Seasonal Inventory Leverage*

Kevin Aretz†
University of Manchester
kevin.aretz@manchester.ac.uk

Hening Liu‡
University of Manchester
hening.liu@manchester.ac.uk

Kevin Schneider✠
University of Cambridge
k.schneider@jbs.cam.ac.uk

This Version: Monday 11th March, 2024

Abstract

We develop a neoclassical theory of a firm with seasonal output prices and inventory building. Our theory suggests that seasonal firms optimally build up output inventories toward their high-price seasons, prepaying quasi-fixed production costs and delevering themselves. In turn, the delevering creates inverse seasonality in their expected returns. Crucially, higher inventory holding costs reduce inventory building, making the expected return seasonality disappear. Supporting our theory, our empirical work reveals that seasonal firms build up inventories toward their high-sales seasons. Also, high seasonal inventory leverage predicts high stock returns and helps explain several seasonal and allegedly-non-seasonal stock anomalies.

Keywords: Asset pricing, real options, seasonalities, seasonal inventory leverage, output inventories.

JEL classification: G11, G12, G13, G14.

†Corresponding author, Accounting and Finance Division, Booth Street West, Manchester M15 6PB, UK, tel.: +44(0)161 275 0122, fax.: +44(0)161 275 4023, URL address: www.kevin-aretz.com.

‡Accounting and Finance Division, Booth Street West, Manchester M15 6PB, UK, tel.: +44(0)161 275 4492, fax.: +44(0)161 275 4023, URL address: www.sites.google.com/site/liuhn2008.

✠Cambridge Endowment for Research in Finance (CERF), Cambridge Judge Business School, Trumpington Street, Cambridge CB2 1AG, UK, URL address: www.sites.google.com/view/kevinschneider.

*We are indebted to Michael Brennan, Sungjun Choi, Ilan Cooper, Alexandre Corhay, Adlai Fisher, Ron Giammarino, Hassan Ilyas, Yamil Kaba, Gaurav Kankanhalli, Matti Keloharju, Nicos Koussis, Bart Lambrecht, Matthias Lassak, Peter Nyberg, Juliusz Radwanski, Mark Shackleton, Greg Tindall, and Paul Zimmermann, seminar participants at Cardiff Business School, and conference participants at the 2021 British Academy of Management (BAM) Conference (online), the 2021 International Young Finance Scholars Conference (online), the 2021 New Zealand Finance Meeting (online), the 2022 Eastern Finance Association Annual Meeting (Washington DC), the 2022 French Finance Association International Conference (Saint Malo), the 2022 European Financial Management Association (EFMA) Annual Meeting (Rome), the 2022 Financial Management Association (FMA) European Conference (Lyon), the 2023 Financial Intermediation Research Society (FIRS) Conference (Vancouver), the 2023 Annual Corporate Finance Conference (Exeter), the 2024 Frontiers of Factor Investing Conference (Lancaster), and the 2024 CAIR Thematic Talk about Investment Decisions (Manchester) for many innovative and constructive suggestions.

Conflict-of-Interest Disclosure Statement:

Kevin Aretz: I have nothing to disclose.

Hening Liu: I have nothing to disclose.

Kevin Schneider: I have nothing to disclose.

1 Introduction

Many empirical studies discover important seasonalities in stock, accounting, and macroeconomic variables. Citing several recent examples, Chang et al. (2017) and Hartzmark and Solomon (2018) find seasonalities in corporate earnings. Heston and Sadka (2008, 2010) and Keloharju et al. (2016, 2021) detect seasonalities in stock returns, showing that a stock's past same-calendar-month return positively predicts its future return ("same-calendar-month anomaly"). Ogden (2003) observes seasonalities in aggregate production, investment, and stock market capitalization. Interestingly, Grullon et al. (2020) connect the seasonalities in stock and accounting data, showing that stocks tend to underperform in their high sales quarters and vice versa ("seasonal sales anomaly"). Yet, despite this abundance of empirical evidence, we lack theoretical studies shedding light on the seasonalities and their connections.

In this paper, we make a first step toward closing this gap in the literature by developing a neoclassical theory of a firm with a seasonal (exogenous) state variable and output-inventory building. Specifically, we study a firm able to produce, store in inventory, and later sell a unique output good at a seasonal stochastic price. Assuming no or linear inventory holding costs, we show that the firm's optimal policy is to build up output inventories over some period before its high price season, to sell those in that season. A crucial implication is that the firm prepays quasi-fixed production costs associated with output kept in-house over that period, gradually lowering its "seasonal inventory leverage" and thus expected return.[1] Strikingly, higher inventory holding costs reduce inventory building, making the endogenous seasonalities disappear. In line with our theoretical work, we empirically show that seasonal firms build up output inventories

[1]Akin to the operating leverage concept, a higher seasonal inventory leverage (induced through *lower* inventory holdings) implies that a fixed percentage increase in sales generates a more positive fixed percentage increase in firm value and vice versa, simply because the present value of inventory holdings are added to firm value.

toward their high-sales seasons. Also, seasonal inventory leverage is positively related to future stock returns and sheds new light on several well-known stock anomalies.

In our theoretical work, we investigate a toy model and a real options model of an all-equity-financed firm owning capacity able to produce a unique output good ("assets-in-place") and operating over a finite continuous horizon. We innovate upon similar models in the literature by, first, exposing the firm to seasonality in the stochastic price at which it sells output and, second, allowing it to hold output in inventory. To be specific, we model the output price as a *generalized* geometric Brownian motion whose drift term contains an additive sine function yielding a seasonal output-price drift.[2] Moreover, while the firm always produces at full capacity since it incurs only fixed (but no variable) production costs, we allow it to sell its concurrent production at a later time by either paying no (toy model) or a linear (real options model) inventory holding cost. Also important, while the toy model assumes that the firm has to precommit on when it will sell its concurrent production, the real options model allows the firm to optimally update its selling dates for output held in inventory as uncertainty resolves. The upshot is that both our models contain three state variables, namely the current time, output held in inventory, and output price, and one set of choice variables, namely the optimal selling dates.

We start our empirical work by testing our theory's prediction that seasonal firms build up output inventories toward their high-demand seasons. Toward that goal, we first turn to the Manufacturers' Shipments, Inventories, and Orders (MSIO) surveys, containing unique monthly data for several manufacturing industries (see, e.g., Jones and Tüzel (2013)). The data show that while the surveyed firms are exposed to strong seasonalities in their new orders, they swiftly react to these, with them shipping out the lion share of new orders in high-order months within those months. In line with our theory, one reason is that those firms hold large work-in-progress and

[2]Among others, Geman and Nguyen (2005) employ trigonometric functions in the drift of diffusion processes to model seasonal patterns in commodity prices and to derive futures prices.

finished-good but small raw material inventories at the start of high-order months, suggesting they actively build up output inventories toward these months. Running a similar albeit cruder analysis over a more diverse set of industries using Compustat data, we further show that the tendency to build up output inventories toward high-demand seasons is significantly greater for (i) firms with more seasonal sales and (ii) industries expected to have lower (e.g., consumer durables) than higher (e.g., services or healthcare) inventory holding costs. In addition, we reveal that seasonal firms lower their cash balances and raise their short-term liabilities over their inventory building periods as well as raise their payouts over their high sales periods.

Having confirmed that our sample firms engage in optimal inventory building in response to seasonalities, we next empirically assess our theory's prediction that firms with amplified inventory holdings relative to their annual average holdings have prepaid quasi-fixed costs and thus have low seasonal inventory leverage and expected returns. In the spirit of Grullon et al. (2020), we thus calculate the historical fraction of a firm's inventory holdings at the *start* of the current fiscal quarter relative to its average holdings over the fiscal year (*QInv*). *QInv* is an *inverse* measure for seasonal inventory leverage and a high value signals a delevered firm (since output kept in inventory *adds* to the firm's value). In line with our theory, our empirical work shows that *QInv* is strongly negatively priced. Our value-weighted portfolio sorts, for example, show that the top-minus-bottom *QInv* tercile spread portfolio yields a mean monthly return of –0.39% (t-statistic: –4.39). Remarkably, the mean spread return is close-to-unaffected by the Fama and French (2016) six-factor-model (–0.35%) or the Hou et al. (2021) augmented q-theory (–0.38%) factors but is more negative in the seasonal-stock (–0.56%) than the non-seasonal-stock (–0.26%) subsample. Also, while our Fama and MacBeth (FM; 1973) regressions confirm our portfolio results, they further validate our theory's prediction that the *QInv* premium is stronger for firms with higher historical quarterly inventories-to-quarterly sales ratios.

We finally study whether seasonal inventory leverage sheds new light on several seasonal and allegedly-non-seasonal stock anomalies. Specifically, while, as we already said, Grullon et al. (2020) show that stocks tend to earn low (high) returns in their high (low) sales quarters, our theory predicts that firms with high (low) abnormal inventory holdings — and thus low (high) seasonal inventory leverage — at the start of those quarters ("inventory builders") drive the low (high) returns. In the same vein, if a firm's high (low) return months identify its low (high) sales quarters, then seasonal inventory leverage should also condition Heston and Sadka's (2008, 2010) same-calendar-month anomaly. Moving to allegedly-non-seasonal anomalies, Heston and Sadka (2008) show that Jegadeesh and Titman's (1993) twelve-month momentum effect is almost entirely driven by the twelve-month-ago return, so that the momentum and same-calendar-month anomalies greatly overlap (see also Novy-Marx (2012)). Finally, Hou et al.'s (2015) past-quarter ROE may also crudely distinguish between a firm's high and low sales seasons. Given that, seasonal inventory leverage may also help to explain those anomalies.

We rely on triple portfolio sorts to evaluate whether seasonal inventory leverage adds to or even explains the anomalies by contrasting standard anomaly-variable spread portfolios with *seasonal-inventory-leverage-neutral* spread portfolios exclusively formed from stocks not building up output inventories toward their high sales seasons. Remarkably, our evidence shows that each single anomaly is stronger in the seasonal than in the non-seasonal stock subsample, a first piece of evidence linking the anomalies to seasonality. More importantly, it further reveals that the seasonal-inventory-leverage-neutral spread portfolios yield mean returns significantly closer to zero than the standard spread portfolios. While the standard *QSales* spread portfolio, for example, produces a highly significant mean monthly return of –0.22% (t-statistic: –2.17), the corresponding number for the seasonal-inventory-leverage-neutral portfolio is an insignificant –0.10% (t-statistic: –0.80), a whopping 55% increase. In comparison, the seasonal-inventory-leverage-neutral same-calendar-month, momentum, and quarterly ROE spread portfolios yield no-longer-significant mean returns 20%, 35%, and 43% closer to zero, respectively.

Our work adds to a large literature studying the pricing of leverage coming from a firm's operational characteristics. On the theoretical front, Rubinstein (1973), Carlson et al. (2004), Cooper (2006), and Lambrecht et al. (2016), among others, show that operating leverage induced through quasi-fixed production costs inflates expected firm returns. On the empirical front, Mandelker and Rhee (1984), Garcia-Feijóo and Jorgensen (2010), and Novy-Marx (2011) support the positive pricing of operating leverage and link it to value anomalies. Conversely, Gu et al. (2018) and Kogan et al. (2023b) show that the pricing of operating leverage depends on investment flexibility and variable production costs. While these studies look into the effects of time-invariant leverage induced through quasi-fixed production costs, we study those of time-variation in a firm's leverage over the fiscal year coming from its optimal seasonality-induced inventory holdings. We also contribute to studies looking into the pricing of inventory variables, again however focusing on the effects of dynamic (rather than static) inventory variables.[3]

We also contribute to a literature on recent seasonalities in firm-level stock and accounting data. Chang et al. (2017) show that the announcement of high (low) earnings in fiscal quarters with historically high (low) earnings produces high (low) abnormal returns, presumably because investors and analysts overweight the information in the most recent 2-3 earnings announcements but neglect seasonal patterns in earnings. As already said, Heston and Sadka (2008, 2010) and Keloharju et al. (2016, 2021) report that the past same-calendar-month return positively prices stocks. While Keloharju et al. (2021) posit that these results come from temporary mis-

[3]Thomas and Zhang (2002) show that annual inventory changes negatively price stocks. Using dynamic investment models, Belo and Lin (2012), Jones and Tüzel (2013), and Chen (2016) rationalize that finding by showing that firms with larger inventory investments have lower adjustment costs, making them more flexible. Different from us, these papers focus on *input* inventory holdings, while we focus on *output* inventory holdings. Similarly, Chen et al. (2005) and Alan et al. (2014) find that absolute inventory holdings also negatively price stocks, possibly because investors only slowly incorporate inventory information from financial reports into their valuations.

pricing effects, Hirshleifer et al. (2020) suggest that they are driven by seasonalities in investors' mood. As also already said, Grullon et al. (2020) find that stocks tend to underperform in their high sales quarters and vice versa, linking their findings to firms' investment behavior, their financial leverage, and investor inattention. We add to these studies by offering evidence that seasonal inventory leverage not only underlies the seasonal sales and same-calendar-month effects but also the (intrinsically related) momentum and quarterly ROE effects.

We finally also add to the real options asset pricing literature. Differentiating ourselves from other models in that literature (see, e.g., Kogan and Papanikolaou (2013, 2014), Hackbarth and Johnson (2015), Aretz and Pope (2018), Kogan et al. (2023a), and others), we consider a firm exposed to seasonalities in its main exogenous state variable, the stochastic price at which it sells its output. As a further difference, we also allow the firm to store produced output in inventory to sell it later, disentangling the production and selling of output.

We organize our paper as follows. In Section 2, we develop a toy and real options model with seasonal output prices and optimal inventory building. In Section 3, we verify the models' predictions that (i) firms with seasonal demand and low inventory holding costs build up output inventories toward their high-demand season, and that (ii) seasonal inventory leverage induced through optimal inventory building positively prices stock returns. In Section 4, we assess whether seasonal inventory leverage sheds new light on stock anomalies. Section 5 sums up and concludes. We relay mathematical proofs as well as supplementary theoretical and empirical material to our main paper appendix and Internet Appendix.

2 Theoretical Analysis

In this section, we develop a toy model and a real options model of a firm with output price seasonality and output inventory building. The key advantage of the toy model is that it allows us to derive the firm's optimal policies and values in quasi-closed-form, shedding more light on our models' main mechanisms. We start with stating the toy model's assumptions and then derive its quasi-closed-form solution. We next discuss how the real options model differs from the toy model and then explain how we numerically solve the real options model. We finally elaborate on the main corporate finance and asset pricing implications of our models.[4]

2.1 The Toy Model

2.1.1 Modeling Assumptions

In our toy model, we study an all-equity-financed firm operating over a continuous time horizon indexed by t and ending at time T. The firm owns $\bar{K}$ capacity units, with each capacity increment enabling it to produce one output increment per time unit. Since the firm only incurs fixed but no variable production costs of $f\bar{K}$, where $f > 0$ is a fixed parameter, it always optimally produces at full capacity. The firm immediately shifts its produced output into inventory, with the stock of output in inventory at time t equal to I_t. In the toy (but not the real options) model, we assume that the firm does not incur costs from holding output in inventory.

The firm can costlessly shift an amount of output $S_t \in [0, I_t]$ out of inventory and sell it at the prevailing output price P_t at future time t. Importantly, however, the toy model assumes that the firm has to precommit on the time at which it will sell its concurrently produced output,

[4] In Section IA.1 of our Internet Appendix, we also look into an extended version of our real options model in which the firm holds a growth option allowing it to discretely raise its capacity-in-place at some fixed cost. The extended model shows that our theoretical conclusions are robust to adding that option.

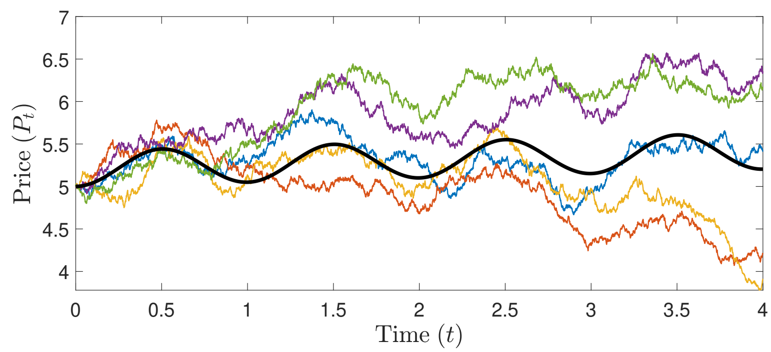

Figure 1: The figure plots five sample paths (non-black lines) and the time-0 conditional expectation (black line) of seasonal stochastic process (1) under the real world measure over the period from $t = 0$ to 4. We describe the stochastic process parameter values used to create the figure in Section 2.1.1.

allowing us to solve the model in quasi-closed-form. We assume that the demand for the firm's output is seasonal, so that P_t obeys a *generalized* geometric Brownian motion (GBM) whose drift term displays seasonal variations due to an additive sine function component:

$$dP_t = (\alpha + \kappa \sin(\eta t))P_t dt + \sigma P_t dB_t, \tag{1}$$

where the constant $\alpha > 0$ is the linear time trend, the constant $\kappa \geq 0$ controls the amplitude of seasonal fluctuations, the constant $\eta > 0$ governs the length of a seasonal cycle,[5] the constant $\sigma > 0$ is volatility, and B_t is a Brownian motion. Setting $\kappa = 0$, stochastic process (1) collapses to a standard GBM, as used in other asset pricing models. Assuming an initial output price value (P_0) of five, an annualized linear time trend (α) of 1%, a seasonal strength (κ) of 0.25,[6] a periodicity (η) of 2π, and an annualized volatility (σ) of 5%, Figure 1 plots five sample paths (non-black lines) plus the time-0 expectation (black line) of stochastic process (1).

[5]The length of a seasonal cycle is $2\pi/\eta$. Since we interpret a seasonal cycle as a year, we always set $\eta = 2\pi$.

[6]To guide our choice of parameter values, we use a maximum likelihood estimation to fit stochastic process (1) to the quarterly sales growth of seasonal firms with at least five years of data, where we classify a firm as seasonal if its average *Seasonality* value exceeds the full-sample median (see Section 3 for more details about our sample data and variable definitions). Doing so, we obtain an average κ estimate of about 0.22.

It follows that we can write the firm's net operating profits per time unit, Π_t, as:

$$\Pi_t = P_t S_t - f\bar{K}, \tag{2}$$

while we can write the law of motion for the output in inventory, I_t, as:

$$\mathrm{d}I_t = (\bar{K} - S_t)\mathrm{d}t. \tag{3}$$

Consistently choosing $S_t = \bar{K}$, the firm always sells all of its concurrently produced output instantaneously, as in other asset pricing models in the extant literature.

Overall, the model's state space consists of the current time t (due to the seasonality and the finite horizon), the output price P_t, and the output in inventory I_t. Conversely, the only choice variables are the amounts of output to be sold out of inventory in the future, S_t.

2.1.2 Optimal Selling Policy

We assume that the firm chooses its optimal selling policy by maximizing its value. To determine how the firm should optimally allocate its *current* production over the period from current time t to T, we notice that this decision does not affect the firm's benefits and costs derived from output produced in the future. The reason is that the current allocations neither affect future output prices nor production costs. Thus, the firm can make decisions about its current production in isolation of those about its future production. In addition, we notice that if it were optimal to sell a non-zero amount of concurrently produced output at future time $\tau^* \in [t, T]$, it would also be optimal to sell all other such output at that time, simply because the present value from selling an output unit does not differ across output units. As a result, the firm either sells

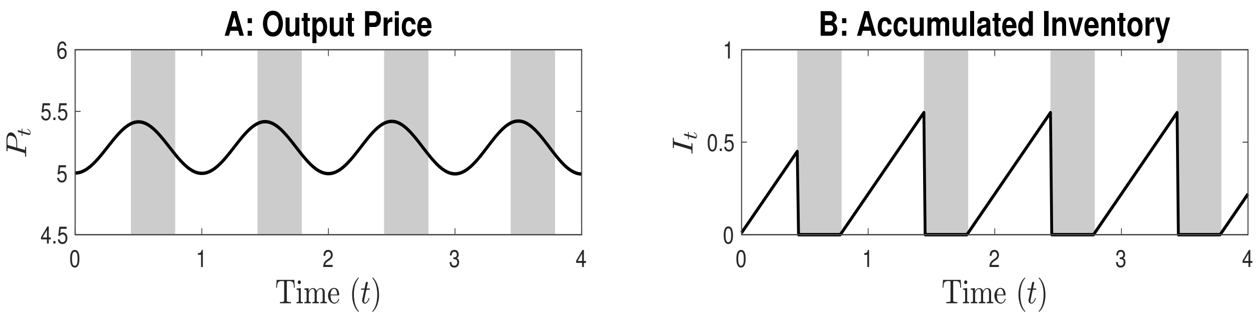

Figure 2: The figure plots the output price P_t (Panel A) and the firm's inventory holdings I_t (Panel B) over the period from $t = 0$ to 4. While the white areas in each subpanel are optimal inventory building periods, the gray-shaded areas are optimal instantaneous selling periods. We describe the parameter values in Section 2.1.2.

no ($S_{\tau^*} = 0$) or all available ($S_{\tau^*} = \bar{K}$) output at future time τ^* ("bang-bang solution;" see, e.g., Majd and Pindyck (1987) and Dixit and Pindyck (1994) for other examples).

To identify the optimal selling date τ^* for concurrent production, the firm can maximize the present value from selling one single output unit at future time u. More rigorously, the firm can solve the following constrained maximization problem:

$$\max_{u\in[t,T]} \mathbb{E}_t^{\mathbb{Q}}[P_u]e^{-r(u-t)} = \max_{u\in[t,T]} P_t \exp\left(-(\mu-\alpha)(u-t) + \frac{\kappa}{\eta}\left(\cos(\eta t) - \cos(\eta u)\right)\right), \tag{4}$$

where $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ is the time-t expectation under the equivalent martingale measure, r is the risk-free rate, and μ is the constant expected return of a mimicking portfolio perfectly correlated with the output price. For future reference, we define $\mathscr{E}_{t,u} \equiv \exp\left(-(\mu-\alpha)(u-t) + \frac{\kappa}{\eta}\left(\cos(\eta t) - \cos(\eta u)\right)\right)$ for $u > t$ and else zero. Taking the partial derivative of Equation (4) with respect to u and setting to zero, we find the local maximums of the objective function to be at $\tau^* = \frac{\pi}{\eta} - \frac{1}{\eta}\sin^{-1}\left(\frac{\mu-\alpha}{\kappa}\right) + \frac{2\pi}{\eta}n$ for $n \in \{\ldots,-2,-1,0,1,2,\ldots\}$. As the local maximums decline with time (due to discounting), we then pick as global maximum the value closest to but larger than t. We stress that in the absence of output price seasonality (i.e., when $\kappa = 0$ and/or $\eta = 0$), we would always find $\tau^* = t$, so that, as in other models, the firm would always instantaneously sell its concurrent production.

Figure 2 aids in better understanding the firm's optimal selling policy. In Panel A, we plot an output price path which is identical to its expectation, to better highlight the implications of seasonality. We rely on the same stochastic process parameter values as in Section 2.1.1, and we assume an expected output-price mimicking portfolio return (μ) of 8% per annum. The white areas indicate periods in which the firm optimally delays selling its output ($\tau^* > t$), whereas the gray areas indicate periods in which it optimally instantaneously sells its output ($\tau^* = t$). In turn, Panel B plots the firm's inventory holdings, showing that those holdings linearly rise over inventory building periods and remain at zero over instantaneous selling periods. Intuitively, the firm builds up output inventories over inventory building periods because the present value of selling output at the next seasonal high exceeds the current output price. Conversely, the firm instantaneously sells it concurrently produced output over some period after the prior seasonal high since the current output price exceeds that same present value.

2.1.3 Firm Valuation and Expected Firm Return

We now derive the firm's value and expected excess return. Starting with the firm's value, W_t, we first use an integral to compute the present value of output produced before the current time t but still optimally held in inventory at that same time, W_t^i, as:

$$W_t^i = \int_0^t P_t \mathscr{E}_{t,\tau_s^*} \bar{K} \, \mathrm{d}s = P_t \bar{K} \int_0^t \mathscr{E}_{t,\tau_s^*} \mathrm{d}s, \tag{5}$$

where τ_s^* is the optimal selling time chosen at time $s < t$, and we again impose $\mathscr{E}_{t,\tau_s^*} = 0$ if $\tau_s^* < t$ to ensure that we do not consider output already sold in the past. We next use another integral to compute the present value of output produced after time t, W_t^p:

$$W_t^p = \int_t^T P_t \mathcal{E}_{t,\tau_s^*} \bar{K} \mathrm{d}s = P_t \bar{K} \int_t^T \mathcal{E}_{t,\tau_s^*} \mathrm{d}s, \tag{6}$$

where τ_s^* is now the firm's expectation of the optimal selling time for output produced at future time $s \geq t$, determined from the law of iterated expectations.[7] The present value of the firm's future fixed production costs, C_t^p, is equal to:

$$C_t^p = \int_t^T f \bar{K} e^{-r(s-t)} \mathrm{d}s = \frac{f \bar{K}}{r} \left(1 - e^{-r(T-t)}\right). \tag{7}$$

Proposition 1 finally combines the three value components into the total value of the firm:

Proposition 1. *The firm's value at time t, $W_t = W(t, P_t)$, is:*

$$W_t = W_t^i + W_t^p - C_t^p = P_t \bar{K} \int_0^T \mathcal{E}_{t,\tau_u^*} \mathrm{d}u - \frac{f \bar{K}}{r} \left(1 - e^{-r(T-t)}\right). \tag{8}$$

We next compute the firm's conditional expected excess return. Since the output price, P_t, is the only priced stochastic variable, that expected return is linear in the firm's elasticity (i.e., the partial derivative of firm value with respect to output price times the ratio of output price to firm value), as in other asset pricing models in the literature. Proposition 2 formalizes this insight and also offers a linear decomposition of the firm's elasticity:

Proposition 2. *The firm's conditional expected excess return per time unit is:*

$$\mathbb{E}[r_W] - r = \beta_t (\mu - r), \tag{9}$$

[7] In words, the firm's best estimate at current time t of its best estimate at future time $s \geq t$ of the present value of selling an output unit net of inventory holding costs at some later time $u \geq s$ is simply its best estimate of that net present value at time t. Given that, the firm can simply maximize over the estimate at time t.

where the elasticity $\beta_t = \frac{\partial W_t / W_t}{\partial P_t / P_t}$ *captures the responsiveness of the firm's value to its output price, and* μ *is the constant expected return of a mimicking portfolio whose price changes perfectly correlate with output price changes. The elasticity* β_t *can be written as:*

$$\beta_t = 1 + \frac{C_t^p}{W_t} = 1 + \frac{C_t^p}{W_t^i + W_t^p - C_t^p}. \tag{10}$$

Equation (10) suggests that the firm's elasticity β_t (and thus its expected excess return) vary seasonally over time. To understand that, we first notice that C_t^p, the present value of future production costs, is a deterministically declining function over time. We next realize that both W_t^i and W_t^p, the present values of selling already-produced and still-to-be-produced output, respectively, rise with the firm approaching its next high-output-price season. The reasons are that, first, the firm's future sales revenues are less strongly discounted and, second, that the firm holds more output in inventory the closer it is to that season. The upshot is that the firm's optimal inventory building toward its next high-output-price season gradually *levers down* (or: delevers) the firm, inducing its expected excess return to fall over that period. In contrast, the static operating leverage coming from quasi-fixed production costs in, for example, Carlson et al. (2004) and Cooper (2006) *lever up* the firm, simply because the present value of those costs are subtracted from (and not added to) the firm's value (see again Equation (10)).

2.2 The Real Options Model

We next introduce the real options model. The real options model innovates over the toy model by, first, assuming that the firm incurs a cost from holding inventories and, second, by no longer requiring the firm to precommit on the time at which it will sell its concurrently produced output. Rather, the firm now owns the real option to optimally update that time.

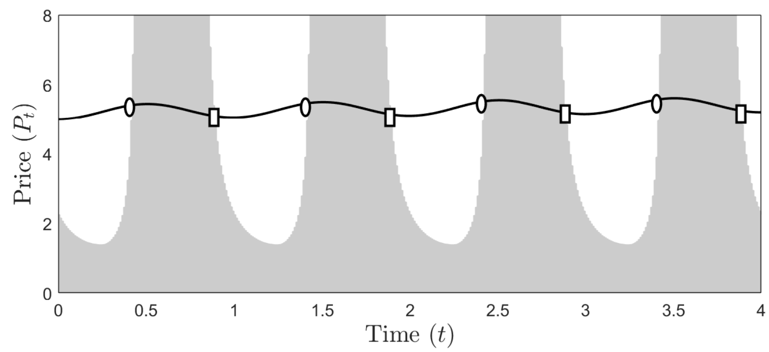

Figure 3: The figure plots the firm's optimal inventory building (white areas) and instantaneous selling (gray areas) regions on the output price P_t-time t plane. We describe the parameter values used to create the figure in the second paragraph of Section 2.1.1 and in the last paragraph of Section 2.1.2.

2.2.1 Optimal Selling Policy and Firm Valuation

We start with valuing the output held in inventory. As each output increment in inventory is again (just like in the toy model) identical to all others, we can find the value of an arbitrary output unit, $V(t, P_t)$, from the following standard Bellman equation:

$$V(t, P_t) = -c_I \mathrm{d}t + \mathbb{E}_t^{\mathbb{Q}}[V(t + \mathrm{d}t, P_t + \mathrm{d}P_t)]\, e^{-r\mathrm{d}t}, \tag{11}$$

where $c_I > 0$ is the cost of holding one output unit in inventory for one time unit.

Using Itô's Lemma, taking the equivalent martingale measure expectation, and re-arranging, we can rewrite Equation (11) as the partial differential equation (PDE):

$$\frac{1}{2}\sigma^2 P^2 \frac{\partial^2 V}{\partial P^2} + (r - \delta_t) P \frac{\partial V}{\partial P} - rV + \frac{\partial V}{\partial t} - c_I = 0, \tag{12}$$

which the output unit value, $V(t, P_t)$, must obey subject to the boundary condition that $V(t, P_t) = P_t$ at the optimal selling time. Due to the early exercise decision, we cannot solve PDE (12) in closed-form and resort to a finite difference approach. See Appendix A.2 for details.

Using the same parameter values as in Sections 2.1.1, 2.1.2, and 2.1.3 and assuming $c_I = 0.25$, Figure 3 displays the firm's optimal inventory building and selling policy in output price P_t-time t space. As long as the output price-time combination is within one of the white areas in the figure, the firm optimally produces to exclusively build up its output inventories. Upon crossing into a gray area (which happens at the circles for the exemplary sample path), the firm then sells off its entire output in inventory and starts producing to instantaneously sell until it crosses back into a white area (which happens at the rectangles). As in the toy model, the firm thus again follows a bang-bang strategy. Interestingly, if the output price becomes sufficiently low (the threshold is about 1.7 in the figure), it is never optimal for the firm to pay inventory holding costs, and it thus consistently sells all of its concurrent production instantaneously.

Relying on the value of a single output unit, $V(t,P_t)$, Proposition 3 offers the real-option-model implied value of a firm with I_t output units in inventory at time t:

Proposition 3. *The firm's value at time t, $W_t = W(t,P_t,I_t)$, is:*

$$W_t = W_t^i + W_t^p - C_t^p = I_t V(t,P_t) + \bar{K}\int_t^T e^{-r(u-t)}\mathbb{E}_t^{\mathbb{Q}}[V(u,P_u)]\mathrm{d}u - \frac{f\bar{K}}{r}\left(1-e^{-r(T-t)}\right), \tag{13}$$

where $W_t^i = I_t V(t,P_t)$, $W_t^p = \bar{K}\int_t^T e^{-r(u-t)}\mathbb{E}_t^{\mathbb{Q}}[V(u,P_u)]\mathrm{d}u$, C_t^p is defined as before, and the future output unit values $V(u,P_u)$ are given by the finite difference grid.

As the distribution of P_t under the equivalent martingale measure is log-normal, we can use the law of the unconscious statistician to (numerically) compute $\mathbb{E}_t^{\mathbb{Q}}[V(u,P_u)]$.

Finally, we can again use Equation (9) to compute the firm's expected return.

2.3 Model Conclusions

We finally use a simulation exercise to determine the main corporate finance and asset pricing implications of our two models. Since the models, however, produce similar implications, we exclusively focus on the real options model in this section, only drawing on the toy model to interpret our results. Notwithstanding, we offer the results from the toy model in Appendix B at the end of this paper. In each simulation, we look into the period from $t = 0$ to 4. To bring the effects of seasonality into sharper focus, we consistently let the output price evolve according to its drift term (i.e., the solid black line in Figure 1).[8] We use the same basecase parameters as in Figures 1 to 3. More specifically, we choose $P_0 = 5.00$, $\alpha = 0.01$, $\kappa = 0.25$, $\eta = 2\pi$, $\sigma = 0.05$, $\mu = 0.08$, $r = 0.01$, $c_I = 0.25$, $\bar{K} = 1.00$, $f = 1.00$, and $T = 40$. To study how the inventory holding costs modulate our conclusions, we later also consider $c_I = 0.00$, 0.66, and $+\infty$.

In Figure 4, we plot the evolution of the output price P_t (Panel A; repeated for convenience), the amount of output the firm holds in inventory I_t (Panel B), the firm's inventory holdings value W_t^i (Panel C), its continuation value W_t^p (Panel D), its total value W_t (Panel E), and its expected excess return $\mathbb{E}[r_W] - r$ (Panel F) over time. The white (gray) shaded areas indicate optimal inventory-building (instantaneous selling) periods. Spurred by the seasonal output price in Panel A, Panel B demonstrates that the firm produces to exclusively build up output inventories over some period before its optimal selling time τ^* (the start of each gray-shaded area). Reaching time τ^*, the firm sells off its entire output in inventory and starts producing to exclusively instantaneously sell over some period. As the firm keeps the value-added generated over inventory-building periods in-house, Panels C and D establish that the value of its inventory holdings W_t^i — but not its continuation value W_t^p — markedly rises over such periods. In turn,

[8]To be clear, while the firm expects the output price to evolve stochastically according to the seasonal GBM (1), it is just a "coincidence" that the actual realization is in agreement with the expectation. We rely on that coincidence to more clearly establish how the firm responds to seasonality in its output price.

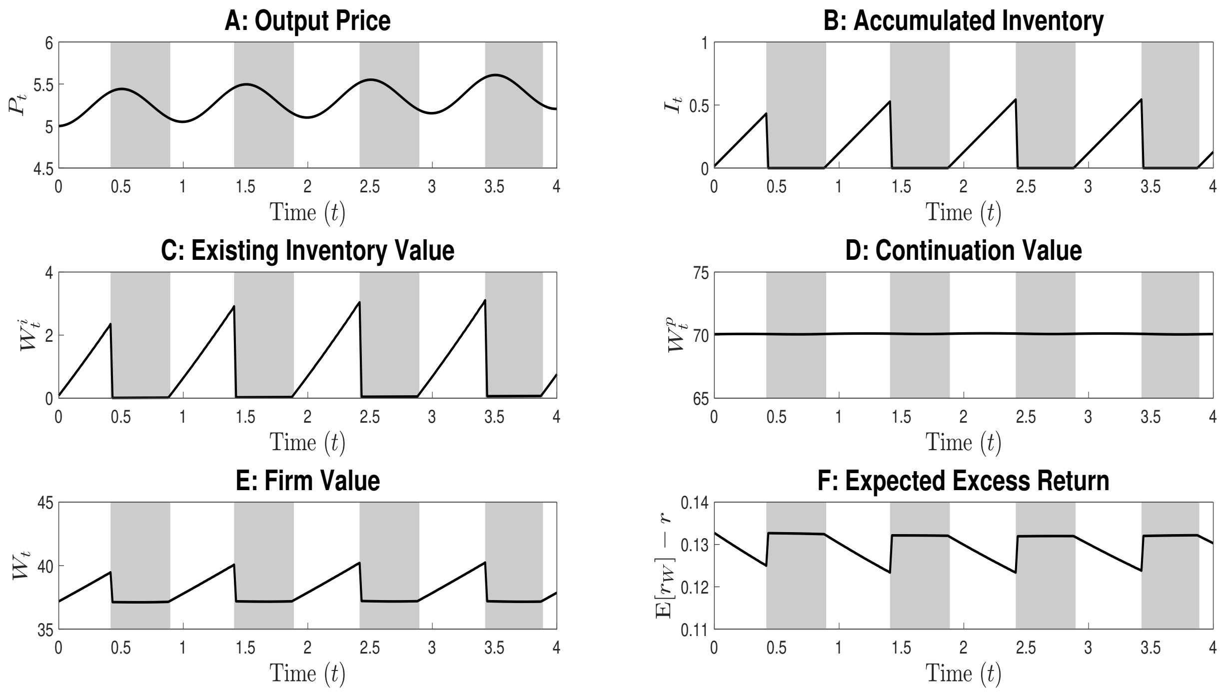

Figure 4: The figure shows the results from our benchmark simulation exercise. While Panel A displays the evolution of the output price, P_t, over the period from $t = 0$ to 4, Panels B to F display how the firm's inventory holdings I_t, inventory holdings value W_t^i, continuation value W_t^p, total value W_t, and conditional expected excess return, $\mathbb{E}[r_W] - r$, vary over that period, respectively. To focus on seasonal patterns, we remove a linear trend from the inventory holding, continuation, and total firm value as well as the conditional expected excess return. While the white areas in each subpanel are optimal inventory-building periods, the gray-shaded areas are optimal instantaneous-selling periods. We describe the parameter values in the first paragraph of Section 2.3.

Panel E reveals that the seasonality in the firm's inventory holdings value maps directly into similar seasonality in its total value W_t. Finally, Panel F confirms that the firm's value increases induced through its optimal inventory building delever it over those same periods, gradually lowering its expected return toward time τ^*. As suggested by Proposition 2, the delevering over the optimal inventory holding periods occurs because the firm's elasticity and thus expected return are inversely related to the present value of its inventory holdings (W_t^i).

Our simulation results so far insinuate that optimal inventory building delevers the firm, yielding endogenous seasonality in the firm's expected return inversely related to the exogenous seasonality in its output price. To further support the claim that optimal inventory building is

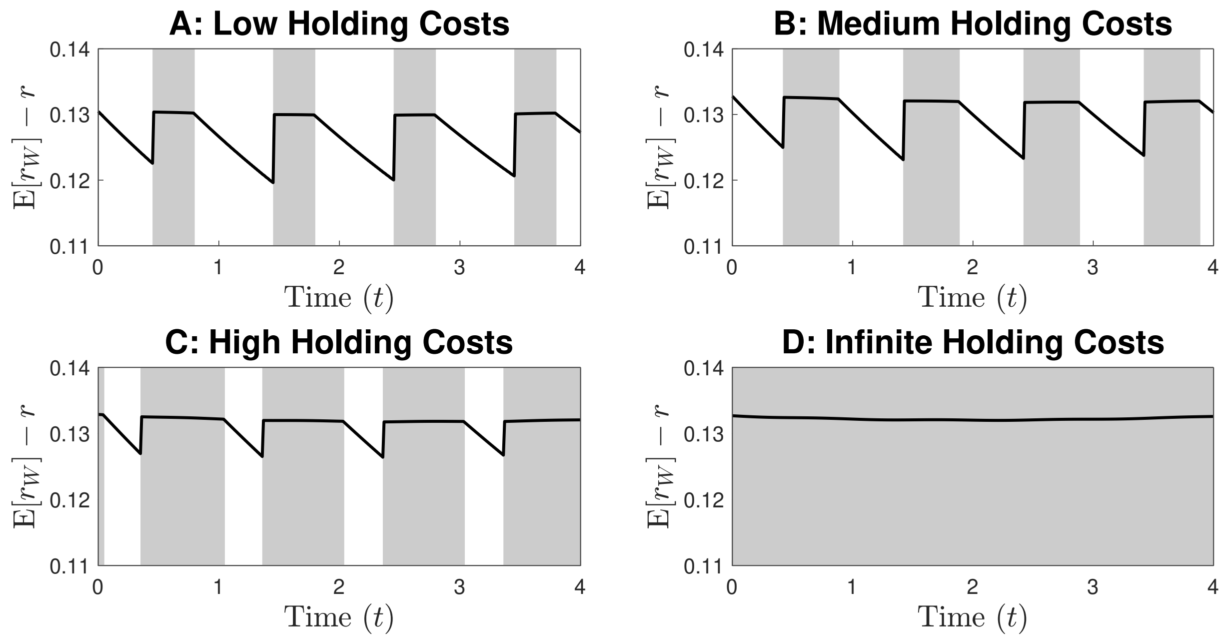

Figure 5: The figure gives comparative statics for the expected excess return. While Panel A displays the evolution of that expected return under an inventory holding cost, c_I, of 0.00, Panels B to D display it under a cost of 0.25, 0.66, and $+\infty$, respectively. To focus on seasonal patterns, we remove a linear trend from the expected returns. We describe the parameter values in the first paragraph of Section 2.3.

the mechanism underlying those conclusions, we next repeat the simulations varying the ease with which the firm can build up inventories by setting the inventory holding cost, c_I, to 0.00, 0.25 (for comparison), 0.66, or $+\infty$. Using those values in Panels A to D, respectively, Figure 5 plots the evolution of the expected excess return from time $t = 0$ to 4. Contrasting the panels, the figure confirms that lower inventory holding costs raise (lower) the duration of the inventory building (instantaneous selling) periods, in turn amplifying the endogenous seasonality in the expected return. Crucially, when the firm finds it prohibitively costly to engage in any inventory building, the expected return does not display seasonalities at all (see Panel D).

It is instructive to consider a version of our toy model in which the firm must instantaneously sell its output to better understand why optimal inventory building is a necessary condition for output price seasonality to translate into inverse expected return seasonality. In that version, the present value of the firm's futures sales revenue, W_t^p, is equal to:

$$W_t^p = P_t \bar{K} \int_t^T e^{-(\mu-\alpha)(u-t)+\frac{\kappa}{\eta}(\cos(\eta t)-\cos(\eta u))} \mathrm{d}u. \tag{14}$$

While it is obvious that P_t evolves seasonally according to GBM (1), the $\frac{\kappa}{\eta}(\cos(\eta t) - \cos(\eta u))$ term inside the integral evolves exactly inverse-seasonally, reaching its seasonal high when P_t reaches its seasonal low, and vice versa. The intuition is that when P_t is at its seasonal high (low) its expected seasonal drift is lowest (highest). In turn, the present value W_t^p displays almost no seasonalities and, since the elasticity β_t continues to be inversely linear in $W_t = W_t^p - C_t^p$ (as in Proposition 2), neither does the expected excess return of the firm, $\mathbb{E}[r_W] - r$.

Overall, this section offers theoretical evidence that in a world with seasonal exogenous output prices and inventory building, firms with lower inventory holding costs optimally build up output inventories over longer periods before their high-price seasons. As a result, they also produce stronger inverse endogenous expected-return seasonalities, induced through their inventory holdings gradually delevering them over their inventory building periods.

3 Main Empirical Evidence

In this section, we empirically evaluate the main predictions of our theory in Section 2. To do so, we first introduce our variable definitions. We next state our data sources and offer descriptive statistics. We then test the prediction that seasonal firms with low inventory holding costs build up output inventories toward their high-demand season. We finally test the prediction that high seasonal inventory leverage predicts high future stock returns in the cross-section.

3.1 Main Variable Definitions

Our theoretical work suggests that firms with seasonal demand and low inventory holding costs actively build up output inventories toward their high-demand seasons, gradually delevering themselves and inducing their expected returns to fall. To capture seasonal swings in a firm's inventory holdings, we compute abnormal quarterly inventory holdings, $QInv_{i,q,y}$, as:

$$QInv_{i,q,y} = \sum_{j=2}^{3} \left(\frac{QuarterlyInventory_{i,q-1,y-j}}{MeanQuarterlyInventory_{i,y-j}} \right) /2, \tag{15}$$

where $QuarterlyInventory_{i,q,y}$ are the total inventory holdings of firm i at the end of quarter q in fiscal year y and $MeanQuarterlyInventory_{i,y}$ are its average quarterly total inventory holdings over fiscal year y. Intuitively, *QInv* captures the mean deviation of the firm's total inventory holdings at the end of the prior quarter (and thus at the start of the current) from its average ("normal") holdings over the fiscal year, calculated over the fiscal years ending in calendar years $t-1$ and $t-2$. As a result, *QInv* is an inverse proxy for seasonal inventory leverage.[9]

In addition to *QInv*, we also calculate an equivalent variable, *QFWInv*, using the sum of work-in-progress and finished-good (rather than total) inventories. While *QFWInv* is only available for a small subset of our sample data, we use it to verify *QInv*. Furthermore, since *QInv* captures only the magnitude of a firm's inventory holdings at the start of some quarter relative to its normal holdings but not a firm's general propensity to store output in inventory, we also compute *InvToSales*, defined as the average fraction of the quarterly inventories-to-quarterly sales ratio over the fiscal years ending in calendar years $t-1$ and $t-2$.

[9] In line with our theory, our empirical work assumes that a firm's output-price seasonality and ability to build up output inventories are highly persistent, so that past values of *QInv* can be used to predict future values. In Section IA.2 in the Internet Appendix, we offer evidence verifying that assumption.

As our theory predicts that its main implications are stronger for seasonal firms (i.e., firms with higher κ values in Section 2), we follow Grullon et al.'s (2020) strategy to distinguish between seasonal and non-seasonal firms. In particular, we compute $QSales_{i,y,q}$, the average historical contribution of a firm's sales over the current quarter to its total annual sales, as:

$$QSales_{i,q,y} = \sum_{j=2}^{3} \left(\frac{QuarterlySales_{i,q,y-j}}{AnnualSales_{i,y-j}} \right) / 2, \tag{16}$$

where $QuarterlySales_{i,q,y}$ are the quarterly sales of firm i over quarter q in fiscal year y and $AnnualSales_{i,y}$ are the annual sales of that firm over fiscal year y. We subsequently calculate the standard deviation of *QSales* over the fiscal year ("*Seasonality*"), classifying firms with a value above (below) the cross-sectional median as seasonal (non-seasonal) firms.[10]

In our supplementary industry-level analysis, we calculate our variables in the same spirit as in Equations (15) and (16), except that we compute contemporaneous (and not historical average) fractions at the monthly (and not quarterly) frequency. Using an industry's new orders as example for a flow variable, we compute $MNewOrders_{i,m,y}$ as:

$$MNewOrders_{i,m,y} = \frac{MonthlyNewOrders_{i,m,y}}{AnnualNewOrders_{i,y}}, \tag{17}$$

where $MonthlyNewOrders_{i,m,y}$ are the new orders of industry i in month m of calendar year y and $AnnualNewOrders_{i,y}$ are the total new orders of that industry over year y. We calculate the flow variables derived from either monthly shipment values (*MShipments*) or monthly unfilled orders (*MUnfilledOrders*) in an analogous manner. Conversely, using total inventories

[10] Our variable definitions ensure that real investors could have calculated the variables in real-time. We form historical averages over two fiscal years to mitigate noise. Taking averages over longer past periods does not change our results. Neither does excluding fiscal years for which quarterly sales do not add up to annual sales.

as example for a stock variable, we compute $MInv_{i,m,y}$ as:

$$MInv_{i,m,y} = \frac{MonthlyInventories_{i,m,y}}{MeanMonthlyInventories_{i,y}}, \tag{18}$$

where $MonthlyInventories_{i,m,y}$ are the total inventory holdings of industry i at the end of month m in calendar year y and $MeanMonthlyInventories_{i,y}$ are the average monthly total inventory holdings of that industry over year y. We finally also calculate three more granular abnormal inventory holding variables, in particular, raw material (*MRMInv*), work-in-progress (*MWIPInv*) and finished-goods (*MFGInv*) inventory holdings, in an analogous manner.

We use the following standard control variables. In our single-firm panel regressions, we control for *Assets*, *QROE*, and *Investment*. Conversely, our portfolio sorts control for either Fama and French's (2016) six-factor-model or Hou et al.'s (2021) augmented q-theory factors. Finally, our single-stock FM regressions control for *QCash*, *QDebt*, *InvToSales*, *MarketBeta*, *MarketSize*, *BookToMarket*, *Momentum*, *Investment*, and *Profitability*. We offer more details about the calculation of our variables in Table C1 in Appendix C at the end of this paper.

3.2 Data Sources and Descriptive Statistics

In our firm-level analysis, we use market data from CRSP, accounting data from Compustat, and factor model data from Ken French's and Lu Zhang's websites. We study common stocks traded on the NYSE, Amex, or NASDAQ excluding financial and utility stocks. We replace a stock's return with its delisting return when the delisting return is available. We exclude observations for which quarterly sales and/or quarterly inventory holdings are negative as well as those for which these variables are not available for the entire fiscal year. In our FM regressions, we further exclude stocks with a market capitalization below the first NYSE quintile or $5 at the end of the prior month. We winsorize all variables except the stock return at the 0.5th and 99.5th percentiles per month. Our sample period is January 1979 to December 2022.[11]

Panel A of Table 1 gives descriptive statistics for *QInv*, *QSales*, *Seasonality*, *InvToSales*, and *QFWInv*. Except for the number of observations, we calculate each statistic first by sample month and then average over time. The panel shows that our main variables are available for more than one million observations (see column (1)). The only exception is *QFWInv*, our abnormal inventory variable computed from finished-good and work-in-progress inventories, which is available for only about 180k observations. The panel further reveals that while *QInv* (*QSales*) necessarily produces a mean value of 1.00 (0.25; see column (2)), *QInv* is about 3-4 times more volatile than *QSales* (see column (3)). Also strikingly, while, in contrast to *QFWInv*, *QInv* also includes raw-material inventories, the two variables are extremely similar, not only in terms of their moments but also percentiles. Looking into the *Seasonality* and *InvToSales* statistics, we find that firms vary markedly in the extent to which single quarters make up the lion share of their annual sales and the extent to which they hold output in inventory.

TABLE 1 ABOUT HERE.

In the industry-level analysis, we use *non-seasonally-adjusted* data on the values of orders, shipments, and inventory holdings of seven major industries from the Manufacturers' Shipments, Inventories, and Orders (MSOI) database made available by the Census Bureau of the United States.[12] Our sample period is January 1993 to December 2021.

Panel B of Table 1 gives descriptive statistics for the industry-level variables. The panel shows that these variables are available for about 2,500 observations. It also reveals that the aggregation

[11]The start of our sample period is dictated by the availability of quarterly inventory data in Compustat.

[12]To be specific, we look into the primary metals (31S), fabricated metal products (32S), machinery (33S), computer and electronics (34S), electrical equipment, appliances, and components (35S), transportation equipment (36S), and furniture and related products (37S) industries. The reason is that the three more granular (raw material, work-in-progress, and finished-good) inventory variables are only available for those industries.

to the industry level slightly attenuates the seasonalities in them (with, e.g., the volatility of *MNewOrders* being less than one-third of the volatility of *QSales* (see column (3)).

3.3 Seasonal Demand and Inventory Building Policies

We first test our theory's prediction that firms with seasonal demand and low inventory holding costs optimally build up output inventories toward their high-demand season, to better serve demand in that season. We start off with using industry-level data from the MSOI database to test that prediction. Two key advantages of the MSOI data are that they allow us to (i) directly observe demand in the form of new orders, and to (ii) decompose total inventory holdings into raw material, work-in-progress, and finished-good holdings. To be more specific, we conduct the following panel regression with fixed effects on the industry-level data:

$$MVariable_{i,m,y} = \boldsymbol{\beta}'\mathbf{MVariables} + \alpha_i + \epsilon_{i,m,y}, \tag{19}$$

where $MVariable \in \{MNewOrders,\ MShipmentValues,\ MUnfilledOrders\}$, **MVariables** is a vector of contemporaneous and/or lagged values of *MNewOrders*, *MInv*, *MFGInv*, *MWIPInv*, and *MRMInv*, $\boldsymbol{\beta}$ is a parameter vector, α is an industry fixed effect, and ϵ is the residual. We cluster the standard errors from the panel regression at the major-industry level.[13,14]

[13]We do not include time fixed effects since if the seven industries received most of their new orders in the same or similar calendar months, then the time fixed effects would subsume the seasonalities in demand and the other variables. Notwithstanding, the inclusion of time fixed effects only mildly changes our results.

[14]We can show that whenever we regress *MNewOrders*, *MShipments*, or *MUnfilledOrders* on one of the inventory variables (*MInv*, *MFGInv*, *MWIPInv*, or *MRMInv*), we would find slope coefficients of 1/12 on those variables if firms preproduced and held in inventory at the start of the relevant month *all* output ordered, shipped, or unfilled over month t (and sales revenues were a constant multiple of historical costs). The technical reason is that *MNewOrders*, *MShipments*, and *MUnfilledOrders* are scaled by their annual sums, whereas the inventory variables are scaled by their annual means. To ease interpretation, we thus multiply the slope coefficients of the inventory variables by twelve, implying that a coefficient of one signals "perfect inventory building."

Table 2 reports the results from estimating regression (19). The plain numbers are coefficient estimates, whereas those in square brackets are t-statistics. While the dependent variable is *MNewOrders* in columns (1) and (2), it is *MUnfilledOrders* in column (5), and *MShipments* in the remaining columns. Moreover, while "lag x" after an independent variable's name indicates that the variable is lagged by x months, "lag x-y" indicates that the variable is the sum of its lagged values over lags x to y. The first two columns confirm that the share of an industry's annual new orders attributable to the current month is highly predictable and seasonal. To be specific, while column (1) shows that the shares over the past twelve months capture 46% of the variations in the share of the current month, column (2) reveals that the twelve-month-ago *same-calendar-month* share captures 43% on its own (estimate: 0.65, t-statistic: 6.14). The upshot is that if an industry received a disproportionally small (large) number of new orders in some month over the previous year, it is highly likely that it will also receive a disproportionally small (large) number of new orders in that same month over the current year.

TABLE 2 ABOUT HERE.

Notwithstanding, columns (3) to (5) suggest that the surveyed firms deal swiftly with seasonal swings in demand. To wit, a higher new-order share in the current month translates into a significantly higher shipment share in the same month (estimate: 0.69; t-statistic: 11.46) in column (3) but not in later months in column (4). Moreover, that same higher new-order share has a much smaller effect on the same-month unfilled-order share (estimate: 0.13; t-statistic: 2.23) in column (5). Finally, the remaining columns indicate that inventory building is at least partially responsible for the swift reactions. In particular, they show that the surveyed firms hold larger (smaller) than usual inventories at the start of their high (low) shipment-share months (estimate: 0.93; t-statistic: 3.66) in column (6) but not one month earlier or later in column (7). Decomposing total inventory holdings into raw material, work-in-progress, and finished-good holdings, column (8) reveals that the higher than usual inventories at the start of high shipment months mostly come from higher finished-good (estimate: 0.69; t-statistic: 4.53) and work-in-progress (estimate: 0.58; t-statistic: 3.87) inventories but, in an important contrast, smaller raw material (estimate: –0.33; t-statistic: –2.60) inventories.

While the MSOI survey data support several important predictions of our theory, they do not allow us to test the predictions that (i) seasonal firms build up more output inventories toward their high-demand seasons than non-seasonal firms, and that (ii) inventory holding costs condition the ability to build up output inventories toward those seasons. The reason is that the survey firms all manufacture relatively homogeneous physical goods (as, e.g., metal, machinery, electrical equipment, etc.) with a likely similar demand seasonality. Given that, we now also conduct a cruder analysis based on firm-level Compustat data over a greater variety of industries. To be specific, we execute the following panel regressions with fixed effects:

$$QSales_{i,q,y} = \beta\, QInv_{i,q,y} + \boldsymbol{\eta}'\mathbf{Control}_{i,q,y} + \alpha_i + \epsilon_{i,q,y}, \tag{20}$$

and

$$\begin{aligned} QSales_{i,q,y} \quad = \quad & \beta\, QInv_{i,q,y} + \gamma(QInv_{i,q,y} \times DSeasonality_{i,q,y}) \\ & + \delta\, DSeasonality_{i,q,y} + \boldsymbol{\eta}'\mathbf{Control}_{i,q,y} + \alpha_i + \epsilon_{i,q,y}, \end{aligned} \tag{21}$$

where **Control** = [*Assets, QROE, Investments*]′, *DSeasonality* is a dummy variable equal to one of *Seasonality* is above its median and else zero, β, γ, and δ are parameters, $\boldsymbol{\eta}$ is a parameter vector, α_i is a firm-fixed effect, and $\epsilon_{i,q,y}$ is the residual. We cluster standard errors at the firm level.[15]

[15] Consistent with footnote (14), we multiply the slope coefficients of the inventory, cash, and short-term debt variables and their interactions in Tables 3, 4, and 5 with four, so that a unit coefficient indicates that firms perfectly build up inventories and short-term debt and run down cash toward their high-sales quarters.

Table 3 reports the results from estimating regressions (20) and (21) in the odd and even-numbered columns, respectively. While we run the estimations in columns (1) and (2) over our full sample, we run those in columns (3) and (4), (5) and (6), (7) and (8), and (9) and (10) over the durable consumer goods, manufacturing, services, and healthcare industry subsamples, respectively.[16] Plain numbers are estimates and those in square brackets t-statistics. Column (1) shows that the average sample firm builds up inventories toward its high-sales quarter (*QInv* estimate: 0.15; t-statistic: 17.66). Crucially, however, column (2) demonstrates that the tendency to build up inventories is almost exclusively driven by seasonal firms. Specifically, while non-seasonal firms yield a *QInv* estimate of 0.00 (t-statistic: 2.41), the corresponding number for seasonal firms is 0.21 (t-statistic of interaction: 17.26). Also reassuringly, columns (3) to (10) demonstrate that seasonal firms with plausibly lower inventory holding costs (as those from the durable-good or manufacturing industry) build up more inventories than those seasonal firms with plausibly higher such costs (as those from the service or healthcare industry).

TABLE 3 ABOUT HERE.

In Table 4, we verify that our evidence in Table 3 is robust to us using *QFWInv* (i.e., abnormal work-in-progress plus finished-good inventories) rather than *QInv* (i.e., abnormal total inventories). Doing so is important since our theory in Section 2 makes predictions about finished-good (and not raw-material) inventories. To that end, the tables gives the results from reestimating regressions (20) and (21) on the full sample (columns (1) and (2); for comparison); on the sub-sample of firms with non-missing work-in-progress and finished-good inventory values ((3) and (4)); and on that subsample but with *QInv* replaced with *QFWInv* ((5) and (6)). Contrasting the columns, we find that *QInv* yields results close-to-identical to *QFWInv*. While the estimate

[16] We do not focus on a comprehensive but rather a highly selective sample of industries with plausibly either relatively low (the former two industries) or high (the latter two) inventory holding costs.

of *QInv* is, for example, 0.22 (*t*-statistic: 9.29) in regression (20) in column (3), the corresponding estimate of *QFWInv* in column (5) is a similar 0.18 (*t*-statistic: 9.82).

TABLE 4 ABOUT HERE.

We finally also examine how seasonal firms finance building up output inventories and when they distribute sales revenues to equityholders. While our theory in Section 2 assumes that those firms continuously raise external equity for financing purposes, and that they instantaneously distribute sales revenues to equityholders, it is likely that our sample firms rely on a greater mix of financing sources and smooth their payouts over time. To study firms' financing sources, we use abnormal quarterly cash holdings, *QCash*, and abnormal quarterly short-term debt holdings, *QDebt*, both computed analogously to *QInv* but with total inventory holdings replaced with cash reserves and current liabilities, respectively.[17] To study their payout policies, we look into abnormal quarterly dividends, *QDiv*, computed analogously to *QSales* but with sales replaced with dividends. See Table C1 in Appendix C at the end of this paper for more details.

Table 5 reports the results from reestimating regressions (20) and (21) using *QCash* (columns (1) to (2)), *QDebt* ((3) to (4)), or *QDiv* ((5) to (6)) as dependent variable. While the odd-numbered columns focus on regression (20), the even-numbered concentrate on regression (21). Moreover, plain numbers are again estimates and those in square brackets *t*-statistics. The table suggests that firms both deplete their cash reserves and take out short-term debt to finance building up output inventories. While column (1), for example, shows that the average sample firm holds smaller cash reserves before a high sales quarter (*t*-statistic: –24.90), column (2) reveals that this effect is much stronger for seasonal firms (*t*-statistic of interaction: –19.87). The table further

[17] We notice that we focus on all current liabilities of a firm (rather than only its short-term debt) since the firm could also finance the build-up of output inventories using trade credit, payable wages and salaries, or deferred taxes. Despite that, we find similar results when we exclusively look into short-term debt financing.

indicates that firms raise their dividends in high sales quarters (t-statistic: 11.51), with the effect again much stronger for seasonal firms (t-statistic of interaction: 15.19).

TABLE 5 ABOUT HERE.

Overall, this section uses both granular industry as well as coarser firm-level data to support our theory's prediction that seasonal firms build up output inventories toward their high demand seasons, to better accommodate demand in those seasons.

3.4 The Pricing of Seasonal Inventory Leverage

As Section 3.3 shows that our sample firms build up output inventories toward their high-demand seasons, and as our theory suggests that doing so delevers firms, we would now naturally expect that seasonal inventory leverage is positively related to the cross-section of future stock returns. To test this prediction, we start off with univariate portfolio sorts. At the end of each month $t-1$, we thus sort our sample stocks into portfolios according to the terciles of the *QInv* distribution at that time. We value-weight the portfolios and form a spread portfolio long the highest *QInv* portfolio and short the lowest. We hold all portfolios over month t. Alternatively, we first split our sample stocks according to the median of the *Seasonality* distribution at the end of each month $t-1$, referring to above (below) median stocks as seasonal (non-seasonal) stocks. We then independently sort into *QInv* portfolios just like before, using the intersection of the two sorts to form double-sorted portfolios. To adjust for systematic risk, we regress the spread portfolio returns on the Fama and French (2016) and Hou et al. (2021) factors.

Table 6 presents the univariate portfolio sort results. Panels A, B, and C focus on the full sample and the seasonal and non-seasonal stock subsamples, respectively. Plain numbers are monthly excess returns (in %; "Return^e") or portfolio characteristics, whereas the numbers in square brackets are Newey and West (1987) *t*-statistics with a twelve-month lag length. The portfolio characteristics are the average number of stocks, the time-series means of the cross-sectional *QSales* and *MarketSize* means per portfolio, and the portfolio Sharpe ratio.

TABLE 6 ABOUT HERE.

The table confirms that seasonal inventory leverage is positively related to future stock returns. Specifically, the full-sample results in Panel A show that while the bottom *QInv* tercile (containing the stocks with the highest seasonal inventory leverage) yields a mean excess return of 0.90% per month (*t*-statistic: 4.17) in column (1), the corresponding number for the top tercile (containing the stocks with the lowest seasonal inventory leverage) in column (3) is 0.51% (*t*-statistic: 2.41). The upshot is that the spread portfolio yields a significant mean excess return of –0.39% (*t*-statistic: –4.39) in column (4). In line with our theoretical priors, Panels B and C reveal that the mean spread return is an even more significantly negative –0.56% (*t*-statistic: –4.65) in the seasonal — but an only weakly significant –0.26% (*t*-statistic: –2.51) in the non-seasonal — stock subsample. Finally, columns (5) and (6) show that adjusting for systematic risk has close-to-no effects on the pricing of seasonal inventory leverage.[18]

Looking at the portfolio characteristics, the table suggests that the portfolios include a large number of stocks and are are thus well diversified. Crucially, however, the extreme seasonal (non-seasonal) *QInv* portfolios in Panel B (C) feature a disproportionally large (small) number of stocks, in line with seasonal stocks being more prone to build up inventories. In accordance, the *QSales* means reveal that seasonal (but not non-seasonal) stocks tend to hold larger inventories at the start of their high-sales quarters. Finally, the Sharpe ratio drops over the *QInv* portfolios in Panels A and B, indicating that low seasonal-inventory-leverage firms are less risky.

[18]In Section IA.3 of our Internet Appendix, we offer evidence that our univariate *QInv* portfolio sort results are robust to (i) using equal (rather than) value-weighted portfolios; (ii) relying on decile (rather than tercile) breakpoints; and (iii) examining delevered asset (rather than equity) returns.

In Table 7, we switch to FM regressions to verify that our portfolio sort results are robust to reasonable methodological changes. In the regressions, we project the single-stock excess return over month t on a rank variable formed from *QInv* and control variables measured at the start of that month.[19] As controls, we employ the rank of *QCash*, the rank of *QDebt*, *InvToSales*, *MarketBeta*, *MarketSize*, *BookToMarket*, *Momentum*, *Investment*, and *Profitability*.[20] While we run the regression in column (1) over our entire sample, we estimate those in columns (2) and (3) ((4) and (5)) on the seasonal and non-seasonal stock subsamples (subsamples of stocks with *InvToSales* values above (firms relying a lot on output inventories) or below (firms not doing so) the cross-sectional median), all respectively. As before, plain numbers are monthly premium estimates (in %), while those in square brackets are Newey and West (1987) t-statistics.

TABLE 7 ABOUT HERE.

The table confirms that the FM regressions yield results in agreement with the portfolio sorts. In particular, column (1) shows that the full sample produces a highly significant negative *QInv* rank premium of –0.33% per month (t-statistic: –4.57). Conversely, columns (2) and (3) corroborate that the same premium is a more negative –0.40% (t-statistic: –4.00) in the seasonal — but a weaker –0.24% (t-statistic: –2.58) in the non-seasonal — stock subsample. Columns (4) and (5) finally indicate that while the premium is a negative and significant –0.54% (t-statistic: –4.56) in the subsample of stocks relying more on inventories, the corresponding number for the other subsample is a less negative and significant –0.20% (t-statistic: –3.50).

We finally offer some evidence supporting our theory's prediction that the pricing of seasonal inventory leverage arises through systematic risk. One obstacle in doing so is that our models

[19] We use the rank variable since *QInv* is highly leptokurtic, possibly leading to concern that extreme *QInv* values could drive our estimates. Despite that, using *QInv* instead of its rank does not change our conclusions.

[20] In Section IA.4 of our Internet Appendix, we show that our FM regression results are robust to also controlling for standard inventory policies and operating flexibility. Specifically, also controlling for annual inventory growth, inventory productivity, and firm flexibility only marginally affects the negative pricing of *QInv*.

do not specify the underlying sources of priced risk (i.e., the factors determining the expected return of the demand-mimicking portfolio μ). Following the literature (e.g., Carlson et al. (2010) and Hackbarth and Johnson (2015)), we thus assume that market risk is one of those sources of risk. To establish that market risk rises in seasonal inventory leverage, Table 8 reports the results from standard (columns (1) and (2)) and weighted ((3) to (6)) least-squares panel regressions of single-stock excess returns on the contemporaneous excess market return, an interaction between market return and *QInv*, and industry fixed effects (only columns (2), (4), (6)):

$$R_{i,t} = \alpha + \left(\beta_0 + \beta_1 QInv_{i,q,y}\right) MKT_t + a_j + \epsilon_{i,t}, \tag{22}$$

where α, β_0, and β_1 are parameters, a_j is the industry-fixed effect, and $\epsilon_{i,t}$ is the residual. While the regressions in columns (3) and (4) use the gross return over the past month as weight, those in columns (5) and (6) use market size at the end of that month. As always, plain numbers are premium estimates, while those in square brackets are *t*-statistics.

TABLE 8 ABOUT HERE.

The table suggests that systematic risk is at least partially behind the pricing of *QInv*. In particular, we consistently find a significantly negative estimate on the interaction between the excess market return and *QInv*, implying that market risk increases with seasonal inventory leverage. The least-squares estimation with industry fixed effects in column (2), for example, reveals that a one-unit rise in *QInv* induces market risk to drop by 0.10 (*t*-statistics: –2.53).

Overall, this section shows that, in line with our theory, seasonal inventory leverage is positively related to future stock returns in portfolio sorts as well as FM regressions. In further agreement with our theory, it also links that positive relation to systematic market risk.

4 Implications for Stock Anomalies

In this section, we study whether seasonal inventory leverage contributes to or even explains several recent seasonal and allegedly-non-seasonal stock anomalies. To do so, we first introduce the anomalies and outline why they are likely to be related to seasonal inventory leverage. We next conduct independent triple portfolio sorts investigating the extent to which seasonal inventory leverage captures the abnormal profitability of the anomalies. We offer more details about the definitions of the anomaly variables in Table C1 in Appendix C.

4.1 The Anomalies and their Relations with Our Theory

We examine whether seasonal inventory leverage adds to or explains Grullon et al.'s (2020) seasonal sales, Heston and Sadka's (2008, 2010) and Keloharju et al.'s (2016, 2021) same-calendar-month, Jegadeesh and Titman's (1993, 2001) momentum, and Hou et al.'s (2015) quarterly ROE anomalies. To recap, the seasonal sales anomaly is that stocks tend to earn low (high) returns in their high (low) sales quarters, while the same-calendar-month anomaly is that stocks tend to earn high (low) returns in their past high (low) return same-calendar-months. Assuming that a high sales quarter predicts a low-return month and vice versa, our theory suggests that the anomalies arise because firms with low inventory holding costs hold large (small) output inventories at the start of their high (low) sales quarters/past low (high) return same-calendar-months, translating into low (high) seasonal inventory leverage. Thus, high *QInv* stocks should drive the low mean returns of high-seasonal sales/low same-calendar-month return stocks, whereas low *QInv* stocks should drive the high mean returns of their counterparts.

Conversely, the momentum anomaly is that stocks with high returns over months $t-2$ to $t-12$ ("winners") tend to earn higher returns than stocks with low returns ("losers") over that same period. Crucially, Heston and Sadka (2008) find that the anomaly is almost exclusively driven by the twelve-month-ago (i.e., month $t-12$) return, closely linking it to their same-calendar-month anomaly. As a result, low *QInv* stocks should drive the high mean returns of winners, whereas high *QInv* stocks should drive the low mean returns of losers. Finally, the quarterly ROE anomaly is that stocks with high past-quarter ROE values tend to earn higher returns than those with low values. As a high (low) ROE must reflect seasonality in sales and profits, it crudely identifies the high and low sales quarters of seasonal firms. Yet, if a firm's last quarter is a high (low) sales quarter, its current quarter is more likely than not a low (high) sales quarter. Thus, low *QInv* stocks should drive the high mean returns of high quarterly ROE stocks, whereas high *QInv* stocks should drive the low mean returns of their counterparts.

4.2 Seasonal Inventory Leverage and the Anomalies

We use triple portfolio sorts to evaluate whether seasonal inventory leverage adds to or explains the anomalies in Section 4.1. At the end of each month $t-1$, we thus separately sort our sample stocks into portfolios according to the median of the *Seasonality* distribution, the terciles of the anomaly variable distribution, and the terciles of the *QInv* distribution at that time. We use the intersection of the three sets of univariate portfolios to create tripe-sorted portfolios. We value-weight the portfolios and hold them over month t. We then consider either all stocks, only those in the top *Seasonality* portfolio, or only those in the bottom *Seasonality* portfolio. In line with the literature, we next create *standard* anomaly spread portfolios by longing the relevant stocks in the top anomaly-variable portfolio and shorting those in the bottom.

To study the effect of seasonal inventory leverage, we contrast the standard anomaly spread portfolios with *seasonal-inventory-leverage-neutral* spread portfolios. While we would typically form the neutral spread portfolios by averaging over all anomaly-variable spread portfolios within each *QInv* portfolio, our theory cautions us against doing so. The reason is that, according to our theory, we should never observe stocks with abnormally high (low) inventories at the start of their low (high) sales season. As a result, there should be few stocks in the corresponding portfolios (e.g., the top *QSales*-bottom *QInv* portfolio), and those stocks should likely have mismeasured *QInv* values. To take our theory seriously and to guard against this issue, we thus choose as seasonal-inventory-leverage-neutral spread portfolios the anomaly-variable spread portfolios formed from only those stocks in the middle *QInv* portfolio. We can interpret these stocks as belonging to the non-inventory-builders with a flat seasonal inventory leverage over the calendar year in our theory in Section 2 (recall Panel D in Figure 5).[21]

Table 9 reports the mean returns (columns (1) and (2)), augmented-q-theory alphas ((3) and (4)), and six-factor-model alphas ((5) and (6)) of spread portfolios formed from *QSales* (Panel A), the past same-calendar-month return (Panel B), the momentum return (Panel C), and the quarterly ROE (Panel D). While columns (1), (3), and (5) focus on the standard spread portfolios, columns (2), (4), and (6) look into the seasonal-inventory-leverage-neutral ones. Conversely, while the first row-pair in each panel considers all stocks, the second (third) explores seasonal (non-seasonal) stocks. As always, plain numbers are mean monthly excess returns or alphas, whereas those in square brackets are Newey and West (1987) t-statistics.

TABLE 9 ABOUT HERE.

The table suggests that all the anomalies still exist in our sample data. In particular, the first row-pairs in column (1) of Panels A to D show that the standard spread portfolios formed from

[21] One concern with that strategy could be that our seasonal-inventory-leverage-neutral spread portfolios contain fewer stocks and are thus less well diversified than the standard spread portfolios, perhaps making it more likely that their mean returns are insignificant. To mitigate that concern, Section IA.5 in our Internet Appendix repeats the construction of the standard spread portfolios analyzed in this section, this time, however, only including the one-third largest stocks in the short and long legs of the standard portfolios to ensure that those legs contain an equal number of stocks across the two types of portfolios. Remarkably, the section shows that doing so has hardly any effect on the mean returns and t-statistics of the standard spread portfolios.

QSales, the past same-calendar-month return, the momentum return, and the quarterly ROE yield mean monthly returns of –0.22%, 0.51%, 0.52%, and 0.42% (t-statistics: –2.17, 3.03, 2.29, and 2.08), all respectively.[22] In agreement with the literature, the same row-pairs in columns (3) and (5), however, demonstrate that while both the augmented-q-theory and the six-factor model price the momentum and quarterly-ROE spread portfolios, only the six-factor model prices the same-calendar-month spread portfolio (likely due to its inclusion of a momentum factor). Finally, neither of the two models is able to price the *QSales* spread portfolio.

More importantly, the table further shows that each anomaly is stronger in the seasonal than in the non-seasonal stock subsample, including the allegedly-non-seasonal momentum and quarterly ROE anomalies. Contrasting the second and third row-pairs in Panels A to D, the *QSales*, same-calendar-month return, momentum, and quarterly ROE spread portfolios, for example, yield highly significant mean monthly returns of –0.34%, 0.68%, 0.87%, and 0.54% in the seasonal but largely insignificant mean returns of –0.03%, 0.47%, 0.26%, and 0.26% in the non-seasonal subsample, all respectively. As our theory in Section 2 suggests that seasonality in demand and thus sales is a necessary condition for inverse seasonality in firms' expected returns to explain the anomalies, our evidence in this paragraph is a first piece of evidence supporting our seasonal-inventory-leverage-based explanation for the anomalies.

The seasonal-inventory-leverage-neutral spread portfolios in columns (2), (4), and (6) of Table 9 offer more direct evidence that seasonal inventory leverage is, at least partially, behind the anomalies. To see that, consider the *QSales* anomaly in Panel A. Forming the *QSales* spread portfolio from only stocks in the middle *QInv* portfolio (i.e., from stocks not building up inventories toward their high sales season and thus with a flat seasonal inventory leverage over

[22]The low t-statistic for the momentum anomaly is a result of our data filters. To show this, we sort stocks into decile portfolios based on the three anomalies without requiring non-missing data for *QInv* and *Seasonality*. The resulting mean monthly spread returns for the same-calendar-month, momentum, and quarterly ROE anomalies are 0.98%, 1.65%, and 0.82% (t-statistics: 3.80, 3.99, and 2.58), all respectively.

time), the mean monthly spread portfolio return rises from a significant –0.22% (t-statistic: –2.17) to an insignificant –0.10% (t-statistic: –0.80), a whopping 55% rise. Similarly, forming the same-calendar-month return, momentum, and quarterly ROE spread portfolios from only stocks in that same *QInv* portfolio in Panels B to D, their mean returns drop from 0.51% to 0.41% (t-statistics: 3.03 vs. 2.06); 0.52% to 0.34% (t-statistics: 2.29 vs. 1.29); and 0.42% to 0.24% (t-statistics: 2.08 vs. 1.05), translating into 20%, 35%, and 43% declines, all respectively.

Overall, this section offers evidence that seasonal inventory leverage significantly adds to the seasonal sales, same-calendar-month, momentum, and quarterly ROE anomalies.

5 Concluding Remarks

We develop a neoclassical theory of a firm with output price seasonality and optimal inventory building. Our theory predicts that a seasonal firm optimally builds up output inventories toward its high-price season, to sell more output in that season. An implication is that the firm prepays quasi-fixed production costs associated with output kept in-house, lowering its seasonal inventory leverage and expected return toward its high-price season. In agreement, our empirical work shows that seasonal firms hold amplified output inventories at the start of their high-sales months or quarters. Moreover, these holdings (which inversely proxy for seasonal inventory leverage) negatively price stocks. Our empirical work finally reveals that seasonal inventory leverage partially explains several seasonal and allegedly-non-seasonal stock anomalies, such as the seasonal sales, same-calendar-month, momentum, and quarterly ROE anomalies.

References

Alan, Yasin, George P. Gao, and Vishal Gaur, 2014, Does inventory productivity predict future stock returns? A retailing industry perspective, *Management Science* 60, 2416–2434.

Aretz, Kevin and Peter F. Pope, 2018, Real options models of the firm, capacity overhang, and the cross section of stock returns, *Journal of Finance* 73, 1363–1415.

Belo, Frederico and Xiaoji Lin, 2012, The inventory growth spread, *Review of Financial Studies* 25, 278–313.

Carlson, Murray, Adlai Fisher, and Ron Giammarino, 2004, Corporate investment and asset price dynamics: Implications for the cross-section of returns, *Journal of Finance* 59, 2577–2603.

Carlson, Murray, Adlai Fisher, and Ron Giammarino, 2010, SEO risk dynamics, *Review of Financial Studies* 23, 4026–4077.

Chang, Tom Y., Samuel M. Hartzmark, David H. Solomon, and Eugene F. Soltes, 2017, Being surprised by the unsurprising: Earnings seasonality and stock returns, *Review of Financial Studies* 30, 281–323.

Chen, Hong, Murray Z. Frank, and Owen Q. Wu, 2005, What actually happened to the inventories of American companies between 1981 and 2000?, *Management Science* 51, 1015–1031.

Chen, Zhanhui, 2016, Time-to-produce, inventory, and asset prices, *Journal of Financial Economics* 120, 330–345.

Cooper, Ilan, 2006, Asset pricing implications of nonconvex adjustment costs and irreversibility of investment, *Journal of Finance* 61, 139–170.

Dixit, Avinash K. and Robert S. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ.

Fama, Eugene F. and Kenneth R. French, 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.

Fama, Eugene F. and Kenneth R. French, 2016, Dissecting anomalies with a five-factor model, *Review of Financial Studies* 29, 69–103.

Fama, Eugene F. and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Garcia-Feijóo, Luis and Randy D. Jorgensen, 2010, Can operating leverage be the cause of the value premium?, *Financial Management* 39, 1127–1154.

Geman, Hélyette and Vu-Nhat Nguyen, 2005, Soybean inventory and forward curve dynamics, *Management Science* 51, 1076–1091.

Grullon, Gustavo, Yamil Kaba, and Alexander Núñez, 2020, When low beats high: Riding the sales seasonality premium, *Journal of Financial Economics* 138, 572–591.

Gu, Lifeng, Dirk Hackbarth, and Timothy Johnson, 2018, Inflexibility and stock returns, *Review of Financial Studies* 31, 278–321.

Hackbarth, Dirk and Timothy Johnson, 2015, Real options and risk dynamics, *Review of Economic Studies* 82, 1449–1482.

Hartzmark, Samuel M. and David H. Solomon, 2018, Recurring firm events and predictable returns: The within-firm time series, *Annual Review of Financial Economics* 10, 499–517.

Heston, Steven L. and Ronnie Sadka, 2008, Seasonality in the cross-section of stock returns, *Journal of Financial Economics* 87, 418–445.

Heston, Steven L. and Ronnie Sadka, 2010, Seasonality in the cross section of stock returns: The international evidence, *Journal of Financial and Quantitative Analysis* 45, 1133–1160.

Hirshleifer, David, Danling Jiang, and Yuting M. DiGiovanni, 2020, Mood beta and seasonalities in stock returns, *Journal of Financial Economics* 137, 272–295.

Hou, Kewei, Haitao Mo, Chen Xue, and Lu Zhang, 2021, An augmented q-factor model with expected growth, *Review of Finance* 25, 1–41.

Hou, Kewei, Chen Xue, and Lu Zhang, 2015, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650–705.

Jegadeesh, Narasimhan and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65–91.

Jegadeesh, Narasimhan and Sheridan Titman, 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *Journal of Finance* 56, 699–720.

Jones, Christopher S. and Şelale Tüzel, 2013, Inventory investment and the cost of capital, *Journal of Financial Economics* 107, 557–579.

Keloharju, Matti, Juhani T. Linnainmaa, and Peter M. Nyberg, 2016, Return seasonalities, *Journal of Finance* 71, 1557–1590.

Keloharju, Matti, Juhani T. Linnainmaa, and Peter M. Nyberg, 2021, Are return seasonalities due to risk or mispricing? Evidence from seasonal reversals, *Journal of Financial Economics* 139, 138–161.

Kogan, Leonid, Jun Li, and Harold H. Zhang, 2023a, Operating hedge and gross profitability premium, *Journal of Finance* 78, 3387–3422.

Kogan, Leonid, Jun Li, Harold H. Zhang, and Yifan Zhu, 2023b, Operating leverage and asset pricing puzzles, *Working Paper*.

Kogan, Leonid and Dimitris Papanikolaou, 2013, Firm characteristics and stock returns: The role of investment-specific shocks, *Review of Financial Studies* 26, 2718–2759.

Kogan, Leonid and Dimitris Papanikolaou, 2014, Growth opportunities, technology shocks, and asset prices, *Journal of Finance* 69, 675–718.

Lambrecht, Bart M., Grzegorz Pawlina, and Joao C. A. Teixeira, 2016, Making, buying, and concurrent sourcing: Implications for operating leverage and stock beta, *Review of Finance* 20, 1013–1043.

Lewellen, Jonathan and Stefan Nagel, 2006, The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Majd, Saman and Robert S. Pindyck, 1987, Time to build, option value, and investment decisions, *Journal of Financial Economics* 18, 7–27.

Mandelker, Gershon N. and S. Ghon Rhee, 1984, The impact of the degrees of operating and financial leverage on systematic risk of common stock, *Journal of Financial and Quantitative Analysis* 19, 45–57.

Newey, Whitney K. and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Novy-Marx, Robert, 2011, Operating leverage, *Review of Finance* 15, 103–134.

Novy-Marx, Robert, 2012, Is momentum really momentum?, *Journal of Financial Economics* 103, 429–453.

Ogden, Joseph P., 2003, The calendar structure of risk and expected returns on stocks and bonds, *Journal of Financial Economics* 70, 29–67.

Petersen, Mitchell A., 2008, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of Financial Studies* 22, 435–480.

Rubinstein, Mark, 1973, A mean-variance synthesis of corporate financial theory, *Journal of Finance* 28, 167–181.

Thomas, Jacob K. and Huai Zhang, 2002, Inventory changes and future returns, *Review of Accounting Studies* 7, 163–187.

Table 1: Descriptive Statistics

The table presents descriptive statistics for our main firm-level (Panel A) and industry-level (Panel B) variables. The descriptive statistics include the total number of observations ("#"), the mean, standard deviation ("Std.Dev."), and several percentiles. Except for the number of observations, we calculate the table entries first by sample month and then average over our sample period. See Table C1 in the appendix of this paper for more details about variable construction.

			Std	Percentiles				
	#	Mean	Dev	1	25	50	75	99
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Firm-Level Variables								
QInv	1,132,788	1.00	0.14	0.56	0.95	1.00	1.05	1.46
QSales	1,132,788	0.25	0.04	0.15	0.23	0.25	0.26	0.38
Seasonality	1,132,788	0.03	0.03	0.00	0.01	0.02	0.04	0.17
InvToSales	1,132,788	0.70	1.07	0.02	0.29	0.55	0.83	3.75
QFWInv	178,992	1.00	0.14	0.59	0.95	1.00	1.05	1.46
Panel B: Industry-Level Variables								
MNewOrders	2,436	0.08	0.01	0.08	0.08	0.08	0.09	0.09
MShipments	2,436	0.08	0.00	0.08	0.08	0.08	0.09	0.09
MUnfilled Orders	2,436	0.08	0.00	0.08	0.08	0.08	0.09	0.09
MInv	2,436	1.00	0.02	0.97	0.99	1.00	1.01	1.03
MFGInv	2,436	1.00	0.03	0.96	0.98	1.00	1.02	1.04
MWIPInv	2,436	1.00	0.02	0.96	0.98	1.00	1.02	1.04
MRMInv	2,436	1.00	0.02	0.97	0.98	1.00	1.02	1.03

Table 2: Manufacturers' Shipments, Inventories, & Orders Survey Regressions
The table presents the results from panel regressions of an industry's monthly new orders (columns (1) and (2)), shipment values ((3), (4), and (6) to (8)), and unfilled orders ((5)) as fractions of the annual totals on various combinations of the contemporaneous and/or lagged new-order fractions, the contemporaneous and/or lagged total-inventory fractions, and the one-month lagged finished good, work-in-progress, and raw material inventory fractions as well as industry fixed effects. Plain numbers are estimates, while the numbers in square brackets are Petersen (2008) t-statistics clustered at the industry level. The regression industries are the primary metals (31S), fabricated metal products (32S), machinery (33S), computers and electronic products (34S), electrical equipment, appliances, and components (35S), transportation equipment (36S), and furniture and related products (37S) industries. See Table C1 in the appendix of this paper for more details about variable construction.

	Dependent Variable:							
	MNew Orders	MNew Orders	MShipments	MShipments	MUnfilled Orders	MShipments	MShipments	MShipments
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
MNewOrders			0.69	0.69	0.13			
			[11.46]	[11.00]	[2.23]			
MNewOrders (Lag 1)	−0.02			0.00				
	[−0.34]			[−0.12]				
MNewOrders (Lag 2-4)	−0.04			0.01				
	[−1.02]			[0.55]				
MNewOrders (Lag 5-11)	−0.08							
	[−9.94]							
MNewOrders (Lag 12)	0.61	0.65						
	[6.24]	[6.14]						
MInv							−0.56	
							[−1.38]	
MInv (Lag 1)						0.93	1.54	
						[3.66]	[3.00]	
MInv (Lag 2-4)							−0.22	
							[−4.23]	
MFGInv (Lag 1)								0.69
								[4.53]
MWIPInv (Lag 1)								0.58
								[3.87]
MRMInv (Lag 1)								−0.33
								[−2.60]
Industry FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SE Cluster	Ind.	Ind.	Ind.	Ind.	Ind.	Ind.	Ind.	Ind.
R-Squared	0.46	0.43	0.66	0.67	0.09	0.11	0.16	0.18
Obs.	2,436	2,436	2,436	2,436	2,436	2,436	2,436	2,436

Table 3: Compustat Regressions of Abnormal Seasonal Sales on Lagged Abnormal Seasonal Inventories

The table presents the results from panel regressions of a firm's quarterly sales as fraction of its annual total (*QSales*) on combinations of its start-of-quarter abnormal total inventory fraction (*QInv*), an interaction between the inventory fraction and the seasonality dummy variable (*QInv* × *DSeasonality*), the seasonality dummy variable (*DSeasonality*), controls, and firm fixed effects. While we run the regressions in columns (1) and (2) over our full sample, we run those in columns (3) and (4); (5) and (6); (7) and (8); and (9) and (10) over durable consumer-goods, manufacturing goods, services, and healthcare industry subsamples, respectively. The control variables are log total assets (*Assets*), the quarterly ROE (*QROE*), and asset growth (*Investment*). Plain numbers are estimates, while the numbers in square brackets are Petersen (2008) t-statistics clustered at the firm level. See Tables C1 and C2 in the appendix of this paper for more details about analysis variable and industry-subsample definitions, respectively.

	Dependent Variable = *QSales*									
	Full Sample		Durable Consumer Goods		Manufacturing Goods		Services		Healthcare	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
QInv	0.15	0.01	0.51	0.06	0.18	0.04	0.03	−0.01	0.02	−0.01
	[17.66]	[2.41]	[9.38]	[5.31]	[8.03]	[5.69]	[2.89]	[−4.09]	[1.30]	[−1.72]
QInv × *DSeasonality*		0.21		0.58		0.22		0.07		0.04
		[17.26]		[8.78]		[6.55]		[4.06]		[1.95]
DSeasonality		−0.05		−0.14		−0.06		−0.02		−0.01
		[−16.91]		[−8.76]		[−6.50]		[−3.66]		[−1.52]
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SE Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
R-Squared	0.03	0.03	0.13	0.16	0.03	0.04	0.01	0.01	0.01	0.01
Obs. (in 1,000s)	854	854	82	82	196	196	83	83	81	81

Table 4: Compustat Regressions of Abnormal Seasonal Sales on Lagged Abnormal Seasonal Total or Finished Good and Work-in-Progress Inventories

The table presents the results from panel regressions of a firm's quarterly sales as fraction of its annual total (*QSales*) on combinations of its start-of-quarter abnormal inventory fraction, an interaction between the inventory fraction and the seasonality dummy variable, the seasonality dummy variable (*DSeasonality*), controls, and firm fixed effects. While we use the abnormal *total* inventory fraction (*QInv*) in columns (1) to (4), we use the abnormal *finished-good plus work-in-progress* inventory fraction (*QFWInv*) in columns (5) and (6). Conversely, while we run the regressions in columns (1) and (2) over the full sample, we run those in columns (3) to (6) over the subsample of observations with non-missing finished-good and work-in-progress inventory values. The control variables are log total assets (*Assets*), the quarterly ROE (*QROE*), and asset growth (*Investment*). Plain numbers are estimates, while the numbers in square brackets are Petersen (2008) t-statistics clustered at the firm level. See Table C1 in the appendix of this paper for more details about the analysis variable definitions.

	Dependent Variable = *QSales*					
	Full Sample (for comparison)		Non-Missing-*QFWInv* Subsample			
	(1)	(2)	(3)	(4)	(5)	(6)
QInv	0.15	0.01	0.22	0.03		
	[17.66]	[2.43]	[9.29]	[4.50]		
QInv × *DSeasonality*		0.21		0.23		
		[17.25]		[7.80]		
QFWInv					0.18	0.03
					[9.82]	[6.04]
QFWInv × *DSeasonality*						0.19
						[8.00]
DSeasonality		−0.05		−0.06		−0.05
		[−16.91]		[−7.72]		[−7.89]
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
SE Cluster	Firm	Firm	Firm	Firm	Firm	Firm
R-Squared	0.03	0.03	0.04	0.05	0.04	0.04
Obs. (in 1,000s)	854	854	146	146	146	146

Table 5: Compustat Regressions of Abnormal Seasonal Sales on Lagged Abnormal Cash and Short-term Debt and Contemporaneous Abnormal Dividends

The table presents the results from panel regressions of a firm's quarterly sales as fraction of its annual total (*QSales*) on combinations of its start-of-quarter abnormal cash (*QCash*) and debt (*QDebt*) fractions, its quarterly dividends as fraction of its annual total (*QDiv*), interactions between the *QCash*, *QDebt*, and *QDiv* variables and our seasonality dummy (*DSeasonality*), *DSeasonality* on its own, control variables, and firm fixed effects. The control variables are log total assets (*Assets*), the quarterly ROE (*QROE*), and asset growth (*Investment*). Plain numbers are estimates, while the numbers in square brackets are Petersen (2008) t-statistics clustered at the firm level. See Table C1 in the appendix of this paper for more details about the definitions of the analysis variables.

	Dependent Variable = *QSales*					
	(1)	(2)	(3)	(4)	(5)	(6)
QCash	−0.07	−0.02				
	[−24.90]	[−25.20]				
QCash × *DSeasonality*		−0.09				
		[−19.87]				
QDebt			0.10	−0.03		
			[11.51]	[−12.64]		
QDebt × *DSeasonality*				0.18		
				[15.19]		
QDiv					0.07	0.03
					[24.21]	[27.30]
QDiv × *DSeasonality*						0.09
						[15.05]
DSeasonality		0.02		−0.04		−0.02
		[20.73]		[−14.87]		[−14.87]
Firm FEs	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
SE Cluster	Firm	Firm	Firm	Firm	Firm	Firm
R-Squared	0.03	0.04	0.01	0.02	0.10	0.13
Obs. (in 1,000s)	847	847	832	832	289	289

Table 6: Univariate Portfolios Sorted on Abnormal Inventories

The table presents the results of portfolios univariately sorted on *QInv*. At the end of each month $t-1$, we first choose either all stocks (Panel A), those with a *Seasonality* value above the median (Panel B, "seasonal stocks"), or those with a *Seasonality* value below that median (Panel C, "non-seasonal stocks"). We next sort the chosen stocks into portfolios according to the unconditional terciles of the *QInv* distribution at that time. We value-weight the portfolios and hold them over month t. We also form spread portfolios long the top and short the bottom value-weighted portfolios ("High – Low"). In columns (1) to (3), we report the mean excess returns (Return^e), mean number of stocks (# Stocks), mean cross-sectional *QSales* averages (*QSales*), mean cross-sectional *MarketSize* averages (*MarketSize*), and the annualized Sharpe ratios (mean excess return over volatility) of the portfolios. Conversely, columns (4) to (6) report the mean excess returns, q^5-model alphas, and FF6-model alphas of the spread portfolios. Plain numbers are estimates, while those in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. See Table C1 in the appendix of this paper for more details about variable definitions.

	QInv			High – Low		
	Low	Medium	High	Mean	q^5-Alpha	FF6-Alpha
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Return^e	0.90	0.70	0.51	−0.39	−0.35	−0.38
	[4.17]	[3.85]	[2.41]	[−4.39]	[−4.01]	[−4.46]
# Stocks	708	708	729			
QSales	0.24	0.25	0.26			
MarketSize	9.22	9.40	9.24			
Sharpe (p.a.)	0.63	0.58	0.36			
Panel B: Seasonal Subsample						
Return^e	0.98	0.61	0.42	−0.56	−0.44	−0.48
	[4.32]	[2.60]	[1.78]	[−4.65]	[−3.69]	[−4.17]
# Stocks	409	267	396			
QSales	0.24	0.25	0.26			
MarketSize	9.05	9.13	9.06			
Sharpe (p.a.)	0.67	0.40	0.27			
Panel C: Non-Seasonal Subsample						
Return^e	0.84	0.78	0.58	−0.26	−0.30	−0.31
	[3.90]	[4.53]	[2.95]	[−2.51]	[−2.69]	[−2.74]
# Stocks	300	441	332			
QSales	0.25	0.25	0.25			
MarketSize	9.29	9.46	9.29			
Sharpe (p.a.)	0.60	0.70	0.46			

Table 7: Fama-MacBeth Regressions of Stock Returns on Abnormal Inventories
The table presents the results from Fama-MacBeth (1973) regressions of single-stock returns over month t on *QInv*, *QCash*, and *QDebt* rank variables, *InvToSales*, *MarketBeta*, *MarketSize*, *BookToMarket*, *Momentum*, *Investment*, and *Profitability*, all measured exclusively using data until the start of that month. While we run the regression on the full sample in column (1), we run it on only those firms with a *Seasonality* value above or below the cross-sectional median at the end of the prior June in columns (2) and (3), respectively. Conversely, we run it on only those firms with a total quarterly inventories-to-quarterly annual sales ratio value above or below the cross-sectional median at the end of the prior June in columns (4) and (5), respectively. The plain numbers are monthly premium estimates, in percent. The numbers in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. See Table C1 in the appendix of this paper for more details about variable definitions.

	Full Sample	Seasonal Subsample	Non-Seasonal Subsample	Large Inventories Subsample	Small Inventories Subsample
	(1)	(2)	(3)	(4)	(5)
QInv (Rank)	−0.33	−0.40	−0.24	−0.54	−0.20
	[−4.57]	[−4.00]	[−2.58]	[−4.56]	[−2.50]
QCash (Rank)	0.19	0.38	−0.03	0.13	0.20
	[2.76]	[3.65]	[−0.48]	[1.54]	[2.17]
QDebt (Rank)	0.00	−0.08	0.12	0.18	−0.12
	[−0.02]	[−0.80]	[1.53]	[1.89]	[−1.27]
InvToSales	−0.07	−0.09	−0.04	−0.14	0.25
	[−1.33]	[−1.41]	[−0.53]	[−1.68]	[1.45]
MarketBeta	−0.07	−0.16	0.02	−0.04	−0.10
	[−0.51]	[−1.10]	[0.12]	[−0.28]	[−0.70]
MarketSize	0.00	−0.01	0.01	0.04	−0.02
	[0.08]	[−0.23]	[0.28]	[1.10]	[−0.61]
BookToMarket	0.06	0.00	0.08	0.08	0.05
	[0.72]	[−0.05]	[0.96]	[0.90]	[0.60]
Momentum	0.64	0.81	0.46	0.73	0.57
	[2.73]	[3.27]	[1.97]	[3.05]	[2.27]
Investment	−0.42	−0.49	−0.19	−0.51	−0.31
	[−3.30]	[−3.28]	[−1.13]	[−3.23]	[−2.04]
Profitability	0.32	0.31	0.28	0.37	0.31
	[2.17]	[1.90]	[1.62]	[2.09]	[1.93]
Constant	0.68	0.77	0.60	0.47	0.79
	[1.91]	[1.97]	[1.64]	[1.24]	[2.04]

Table 8: Regressions of Stock Returns on Market Returns Interacted with Abnormal Inventories
The table presents the results from ordinary (columns (1) and (2)) and weighted (columns (3) to (6)) least-squares panel regressions of single-stock returns over month t on the contemporaneous market return minus the risk-free rate of return ("excess market return;" $MktReturn^e$) and that same market return interacted with start-of-quarter abnormal inventories ($MktReturn^e \times QInv$). While the regressions in the even-numbered columns do not include fixed effects, those in the odd-numbered columns include fixed effects based on Kenneth French's 49 industry classification scheme. Also, while the regressions in columns (3) and (4) use a stock's gross return over the last month as weight, those in columns (5) and (6) use its stock capitalization at the end of that month. Plain numbers are estimates, while the numbers in square brackets are Petersen (2008) t-statistics clustered at the firm level. See Table C1 in the appendix of this paper for more details about variable definitions.

	(1)	(2)	(3)	(4)	(5)	(6)
$MktReturn^e$	1.25	1.25	1.21	1.21	1.18	1.18
	[32.54]	[31.32]	[31.77]	[30.69]	[21.39]	[21.21]
$MktReturn^e \times QInv$	−0.10	−0.10	−0.08	−0.09	−0.11	−0.12
	[−2.57]	[−2.53]	[−2.15]	[−2.21]	[−2.07]	[−2.21]
Industry FEs	No	Yes	No	Yes	No	Yes
Technique	OLS	OLS	WLS	WLS	WLS	WLS
Weight	–	–	Return	Return	Size	Size
SE Cluster	Firm	Firm	Firm	Firm	Firm	Firm
R-Squared	0.19	0.18	0.19	0.19	0.23	0.23
Obs. (in 1,000s)	517	516	508	507	517	516

Table 9: Standard and Seasonal-Inventory-Leverage-Neutral Anomaly Spread Portfolios
In this table, we present the mean returns (columns (1) and (2)), augmented-q-theory-model alphas ("q^5;" (3) and (4)), and six-factor-model alphas ("FF6;" (5) and (6)) of seasonal sales (Panel A), same-calendar-month return (Panel B), momentum (Panel C), and quarterly ROE (Panel D) spread portfolios. At the end of each month $t-1$, we form the spread portfolios from independently sorting our sample stocks into portfolios based on the median of *Seasonality*, the terciles of the anomaly variable, and the terciles of *QInv* at that time. We value-weight the portfolios and hold them over month t. We then consider either all stocks ("All Stocks"), only those in the top *Seasonality* portfolio ("Seasonal Stocks"), or only those in the bottom Seasonality portfolio ("Non-Seasonal Stocks"). We finally form standard anomaly spread portfolios through longing the top-anomaly-variable and shorting the bottom portfolios (columns (1), (3), and (5)), while we form *seasonal-inventory-leverage-neutral* spread portfolios through longing the top-anomaly-variable and shorting the bottom portfolios within the middle *QInv* portfolio ((2), (4), and (6)). Plain numbers are estimates, while those in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. See Table C1 in the appendix of this paper for more details about the anomaly variable definitions.

	Mean Spread Return		q^5-Spread Alpha		FF6-Spread Alpha	
	Standard	*QInv* Neutral	Standard	*QInv* Neutral	Standard	*QInv* Neutral
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Seasonal Sales Anomaly						
All Stocks	−0.22	−0.10	−0.12	0.02	−0.18	−0.04
	[−2.17]	[−0.80]	[−1.00]	[0.12]	[−1.63]	[−0.26]
Seasonal Stocks	−0.34	−0.16	−0.24	−0.04	−0.28	−0.03
	[−2.62]	[−0.96]	[−1.52]	[−0.20]	[−2.07]	[−0.20]
Non-Seasonal Stocks	−0.03	−0.02	0.07	0.10	0.05	0.03
	[−0.27]	[−0.11]	[0.54]	[0.55]	[0.41]	[0.18]
Panel B: Same-Calendar-Month Anomaly						
All Stocks	0.51	0.41	0.50	0.45	0.51	0.45
	[3.03]	[2.06]	[2.93]	[2.06]	[2.96]	[2.11]
Seasonal Stocks	0.68	0.65	0.58	0.66	0.63	0.61
	[3.32]	[2.42]	[2.49]	[1.75]	[2.81]	[1.81]
Non-Seasonal Stocks	0.47	0.41	0.53	0.43	0.46	0.44
	[2.94]	[2.35]	[3.06]	[2.45]	[2.75]	[2.36]

(continued on next page)

Table 9: Standard and Seasonal-Inventory-Leverage-Neutral Anomaly Spread Portfolios (cont.)

	Mean Spread Return		q^5-Spread Alpha		FF6-Spread Alpha	
	Standard	*QInv* Neutral	Standard	*QInv* Neutral	Standard	*QInv* Neutral
	(1)	(2)	(3)	(4)	(5)	(6)
Panel C: Momentum Anomaly						
All Stocks	0.52	0.34	−0.10	−0.09	−0.05	−0.08
	[2.29]	[1.29]	[−0.41]	[−0.38]	[−0.35]	[−0.43]
Seasonal Stocks	0.87	0.93	0.08	0.24	0.23	0.37
	[3.35]	[2.77]	[0.33]	[0.77]	[1.47]	[1.52]
Non-Seasonal Stocks	0.26	0.11	−0.14	−0.15	−0.17	−0.22
	[1.10]	[0.44]	[−0.56]	[−0.60]	[−1.27]	[−1.25]
Panel D: Quarterly-ROE Anomaly						
All Stocks	0.42	0.24	−0.43	−0.42	−0.02	0.06
	[2.08]	[1.05]	[−2.14]	[−2.12]	[−0.10]	[0.40]
Seasonal Stocks	0.54	0.38	0.00	−0.28	0.26	0.13
	[2.63]	[1.33]	[0.03]	[−0.90]	[2.08]	[0.50]
Non-Seasonal Stocks	0.26	0.17	−0.24	−0.24	0.01	0.05
	[1.15]	[0.69]	[−1.67]	[−1.22]	[0.09]	[0.27]

A Model Solutions

In this appendix, we provide additional information about our theoretical models. In particular, we first describe how we identify the inventory building and instantaneous selling regions in the toy model. We next outline how we numerically solve the real options model.

A.1 Derivation of the Optimal Selling Date in the Toy Model

In Equation (1) in our main text, we conjecture that the output price, P_t, in the toy and real options models evolves according to the generalized geometric Brownian motion (GBM):

$$\mathrm{d}P_t = (\alpha + \kappa \sin(\eta t))P_t \mathrm{d}t + \sigma P_t \mathrm{d}B_t. \tag{A1}$$

We now let μ be the expected return of a portfolio perfectly replicating the stochastic variations in the output price and thus reflecting its systematic risk. In line with Dixit and Pindyck (1994), we then define the "expected-return shortfall" of the output price as $\delta_t = \mu - \frac{1}{P_t \mathrm{d}t}\mathbb{E}[\mathrm{d}P_t] = \mu - \alpha - \kappa \sin(\eta t)$, which is a sinusoidal function of time. Intuitively, we can interpret δ_t as the opportunity cost of selling output immediately rather than in the future. Using the definition for the expected-return shortfall δ_t, we can rewrite the output price differential in Equation (A1) as:

$$\mathrm{d}P_t = (\mu - \delta_t)P_t \mathrm{d}t + \sigma P_t \mathrm{d}B_t, \tag{A2}$$

whose closed-form solution is equal to:

$$P_t = P_0 \exp\left(\int_0^t (\mu - \delta_u)\mathrm{d}u - \frac{1}{2}\sigma^2 t + \sigma B_t\right) \tag{A3}$$

$$= P_0 \exp\left(\left(\alpha - \frac{1}{2}\sigma^2\right) t + \frac{\kappa}{\eta}\big(1 - \cos(\eta t)\big) + \sigma B_t\right). \tag{A4}$$

Conversely, switching to the equivalent martingale measure $\mathbb{Q}$ under which the instantaneous drift changes from $\mu - \delta_t$ to $r - \delta_t$, the explicit solution for the output price becomes:

$$P_t = P_0 \exp\left(\int_0^t (r - \delta_u)\mathrm{d}u - \frac{1}{2}\sigma^2 t + \sigma B_t^{\mathbb{Q}}\right) \tag{A5}$$

$$= P_0 \exp\left(\left(r - \mu + \alpha - \frac{1}{2}\sigma^2\right) t + \frac{\kappa}{\eta}\big(1 - \cos(\eta t)\big) + \sigma B_t^{\mathbb{Q}}\right), \tag{A6}$$

where $B_t^{\mathbb{Q}}$ is a standard Brownian motion under the equivalent martingale measure.

Using Equation (A6), the conditional expectation of the time-t output price taken under the equivalent martingale measure $\mathbb{Q}$ at time s is equal to:

$$\mathbb{E}_s^{\mathbb{Q}}[P_t] = P_s \exp\left((r - \mu + \alpha)(t - s) + \frac{\kappa}{\eta}\big(\cos(\eta s) - \cos(\eta t)\big)\right), \tag{A7}$$

so that we can write the first-order condition for maximization problem (4) in our main text as:

$$P_t \exp\left(-(\mu - \alpha)(t^* - t) - \frac{\kappa}{\eta}\big(\cos(\eta t^*) - \cos(\eta t)\big)\right)(\mu - \alpha - \kappa \sin(\eta t^*)) = 0, \tag{A8}$$

which holds when $t^* = \frac{\pi}{\eta} - \frac{1}{\eta}\sin^{-1}\left(\frac{\mu - \alpha}{\kappa}\right) + \frac{2\pi}{\eta} n$ for $n \in \{\ldots, -1, 0, 1, \ldots\}$ and is equivalent to $\delta_{t^*} = 0$.

A.2 Numerical Solution for the Real Options Model

We next discuss how we use the explicit finite difference method to solve the real options model introduced in Section 2.2 of our main paper. As shown in Equation (12) of that paper, the value of an output unit in inventory in that model must satisfy the linear PDE:

$$rV = -c_I + \frac{\partial V}{\partial t} + (r - \delta_t)P\frac{\partial V}{\partial P} + \frac{1}{2}\sigma^2 P^2 \frac{\partial^2 V}{\partial P^2}. \tag{A9}$$

Using the log-transformation $p_t = \ln(P_t)$ to improve numerical stability, we can simplify PDE (A9) to:

$$rV = -c_I + \frac{\partial V}{\partial t} + \left(r - \delta_t - \frac{1}{2}\sigma^2\right)\frac{\partial V}{\partial p} + \frac{1}{2}\sigma^2 \frac{\partial^2 V}{\partial p^2}. \tag{A10}$$

Consider a grid (t_i, p_j) of time and the log output price on which we approximate the partial derivatives of the PDE using the forward and central finite differences:

$$\frac{\partial V}{\partial t} = \frac{V_{i+1,j} - V_{i,j}}{\Delta t} + \mathcal{O}(\Delta t), \tag{A11}$$

$$\frac{\partial V}{\partial p} = \frac{V_{i+1,j+1} - V_{i+1,j-1}}{2\Delta p} + \mathcal{O}(\Delta p^2), \tag{A12}$$

$$\frac{\partial^2 V}{\partial p^2} = \frac{V_{i+1,j+1} - 2V_{i+1,j} + V_{i+1,j-1}}{(\Delta p)^2} + \mathcal{O}(\Delta p^2). \tag{A13}$$

Using these approximations, the PDE turns into the explicit finite difference scheme:

$$V_{i,j} = a_i V_{i+1,j-1} + b V_{i+1,j} + c_i V_{i+1,j+1} + d, \tag{A14}$$

where:

$$a_i = \frac{1}{1 + r\Delta t}\left(-\left(r - \delta_{t_i} - \frac{1}{2}\sigma^2\right)\frac{\Delta t}{2\Delta p} + \frac{1}{2}\sigma^2\frac{\Delta t}{(\Delta p)^2}\right), \tag{A15}$$

$$b = \frac{1}{1 + r\Delta t}\left(1 - \sigma^2\frac{\Delta t}{(\Delta p)^2}\right), \tag{A16}$$

$$c_i = \frac{1}{1 + r\Delta t}\left(\left(r - \delta_{t_i} - \frac{1}{2}\sigma^2\right)\frac{\Delta t}{2\Delta p} + \frac{1}{2}\sigma^2\frac{\Delta t}{(\Delta p)^2}\right), \tag{A17}$$

$$d = \frac{-c_I \Delta t}{1 + r\Delta t}. \tag{A18}$$

We impose standard boundary conditions. When $P_t \to 0$, the output price approaches an absorbing boundary such that $V_{i,j_{\min}} = 0$. When $P_t \to \infty$, we use the solution with pre-commitment by setting $V_{i,j_{\max}} = \mathbb{E}^{\mathbb{Q}}_{t_i}[P_{t^*}|P_{t_i} = P_{j_{\max}}]e^{-r(t^* - t_i)} - C_I(t_i, t^*)$, where $C_I(s,t) = \int_s^t e^{-r(u-s)}c_I \mathrm{d}u = \frac{c_I}{r}\left(1 - e^{-r(t-s)}\right)$ is the holding cost of an inventory unit from s to t and t^* is the next optimal selling point between t_i and T identified as in Section A.1 of this appendix after accounting for the presence of inventory holding costs. When $t \to T$ (the liquidation time), the firm immediately sells all remaining output, implying that $V_{i_{\max},j} = P_j$. We incorporate the free boundary condition $V_{i^*,j} = P_j$ by setting:

$$V_{i,j} = \max\{a_i V_{i+1,j-1} + b V_{i+1,j} + c_i V_{i+1,j+1} + d, P_j\}. \tag{A19}$$

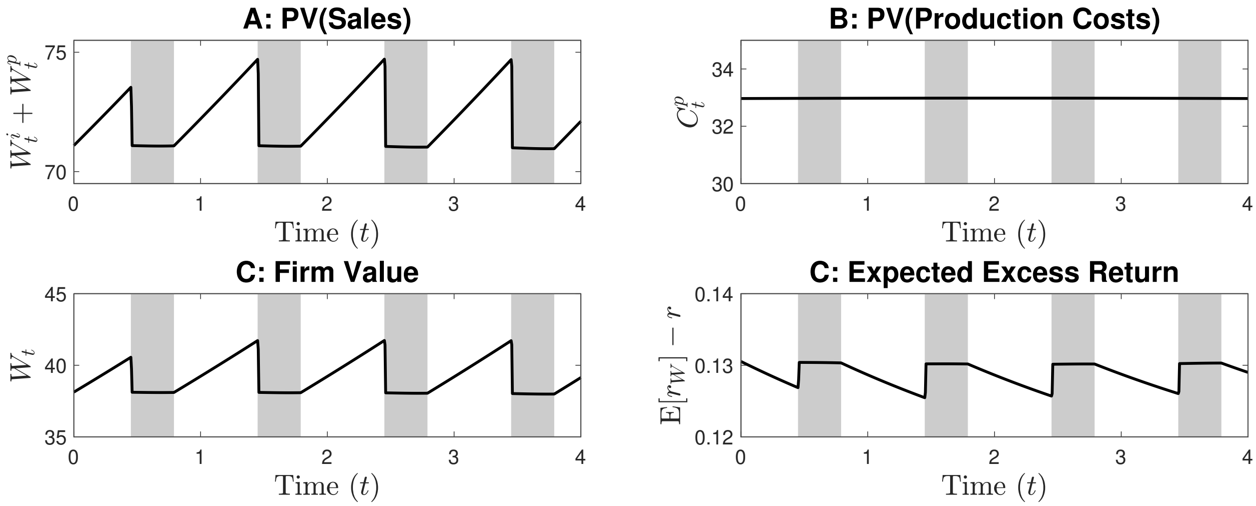

Figure B1: The figure plots the present value of the firm's future sales revenue $W_t^i + W_t^p$ (Panel A), the present value of its future fixed production costs C_t^p (Panel B), its total value W_t (Panel C), and its expected excess return (Panel D) over the period from $t = 0$ to 4. To focus on seasonal patterns, we remove a linear trend before plotting. While the white areas in each subpanel are optimal inventory building periods, the gray-shaded areas are optimal instantaneous selling periods. We describe the parameter values in Section 2.1.3 in our main paper.

B Toy Model Simulation Exercise

In this section, we offer the results from simulations of the toy model analogous to those of the real options model in Section 2.3. Using the same basecase parameters as in Section 2.3, Figure B1 plots the present value of firm's future sales revenue $W_t^i + W_t^p$ (Panel A), the present value of its future fixed production costs C_t^p (Panel B), its total value $W_t = W_t^i + W_t^p - C_t^p$ (Panel C), and its expected excess return (Panel D) over time. The figure yields conclusions in complete agreement with the corresponding real-options-model figure, Figure 4 in our main text. More specifically, Panel A reveals that the present value of the firm's future sales revenues rises over optimal inventory building periods. In contrast, since the firm pays constant production costs at each time, Panel B demonstrates that the present value of those costs are flat over time. Combining the two present values into the firm's total value, Panel C confirms that the total value inherits the seasonality from the present value of its future sales revenue. Finally, in line with the real options model, Panel D documents that the seasonality in the firm's total value translates into inverse seasonality in its expected excess return, simply because the firm's elasticity is linear in the inverse of its total value (see Equation (9) in our main text).

C Variable Definitions

Table C1: Variable Definitions

The table presents the definitions of our analysis variables. While Panel A focuses on our main firm-level variables, Panels B, C, and D focus on our main industry-level, anomaly, and control variables, respectively. We update the variables indexed by "M" ("Q") ["A"] on a monthly (quarterly) [annual] basis and use their values to condition returns over month $t+1$ (month $t+1$) [the period from July of year t to June of year $t+1$]. We show the data-provider mnemonics of the CRSP and Compustat variables in parentheses.

Variable Name	Variable Definition
	Panel A: Main Firm-Level Variables
QInv (Q)	Ratio of quarterly total inventories (invtq) at the start of the current quarter to mean quarterly inventories over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.
QSales (Q)	Ratio of quarterly sales (saleq) over the current quarter to total sales over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$ (see Grullon et al. (2020)).
Seasonality (A)	Standard deviation of the four *QSales* values for the current fiscal year.
InvToSales (A)	Ratio of the mean of quarterly total inventories (invtq) over the fiscal year to the mean of quarterly sales (saleq) over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.
QFWInv (Q)	Ratio of the sum of quarterly finished-good and work-in-progress inventories (invfgq+invwipq) at the start of the current quarter to the mean of such inventories over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.
QCash (Q)	Ratio of quarterly cash and short-term investments (cheq) at the start of the current quarter to the mean of such cash holdings over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.
QDebt (Q)	Ratio of quarterly current liabilities (lctq) at the start of the current quarter to the mean of such liabilities over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.
QDiv (Q)	Ratio of quarterly cash dividends (dvy) over the current quarter to total dividends over the fiscal year, averaged over the two fiscal years ending in calendar years $t-1$ and $t-2$.

(continued on next page)

Table C1: Variable Definitions (cont.)

Variable Name	Variable Definition
	Panel B: Main Industry-Level Variables
MNewOrders (M)	Ratio of new output-good orders received over the current month to total new output-good orders over the calendar year.
MShipments (M)	Ratio of the value of output goods shipped out over the current month to total value of shipments of output goods over the calendar year.
MUnfilledOrders (M)	Ratio of the unfilled output-good orders over the current month to total unfilled output-good orders over the calendar year.
MInv (M)	Ratio of total inventories at the end of the current month to mean total inventories over the calendar year.
MFGInv (M)	Ratio of finished-good inventories at the end of the current month to mean finished-good inventories over the calendar year.
MWIPInv (M)	Ratio of work-in-progress inventories at the end of the current month to mean work-in-progress inventories over the calendar year.
MRMInv (M)	Ratio of raw-material inventories at the end of the current month to mean raw-material inventories over the calendar year.
	Panel C: (Additional) Anomaly Variables
SameCMReturn (M)	The stock return (ret) over the same calendar month as the current, averaged over calendar years $t-1$ to $t-7$ (see Heston and Sadka (2008)).
Momentum (M)	Log of the gross monthly stock return (ret) compounded over months $t-2$ to $t-12$ (see Jegadeesh and Titman (1993)).
QROE (M)	The ratio of quarterly income before extraordinary items (ibq) to the quarterly book value of equity, where the book value of equity is shareholders' equity (seqq, ceqq+pstkq, or atq–ltq, in that order of availability) plus deferred taxes and investment tax credits (txditcq, zero if missing) minus the book value of preferred stock (pstkrq, zero if missing). We take quarterly income from the latest earnings announcement date (rdq) and all other variables from the earnings announcement date before. We set the variable to missing if the earnings announcement date is more than six months after the earnings date (see Hou et al. (2015)).

(continued on next page)

Table C1: Variable Definitions (cont.)

Variable Name	Variable Definition
	Panel D: (Additional) Control Variables
MarketBeta (M)	Sum of slope coefficients from a stock-level regression of excess stock returns (ret) on current, one-day lagged, and the sum of two-, three-, and four-day lagged excess market returns, where the regression is run using daily data over the prior 12 months. We require that the regression is run on at least 200 observations (see Lewellen and Nagel (2006)).
MarketSize (A)	Log of the product of the stock price (abs(prc)) and common shares outstanding (shrout) at the end of the prior calendar year (in millions).
BookToMarket (A)	Log of the ratio of the book value of equity to the market value of equity (abs(prc) times shrout), where the book value of equity equals stockholder's equity (seq) plus deferred taxes (txditc) plus investment tax credits (itcb, zero if missing) minus preferred stock (pstkrv, pstkl, pstk, or zero, in that order of availability). We take all the variables from the fiscal year-end in calendar year $t-1$ (see Fama and French (2015)).
Investment (A)	Log of the gross percentage change in total assets (at) over the fiscal year ending in calendar year $t-1$ (see Fama and French (2015)).
Profitability (A)	Ratio of sales (sale) net of costs of goods sold (cogs), selling, general, and administrative expenses (xsga, zero if missing), and interest expenses (xint, zero if missing) to the book value of equity, which equals stockholder's equity (seq) plus deferred taxes (txditc) plus investment tax credits (itcb, zero if missing) minus preferred stock (pstkrv, pstkl, pstk, or zero, in that order of availability). We take all the variables from the fiscal year ending in calendar year $t-1$ (see Fama and French (2015)).
Assets (A)	Log of total assets (at) from the fiscal year-end in calendar year $t-1$.

Table C2: Industry Classifications

The table presents the Standard Industrial Classification (SIC) codes used in Table 3.

Label	SIC Codes
Non-Perishable Consumer Goods	2200-2399, 2500-2519, 2590-2599, 2700-2749, 2770-2799, 3100-3199, 3630-3659, 3710-3711, 3714, 3716, 3750-3751, 3792, 3900-3939, 3940-3989, 3990-3999.
Manufacturing Goods	2520-2589, 2600-2699, 2750-2769, 2800-2829, 2840-2899, 3000-3099, 3200-3569, 3580-3621, 3623-3629, 3700-3709, 3712-3713, 3715, 3717-3749, 3752-3791, 3793-3799, 3860-3899.
Services	7000-8999.
Healthcare	2830-2839, 3693-3693, 3840-3859, 8000-8099.

Internet Appendix:

Seasonal Inventory Leverage

AUTHOR 1, AUTHOR 2, and AUTHOR 3

In this Internet Appendix, we offer supplementary theoretical derivations and empirical evidence on the corporate finance and asset pricing implications of seasonal inventory leverage induced through optimal inventory building. Section IA.1 starts with extending the real options model in Section 2.2 of our main paper by awarding a growth option to the firm, to verify that seasonal investment policies do not alter our main conclusions. Section IA.2 confirms that our main analysis variable, *QInv*, is persistent over time, allowing us to its past values to proxy for its future values (Table IA.1). Section IA.3 reveals that our univariate portfolio sorts in Section 3.4 of our main paper are robust to (i) using equal rather than value weights (Table IA.2); (ii) relying on decile rather than tercile breakpoints; and (iii) evaluating delevered asset rather than equity returns. In Section IA.4, we confirm that *QInv* continues to be significantly negatively priced even after adding further controls related to firms' inventory policies as well as their operating flexibility (Table IA.3). In Section IA.5, we finally repeat our calculations of the mean returns and alphas of the standard and seasonal-inventory-leverage-neutral anomaly spread portfolios in Section 4 of our main paper, this time, however, ensuring that the long (alternatively: short) legs of the standard and seasonal-inventory-leverage-neutral portfolios contain an identical number of stocks and are thus equally well diversified (Table IA.4).

IA.1 Growth Option Extension

We begin with endowing the firm in our real options model in Section 2.2 of our main paper with a growth option allowing it to discretely and irreversibly raise its production capacity by $\psi\bar{K}$ additional units at an investment cost of k, where $\psi > 0$ and $k > 0$ are fixed parameters (see, e.g., McDonald and Siegel (1986) and Dixit and Pindyck (1994)). To do so, we follow the methodology outlined in Appendix A.2 of our main paper to value the additional units, $U(t, P_t)$. We can easily do so since our model assumes constant returns to scale, no effect of quantity on price, and no increasing marginal production costs. The upshot is that we only have to change $\bar{K}$ to $\psi\bar{K}$ in the finite difference grid introduced in that appendix and that we can then calculate $U(t, P_t)$ from applying Equation (13) in our main paper to the outcomes obtained from solving the finite difference grid.

We next value the growth option on the $\psi\bar{K}$ additional production units. Since exercising that option yields production units which have not yet produced output, the value of the option does not depend on output in inventory I_t, allowing us to write it as $F = F(t, P_t, I_t = 0) = F(t, P_t)$. As well known from the literature, the firm optimally exercises the option whenever the output price P_t rises above the time-varying threshold P_t^*. Before the optimal exercise, the value of the growth option must satisfy the following generalized Black-Scholes (1973) PDE:

$$\frac{\partial F}{\partial t} + \left(r - \delta_t - \frac{1}{2}\sigma^2\right)\frac{\partial F}{\partial p} + \frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial p^2} - rF = 0, \tag{IA1}$$

which is identical to PDE (A10) except for the inventory holding cost c_I on the left-hand side. We can thus use the familiar explicit finite difference scheme (excluding the right-hand side constant):

$$F_{i,j} = a_i F_{i+1,j-1} + b F_{i+1,j} + c_i F_{i+1,j+1}, \tag{IA2}$$

where the coefficients a_i, b, and c_i are as in Appendix A.2 in our main paper. We impose the usual boundary conditions for American call options, $F(t_i, P_{\min}) = 0$ and $F(t_i, P_{\max}) = U(t_i, P_{\max}) - k$, as well as the terminal boundary condition $F(t_{\max}, P_j) = \max\{U(t_{\max}, P_j) - k, 0\}$. We incorporate the early

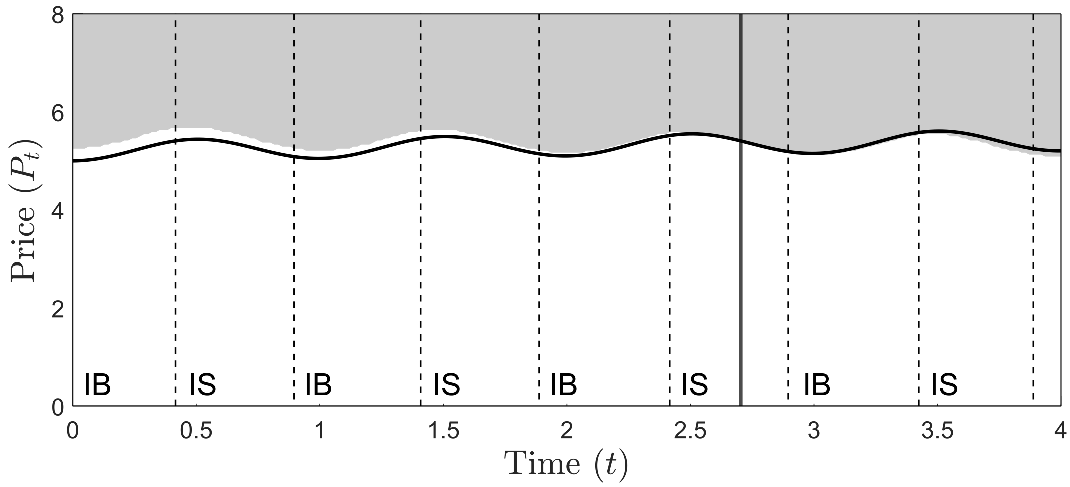

Figure IA.1: The figure plots the optimal output price threshold P_t^* above which the firm exercises its growth option (the lower boundary of the gray area), the optimal growth-option exercise region (the gray area), and the optimal continuation region (the white area) over the period from $t = 0$ to 4. We add an output price trajectory as solid black line. While the solid vertical line indicates the optimal exercise of the growth option, the broken lines separate the optimal inventory-building (IB) and instantaneous-selling (IS) periods. We describe the parameter values in Section IA.1.

exercise boundary condition $F(t, P_t^*) = U(t, P_t^*) - k$ by relying on:

$$F_{i,j} = \max\{a_i F_{i+1,j-1} + b F_{i+1,j} + c_i F_{i+1,j+1}, U(t_i, P_j) - k\}. \tag{IA3}$$

In our growth option extension, the value of the firm, $W(t, P_t, I_t)$, is then the sum of its production capacity-in-place, $V(t, P_t, I_t)$, valued as described in Appendix A.2, and its growth option, $F(t, P_t)$:

$$W(t, P_t, I_t) = V(t, P_t, I_t) + F(t, P_t), \tag{IA4}$$

whereas its conditional expected excess return, $\mathbb{E}[r_W] - r$, is the value-weighted average of the expected returns of the production capacity-in-place and the growth option:

$$\mathbb{E}[r_W] - r = \left(\frac{V(t, P_t, I_t)}{W(t, P_t, I_t)}\beta_V + \frac{F(t, P_t)}{W(t, P_t, I_t)}\beta_F\right)(\mu - r), \tag{IA5}$$

where $\beta_V = \frac{\partial V/V}{\partial P/P}$ and $\beta_F = \frac{\partial F/F}{\partial P/P}$ are the sensitivities of the value components to the output price.

Using the same parameter values as in Section 2.1.2 of our main paper, an investment cost, k, of 10, and an increase in the firm's production capacity upon it exercising its growth option, ψ, of 25%, Figure IA.1 plots the firm's optimal growth-option exercise strategy over the period from $t = 0$ to 4. To be more specific, the lower boundary of the gray area in the figure is the optimal output price threshold above which the firm exercises its growth option, P_t^*, the gray area is the optimal exercise region, and the white area is the optimal continuation region. Conversely, the solid black line is the output price trajectory also used in the main paper (see Figure 1 in that paper). The solid vertical line indicates the exercise of the growth option, whereas the broken vertical lines separate the optimal inventory-building (IB) and instantaneous-selling (IS) periods. The figure demonstrates that seasonality in the output price translates into seasonality in the optimal growth-option exercise threshold. Interestingly, however, the firm tends to invest *after* the seasonal output-price peak and closer toward the end of its instantaneous selling period. The reason is that adding additional production capacity during a low output price season allows the firm to build up greater output inventories toward its next high output price season and, as a result, to generate even higher sales revenues in that latter season.

We next study the implications of adding the growth option for our seasonal inventory leverage conclusions. To do so, Figure IA.2 plots the output price realization (Panel A), accumulated inventory (Panel B), the value and expected return of the capacity-in-place (Panels C and D), the value and expected return of the growth option (Panels E and F), and the value and expected return of the firm (Panels G and H) from time $t = 0$ to 4, respectively. As before, the vertical line indicates the growth option exercise, and the white (gray) shaded areas the optimal inventory-building (instantaneous-selling) regions. Figure IA.2 yields conclusions largely consistent with those from Figure 4 in our main paper, at least some time before the exercise of the growth option. To wit, Panels C and D show that both the value and expected return of the capacity-in-place vary seasonally, with the value (expected return) rising (falling) as the firm approaches its high-output-price season. In contrast, Panels E and F reveal that while the value of the growth option gradually and monotonically increases over time, its expected return moves inverse-seasonally to that of the capacity-in-place, mirroring the seasonal

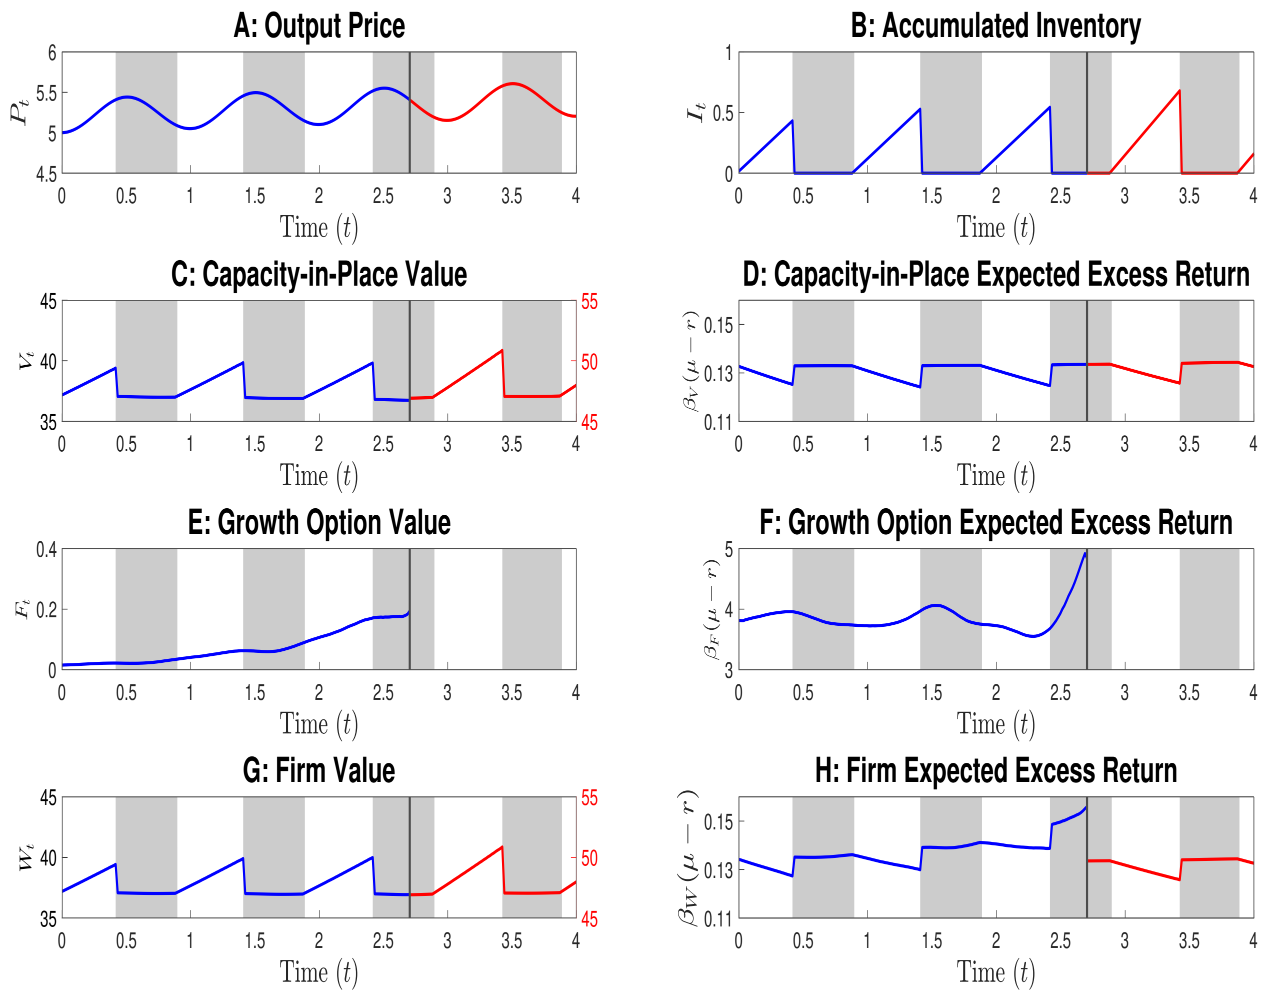

Figure IA.2: The figure plots the output price realization P_t (Panel A), accumulated inventory I_t (Panel B), capacity-in-place value V_t and expected excess return $\beta_V(\mu - r)$ (Panels C and D), growth option value F_t and expected excess return $\beta_F(\mu - r)$ (Panels E and F), and firm value W_t and expected excess return $\beta_W(\mu - r)$ (Panels G and H) over the period from $t = 0$ to 4, all respectively. The vertical line indicates the growth option exercise. The gray areas indicate the periods over which the firm acts as an instantaneous seller. To focus on seasonal patterns, we remove a linear trend from the firm value and expected return in Panels C, D, G, and H. We describe the basecase parameter values in Section IA.1.

pattern in the optimal exercise threshold.[23] Despite that, Panels G and H confirm that, combining the firm's capacity-in-place and growth option into its total assets, the firm's value and expected return

[23]The reason is that the seasonality in the output price induces the option to gradually move deeper out-of-(in)-the-money over inventory-building (instantaneous-selling) periods, leading its expected return to rise (fall).

display seasonal patterns in complete accordance with those of the capacity-in-place, simply because the capacity-in-place makes up the lion share of the firm's value.

Notwithstanding, an interesting twist occurs shortly before the growth option is exercised. To wit, as the output price P_t draws close to the optimal growth-option exercise threshold P_t^*, Panels E and F illustrate that both the value and the expected return of the growth option shoot up. In turn, Panel H demonstrates that the rapid increase in the value and expected return of the growth option can drag up the expected firm return, with the expected firm return becoming flatter (rising) over optimal inventory building (instantaneous selling) periods. That said, the panel finally reveals that the expected firm return immediately returns to its seasonal pattern observed in the main paper or long before the growth option's exercise directly after that option's exercise.

IA.2 Persistence in Seasonalities and Inventory Building

In our theoretical work in Section 2 of our main paper, we assume that the seasonality in a firm's output price and its ability to build up output inventories do not vary over time, allowing us to use past *QInv* values to proxy for future values in our empirical tests. While Grullon et al. (2020) show that a firm's output price seasonality, as measured using *QSales*, is indeed highly persistent, we next establish that its ability to build up output inventories, as measured using *QInv*, is so, too. To do so, Table IA.1 shows the proportions of all (Panel A), seasonal (Panel B), and non-seasonal (Panel C) stocks currently in the low, medium, or high *QInv* tercile portfolio but in the same or another portfolio one (first subpanel), two (second), and five (third) years later. The table confirms that *QInv* is highly persistent, especially in the seasonal-firm subsample. Panel B.1, for example, reports that about 70% of seasonal firms currently in the bottom (alternatively: top) portfolio are in that same portfolio one year later. Even more strikingly, Panel B.3 shows that, even after five years, these two proportions only drop down to about 57%. In contrast, Panel C.3 reveals that the proportions of non-seasonal firms still in the top or bottom *QInv* portfolio after five years are a mere about 35% and 37%, respectively. Overall,

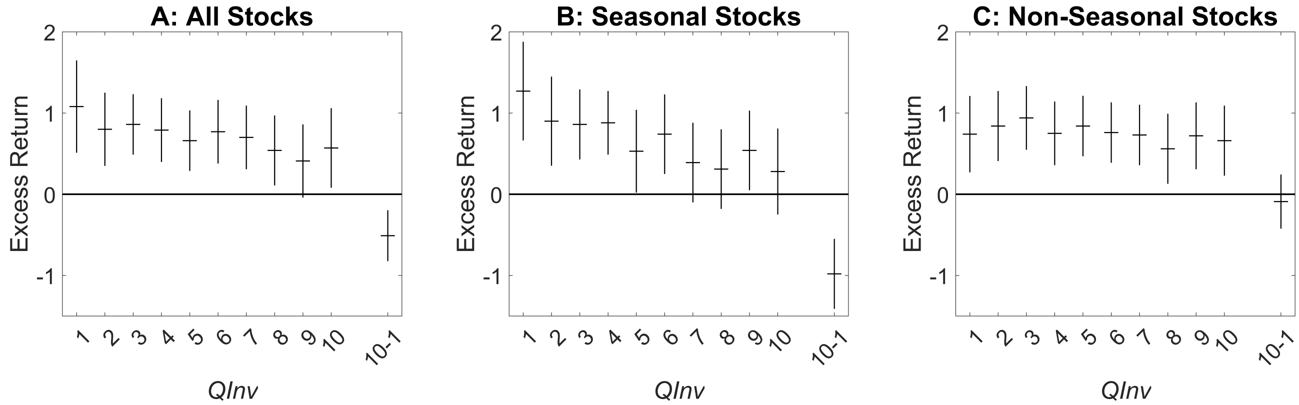

Figure IA.3: The figure plots the mean excess returns of decile portfolios univariately sorted on *QInv* and the high-minus-low spread portfolio ("10-1;" horizontal lines) plus their Newey-West (1987) 95% confidence intervals (vertical lines). Panel A focuses on all stocks, Panel B on seasonal stocks, and Panel C on non-seasonal stocks.

our migration tests offer strong evidence that the ability of firms to build up output inventories to accommodate seasonalities in their output prices does not change much over time.

TABLE IA.1 ABOUT HERE.

IA.3 Portfolio Sort Robustness Tests

We next present robustness tests for our univariate portfolio sorts based on *QInv* in Section 3.4 of our main paper. To begin with, Table IA.2 repeats the portfolio sorts in Table 6 of our main paper using equal (rather than value) weights. Consistent with our main empirical evidence, the table suggests that mean returns and alphas also decline over the equally-weighted *QInv* portfolios, especially when we form those portfolios from only seasonal stocks. Despite that, we notice that *QInv* more negatively prices the value rather than equally weighted portfolios, in line with Grullon et al.'s (2020) evidence that their seasonal sales anomaly is also stronger in value-weighted portfolios.

TABLE IA.2 ABOUT HERE.

In Figure IA.3, we plot the results from repeating the univariate *QInv* sorts using decile (rather than tercile) breakpoints. Panel A focuses on all stocks, Panel B on seasonal stocks, and Panel C on

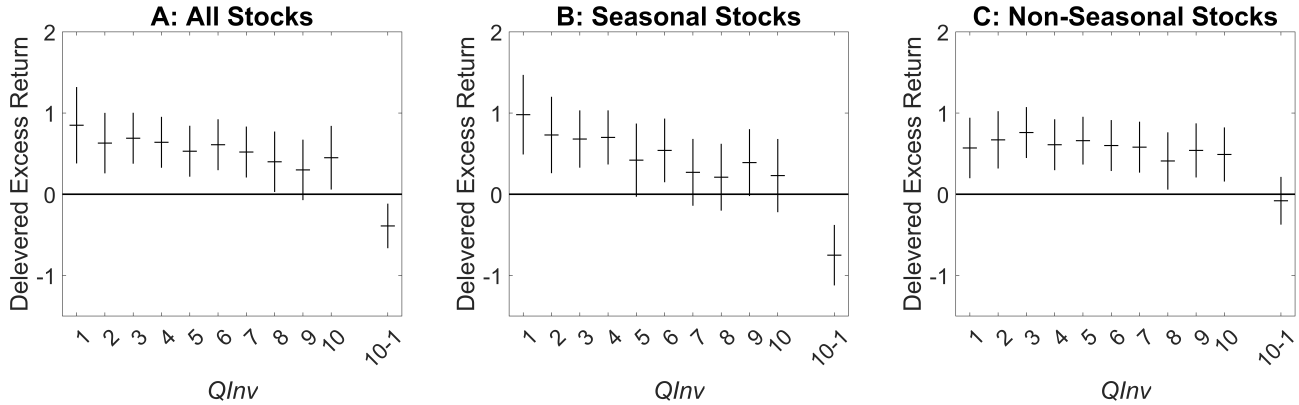

Figure IA.4: The figure plots the mean excess delevered returns of decile portfolios univariately sorted on *QInv* and the high-minus-low spread portfolios ("10-1;" horizontal lines) plus their Newey-West (1987) 95% confidence intervals (vertical lines). Panel A focuses on all stocks, Panel B on seasonal stocks, and Panel C on non-seasonal stocks.

non-seasonal stocks. While the horizontal lines give the mean excess returns of the decile portfolios and the top-minus-bottom spread portfolios, the vertical lines report their 95% confidence intervals derived from Newey and West (1987) standard errors with a twelve-month lag length. The figure suggests that the mean excess returns decline almost monotonically over the portfolios, at least in case of the full sample and the seasonal-stock subsample. In particular, while the mean return of the spread portfolio is –0.51% (t-statistic: –3.12) per month in the full sample, it is a more negative –0.98% (t-statistic: –4.40) in the seasonal-stock subsample. Conversely, that same spread portfolio does not yield a significant mean return in the non-seasonal stock subsample. Adjusting for standard risk factors, the spread portfolio alphas are always close to their mean returns (not shown).

Using the same decile sorts as in Figure IA.3, Figure IA.4 reports the mean excess *delevered* returns of the portfolios plus their 95% confidence intervals. We look into delevered returns since our theory abstracts from financial leverage and thus focuses on asset (and not stock) returns. We follow Doshi et al. (2019) in calculating delevered returns, multiplying the original returns with one minus the stock's financial leverage.[24] The figure shows that *QInv* is also significantly negatively related to future asset returns, with a monthly premium of –0.39% (t-statistic: –2.72). Also, the premium is again an even more significantly negative –0.75% (t-statistic: –3.88) in the seasonal but an only insignificant

[24]Financial leverage is the ratio of total liabilities to the sum of market size and total liabilities, where market size is from the end of the prior calendar year and total liabilities from the fiscal year ending in that calendar year.

–0.08% (t-statistic: –0.52) in the non-seasonal stock subsample. Overall, Figure IA.4 suggests that the pricing of *QInv* in stocks is unlikely to arise from firms' dynamic leverage choices.

IA.4 FM Regression Robustness Tests

In Table IA.3, we document that our FM regressions of single-stock returns on *QInv* and control variables in Section 3.4 of our main paper are robust to adding further control variables proxying for firms' inventory policies and operating flexibility. To be more specific, we add the annual change in inventory *InvGrowth* (Belo and Lin (2012)), inventory turnover *InvTurnover* (Alan et al. (2014)), and asset inflexibility *Inflex* (Gu et al. (2018)). The table suggests that, even in the presence of the additional controls, *QInv* continues to be negatively related to future stock returns in the full sample, the seasonal subsample, and the large inventories subsample, in line with our theory and the results in Table 7 of our main paper. As *InvGrowth* and *InvTurnover* are calculated from annual accounting data and *Inflex* from a time-series of quarterly accounting variables, it is perhaps unsurprising that their inclusion in our regression models cannot subsume the pricing power of *QInv*.

TABLE IA.3 ABOUT HERE.

IA.5 Equally-Well-Diversified Standard and Seasonal-Inventory-Leverage-Neutral Anomaly Spread Portfolios

In Section 4.2 of our main paper, we form the standard anomaly-variable spread portfolios from the entire, seasonal, or non-seasonal stock universe, while we form the seasonal-inventory-leverage-adjusted anomaly-variable spread portfolios only from the corresponding stocks within the middle *QInv* portfolio. We do so since both our theory and intuition dictate that it makes no sense for seasonal firms to hold deflated (inflated) output inventories at the start of their high (low) sales seasons. In accordance, we find that the, for example, high *QSales*-low *QInv* portfolio is indeed sparsely populated, presumably mostly with firms with mismeasured *QInv* values. A concern with our remedy, however, could be that the seasonal-inventory-leverage-adjusted spread portfolios are less well diversified than the standard spread portfolios, with both their short and long legs containing only one-third of the stocks relative to the short and long legs of the standard spread portfolios. In turn, the weaker diversification could raise the volatility of the seasonal-inventory-leverage-adjusted spread portfolios, making it more likely that their mean excess returns are only insignificantly different from zero.

To put the standard and seasonal-inventory-leverage-adjusted spread portfolios on equal footing, Table IA.4 repeats the portfolio sorts in Table 9 of our main paper, this time, however, ensuring that the two types of portfolios contain an identical number of stocks in their short and long legs. To that end, we use only the 33% largest stocks (in terms of market capitalization) in both the long and short leg of the standard spread portfolios. In contrast, we use the same stocks as in Table 9 in the short and long legs of the seasonal-inventory-leverage-adjusted spread portfolios. The table suggests that the updated standard spread portfolios yield mean excess returns and *t*-statistics almost identical to the original ones. While the original full-sample standard *QSales* spread portfolio, for example, yields a mean monthly return of –0.22% (*t*-statistic: –2.17), the corresponding number for the updated portfolio is –0.24% (*t*-statistic: –2.32). Given the close alignment between the original and updated standard spread portfolio results, we have decided, for the sake of simplicity, to continue to report the better diversified standard spread portfolio results in our main paper.

TABLE IA.4 ABOUT HERE.

References

Alan, Yasin, George P. Gao, and Vishal Gaur, 2014, Does inventory productivity predict future stock returns? A retailing industry perspective, *Management Science* 60, 2416–2434.

Belo, Frederico and Xiaoji Lin, 2012, The inventory growth spread, *Review of Financial Studies* 25, 278–313.

Black, Fischer and Myron C. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Dixit, Avinash K. and Robert S. Pindyck, 1994, *Investment under Uncertainty*, Princeton University Press, Princeton, NJ.

Doshi, Hitesh, Kris Jacobs, Praveen Kumar, and Ramon Rabinovitch, 2019, Leverage and the cross-section of equity returns, *Journal of Finance* 74, 1431–1471.

Fama, Eugene F. and James D. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

Grullon, Gustavo, Yamil Kaba, and Alexander Núñez, 2020, When low beats high: Riding the sales seasonality premium, *Journal of Financial Economics* 138, 572–591.

Gu, Lifeng, Dirk Hackbarth, and Timothy Johnson, 2018, Inflexibility and stock returns, *Review of Financial Studies* 31, 278–321.

McDonald, Robert L. and Daniel R. Siegel, 1986, The value of waiting to invest, *Quarterly Journal of Economics* 101, 707–727.

Newey, Whitney K. and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Table IA.1: Portfolio Migration Matrixes

The table gives the average portfolio migration matrixes of univariate tercile portfolios sorted based on *QInv*. We report those matrixes for the full sample (Panel A), the seasonal-stock subsample (Panel B), and the non-seasonal-stock subsample (Panel C). Each row of the table offers the proportions of stocks currently in the low (first row in each subpanel), medium (second), or high (third) *QInv* portfolio but in the low (first column), medium (second), or high (third) portfolio one year (subpanel 1), two years (subpanel 2), or five years (subpanel 3) later, averaged over our sample period.

	QInv		
	Future Low	Future Medium	Future High
	(1)	(2)	(3)
	Panel A: Full Sample		
	Panel A.1: One Year		
Current Low	0.64	0.24	0.12
Current Medium	0.22	0.54	0.24
Current High	0.11	0.24	0.64
	Panel A.2: Two Years		
Current Low	0.46	0.29	0.24
Current Medium	0.26	0.46	0.28
Current High	0.23	0.30	0.47
	Panel A.3: Five Years		
Current Low	0.43	0.32	0.24
Current Medium	0.25	0.47	0.28
Current High	0.23	0.33	0.44
	Panel B: Seasonal Subsample		
	Panel B.1: One Year		
Current Low	0.70	0.18	0.12
Current Medium	0.27	0.46	0.27
Current High	0.12	0.18	0.70
	Panel B.2: Two Years		
Current Low	0.56	0.22	0.22
Current Medium	0.31	0.38	0.31
Current High	0.22	0.21	0.57
	Panel B.3: Five Years		
Current Low	0.57	0.23	0.20
Current Medium	0.31	0.39	0.30
Current High	0.21	0.22	0.57

(continued on next page)

Table IA.1: Portfolio Migration Matrixes (cont.)

	QInv		
	Future Low	Future Medium	Future High
	(1)	(2)	(3)
Panel C: Non-Seasonal Subsample			
Panel C.1: One Year			
Current Low	0.59	0.32	0.10
Current Medium	0.18	0.62	0.20
Current High	0.09	0.31	0.59
Panel C.2: Two Years			
Current Low	0.39	0.39	0.21
Current Medium	0.21	0.55	0.24
Current High	0.20	0.40	0.40
Panel C.3: Five Years			
Current Low	0.35	0.44	0.21
Current Medium	0.20	0.57	0.23
Current High	0.19	0.44	0.37

Table IA.2: Equally-Weighted Univariate Portfolios Sorted on Abnormal Inventories
The table presents the results of portfolios univariately sorted on *QInv*. At the end of each month $t-1$, we first choose either all stocks (Panel A), those with a *Seasonality* value above the median (Panel B, "seasonal stocks"), or those with a *Seasonality* value below that median (Panel C, "non-seasonal stocks"). We next sort the stocks into portfolios according to the unconditional terciles of the *QInv* distribution at that time. We equally-weight the portfolios and hold them over month t. We also form spread portfolios long the top and short the bottom equally-weighted portfolios ("High – Low"). In columns (1) to (3), we report the mean excess returns (Return^e), the mean number of stocks (# Stocks), the mean cross-sectional *QSales* averages (*QSales*), the mean cross-sectional *MarketSize* averages (*MarketSize*), and the annualized Sharpe ratios (mean excess returns over volatilities) of the portfolios. Conversely, columns (4) to (6) report the mean excess returns, q^5-model alphas, and FF6-model alphas of the spread portfolios. Plain numbers are estimates, while those in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. See Table C1 in the appendix of our main paper for more details about variable definitions.

	QInv			High – Low		
	Low	Medium	High	Mean	q^5-Alpha	FF6-Alpha
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Full Sample						
Return^e	1.05	0.97	0.85	−0.19	−0.16	−0.19
	[3.64]	[3.74]	[2.94]	[−3.52]	[−2.47]	[−3.27]
# Stocks	708	708	729			
QSales	0.24	0.25	0.26			
MarketSize	5.28	5.70	5.25			
Sharpe (p.a.)	0.56	0.58	0.45			
Panel B: Seasonal Subsample						
Return^e	0.97	0.91	0.74	−0.23	−0.19	−0.22
	[3.12]	[3.10]	[2.39]	[−3.48]	[−2.42]	[−3.09]
# Stocks	409	267	396			
QSales	0.24	0.25	0.26			
MarketSize	5.07	5.22	4.99			
Sharpe (p.a.)	0.48	0.48	0.37			
Panel C: Non-Seasonal Subsample						
Return^e	1.14	1.00	0.99	−0.16	−0.15	−0.18
	[4.39]	[4.12]	[3.67]	[−2.43]	[−2.12]	[−2.61]
# Stocks	300	441	332			
QSales	0.25	0.25	0.25			
MarketSize	5.56	5.98	5.55			
Sharpe (p.a.)	0.68	0.64	0.57			

Table IA.3: Fama-MacBeth Regressions of Stock Returns on Abnormal Inventories Also Controlling for Other Optimal Inventory Policy and Operating Flexibility Variables

The table presents the results from Fama-MacBeth (1973) regressions of single-stock returns over month t on *QInv*, *QCash*, and *QDebt* rank variables, *InvToSales*, *InvGrowth*, *InvTurnover*, *Inflex*, and the standard FF6 controls, *MarketBeta*, *MarketSize*, *BookToMarket*, *Momentum*, *Investment*, and *Profitability*, all measured using data until the start of that month. While we run the regression on the full sample in column (1), we run it on only those firms with a *Seasonality* value above or below the cross-sectional median at the end of the prior June in columns (2) and (3), respectively. Conversely, we run it on only those firms with a total quarterly inventories-to-quarterly annual sales ratio value above or below the cross-sectional median at the end of the prior June in columns (4) and (5), respectively. The plain numbers are monthly premium estimates, in percent. The numbers in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. For the sake of brevity, we do not report the effects of the FF6 controls. See Table C1 in the appendix of our main paper and Section IA.4 in this Internet Appendix for more details about variable definitions.

	Full Sample	Seasonal Subsample	Non-Seasonal Subsample	Large Inventories Subsample	Small Inventories Subsample
	(1)	(2)	(3)	(4)	(5)
QInv (Rank)	−0.34	−0.41	−0.29	−0.59	−0.19
	[−4.32]	[−3.79]	[−2.60]	[−4.52]	[−2.12]
QCash (Rank)	0.16	0.29	−0.01	0.08	0.17
	[2.32]	[2.87]	[−0.19]	[1.01]	[1.71]
QDebt (Rank)	0.01	−0.05	0.13	0.18	−0.10
	[0.18]	[−0.43]	[1.58]	[1.84]	[−0.94]
InvToSales	−0.10	−0.14	−0.08	−0.17	0.23
	[−1.59]	[−1.70]	[−0.91]	[−1.69]	[1.17]
InvGrowth	−0.04	−0.01	−0.02	−0.29	0.03
	[−0.68]	[−0.11]	[−0.21]	[−2.53]	[0.44]
InvTurnover	0.00	0.00	0.00	0.00	0.00
	[−1.88]	[−2.54]	[−0.67]	[0.10]	[−1.13]
Inflex	0.00	0.00	0.00	0.00	0.00
	[0.33]	[−0.37]	[0.29]	[0.63]	[0.24]
Constant	0.77	0.92	0.68	0.58	0.83
	[2.16]	[2.33]	[1.84]	[1.45]	[2.15]
FF6 Controls	Yes	Yes	Yes	Yes	Yes

Table IA.4: Equally-Well-Diversified Standard and *QInv*-Neutral Anomaly Spread Portfolios
In this table, we present the mean returns (columns (1) and (2)), augmented-q-theory-model alphas ("q^5;" (3) and (4)), and six-factor-model alphas ("FF6;" (5) and (6)) of seasonal sales (Panel A), same-calendar-month return (Panel B), momentum (Panel C), and quarterly ROE (Panel D) spread portfolios. At the end of each month $t-1$, we form the spread portfolios from independently sorting our sample stocks into portfolios based on the median of *Seasonality*, the terciles of the anomaly variable, and the terciles of *QInv* at that time. We value-weight the portfolios and hold them over month t. We then consider either all stocks ("All Stocks"), only those in the top *Seasonality* portfolio ("Seasonal Stocks"), or only those in the bottom Seasonality portfolio ("Non-Seasonal Stocks"). We finally form standard anomaly spread portfolios through longing the top-anomaly-variable and shorting the bottom portfolios (columns (1), (3), and (5)), while we form *seasonal-inventory-leverage-neutral* spread portfolios through longing the top-anomaly variable and shorting the bottom portfolios within the middle *QInv* portfolio ((2), (4), and (6)). Importantly, we only consider the 33% largest stocks (in terms of market capitalization) in the short and long legs of the standard portfolios, to ensure a similar diversification across the two types of portfolios. Plain numbers are estimates, while those in square brackets are Newey-West (1987) t-statistics with a twelve-month lag length. See Table C1 in the appendix of this paper for more details about the anomaly variable definitions.

	Mean Spread Return		q^5-Spread Alpha		FF6-Spread Alpha	
	Modified Standard	*QInv* Neutral	Modified Standard	*QInv* Neutral	Modified Standard	*QInv* Neutral
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Seasonal Sales Anomaly						
All Stocks	−0.24	−0.10	−0.15	0.02	−0.20	−0.04
	[−2.32]	[−0.80]	[−1.16]	[0.12]	[−1.80]	[−0.26]
Seasonal Stocks	−0.36	−0.16	−0.27	−0.04	−0.30	−0.03
	[−2.69]	[−0.96]	[−1.60]	[−0.20]	[−2.12]	[−0.20]
Non-Seasonal Stocks	−0.06	−0.02	0.04	0.10	0.02	0.03
	[−0.50]	[−0.11]	[0.29]	[0.55]	[0.19]	[0.18]
Panel B: Same-Calendar-Month Anomaly						
All Stocks	0.49	0.41	0.49	0.45	0.49	0.45
	[2.83]	[2.06]	[2.80]	[2.06]	[2.76]	[2.11]
Seasonal Stocks	0.67	0.65	0.57	0.66	0.61	0.61
	[3.12]	[2.42]	[2.34]	[1.75]	[2.61]	[1.81]
Non-Seasonal Stocks	0.46	0.41	0.53	0.43	0.45	0.44
	[2.81]	[2.35]	[2.99]	[2.45]	[2.62]	[2.36]

(continued on next page)

Table IA.4: Equally-Well-Diversified Standard and *QInv*-Neutral Anomaly Spread Portfolios (cont.)

	Mean Spread Return		q^5-Spread Alpha		FF6-Spread Alpha	
	Modified Standard	*QInv* Neutral	Modified Standard	*QInv* Neutral	Modified Standard	*QInv* Neutral
	(1)	(2)	(3)	(4)	(5)	(6)
Panel C: Momentum Anomaly						
All Stocks	0.50	0.34	−0.12	−0.09	−0.07	−0.08
	[2.19]	[1.29]	[−0.47]	[−0.38]	[−0.46]	[−0.43]
Seasonal Stocks	0.86	0.93	0.10	0.24	0.23	0.37
	[3.29]	[2.77]	[0.38]	[0.77]	[1.40]	[1.52]
Non-Seasonal Stocks	0.24	0.11	−0.16	−0.15	−0.19	−0.22
	[1.01]	[0.44]	[−0.60]	[−0.60]	[−1.31]	[−1.25]
Panel D: Quarterly-ROE Anomaly						
All Stocks	0.37	0.24	−0.16	−0.42	0.09	0.06
	[1.87]	[1.05]	[−1.19]	[−2.12]	[0.71]	[0.40]
Seasonal Stocks	0.47	0.38	−0.04	−0.28	0.20	0.13
	[2.21]	[1.33]	[−0.26]	[−0.90]	[1.34]	[0.50]
Non-Seasonal Stocks	0.24	0.17	−0.25	−0.24	0.01	0.05
	[1.05]	[0.69]	[−1.66]	[−1.22]	[0.05]	[0.27]