Financial Studies* Volume 28, No. 2., pages 391-445. https://econpapers.repec.org/article/ouprfinst/v_3a28_3ay_3a2015_3ai_3a2_3ap_3a391-445..htm.

[Hutt, William HaroldHutt, William Harold1939] Hutt, William Harold (1939). Theory of Idle Resources: A study in definition. London, J. Cape https://cdn.mises.org/The%20Theory%20of%20Idle%20Resources_2.pdf.

[Jones, Christopher S. and Selale TuzelJones, Christopher S. and Selale Tuzel2013] Jones, Christopher S. and Selale Tuzel (2013) "Inventory investment and the cost of capital". *Journal of Financial Economics* vol. 107, issue 3, pages 557-579.. https://www.sciencedirect.com/science/article/abs/pii/S0304405X12001869.

[Kahn J.Kahn J.1992] Kahn J. (1992). "Why Is Production More Volatile than Sales? Theory and Evidence on the Stockout Avoidance Motive for Inventory Holding." *The Quarterly Journal of Economics* Volume 107, No.2, pages 481- 510. https://www.jstor.org/stable/2118479.

[Kahn J. and M. Bils.Kahn J. and M. Bils.2000] Kahn J. and M. Bils. (2000). "What Inventory Behavior Tells Us about Business Cycles." *The American Economic Review* Volume 90, No. 3, 458-481. https://www.aeaweb.org/articles?id=10.1257/aer.90.3.458.

[Kashyap, Lamont, and SteinKashyap, Lamont, and Stein1994] Kashyap, Anil K., Owen A. Lamont, and Jeremy C. Stein (1994). "Credit Conditions and the Cyclical Behavior of Inventories." *The Quarterly Journal of Economics* Volume 109, No. 3, pages 565–592. https://doi.org/10.2307/2118419.

[Khan, Aubhik, and Julia K. Thomas.Khan, Aubhik, and Julia K. Thomas.2007] Khan, Aubhik, and Julia K. Thomas. (2007). "Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies." *The American Economic Review* Volume 97, No. 4, pages 1165-1188. https://www.aeaweb.org/articles?id=10.1257/aer.97.4.1165.

[Kewei Hou, Haitao Mo, Chen Xue, Lu ZhangKewei Hou, Haitao Mo, Chen Xue, Lu Zhang2021] Kewei Hou, Haitao Mo, Chen Xue, Lu Zhang (2021). "An Augmented q-Factor Model with Expected Growth". *Review of Finance* Volume 25, Issue 1, February 2021, pages 1–41. https://academic.oup.com/rof/article/25/1/1/5727769.

[Kogan, Leonid, Jun Li and Harold H. ZhangKogan, Leonid, Jun Li and Harold H. Zhang2023] Kogan, Leonid, Jun Li and Harold H. Zhang (2023). "Operating Hedge and Gross Profitability Premium" *Journal of Finance* Volume. 78, No. 6, December 2023, pages 3387-3422. https://onlinelibrary.wiley.com/doi/full/10.1111/jofi.13275.

[Kogan, Leonid, Jun Li, Harold H. Zhang and Yifan ZhuKogan, Leonid, Jun Li, Harold H. Zhang and Yifan Zhu2020] Kogan, Leonid, Jun Li, Harold H. Zhang and Yifan Zhu (2020). "Operating leverage and Asset pricing puzzles" *Working Paper*. https://www.cafr-sif.com/2021/files/213%20Operating%20Leverage%20and%20Hedging%20A%20Tale%20of%20Two%20Production%20Costs%20for%20Asset%20Pricing.pdf.

[Kydland, Finn E. and Edward C. PrescottKydland, Finn E. and Edward C. Prescott1982] Kydland, Finn E. and Edward C. Prescott (1982). "Time to Build and Aggregate Fluctuations". *Econometrica* Volume 50, No. 6, pages 1345-1370. https://www.jstor.org/stable/1913386.

[Lachmann, Ludwig M.Lachmann, Ludwig M.1936] Lachmann, Ludwig M. (1936). "Commodity Stocks and Equilibrium." *Review of Economic Statistics* Volume 3, pages 230–34. https://www.mises.at/static/literatur/Buch/lachmann-expectations-and-the-meaning-of-institutions.pdf.

[Lachmann, Ludwig M. and F. Snapper.Lachmann, Ludwig M. and F. Snapper.1936] Lachmann, Ludwig M. and F. Snapper. (1938). "Commodity Stocks in the Trade Cycle." *Economica* Volume 5, pages 435–54. https://cdn.mises.org/Commodity%20Stocks%20in%20the%20Trade%20Cycle_2.pdf.

[Lee, Hsiao-Hui, Zhou, Jianer and Po-Hsuan Hsu.Lee, Hsiao-Hui, Zhou, Jianer and Po-Hsuan Hsu.2015] Lee, Hsiao-Hui, Zhou, Jianer and Po-Hsuan Hsu. (2015). "The role of innovation in inventory turnover performance." *Decision Support Systems* Volume 76, 2015, pages 35-44. https://doi.org/10.1080/0022250X.2001.9990249.

[Leve, BaruchLeve, Baruch1974] Leve, Baruch (1974). "On the Association between Operating Leverage and Risk." *Journal of Financial and Quantitative Analysis* Volume. 9, No. 4, pages 627-641. https://econpapers.repec.org/article/cupjfinqa/v_3a9_3ay_3a1974_3ai_3a04_3ap_3a627-641_5f01.htm.

[LovellLovell1961] Lovell, Michael C. (1961). "Manufacturers' Inventories, Sales Expectations, and the Acceleration Principle." *Econometrica* Volume 29, No. 3, pages 293–314. https://www.jstor.org/stable/1909630.

[Lucas, Robert E.Lucas, Robert E.1977] Lucas, Robert E. (1977). "Understanding Business Cycles". Carnegie-Rochester Conference Series on Public Policy. Volume 5, 1977, Pages 7-29 reprinted in Robert E. Lucas, Studies in Business Cycle Theory. Cambridge, Mass.: MIT Press, 1989, pp. 215–39. https://www.sciencedirect.com/science/article/abs/pii/0167223177900021.

[McCall, John J.McCall, John J.1965] McCall, John J. (1965). "The Economics of Information and Optimal Stopping Rules" *The Journal of Business* Volume 38, No. 3., pages 300-317. https://www.jstor.org/stable/2351065.

[McConnell and Pérez-QuirosMcConnell and Pérez-Quiros2000] McConnell, Margaret M. and Gabriel Pérez-Quiros (2000). "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review* Volume 90, No. 5, pages 1464–1476. https://www.aeaweb.org/articles?id=10.1257/aer.90.5.1464.

[Metzler, Lloyd A.Metzler, Lloyd A.1941] Metzler, Lloyd A. (1941). "The Nature and Stability of Inventory Cycles." *Review of Economics and Statistics* Volume 23, No. 3, pages 113-129, reprinted in Robert A. Gordon and Lawrence R. Klein, eds., Readings in Business Cycles. Homewood, Ill.: Irwin, 1965. https://www.jstor.org/stable/1927555.

[Mortensen, DaleMortensen, Dale1986] Mortensen, Dale (1986). "Job search and labor market analysis" in Ashenfelter, O.; Card, D. (eds.) The Handbook of Labor Economics, Vol. 2. Amsterdam: North-Holland. https://www.sciencedirect.com/science/article/abs/pii/S1573446399300250.

[Mortensen, Dale; Pissarides, ChristopherMortensen, Dale; Pissarides, Christopher1994] Mortensen, Dale; Pissarides, Christopher (1994). "Job creation and job destruction in the theory of unemployment" *Review of Economic Studies* Volume 61, No. 3., pages 397–415. http://bibliotecadigital.econ.uba.ar/download/Pe/187963.pdf.

[Novy-Marx, RobertNovy-Marx, Robert2010] Novy-Marx, Robert (2010). "The other side of value: the gross profitability premium" *Journal of Financial Economics* Volume. 108, No. 1, April 2013, pages 1-28. https://www.sciencedirect.com/science/article/abs/pii/S0304405X13000044.

[Opler, Tim, Lee Pankowitz, Rene Stulz, and Rohan WilliamsonOpler, Tim, Lee Pankowitz, Rene Stulz, and Rohan Willian
Opler, Tim, Lee Pankowitz, Rene Stulz, and Rohan Williamson (1999). "The determinants and implications of corporate cash holdings" *Journal of Financial Economics* Volume 52, No. 1., pages 3–46. https://ideas.repec.org/a/eee/jfinec/v52y1999i1p3-46.html.

[Pástor, Ľuboš and Robert F. StambaughĽuboš Pástor and Robert F. Stambaugh2003] Pástor, Ľuboš and Robert F. Stambaugh (2003). "Liquidity Risk and Expected Stock Returns." *Journal of Political Economy* Volume 111, No. 3, pages 642-685. https://www.journals.uchicago.edu/doi/10.1086/374184.

[Peterson, Bruce C., Carpenter, Robert E., Fazzari, Steven M., Kashyap, Anil K. and Benjamin M. FriedmanPeterson, Bruc
Peterson, Bruce C., Carpenter, Robert E., Fazzari, Steven M., Kashyap, Anil K. and Benjamin M.

Friedman (1994). "Inventory Investment, Internal Finance Fluctuations, and the Business Cycle." *Brookings Papers on Economic Activity* Number 2, 1994. https://www.brookings.edu/wp-content/uploads/1994/06/1994b_bpea_carpenter_fazzari_petersen_kashyap_friedman.pdf.

[Ramey V. and K. West.Ramey V. and K. West.1999] Ramey V. and K. West. (1999). Inventories. Chapter 13 in the Handbook of Macroeconomics, Vol. 1, Elsevier Science. https://users.ssc.wisc.edu/~kwest/publications/1990/Inventories.pdf.

[Ramey and VineRamey and Vine2006] Ramey, Valerie A. and Daniel J. Vine (2006). "Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry." *The Quarterly Journal of Economics* Volume 121, No. 1, pages 221–255. https://doi.org/10.1162/qjec.2006.121.1.221.

[SloanSloan1996] Sloan, Richard G. (1996). "Do Stock Prices Fully Reflect Information in Accruals and Cash Flows about Future Earnings?" *The Accounting Review* Volume 71, No. 3, pages 289–315. https://www.jstor.org/stable/248290.

[Stigler, George J.Stigler, George J.1961] Stigler, George J. (1961). "The Economic of Information" *Journal of Political Economy* Volume 69, No. 3., pages 213-225. https://home.uchicago.edu/~vlima/courses/econ200/spring01/stigler.pdf.

[Stigler, George J.Stigler, George J.1962] Stigler, George J. (1962). "Information in the labor market" *Journal of Political Economy* Volume 70, No. 5., pages 94-105. https://www.jstor.org/stable/1829106.

[Thomas and ZhangThomas and Zhang2002] Thomas, Jacob K. and Huai Zhang (2002). "Inventory Changes and Future Returns." *Review of Accounting Studies* Volume 7, Issues 2–3, pages 163–187. https://doi.org/10.1023/A:1020238227921.

[Wang, BaolinWang, Baolin2019] Wang, Baolin (2019). "The Cash Conversion Cycle Spread" *Journal of Financial Economics* Volume 133, No. 2., pages 472-497. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2964330.

[Yogo, MotohiroYogo, Motohiro2016] Yogo, Motohiro (2006). "A Consumption-Based Explanation of Expected Stock Returns". *Journal of Finance* Volume 61, Issue 2, April 2006, pages 539-580. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2006.00848.x.