

An Empirical Investigation of Dynamic Ordering Policies

Chad R. Larson

Bauer College of Business, University of Houston, Houston, Texas 77004, crlarson@uh.edu

Danko Turcic, Fuqiang Zhang

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130
{turcic@wustl.edu, fzhang22@wustl.edu}

An adaptive base stock policy is a well-known tool for managing inventories in nonstationary demand environments. This paper presents empirical tests of this policy using aggregate, firm-level data. First, we extend a single-item adaptive base stock policy in previous literature to a multi-item case. Second, we transform the policy derived for the multi-item case to a regression model that relates firm-level inventory purchases to firm-level sales and changes in sales forecasts. We focus on two research questions: Can the adaptive base stock policy explain cross-sectional ordering behaviors under sales growth? To the extent that the adaptive base stock policy fails to explain ordering behaviors under sales growth, are there frictions that explain such a finding? Our empirical results demonstrate disparities in ordering behaviors between firms experiencing high and moderate sales growth. Contrary to theoretical prediction, this implies that inventory purchases are a function of not only current sales and changes in sales forecast but also past sales growth. As potential explanations for this departure from theoretical prediction, we show that both future demand dynamics and inventory holding risks depend on past sales growth. In addition, we find that firms' inventory holding risks may also be affected by purchasing constraints imposed by supply chain contracts. Our results provide managerial implications for practitioners and inform future theoretical research.

Keywords: inventory policies; stochastic models; forecasting; ARIMA processes; nonstationary demand

History: Received November 6, 2012; accepted August 24, 2014, by Martin Lariviere, operations management.

Published online in *Articles in Advance* March 12, 2015.

1. Introduction

Inventory management problems are ubiquitous in business environments, and a large body of theories has been developed to study optimal inventory policies in various contexts (Zipkin 2000). Formulas are commonly available to compute optimum or near-optimum order quantities at a single stock keeping unit (SKU) level. However, little is known about what these tools imply (if anything) about the behaviors of inventories aggregated at firm or economy levels. Economists, aggregate planners, and empiricists have an interest in aggregate inventories (Rumyantsev and Netessine 2007, Gaur and Kesavan 2009). For example, economists study aggregations of firm inventories and how they contribute to an economy's GDP volatility (Khan and Thomas 2007). Operations planners aggregate inventories across multiple SKUs to evaluate a firm's capacity requirements in staffing, raw materials, and warehousing. Consequently, there is a need to understand the link between inventory theory and real data (Wagner 2002). The theory, however, has been developed predominately at the single-item level, and the majority of inventory data

is available only at the firm level, aggregated across SKUs. This mismatch between theory and data limits the use of theoretical inventory models in empirical research.

Rumyantsev and Netessine (2007) were among the first to deal with *SKU-to-firm* aggregation. Using aggregate inventory data, they examine whether insights from classic, SKU-level models hold at the firm level. Their work complements earlier research from Caplin (1985), who contributes to *firm-to-economy* aggregation. The results in Caplin (1985) imply that if there are multiple firms with dependent demands, and each firm follows an (S, s) policy, then, in expectation, aggregate inventory purchases in each period will always replace aggregate sales from the immediate period. In a later study, Mosser (1991) confirms this finding empirically (see Theorem 3.4 in Caplin 1985 and Equations (3)–(5) in Mosser 1991). All these studies, however, assume that successive demands are independent and identically distributed (i.i.d.) random variables (see p. 1396 in Caplin 1985 and p. 414 in Rumyantsev and Netessine 2007). This assumption is restrictive because it precludes phenomena such as

sales growth and seasonality (e.g., see Johnson and Thompson 1975, Lovejoy 1992, Morton and Pentico 1995, Graves 1999).

At the SKU level, *adaptive base stock policy* is a broadly publicized tool for managing inventories when demands are *not* i.i.d. (for details, see Veinott 1965b, Graves 1999). Extant research, however, does not explore whether the adaptive base stock policy explains changes in inventories observed in real data. This paper fills this gap using two steps. First, we extend the single-item adaptive base stock policy presented in Graves (1999) to a multi-item case, allowing stochastic demands to be dependent across items and time. Following Veinott (1965a), we derive a myopic inventory purchasing rule that minimizes the expected discounted costs over an infinite time horizon. This rule consists of two components: The first replenishes observed sales in the immediate period, and the second adjusts the base-stock level to accommodate changes in demand forecasts. In the second step, we transform the policy derived in the first step to a regression model that relates firm-level inventory purchases to firm-level sales and firm-level changes in sales forecasts. The goal is to answer the following questions: (a) Can the adaptive base stock policy explain cross-sectional ordering behavior under sales growth? (b) To the extent that the adaptive base stock policy fails to explain ordering behavior under sales growth, do frictions explain the finding?

Our primary finding is that firms' inventory purchases are not fully explained by the adaptive base stock policy. The policy suggests that inventory purchases should be a function only of current sales and change in sales forecast. However, we find significant differences in ordering behaviors between firms that experienced high and moderate growth in the prior period. In particular, high-growth firms purchase less inventory than predicted by the benchmark policy, implying that inventory purchases are also a function of past sales growth even after adjusting for forecasts and many control variables.

We further explore this finding by investigating the validity of the assumptions underlying the adaptive base stock policy. First, using a generalized autoregressive conditional heteroskedastic (GARCH) model, we document that the volatility of future demand depends on an information process that includes past demand changes. High-growth firm-quarters have future demand volatility that is, on average, over five times that of moderate-growth firm-quarters. This finding contrasts previous results from Montgomery and Johnson (1976, pp. 217–221), who find that demands of some products follow stationary, moving-average processes. It also contrasts familiar demand processes considered in theoretical literature. From Johnson and Thompson (1975), Lamond and Sobel

(1995) and Graves (1999), demand forecast errors do not depend on the magnitude of the demand process; Song and Zipkin (1996) assume i.i.d. demand with a probable negative shock to mean demand; and from Morton and Pentico (1995), demands are independent but not necessarily distributed identically. In practice, we show that firms face demand that is both autoregressive as Johnson and Thompson (1975) and Graves (1999) consider, and has state-dependent variance as Morton and Pentico (1995) and Song and Zipkin (1996) argue. Second, we find that the inventory purchasing restraint among high-growth firms coincides with the risk of higher inventory holding costs because of an association between higher demand growth and an increased probability of incurring significant inventory obsolescence. We link high realized holding costs to high demand growth, and thus findings suggest that high-growth firms are ordering more conservatively because of state-dependent holding costs. This link also contrasts with theoretical literature that routinely assumes holding costs are linear with stationary coefficients (Zipkin 2000, Porteus 2002).

Lastly, we document that not all high-growth firms reduce inventory purchases to offset increased inventory holding costs that accompany higher growth. Using a unique sample of supply chain contracts, we argue that these high-growth, high-holding-cost firms are much more likely to have entered supply chain contracts that limit purchasing flexibility. This evidence of contractual inflexibility explains why these high-growth, high-holding-cost firms do not adjust their purchasing downwardly similar to other high-growth firms. Results therefore indicate that demand, inventory holding cost, and supply chain frictions exist in a nontheoretical world; these previously unexplored frictions affect firms' ordering behavior and help explain empirical observations that are inconsistent with what the theoretical literature predicts.

The rest of this paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the data and variables. Section 4 develops the hypotheses, and the main results are presented in §5. Section 6 concludes the paper.

2. Previous Research

Behaviors of aggregate inventories have been an area of intense inquiry in empirical inventory literature. Rajagopalan and Malhotra (2001) study U.S. manufacturing inventories from 1961 to 1994, reporting that inventory levels have been decreasing because of advances in inventory theory, developments of information technologies, and adoption of just-in-time manufacturing techniques. Chen et al. (2005, 2007) analyze inventory trends in U.S. public companies between 1981 and 2001, and make similar

observations: median manufacturing inventory levels declined from 96 to 81 days; in the retail and wholesale segments, median inventory levels decreased from 72 to 52 days during the same period. Moreover, they show that public companies with abnormally high inventory levels experienced abnormally low levels of financial returns, but on average, lower inventory levels do not necessarily associate with higher financial returns. Gaur et al. (2005) examine firm-level inventory behaviors at retail firms, and propose a model that explains disparities in inventory turns across companies. Using their model, they demonstrate that a portion of this variation can be explained by gross margin, capital intensity, and sales surprise (i.e., the ratio of sales to expected sales for the year). In addition, they find important links between the behaviors of aggregate inventories and financing by showing that inventory turnover for retail firms correlates positively with capital intensity.

Gaur and Kesavan (2009) extend the model presented by Gaur et al. (2005) to include the effects of firm size and sales growth rate on inventory turnover. Regarding size, they find strong evidence of diminishing returns to scale, and regarding sales growth rate, they find inventory turnover increases with sales growth rate. Cachon and Olivares (2010) examine drivers of finished goods inventory in the U.S. automobile industry, and find that the difference in finished goods inventory at the major auto manufacturers can largely be explained by two factors: number of dealerships in a firm's distribution network and a firm's production flexibility. Although these studies deal with aggregate inventories, they do not necessarily test implications stemming from classic inventory models, one aspect that distinguishes these papers from ours.

Recently, there has been considerable interest in understanding how inventory and supply chain management affect firm profitability, and how potential investors might view firms with investments in inventory. Papers in this area clarify that supply chain glitches have long-lasting, negative effects on firm performance (Hendricks and Singhal 2005a, b), holding excess inventory has a negative effect on stock prices (Hendricks and Singhal 2009), firms reduce inventory when the market discounts high inventory firms (Lai 2006), historic inventory and gross margin contain information useful to forecast sales (Kesavan et al. 2010), and inventory turns move counter-cyclically with macroeconomic shocks (Kesavan and Kushwaha 2011).

In addition, there is a significant body of literature that investigates the relationship between the variability in production and variability in sales. For example, in the operations management literature, Cachon et al. (2007) compare the variances of production and

sales to examine the strength of the bullwhip effect. Related research in economics focuses on the role inventories play in production volatility and business cycles; see Caplin (1985), Kahn (1987), Mosser (1991), and Khan and Thomas (2007) for representative studies. Although the questions we ask in this paper are not necessarily the same, our paper is connected to these studies: We start our analysis by constructing an SKU-to-firm model of aggregate inventory behaviors; this step is similar to Caplin (1985), who constructs a firm-to-economy aggregation model. Using our aggregate model, we construct a regression model that we use in our empirical analysis; this step is analogous to Mosser (1991), who constructs a regression model from the results in Caplin (1985).

Our paper relates most closely to Rummyantsev and Netessine (2007), who test whether a variety of stocking predictions obtained from classic, SKU-level inventory models are proxies for stocking behaviors at the firm level. The study finds that product-level predictions are robust and extend beyond individual products to the aggregate firm level. Note, however, that the study stops short of proposing an SKU-to-firm aggregation model. Using several well-known SKU-level models, the authors formulate a number of hypotheses and test whether they hold for aggregate inventories. Their study, however, does not control for sales growth. Consequently, when it comes to accumulating inventories, firms experiencing high sales growth are assumed to face the same cost-benefit tradeoffs as firms whose sales are declining. Our paper differs from Rummyantsev and Netessine (2007) regarding modeling SKU-to-firm aggregation and how sales growth affects aggregate inventory purchases. In contrast to Rummyantsev and Netessine (2007), we document that the SKU-to-firm aggregate model is not necessarily a proxy for stocking behaviors at the firm level if a firm is experiencing high sales growth. Relative to extant literature, we characterize the behaviors of aggregate inventories in firms that face nonstationary demand environments, and empirically test it. To the extent that the model fails to predict empirical reality, we identify frictions that may be the culprits behind negative findings.

3. Samples and Variables

We examine two primary samples in our analysis: a large and small sample. Our large sample consists of data from the COMPUSTAT fundamentals quarterly file spanning from 1980 to 2008. The small sample uses data from the COMPUSTAT fundamentals annual file and inventory obsolescence data collected from annual reports. The size of our obsolescence sample is restricted by the cost of collecting data, and is thus limited to 2002 through 2004. The primary financial variables of interest are changes in

Table 1 Sample Descriptive Statistics

Panel A: Dependent and independent variables					
Variable	Mean	Std. dev.	P25	Median	P75
$PURCHASES_t/COGS_t$	1.032	0.219	0.958	1.009	1.083
$\Delta SALES_t$	0.039	0.384	-0.039	0.082	0.206
$INVTURN_t$	2.290	3.516	0.692	1.143	2.150
GM_t	0.303	0.307	0.201	0.314	0.453
Panel B: Industry classification					
Two-digit SIC codes	Industry	Firm-quarter observations	Sample percentage		
36	Electronic and other electrical equipment	45,549	12.81		
35	Industrial and commercial machinery	35,398	9.96		
38	Measuring and analyzing instruments	29,478	8.29		
28	Chemicals and allied products	25,862	7.27		
73	Business services	17,881	5.03		
50	Wholesale trade—Durable goods	14,704	4.14		
20	Food and kindred products	11,717	3.30		
34	Fabricated metal productions	10,392	2.92		
37	Transportation equipment	10,349	2.91		
59	Miscellaneous retail	9,432	2.65		
13	Oil and gas extraction	8,716	2.45		
51	Wholesale trade—Nondurable goods	8,403	2.36		
33	Primary metal industries	7,889	2.22		
58	Eating and drinking places	7,693	2.16		
	Other	112,033	31.53		
	Total	355,496	100.00		

Notes. The sample consists of all nonfinancial and utility firm-quarters listed on the NYSE, NASDAQ, or AMEX, CRSP, and the COMPUSTAT fundamentals quarterly file between 1980 and 2008 with inventory greater than 1% of average total assets and available data to calculate change in sales for periods $(t - 1)$, t , and $(t + 1)$, and purchases to cost of goods sold in period t . Sample consists of 355,496 firm-quarter observations. $PURCHASES_t/COGS_t$ is measured as purchases (ending inventory plus cost of goods sold minus beginning inventory) divided by cost of goods sold. $\Delta SALES_t$ is measured as four quarters ago percentage change in sales. $INVTURN_t$ is cost of goods sold scaled by average inventory. GM_t is sales minus cost of goods sold scaled by sales.

sales ($\Delta SALES_t$) and purchases scaled by cost of goods sold ($PURCHASES_t/COGS_t$). Purchases are calculated as cost of goods sold plus ending inventory, minus beginning inventory. We measure quarterly changes in sales as the percentage change in sales from four quarters previous. This controls for the effects of seasonality in quarterly sales. Annual changes are measured as year-over-year percentage changes in sales.¹ We require the availability of COMPUSTAT data for each of the variables for period t and change in sales for periods $(t - 1)$ and $(t + 1)$. Because our interest is examining firms' ordering policies with respect to inventory, we eliminate firms with insignificant inventory levels: fewer than 1% of average total assets. Consistent with other studies, we find that a small number of extreme outliers characterize the

distributions of the financial variables. Therefore, we follow the standard procedure of winsorizing financial variables at the 1st and 99th percentiles of the distribution.

In addition to COMPUSTAT data, we also use analyst sales forecasts from IBES. The final large sample consists of 355,496 firm-quarter observations from 1980 through 2008.² Table 1 provides descriptive statistics and a breakdown of firm-quarter observations by industry using two-digit Standard Industrial Classification (SIC) codes. Panel A provides descriptive statistics for our primary variables of interest and control variables that we employ and discuss in further detail in the following section. $PURCHASES_t/COGS_t$ is approximately 1 for the median firm, implying the median firm purchases the same amount of inventory as it sells each quarter. The mean and median of sales growth, $\Delta SALES_t$,

¹ For the purpose of aligning firms in event time, when using quarterly data, we use fiscal years and quarters as reported by COMPUSTAT. COMPUSTAT records the fiscal year as the year in which the majority of calendar months included in the fiscal year occurred. Because all firms do not have the same fiscal year-ends, this results in some mismatches of firms in calendar time. We replicate quarterly results using only 12/31 fiscal year-end firms, and our inferences are unchanged.

² As in extant literature examining inventory, we eliminate all financial services companies (SIC 6000–6999) and utilities (SIC 4900–4999). For comparability, we also limit the sample to domestic firms (popsrc = D and fic = USA) traded on the NYSE, NASDAQ, or AMEX.

are 3.9% and 8.2%, respectively. Inventory turnover, $INVTURN_t$, is calculated as cost of goods sold scaled by average inventory, and gross margin, GM_t , is calculated as sales minus cost of goods sold, scaled by sales. The variables are symmetrically distributed, though inventory turnover and the ratio of purchases to cost of goods sold are both slightly right-skewed, and sales growth is slightly left-skewed.

For succinctness, panel B only includes industries that represent more than 2% of sample observations. The majority of industries fell into this “other” category, representing 31.53% of sample observations. The most heavily represented industry is the electronic and other electrical equipment industry, comprising 12.81% of sample observations. The next three largest industries, in order, are industrial and commercial machinery, measuring and analyzing instruments, and chemicals and allied products industries.

To create the small sample, we collect inventory obsolescence data from annual reports, Securities and Exchange Commission (SEC) 10-K filings.³ Accounting rules require firms to take an impairment charge to income by reducing the book value of inventory to current market value if its market value drops below its original cost. Thus, accounting rules require that as inventory becomes obsolete, firms reflect the reduction in value in financial statements. Accountants commonly refer to this reduction in value as an inventory write-down or write-off.

Using Morningstar’s 10-K Wizard, we collect data by conducting a keyword search for any occurrence of the word “write” within 10 words of the word “inventory.” We then filter these results using the CIK numbers of all firms that appeared in both the CRSP and COMPUSTAT databases for the years 2001 through 2004 and were traded on a major U.S. exchange.⁴ After filtering, we identify 5,638 filings. The number of filings relative to the final sample illustrates the challenge of using only a keyword search to identify inventory obsolescence events. The majority of firms with a material amount of inventory mention in their annual reports the possibility of inventory write-downs or the accounting requirement to write down inventory if the value declines significantly. Therefore,

identifying write-downs requires us to then individually search and read each 10-K for discussion or documentation of an inventory write-down during the fiscal year.

Finding evidence of an inventory write-down, we collect the inventory write-down amount from the annual report.⁵ After collecting the inventory obsolescence data, we merge them with data from the COMPUSTAT fundamental annual file to create a sample of inventory obsolescence and nonobsolescence firm-years. Our search yields many inventory charges that are insignificant in magnitude and serial in frequency. Because we are interested in obsolescence events, we require that the charge be equal to or larger than 1% of average total assets for inclusion in our sample. In addition, because we empirically observe that many firms take subsequent charges following a first large charge, we limit the sample to the first significant inventory impairment we identify in the search. The final sample of first-time, significant inventory obsolescence firm-years between 2002 and 2004 with available data consist of 290 observations.

In the last set of analyses, we use our small sample in conjunction with a novel sample of supply chain contracts to explore firm supply flexibility. Securities rules require that firms disclose material contracts in filings with the SEC. These material contract disclosures allow us to observe supply chain contracting behaviors directly. Despite this great resource, there are a few limitations. The first is that firms are only required to disclose what are considered material contracts. Materiality is generally defined as contracts that represent over 10% of a firm’s business. The second limitation is that firms may redact information that if disclosed, might divulge proprietary information. This information typically includes selected pricing information and technical specifications. Despite these limitations, these contracts provide a previously untapped resource for empirically testing supply chain relationships. We discuss collection, reading, and coding of these supply chain contracts in detail in §4.2.

4. Hypothesis Development

4.1. Adaptive Replenishment Policy

We begin with the primary research question of whether an adaptive base stock policy explains

³ We search annual financial statements rather than quarterly financial statements for two reasons. First, firms provide more financial details in annual filings. As a result, a quarterly analysis is likely to miss many obsolescence events. Second, collecting data from financial statements is expensive. Using quarterly statements would quadruple the number of financial statements to read to over 20,000.

⁴ CIK stands for “central index key,” and is used on the SEC’s computer systems to identify corporations and individuals who filed disclosures with the SEC. COMPUSTAT also includes this identifier making it a useful identifier for matching COMPUSTAT and SEC data.

⁵ The language in annual reports detailing inventory write-downs is largely uniform. Here is an example of a description from the 2004 10-K annual report of pharmaceutical company, Chiron Corp.: “In 2004, gross profit margins decreased to 47% from 58% in 2003, primarily because of the write-off of our entire inventory of FLUVIRIN vaccine, resulting in a \$91.3 million charge to cost of sales in the third quarter 2004, as well as the fact that there were no FLUVIRIN vaccine sales for the 2004–2005 season” (<http://www.getfilings.com/o0001104659-05-011453.html>).

cross-sectional firm ordering behaviors. The inherent difficulty with testing this question empirically lies in the availability of suitable data. Although nearly all inventory models predict ordering behaviors at the single-item level, most firms carry multiple items (or SKUs), and only report inventory purchases and inventory positions at the firm level, aggregated across SKUs. We close the gap between theory and data by constructing an aggregate model that predicts ordering behaviors at the firm level. We do this in two steps. In the *first step*, we extend the single-item model from Graves (1999) by considering a dynamic, non-stationary, multiproduct inventory model. Demands in the model are random vectors, which are not necessarily independent or distributed identically (for details, see Equations (6) in the appendix). There are no fixed ordering costs and unfilled demand is backlogged. Ordering, holding, and back-order costs are linear with stationary coefficients. We do, however, allow a fixed delivery lead time. Future costs are discounted. The objective is to choose an ordering rule that minimizes expected discounted costs over an infinite time horizon. This rule is given in Proposition 1, which we prove in the appendix.

PROPOSITION 1. *Consider a dynamic inventory problem with n products labeled $1, \dots, n$. Let $\mathbf{q}_t = (q_t^1, q_t^2, \dots, q_t^n)$ be the vector of orders placed during period t for delivery during period $t + L$; $\mathbf{D}_t = (D_t^1, D_t^2, \dots, D_t^n)$ and $\mathbf{F}_t = (F_t^1, F_t^2, \dots, F_t^n)$ are the vectors of period t demands and period t demand forecasts (see Equations (6) and (8) in the appendix). The following ordering rule minimizes expected discounted costs over an infinite time horizon:*

$$\mathbf{q}_t = \mathbf{D}_t + L(\mathbf{F}_{t+1} - \mathbf{F}_t). \quad (1)$$

The ordering rule we present previously is effectively a myopic ordering policy that minimizes the per-period expected costs one lead time into the future. There are two components to the order vector, \mathbf{q}_t , given in Proposition 1: The first replenishes the sales observed in the current period, and the second adjusts the base-stock level to accommodate changes in the demand forecast. Period t order quantity is therefore a linear function of both period t demand and the change in forecasted demand or sales. For $n = 1$, Equation (1) reduces to the single-item ordering rule presented on p. 52 in Graves (1999).

Although Equation (1) describes firm-level ordering behaviors, inventory purchases on the left side, \mathbf{q}_t , are specified in units of inventory across n SKUs. Because we are unable to observe inventory units directly, a *second step* is required to transform Equation (1) into a relation that can be tested empirically using data described in §3. A derivation included in

the appendix reveals that Equation (1) can be transformed into

$$\frac{\text{PURCHASES}_t}{\text{COGS}_t} = b_0 + b_1 \Delta \text{FSALES}_{t+1}, \quad (2)$$

where PURCHASES_t denotes aggregate, firm-level inventory purchases during period t , and COGS_t represents aggregate cost of goods sold during period t , both expressed in dollars. Scaling inventory purchases by cost of goods sold has a straightforward and natural interpretation. Firms whose ratio is equal to one are purchasing just enough inventory to replace inventory sold. Firms with a ratio higher than one are increasing inventory, and firms with a ratio lower than one are depleting inventory levels. $\Delta \text{FSALES}_{t+1}$ is the change in sales forecast from period t to period $t + 1$, scaled by period t sales forecast and both expressed in dollars (i.e., $\Delta \text{FSALES}_{t+1}$ is the percentage change in sales forecast; for a detailed definition, see Equation (15) in the appendix). We estimate b_0 and b_1 as regression coefficients. We interpret b_1 as the response rate, which measures how firms' inventory purchases respond to changes in sales forecast. If the policy in Proposition 1 is predictive, then based on Equation (2), we expect to find that a firm's inventory purchases are a function only of current sales and change in sales forecast; in particular, they should *not* depend on sales growth from prior periods. This leads to our primary hypothesis stated in null form:

MAIN HYPOTHESIS. *A firm's inventory purchases are a function of current sales and change in sales forecast, but not past sales growth.*

To test the hypothesis, we partition the sample into three groups based on sales growth in the previous period, $t - 1$, *HIGHGROWTH*, *LOWGROWTH*, and *MODERATEGROWTH*. *HIGHGROWTH* represents the top quartile, *LOWGROWTH* the bottom quartile, and *MODERATEGROWTH* the middle two quartiles of sales growth firm-quarters during period $t - 1$. Portfolios are formed each quarter based on sample firm-quarter observations.⁶ When discussing results, we focus on the differences between the high-growth

⁶ We calculate quartile cutoffs each period for two reasons. First, they produce consistent positive and negative sales growth cutoffs without resulting in cutoffs that are extreme. We do not use periodic tercile cutoffs because they much more frequently switch between positive and negative sales growth depending on the state of the economy. Despite this, we replicate regression results using both terciles and quintile cutoffs, and find that results are insensitive to this design choice. Second, calculating cutoffs each period eliminates potential look-ahead bias. This design choice does present two potential issues. First, the cutoffs vary by period based on economic growth rates. Second, if cutoffs change significantly, firms around the cutoff borders jump between groups despite an unchanged growth rate. We address this by replicating regression results using one constant set of cutoffs based on the entire sample period. We find that our inferences are unaffected by this research design choice.

and moderate-growth groups. Differences between the low-growth and moderate-growth groups may be of interest, but should be interpreted with caution. The low-growth group, which includes many firms experiencing shrinking sales, is particularly susceptible to survivorship bias. This bias does not extend to the other two groups. The empirical model we use to test the main hypothesis is

$$\begin{aligned}
 &PURCHASES_t/COGS_t \\
 &= b_0 + b_1\Delta FSALES_{t+1} \\
 &\quad + b_2\Delta FSALES_{t+1} \times HIGHGROWTH_{t-1} \\
 &\quad + b_3\Delta FSALES_{t+1} \times LOWGROWTH_{t-1} \\
 &\quad + b_4\Delta AVGINDSALES_{q+1} \\
 &\quad + b_5\Delta FSALES_{t+1} \times INVTURN_t \\
 &\quad + b_6\Delta FSALES_{t+1} \times GM_t + \epsilon_t. \tag{3}
 \end{aligned}$$

The *HIGHGROWTH* and *LOWGROWTH* indicator variables in Equation (3) allow us to test *whether inventory purchases depend on past sales growth*. Because we use quarterly data, seasonality in sales is a concern. For example, some industries naturally experience higher sales in the fourth quarter, and in expectation of the increased sales, they increase inventory. We control for seasonality by using year-over-year changes in forecasted sales. In addition to seasonally adjusting sales growth, we control for disparities across industry by including industry fixed effects and $\Delta AVGINDSALES_{q+1}$, the average percentage change in industry sales from quarter q to $q + 1$ (e.g., the average change in sales for the retail industry from the third to fourth quarters). $\Delta AVGINDSALES_{q+1}$ controls for industry seasonal fluctuations that are not captured by adjusting sales growth seasonally at the firm level. Interaction term $\Delta FSALES_{t+1} \times INVTURN_t$ allows the rate of replacement to change with how quickly firms turnover inventory. $INVTURN_t$ is calculated as cost of goods sold, scaled by average total inventory. We also control for gross margin, GM_t , by combining it with changes in forecasted sales. GM_t is calculated as sales less cost of goods sold, scaled by sales. We do this because everything else equal, firms with higher gross margins are likely to stock higher inventory because stockouts are more costly (Zipkin 2000).

One methodological challenge with testing the hypothesis is that $PURCHASES_t$ and $COGS_t$ in Equation (2) represent values at the beginning of the replenishment cycle, whereas the firm-level data allow us to observe inventories only at the beginning of certain periods (months, quarters, years), which do not necessarily coincide with the beginning of the inventory replenishment cycle. We follow Rumyantsev and Netessine (2007, p. 414), who argue

these snapshots of aggregate inventories can be used to analyze properties derived from classic inventory models because "...these properties apply equally to optimal inventory decisions as well as to observations of inventory positions at random points during the review period."

If Equation (1) is descriptive, we expect b_0 will be approximately 1, consistent with firms replacing their aggregate sales from the current period. We also expect that if aggregate order quantity is not history dependent, then b_2 and b_3 will not be different from zero. If we find that $b_2 < 0$ ($b_2 > 0$) then this suggests aggregate order quantity is a function of past changes in aggregate sales. In such a case, it follows that compared to moderate-growth firms, high-growth and low-growth firms either shrink or expand their aggregate inventories at different rates.

In §5, the empirical evidence *rejects* our primary hypothesis. Instead, we consistently find that b_2 is negative, implying that high-growth firms purchase less of their change in expected future demand than moderate-growth firms do. The remainder of our analysis explores frictions that explain why inventory purchases are a function of past sales growth and why, on average, high-growth firms may exercise more restraint when purchasing inventory.

4.2. Frictions

We continue by examining frictions that in our context include both explicit and implicit costs associated with inventory transactions (e.g., see Stoll 2000). Although we recognize that a number of potentially important frictions might explain why the adaptive base stock policy does not explain cross-sectional, aggregate-firm ordering behaviors under growth, we focus on two assumptions that underlie the inventory model described in §4.1: (i) demand process is an autoregressive integrated moving average (ARIMA) process with stationary parameters; (ii) holding costs are linear with coefficients that do not depend on past sales growth.

In deriving the policy given in Proposition 1, we make direct use of Assumptions (i) and (ii). Hence it would accord with rejection of our primary hypothesis if any of these assumptions fail to hold in empirical practice. We first propose how to test Assumptions (i) and (ii) empirically, and in §5, we present the findings.

4.2.1. Firm-Level Demand Process. Aggregate, firm-level sales are given by $\mathbf{r} \cdot \mathbf{S}_t$, where \mathbf{r} is the vector of selling prices and \mathbf{S}_t is vector of sales quantities. Under a back-order assumption, $\mathbf{S}_t = \mathbf{D}_t$, where $\mathbf{D}_t = (D_t^1, D_t^2, \dots, D_t^n) = (F_t^1 + \epsilon_t^1, F_t^2 + \epsilon_t^2, \dots, F_t^n + \epsilon_t^n)$ is the vector of period t demands (see Equations (6) and (9)). From (6), it follows that $\mathbf{r} \cdot \mathbf{S}_t$ is a normal random variable with a mean of $\mathbf{r} \cdot \mathbf{F}_t$ and variance given by $\text{var}(\mathbf{r} \cdot \boldsymbol{\epsilon}_t) = \sum_{i=1}^n \sum_{j=1}^n \text{cov}(r^i \epsilon_t^i, r^j \epsilon_t^j)$, which suggests

that the variance of aggregate sales should *not* depend on past demand growth.

We begin empirical analysis by providing descriptive evidence of the relationship between sales growth and future sales volatility. We sort firm-quarters with sales growth in the same sales growth percentile into portfolios. For example, if a firm-quarter has sales growth of 10.2% during period t , we sort this firm-quarter into the 10% sales growth portfolio, along with all firm-quarters with sales growth greater than or equal to 10% and less than 11%. Sorting firms into sales growth percentiles holds sales growth variation within each portfolio relatively constant. We calculate the standard deviation of sales growth realizations during period $t + 1$ for each portfolio. If prior period sales growth is not a factor in future sales growth volatility, we do not expect a pattern across prior period sales growth percentiles. If sales growth volatility is increasing in the extremity of prior period sales growth, we expect a U-shaped pattern between prior period sales growth percentiles and the standard deviation of future sales growth.

Although this exercise provides intuitive, descriptive evidence, it is not a formal test, and is subject to limitations. The primary limitation is that firm sales growth volatility is a latent variable. Therefore, we use a generalized autoregressive conditional heteroskedastic (GARCH) model to test the relationship between period $t + 1$ sales growth volatility and period t sales growth. Unlike traditional ordinary least squares regression, which assumes error terms are distributed normally with constant variance, GARCH treats the volatility of error terms as a function of previous period errors and error-term volatility.

GARCH models are time-series models typically estimated at the firm level (Bollerslev 1986). Because we are interested in common variation across firms and are working with panel data, we employ a pooled GARCH model (Cermeno and Grier 2001, 2006). Estimation of a GARCH model requires a considerable time series for each firm. Therefore, for inclusion in this test, we require that firms have a minimum of 40 consecutive quarterly observations.⁷ The final

sample for this test included 2,538 firms and 172,580 firm-quarters. Using maximum-likelihood estimation, the GARCH model estimates the following equations simultaneously:

$$\begin{aligned} \Delta SALES_{t+1} = & \theta + \gamma_1 \Delta SALES_t \\ & + \gamma_2 \Delta SALES_t \times HIGHGROWTH_t \\ & + \gamma_3 \Delta SALES_t \times LOWGROWTH_t \\ & + \sqrt{h_{t+1}} \epsilon_{t+1}, \end{aligned} \quad (4a)$$

$$\begin{aligned} h_{t+1} = & \omega + \alpha (\Delta SALES_t - (\theta + \gamma_1 \Delta SALES_{t-1} \\ & + \gamma_2 \Delta SALES_{t-1} \times HIGHGROWTH_{t-1} \\ & + \gamma_3 \Delta SALES_{t-1} \times LOWGROWTH_{t-1}))^2 \\ & + \beta h_t. \end{aligned} \quad (4b)$$

In Equation (4a), we model sales growth as an AR(1) process, but allow for the AR(1) coefficient to vary based on the current period's sales growth. Including the interaction terms allows us to first identify whether the autoregressive process is stationary, and second whether sales growth volatility is a function of past sales growth after controlling for stronger or weaker mean reversion in high-growth and low-growth sales.⁸ First, if we find that the AR(1) coefficients vary based on sales growth, this suggests the autoregressive process is not stationary. The γ_1 coefficient represents the amount of mean reversion in sales growth from one period to the next. If extreme sales growth mean reverts faster, we expect negative coefficients for γ_2 and γ_3 . This result would suggest mean reversion for sales growth is nonstationary. Equation (4b) models error-term volatility, h_{t+1} , as a function of the squared unexpected sales growth during period t and the volatility of the period t error. If future sales growth volatility is an increasing function of the deviation of the previous period's sales growth from its expected level, we expect a positive coefficient for α . *A positive coefficient suggests the volatility of sales growth is a function of past sales growth.*

by Fama and MacBeth (1973). This test is robust to cross-sectional correlation, and results of this analysis are qualitatively similar to those reported in the pooled analysis.

⁸ One concern is that results are driven by the classification scheme for low-growth and high-growth firms; the causality of the relationship is reversed. High sales growth volatility leads to high-growth and low-growth sales growth. We provide two pieces of evidence that remove this concern. First, as mentioned in the previous footnote, we estimate firm-level GARCH models, and these models do not suffer from this limitation. Second, in untabulated results, we find that inferences remain unchanged if we define *HIGHGROWTH* and *LOWGROWTH* over two periods rather than one; high-growth firms are only those firms that have high-growth during periods $(t - 1)$ and $(t - 2)$.

⁷ As pointed out by Cermeno and Grier (2006), to estimate a variance/covariance matrix robust to cross-sectional correlation, the number of time-series observations must be large relative to the number of firms. As the number of firms becomes large, the number of parameters required for estimating the variance/covariance matrix becomes too large. Because we have a large number of firms and short time series, we assume no cross-sectional correlation and estimate a GARCH model as Cermeno and Grier (2001) do, rather than Cermeno and Grier (2006). In untabulated tests, we estimate firm-level GARCH models and then test the mean GARCH coefficients by scaling the average coefficient by the standard deviation of the firm coefficients, similar to the method used

4.2.2. Firm-Level Holding Cost. Inventory control has long focused on managing types and sources of uncertainty in the demand for products. There are however other important sources of uncertainty that receive scant attention, including inventory obsolescence and consequent inventory write-downs.⁹ Many products are subject to certain or possible obsolescence, such as products in industries with high rates of technical innovation like computers and pharmaceuticals, and products in markets with frequent shifts in consumer tastes, including books, records, perfumes, and some food items. Although both the timing and degree of obsolescence are typically uncertain, extant literature suggests inventory obsolescence is a significant cost of holding inventory (Allen et al. 2013).

To measure obsolescence risk, we apply the general modeling framework used by Song and Zipkin (1996), who model it as a state-dependent holding cost. The state-of-the-world determines whether the firm enters an obsolete state, and the inventory position determines the holding cost magnitude. The firm remains in the nonobsolete state for a random amount of time, and after that, the world jumps to the obsolete state. It has been assumed that the probability of jumping to the obsolete state does not depend on the sales history.

Using firm announcements that a significant amount of a firm's inventory has become obsolete, we propose the following model to estimate the probability that a firm enters the obsolete state-of-the-world during period $t + 1$, given that it is a nonobsolete state during period t :

$$\begin{aligned} \text{Ln}[\text{Pr}(\text{OBSOLETE}_{t+1} = 1) / (1 - \text{Pr}(\text{OBSOLETE}_{t+1} = 1))] \\ = b_0 + b_1 \Delta \text{SALES}_t + b_2 \Delta \text{SALES}_{t-1} + b_3 \text{INVACC}_t \\ + b_4 \text{INVACC}_{t-1} + b_5 \text{INV}_{t-2} + b_6 \text{SIZE}_{t-1} + \epsilon_{t+1}. \end{aligned} \quad (5)$$

We estimate Equation (5) using logistic regression, with standard-error estimates clustered by firms and

⁹Under U.S. Generally Accepted Accounting Principles (GAAP), inventories are recorded at cost. However, if the market value of inventory declines below its original cost for reasons such as obsolescence, a company *must* write down the inventory to the new market value to recognize the loss. Best Buy's financial statements, for example, include the following disclosure: "...merchandise inventories are recorded at the lower of average cost or market..." Any inventory write-down must be reflected as an *expense* (part of cost of goods sold) on the income statement. Thus, if the value of inventory declines, a company incurs a financial loss. These costs can be significant: For example, there is the well-publicized \$2.25 billion inventory write-down at Cisco Systems, which led to a decline in the company's market value from \$430 billion in March of 2000 to \$108 billion in March of 2001; more recently, the digital video recorder (DVR) supplier TiVo posted a net loss of \$17.7 million in its fiscal second quarter of 2007, which the company credited to an inventory write-down.

years.¹⁰ *OBSOLETE* is an indicator that equals one if the firm-year observation had an inventory obsolescence event and zero otherwise. As with our previous tests, ΔSALES is measured as percentage change in sales. Consistent with Allen et al. (2013), we measure *INVACC*, inventory accruals as change in inventory, scaled by average total assets. As an additional control, we include industry fixed effects in the estimation. Industries are measured as two-digit SIC codes. If the probability of obsolete inventory is increasing in prior sales growth, we expect positive coefficients for b_1 or b_2 . The reason we examine both one- and two-year lagged sales growth is that extant accounting research suggests managers can delay accounting recognition of an economic write-down to periods subsequent to the economic devaluation of assets (Vyas 2011). Thus, it would not be surprising to find that b_2 is significantly positive and b_1 insignificant. In addition to sales growth and consistent with research from Allen et al. (2013), we expect that the probability of obsolete inventory is increasing in inventory accruals. Therefore, we also expect to find positive coefficients for b_3 and b_4 .

4.2.3. Gross Margins and Ordering Flexibility. The results we find and present in §5 suggest that the probability of incurring an inventory write-down increases with past sales growth. This finding, however, deserves further investigation. Even with state dependent holding cost parameters, if firms experiencing high growth were to offset increased expected holding costs fully by purchasing lower levels of inventory, then we would not expect to find a relationship between sales growth and inventory obsolescence. Because we document that realized inventory obsolescence is increasing in sales growth, this suggests that some high-growth firms fail to adjust inventory purchases downward to offset an expected increase in inventory obsolescence.

We examine two possible causes of this behavior. First, firms that appear to overorder and thus experience inventory obsolescence might simply have higher marginal costs for potential lost sales. If this were the case, we would expect high-growth inventory obsolescence firms to have higher profit margins than high-growth, nonobsolescence firms. The second explanation we examine is whether inventory obsolescence firms face restrictions in ordering, including minimum order quantities or contractual requirements to maintain minimum inventory. This could be especially true in high-growth industries in which product sourcing is difficult because of competition for upstream resources. In these situations, it seems

¹⁰As an additional robustness test, we also estimated a probit regression. Results were qualitatively similar and all inferences remained unchanged.

reasonable that firms trade-off purchasing flexibility for guaranteed supply.

To examine these two explanations, we compare inventory obsolescence firms to nonobsolescence firms. To control for disparities between these two groups, we create a matched sample. Following previous literature, we use propensity score matching (for an overview, see Caliendo and Kopeinig 2008). Using the previous empirical model of the probability of inventory obsolescence, Equation (5), we match inventory obsolescence firms to control firms based on the probability of inventory obsolescence, and industry and year. This method allows us to control for both observable factors that led to inventory obsolescence, in addition to unobservable factors correlated with time and industry.¹¹

To test the first explanation, we compare the margins of the sample and control firms. Although we cannot observe the marginal profit of sales, we can measure average profit (i.e., gross margins, $(SALES - COGS)/SALES$). We expect that if inventory obsolescence firms overorder to avoid potential lost sales, inventory obsolescence firms should have higher profits on inventory than nonobsolescence firms. Because inventory obsolescence firms might adjust prices downward and thus decrease realized margins artificially during periods of overpurchasing (such as those periods during and immediately prior to inventory obsolescence), we compare not only contemporaneous margins but also margins in the previous two periods.

To test the second explanation, we compare the supply flexibility of the sample and control firms using a novel, small sample of publicly available supply chain contracts. Described previously, public companies are required to disclose in filings with the SEC any material contracts they enter. Firms may attach contracts as exhibits to several SEC filings (e.g., 10-Ks, 10-Qs and 8-Ks). Morningstar's 10-K Wizard identifies these material contracts separately, allowing us to search for material supply chain contracts for all of the inventory obsolescence and matched firms. Similar to our inventory obsolescence search, we identify supply chain contracts using CIK numbers and common supply chain contract terms (e.g., supply, supplier, purchase, and buyer). Because all firms do not disclose contracts (as noted in §3), we limit analysis to only those matched firm pairs for which at least one material supply contract was disclosed for each

firm within the five years prior to inventory obsolescence.¹² We collect supplier contracts for 30 firm-pairs and 10 firm-pairs classified as high-growth firms during period $t - 1$. We acknowledge that these small samples should make readers cautious when generalizing results, but given the paucity of empirical evidence on supply chain contracting and the striking statistical significance of the results, we argue the evidence is compelling. We also expect that this initial work in the area will further interest in studying supply chain contracts empirically.

Because we are first to perform such an analysis, we devise a method for measuring how supply chain contracts influence firm ordering flexibility. Our classification scheme is guided by the supply chain contract literature (Kessinger and Pieper 2005, §7) that defines supply contract clauses that reduce (increase) purchasing flexibility. Clauses that reduce purchasing flexibility impose lower bounds on buyers' order quantities, delivery lead times, inventory levels, and wholesale prices. Clauses that increase purchasing flexibility allow for inventory buybacks and impose upper bounds on buyers' delivery lead times and wholesale prices. Using the matched sample of supply contracts, we conduct a keyword search for any occurrences of the following words: guarantee, forecast, lead time, buyback, order, quantity, inventory, source, support, and joint venture. For example, by searching for the keyword "quantity" in a supply contract between CARBOMEDICS, INC. ("Supplier") and ATS MEDICAL, INC. ("Buyer") included in ATS MEDICAL, INC., March 29, 2001, 10-K filing, we find the following clause:

...

2. *Purchase of Goods.* (a) Buyer agrees to purchase, during the first five Contract Years of this Agreement, at least the minimum quantity of the Components per year specified in Exhibit B attached to and made part of this Agreement. Thereafter the minimum purchase requirement each year will be at least the lower of either the minimum number of Component Sets set forth on Exhibit B or the number of valve sets actually sold and/or disposed of by any means by Buyer.

...

The clause clearly reduced ATS MEDICAL, INC. inventory purchasing flexibility because it might force it to place unwanted orders. Similarly, by searching for the keyword "forecast" in a supply contract between Lucent Technologies, Inc. ("Buyer") and FS

¹¹ Control firm-years must have a propensity score within 20% of sample firm-years and be from the same calendar year. Additionally, firms must be from the same industry. We first match on three-digit SIC codes and select the nearest propensity score. If no suitable match is found, we match on two-digit SIC codes. For 16 cases, we find no suitable matches. Therefore, we eliminate these firms from the matched sample analysis.

¹² Ideally, we would like to ensure that the contract is in force at the point of obsolescence, but contract lengths are commonly redacted. We select five years because in the small subset of situations where contract lengths are not redacted, this is the modal contract length. Because all pairs do not have the same number of disclosed contracts, we code only the first available contract in the search time-period for each matched pair.

Corporation included in FS Corporation (“Supplier”) August 14, 2002, 10-K filing, we found the following:

...
 §44—*Forecasting*. Buyer will provide Supplier with weekly nonbinding product and volume forecasts through the use of forecasting methods agreed to between the parties. Any forecasts provided by Buyer to Supplier do not represent a commitment to purchase and are for planning purposes only.
 ...

The clause increased Lucent’s inventory purchasing flexibility because it facilitated the supplier’s capacity planning and future availability of inventory, allowing Lucent to postpone inventory purchases and alleviate the need for stockpiling.

By reading the supply contract clauses identified by the keyword searches, we create six indicator variables for supply contract clauses that increased purchasing flexibility, and six terms that reduced purchasing flexibility. The following contractual features increase flexibility: (1) *PROTECTION*—the contract terms include wholesale price guarantees, which offer flexibility in timing of purchases; (2) *NONBINDING*—the contract requires that the buyer provide the supplier with a *nonbinding* sales forecast, which facilitates future capacity planning and inventory availability; (3) *TECHNICAL*—the contract guarantees the buyer technical support and advisory services from the supplier, which facilitate inventory management and sales; (4) *NONBINDINGINVCOM*—the contract requires that the buyer provide the supplier with a *nonbinding* order forecast, which facilitates future capacity planning and inventory availability; (5) *RETURNS*—the contract allows for inventory returns or guarantees buybacks; (6) *LEADTIME*—the contract guarantees delivery lead times, which offers flexibility in timing of purchases.

Whereas these contractual features reduce flexibility: (1) *BINDINGFORECAST*—the contract requires that the buyer provide the supplier with a *binding* forecast of future order quantities; (2) *MINORDER*—the contract requires the buyer place a minimum number of orders of specified size during a given period; (3) *MINQUANTITY*—the contract specifies that the buyer purchase a minimum quantity during a given period; (4) *BINDINGINVCOM*—the contract requires the buyer maintain a minimum level of inventory; (5) *SOLESOURCING*—the contract requires the buyer sole-source from the supplier for a particular inventory item; (6) *JOINTVENTURE*—the contract is structured as a joint venture that inhibits supply competition.

We create three aggregate variables to assess the overall flexibility of the contract, *FLEX*, which is the sum of the flexibility-increasing contract variables; *INFLEX*, which is the sum of the flexibility-decreasing

contract variables; and *FLEX* minus *INFLEX*, which is a summary statistic for total contract flexibility. Although simple, the method allows us to quantify frictions in inventory purchases; something not yet quantified using alternative or more sophisticated models. Given the novelty of our method and loss of information in aggregation, we report differences not only in both the three aggregate variables, but also for all 12 of the flexibility variables.

5. Results

5.1. Replenishment Results

To draw reliable inferences from empirical tests of our primary hypothesis, we need to ensure that results are robust regarding the choice of a sales forecasting model. We approach this by embedding three sales-forecast models in the tests. First, we use sales growth from the previous period as a forecast for the next period’s sales growth. This model assumes naively that sales growth this period will persist. Because extant literature suggests that extreme sales growth more quickly mean reverts (Nissim and Penman 2001), we employ a perfect foresight model by replacing our naive expectation with realized sales growth, $\Delta SALES_{t+1}$. In comparison to the naive model, this model assumes the other extreme—managers forecast future sales perfectly. Although managers are unlikely to forecast sales growth exactly, if we assume that forecast errors are, on average, zero, this perfect foresight model is a good expectations model. The last forecast model we use is analyst forecasts of sales in $t+1$.¹³ Prior analyst-forecast literature demonstrates that incentives lead managers to guide analysts’ forecasts to make them accurate (Cotter et al. 2010, Matsumoto 2002). If this is the case, then analyst forecasts are a good measure of managers’ expectations of future sales growth. We scale each of the sales forecast changes, $(FSALES_{t+1} - FSALES_t)$, by the sales forecast, $FSALES_t$, during period t (see the appendix for details). Therefore, the change in sales forecast variables can be viewed as the percentage change in sales forecast.

The results of the tests are presented in Table 2. In addition to control variables and industry fixed effects, we cluster standard errors by quarters. Panel A uses the naive forecast model, $\Delta SALES_t$, and panels B and C use the more sophisticated forecast models, perfect foresight, $\Delta SALES_{t+1}$, and analysts’ sales forecasts, $\Delta FSALES_{t+1}$, respectively. Consistent with expectations, coefficients for all control variables, $\Delta AVGINDSALES_{q+1}$, $\Delta SALES_t \times INVTURN_t$,

¹³ We use the mean analyst forecast closest to the beginning of each quarter if the forecast is within 30 days of the start of the quarter. If no forecast is available in this window, we use the first forecast within 30 days of the start of the quarter, if available.

Table 2 Quarterly Purchases Scaled by Cost of Goods Sold in t Regressed on Forecasted Changes in Sales

Panel A: $\Delta SALES_t$							
Constant	$\Delta SALES_t$	$\Delta SALES_t \times$ $HIGHGROWTH_{t-1}$	$\Delta SALES_t \times$ $LOWGROWTH_{t-1}$	$AVG \Delta INDSALES_{q+1}$	$\Delta SALES_t \times$ $INVTURN_t$	$\Delta SALES_t \times$ GM_t	R^2
0.999*** (0.002)	0.059*** (0.002)			0.453*** (0.015)	-0.005*** (0.0002)	0.019*** (0.002)	0.041
0.999*** (0.002)	0.072*** (0.005)	-0.019*** (0.005)	0.005 (0.005)	0.453*** (0.014)	-0.005*** (0.0002)	0.020*** (0.002)	0.041
Panel B: $\Delta SALES_{t+1}$							
Constant	$\Delta SALES_{t+1}$	$\Delta SALES_{t+1} \times$ $HIGHGROWTH_{t-1}$	$\Delta SALES_{t+1} \times$ $LOWGROWTH_{t-1}$	$AVG \Delta INDSALES_{q+1}$	$\Delta SALES_{t+1} \times$ $INVTURN_t$	$\Delta SALES_{t+1} \times$ GM_t	R^2
0.995*** (0.002)	0.094*** (0.002)			0.452*** (0.014)	-0.007*** (0.0002)	0.022*** (0.002)	0.062
0.995*** (0.002)	0.113*** (0.004)	-0.028*** (0.004)	-0.005 (0.005)	0.452*** (0.014)	-0.007*** (0.0002)	0.024*** (0.002)	0.063
Panel C: Analysts' forecasted $\Delta SALES_{t+1}$							
Constant	$\Delta FSALES_{t+1}$	$\Delta FSALES_{t+1} \times$ $HIGHGROWTH_{t-1}$	$\Delta FSALES_{t+1} \times$ $LOWGROWTH_{t-1}$	$AVG \Delta INDSALES_{q+1}$	$\Delta FSALES_{t+1} \times$ $INVTURN_t$	$\Delta FSALES_{t+1} \times$ GM_t	R^2
0.991*** (0.003)	0.111*** (0.005)			0.450*** (0.032)	-0.009*** (0.0009)	0.038*** (0.006)	0.088
0.993*** (0.004)	0.130*** (0.011)	-0.034*** (0.012)	0.034** (0.016)	0.453*** (0.050)	-0.008*** (0.001)	0.043*** (0.007)	0.091

Notes. The sample consists of all nonfinancial and utility firm-quarters listed on the NYSE, NASDAQ, or AMEX, CRSP, and the COMPUSTAT fundamentals quarterly file between 1980 and 2008 with inventory greater than 1% of average total assets and available data to calculate change in sales for periods $(t - 1)$, t , and $(t + 1)$, and purchases to cost of goods sold in period t . The dependent variable, $PURCHASES_t / COGS_t$, is measured as purchases (ending inventory plus cost of goods sold minus beginning inventory) divided by cost of goods sold. $\Delta SALES_t$ is measured as four quarters ago percentage change in sales. $HIGHGROWTH_t$ and $LOWGROWTH_t$ are indicator variables equal to 1 if a firm ranks in the top or bottom quartile of sales growth in quarter t , respectively, and 0 otherwise. For firm i in quarter q , $AVG \Delta INDSALES_q$ is calculated as the average industry sales growth for firm i 's industry in quarter q . $INVTURN$ is cost of goods sold scaled by average inventory. GM is sales minus cost of goods sold scaled by sales. $\Delta FSALES_{t+1}$ is the analyst forecasted growth in sales measured as the mean analyst forecast of sales in period $t + 1$ at the end of quarter t minus the mean analyst forecast of sales in period $t - 3$ scaled by period $t - 3$ forecast. Analyst sales forecasts are collected from the IBES summary file. All specifications include industry fixed effects and robust standard errors clustered on each quarter.

Significant at the 0.05 level; *significant at the 0.01 level.

and $\Delta SALES_t \times GM_t$, are significant and in predicted directions. However, we also find that the b_2 coefficients are in each panel negative, suggesting *inventory purchases depend on past sales growth*.

To illustrate the results, the first estimation in panel A shows that using our data, we recover a policy consistent with the form of a base stock policy under which firms replace inventory they sold ($Constant = 0.999$) and adjust the base-stock level to accommodate changes in forecasted sales growth ($b_1 = 0.059$). This accords with findings from Mosser (1991), who studies *firm-to-economy* aggregation, and reports that in absence of growth, aggregate inventory purchases in each period always replace the aggregate sales from the immediate period.

In panel A, we find that *HIGHGROWTH* firm quarters behave differently from *MODERATEGROWTH* firm quarters. Because $b_2 < 0$, this implies *observations with high sales growth in the previous quarter, in aggregate, purchase less of their expected future demand than observations with only moderate growth in the previous*

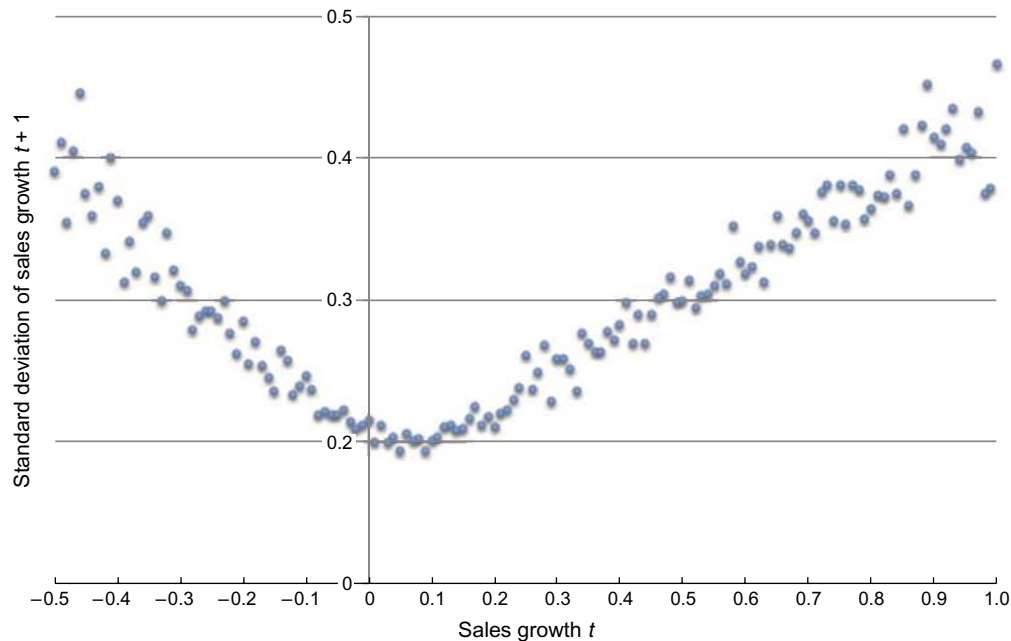
quarter. The results in Table 2, which are robust to several sales forecast specifications, provide evidence consistent with high-sales-growth firms ordering less of their forecast sales changes than moderate-growth firms. Therefore, we reject the primary hypothesis. The following results may help explain why high-growth firms' aggregate ordering behaviors differ from moderate-growth firms'.

5.2. Friction Results

5.2.1. Firm-Level Demand Process.

We begin the exploration by presenting Figure 1, which shows a plot of the relationship between period t sales growth and the standard deviation of period $t + 1$ sales growth for each portfolio between -50% and 100% sales growth.¹⁴ Recall from §4.2 that we sort firm-quarter observations in the same sales growth

¹⁴The pattern holds beyond -50% and 100% sales growth, but for parsimony, we limit the presentation to sales growth between -50% and 100%.

Figure 1 (Color online) Quarterly Sales Growth in t and Sales Growth Volatility in $t + 1$ 

Note. Firm-quarter observations are grouped into portfolios of observations with the same percentage quarterly sales growth in period t (e.g., all firms between 10% and 11% sales growth are grouped in one portfolio).

percentile into the same portfolio. Therefore, the figure, which covers sales growth from -50% to 100% , plots 151 points, one for each sales growth portfolio formed in t . Each portfolio has more than 200 observations, with most portfolios in excess of 500. The y axis represents the standard deviation of sales growth during period $t + 1$ for firms in the same sales growth portfolio during period t , and is calculated as $SD = \sqrt{(1/(N-1)) \sum_{i=1}^N (\Delta Sales_{i,t+1} - \overline{\Delta Sales_{t+1}})^2}$, where N is the number of firms in a sales growth portfolio. If the standard deviation of sales growth in $t + 1$ is unrelated to the sales growth realization during period t , we expect a plot that is flat. What we observe instead is a strong, U-shaped pattern, suggesting that as sales growth becomes more extreme, positive or negative, the volatility of the next period's sales growth increases. We also observe that a significant number of firms realize negative sales growth. In untabulated results, we find that 31.1% of quarterly observations experience decreases in sales.

Figure 1 provides informal evidence concerning sales growth volatility. We present the GARCH model analysis in Table 3. Based on estimations, we observe first that sales growth is not only autoregressive, but the autoregressive coefficient differs depending on prior-period sales growth. The γ coefficients show that for moderate growth, sales growth persists at a rate of 0.773 (γ_1), but at only a rate of 0.707 ($\gamma_1 + \gamma_2$) and 0.594 ($\gamma_1 + \gamma_3$) for high-growth and low-growth firms, respectively, suggesting sales growth is nonstationary. Second, the volatility of sales growth

Table 3 Pooled Generalized Autoregressive Conditional Heteroskedastic Model of Quarterly Sales Growth

Explanatory variable	Estimate (t -statistics)
θ	0.021 (135.41)
$\Delta SALES_t (\gamma_1)$	0.773 (624.52)
$\Delta SALES_t \times HIGHGROWTH_t (\gamma_2)$	-0.066 (-55.20)
$\Delta SALES_t \times LOWGROWTH_t (\gamma_3)$	-0.179 (-93.52)
ω	0.003 (552.82)
$ARCH(\alpha)$	0.566 (465.66)
$GARCH(\beta)$	0.616 (1372.50)

Notes. The sample consists of all nonfinancial and utility firm-quarters listed on the NYSE, NASDAQ, or AMEX, CRSP, and the COMPUSTAT fundamentals quarterly file between 1980 and 2008 with inventory greater than 1% of average total assets and available data to calculate change in sales for periods $(t-1)$, t , and $(t+1)$, and purchases to cost of goods sold in period t . For inclusion in estimation, firms must have the required data for a minimum of 40 consecutive quarters. $\Delta SALES_t$ is measured as the percentage change from current to sales four quarters previously. $HIGHGROWTH_t$ and $LOWGROWTH_t$ include firm-quarters ranking in the top or bottom quartile of sales growth in quarter t , respectively.

Table 4 Summary Statistics for Inventory Obsolescence Firm-Years

Variable	Mean	Std. dev.	P25	Median	P75
Inventory Obsolescence Amount _t (millions)	13.21	45.69	0.98	3.00	1.30
Inventory Obsolescence Amount _t /Average Total Assets _t	-0.037	0.036	-0.041	-0.023	-0.016
SALES _t (millions)	412.96	991.79	29.94	83.13	32.42
ROA _t	-0.154	0.253	-0.253	-0.082	0.201
EWRET _t	-0.215	0.661	-0.582	-0.337	0.002

Notes. The sample consists of 290 firms experiencing a first-time inventory obsolescence of more than 1% of average total assets between calendar years 2002 and 2004. SALES is measured as total sales. ROA is measured as income before extraordinary items scaled by average total assets. EWRET is the buy-hold annual return less the compounded monthly annual equally weighted return. All variables with the exception of returns are winsorized at the 1st and 99th percentiles.

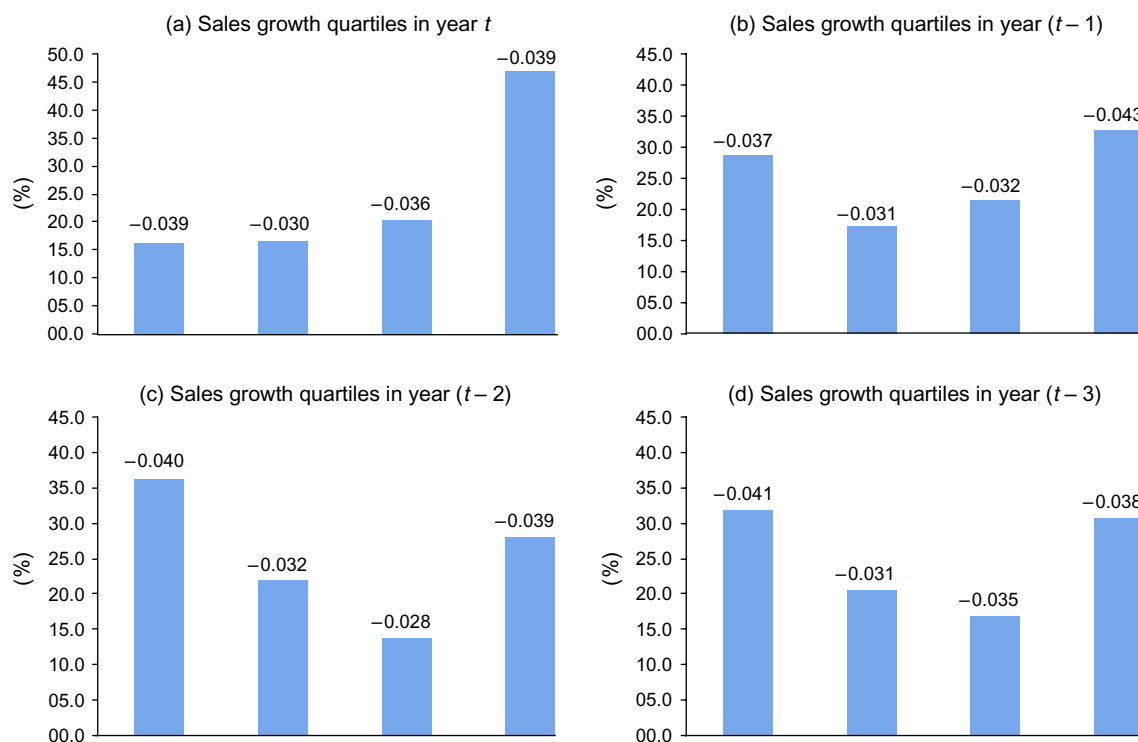
is a function of both sales growth ($\alpha = 0.566$) and prior-period sales growth volatility ($\beta = 0.616$). These results accord with inventory purchases being a function of past sales growth, a finding that potentially explains the rejection of our primary hypothesis.

5.2.2. Firm-Level Holding Cost. The next analysis tests whether the probability of incurring inventory obsolescence is constant or varies as a function of sales growth. If the probability is an increasing function of sales growth, then it explains why firms facing higher sales growth replace inventory at a rate lower than moderate-growth firms do. Because the inventory obsolescence sample is a unique sample collected from annual reports, we provide descriptive statistics for the 290 inventory obsolescence firm-years

in Table 4. The mean (median) inventory obsolescence charge is \$13.2 (\$3.0) million, representing 3.7% (2.3%) of a firm’s mean (median) total assets and suggesting that the magnitude of these write-downs is significant. The mean (median) firm experiences a -15.4% (-0.082%) return on assets, ROA_t, and a -21.5% (-33.7%) equally weighted, market-adjusted return, EWRET_t, in the year of inventory obsolescence, suggesting inventory obsolescence is associated with extreme negative consequences for shareholders.

We begin analysis of inventory obsolescence and sales growth by presenting graphical evidence. Figure 2 presents the percentage of firms falling into each quartile of sales growth in the period of obsolescence, t , and the three periods prior to obsolescence, ($t - 1$),

Figure 2 (Color online) Percentage of Inventory Obsolescence Sample Falling in Each Sales Growth Quartile in Periods ($t - 1$), ($t - 2$), and ($t - 3$) Relative to the Year of Inventory Obsolescence, t



Notes. The entire sample of firms is sorted each year on sales growth and ranked into quartiles. Quartiles are sorted from the highest sales growth on the left to lowest sales growth on the right. In this figure, the number at the top of each bar represents the mean inventory obsolescence cost scaled by average total assets for firms in that sales growth quartile.

Downloaded from informs.org by [128.252.111.87] on 06 March 2016, at 20:52 . For personal use only, all rights reserved.

Table 5 Logistic Regressions of Inventory Obsolescence ($t + 1$) Regressed on Changes in Inventory, Sales, and Controls (2002–2004)

	(A)	(B)	(C)
$\Delta SALES_t$	-0.034 (0.050)		-0.193 (0.152)
$\Delta SALES_{t-1}$	0.213*** (0.039)		0.188*** (0.031)
$INVACC_t$		3.049*** (0.632)	3.440*** (0.603)
$INVACC_{t-1}$		3.769*** (0.865)	3.245*** (0.679)
INV_{t-2}		2.607*** (0.592)	2.643*** (0.659)
$SIZE_{t-1}$	-0.214*** (0.041)	-0.187*** (0.041)	-0.185*** (0.040)
Constant	-14.257*** (0.753)	-15.567*** (0.638)	-15.093*** (0.736)
Observations	6,998	6,468	6,448

Notes. The sample consists of all nonfinancial and utility firm-quarters listed on the NYSE, NASDAQ, or AMEX, CRSP, and the COMPUSTAT fundamentals quarterly file between 2002 and 2004 with inventory greater than 1% of average total assets and available data to calculate change in sales for periods $(t - 1)$, t , and $(t + 1)$, and purchases to cost of goods sold in period t . The dependent variable, *OBSOLETE*, equals 1 for the 290 firms experiencing a first-time inventory obsolescence of more than 1% of average total assets and 0 otherwise. $\Delta SALES$ is measured as the year-over-year percentage change in sales. $INVACC$ is measured as change in ending inventory scaled by average total assets. INV is ending inventory scaled by average total assets. $SIZE$ is measured as the log of average total assets. Regressions include industry fixed effects.

***Significant at the 0.01 level.

$(t - 2)$, and $(t - 3)$. The mean size of obsolescence cost, scaled by average total assets for firms falling into each quartile, is shown at the top of each bar. If there is no relationship between sales growth and obsolete inventory then we expect approximately 25% of firms to fall into each quartile.

Unsurprisingly, during the period of the inventory write-down, the lowest sales growth quartile is extremely overrepresented (46.9%). During the period immediately prior to obsolescence, it appears firms are distributed evenly across quartiles, with the largest percentage of firms in the lowest sales growth quartile (32.8%). During periods $(t - 2)$ and $(t - 3)$, the strongest representation of firms appears in the highest sales growth quartiles (36.2% and 31.9%, respectively). These are the most extreme obsolescence costs because the mean cost is -4.0% and -4.1% of average total assets during periods $(t - 2)$ and $(t - 3)$, respectively. This descriptive evidence accords with a positive relationship between sales growth and inventory obsolescence, especially during periods $(t - 2)$ and $(t - 3)$. We find no relationship during period $t - 1$, according with extant empirical results that managers delay accounting recognition of economic decreases in asset values (Vyas 2011). We now provide formal statistical tests on the relationships between sales growth and inventory obsolescence.

Shown in Table 5, we estimate the relationship between inventory obsolescence and sales growth. Regression (A) presents results for sales growth alone. Regression (B) presents results for just inventory accruals, and regression (C) presents results for both sales growth and inventory accruals. Results are consistent with *inventory obsolescence events being an increasing function of past sales growth*. It appears that sales growth during the period immediately prior to an inventory obsolescence event has no explanatory power, but sales growth two periods prior to an inventory obsolescence event does have explanatory power in the predicted direction.

In regression (B), we document that inventory obsolescence is also increasing in positive inventory accruals. This is unsurprising; if firms were able to decrease inventory prior to it becoming obsolete by selling the inventory for book value, firms would not experience obsolescence.¹⁵

Next, we test whether the purchasing behaviors of high-growth firms differ based on whether firms ex post experience an inventory obsolescence event. We use annual rather than quarterly data because the inventory obsolescence data are collected annually. Because we are interested in providing evidence on the differences in purchasing among high-growth firms, for parsimony, we estimate results only for the naive sales growth forecast, $\Delta SALES_t$. The tests use an identical model as in previous tests of purchasing behaviors (Equation (3) and Table 2), except that we modify the growth interaction terms by adding an additional indicator variable for firm-year observations that ex post experience obsolete inventory.

If high-growth, firm-year observations that experience subsequent obsolete inventory purchase and hold more inventory than other high-growth, firm-year observations, we expect a positive interaction term. We also test whether high-growth, firm-year observations that experience obsolescence purchase and hold more or less inventory prior to inventory becoming obsolete than moderate-growth, firm-year observations using the same model. By examining the sum of coefficients on the main interaction effect ($\Delta SALES_t \times HIGHGROWTH_{t-1}$) and the obsolescence interaction effect ($\Delta SALES_t \times HIGHGROWTH_{t-1} \times OBSOLETE_{t+1}$), we compare high-growth, firm-year observations that subsequently experience obsolescence to moderate-growth, firm-year observations. If firms experiencing high-growth and subsequent obsolescence purchase more inventory than firms

¹⁵ As a robustness test, we also include additional control variables from Francis et al. (1996) in estimation of Equation (5). We control for current and prior-three-years changes in earnings, an indicator variable for negative unexpected earnings, and market-adjusted stock returns. Results are qualitatively similar, and all inferences remain unchanged.

experiencing moderate sales growth, we expect the sum of the coefficients to be positive. If we find that the coefficient sum is not different from zero, this would suggest that the purchasing behaviors of firms experiencing high growth and subsequent obsolescence are indistinguishable from behaviors of firms experiencing moderate growth.

Table 6 presents results of the estimation. Because the sample period differs from the sample in Table 2, regressions (A) and (B) replicate those models with the more limited sample. Consistent with previous results, high-growth firms replenish inventory at a lower rate than moderate-growth firms do. Now we turn to the more extensive specification, regression (C) of Table 6. We are interested in the coefficient of $\Delta SALES_t \times HIGHGROWTH_{t-1} \times OBSOLETE_{t+1}$. The coefficient is

Table 6 Purchases Scaled by Cost of Goods Sold in t Regressed on Change in t Sales, Extreme-Growth Indicators for $(t - 1)$ and Obsolescence Indicators for $(t + 1)$

	PURCHASES _{<i>t</i>} /COGS _{<i>t</i>}		
	(A)	(B)	(C)
$\Delta SALES_t$	0.068*** (0.006)	0.096*** (0.009)	0.096*** (0.009)
$\Delta SALES_t \times HIGHGROWTH_{t-1}$		-0.037*** (0.008)	-0.040*** (0.009)
$\Delta SALES_t \times LOWGROWTH_{t-1}$		-0.034* (0.018)	-0.034* (0.018)
$\Delta SALES_t \times HIGHGROWTH_{t-1} \times OBSOLETE_{t+1}$			0.038** (0.019)
$\Delta SALES_t \times LOWGROWTH_{t-1} \times OBSOLETE_{t+1}$			0.002** (0.010)
$\Delta SALES_t \times INVTURN_t$	0.0002*** (0.000)	0.0002** (0.000)	0.0002** (0.000)
$\Delta SALES_t \times GM_t$	0.053*** (0.003)	0.052*** (0.004)	0.052*** (0.004)
Constant	0.991*** (0.012)	0.992*** (0.011)	0.992*** (0.011)
Observations	8,004	8,004	8,004
R ²	0.170	0.174	0.174

Regression (C): $\Delta SALES_t \times HIGHGROWTH_{t-1} + \Delta SALES_t \times HIGHGROWTH_{t-1} \times OBSOLETE_{t+1} = 0$;
 F-statistic = 0.02, p-value = 0.883

Notes. The sample consists of all nonfinancial and utility firm-years listed on NYSE/NASDAQ/AMEX, CRSP, or COMPUSTAT between 2002 and 2004, with inventory greater than 1% of average total assets and available data to calculate change in sales data for periods $(t - 1)$, t , and $(t + 1)$, and purchases to cost of goods sold during period t . The dependent variable, PURCHASES_{*t*}/COGS_{*t*}, is measured as purchases (ending inventory plus cost of goods sold, minus beginning inventory) divided by cost of goods sold. $\Delta SALES_t$ is measured as the year-over-year percentage change in sales. HIGHGROWTH_{*t-1*} and LOWGROWTH_{*t-1*} are indicators equal to 1 if a firm ranks in the top or bottom quartiles, respectively, of sales growth in year $(t - 1)$, and zero otherwise. OBSOLETE_{*t+1*} is an indicator equal to 1 if a firm experiences a first-time obsolescence event of at least 1% of total assets during period $(t + 1)$, and zero otherwise. INVTURN is average cost of goods sold, scaled by average inventory. GM is sales minus cost of goods sold, scaled by sales. Regressions include industry fixed effects, and standard errors are clustered by year.

*Significant at the 0.10 level; **significant at the 0.05 level; ***significant at the 0.01 level.

positive, suggesting high-growth obsolescence firms replace inventory at a higher rate than other high-growth firms do. The sum of coefficients $\Delta SALES_t \times HIGHGROWTH_{t-1} \times OBSOLETE_{t+1}$ and $\Delta SALES_t \times HIGHGROWTH_{t-1}$ is insignificant and nearly zero (-0.002). An F-test shows that the purchasing behaviors of high-growth, inventory obsolescence firm-year observations cannot be distinguished from the purchasing behaviors of moderate-growth firm-year observations. Results demonstrate that firms with high sales growth and subsequent obsolescence replace inventory at a higher rate than firms experiencing high-growth alone, and the purchasing behaviors of firms with high sales growth and subsequent obsolescence is indistinguishable from the behaviors of firms with moderate sales growth.

5.2.3. Gross Margins and Ordering Flexibility.

Table 7 presents the results from gross margin tests. Recall that we are comparing contemporaneous and two years of lagged gross margins between inventory obsolescence and nonobsolescence firms. Panel A provides evidence from all inventory obsolescence firms, and panel B focuses only on those inventory obsolescence firms that experienced high-growth during period $t - 1$. The three columns present gross margins for both inventory obsolescence and matched firms for the three periods prior to inventory obsolescence. In both panels A and B, we find no evidence that profit margins are higher for inventory obsolescence firms. Contrary to expectations, this suggests inventory obsolescence firms do not, on average, appear to order more inventory to offset higher stockout costs.

Next, using supply contracts filed with the SEC, we test whether inventory obsolescence correlates with

Table 7 Differences in Gross Margin for Inventory Obsolescence and Matched Firms

	GM _{<i>t</i>}	GM _{<i>t-1</i>}	GM _{<i>t-2</i>}
Panel A: All inventory obsolescence firms			
Obsolescence firms	0.353	0.354	0.347
Matched firms	0.372	0.373	0.358
Difference	-0.018	-0.019	-0.011
(<i>t</i> -statistic)	(-0.78)	(-0.73)	(-0.36)
Panel B: High-growth-inventory obsolescence firms			
Obsolescence firms	0.324	0.399	0.396
Matched firms	0.342	0.365	0.361
Difference	-0.018	0.034	0.034
(<i>t</i> -statistic)	(-0.29)	(0.76)	(0.89)

Notes. The sample consists of 274 obsolescence and matched firms between calendar years 2002 and 2004. Obsolescence firms are matched using propensity score matching. The logistic model used for the propensity score matching is $\text{Ln}[\text{Pr}(OBSOLETE_{t+1} = 1)/(1 - \text{Pr}(OBSOLETE_{t+1} = 1))] = \alpha + \beta_1 \Delta SALES_t + \beta_2 \Delta SALES_{t-1} + \beta_3 \Delta INV_t + \beta_4 \Delta INV_{t-1} + \beta_5 INV_{t-2} + \beta_6 SIZE_t + \epsilon_{t+1}$. A matched firm must have a propensity score within 20% of the obsolescence firm and be from the same year and industry. GM is measured as total sales less cost of goods sold scaled by total sales. High growth is measured in period $t - 1$ and inventory obsolescence in period $t + 1$.

Downloaded from informs.org by [128.252.111.87] on 06 March 2016, at 20:52. For personal use only, all rights reserved.

Table 8 Differences in Contract Flexibility for Inventory Obsolescence and Matched Firms

	PROTECTION	NONBINDING	TECHNICAL	NONBINDINGINVCOM	RETURNS	LEADTIME	FLEX
Panel A: Flexible contract terms for all inventory obsolescence firms							
Obsolescence firms	0.80	0.13	0.30	0.07	0.00	0.27	1.57
Matched firms	0.83	0.57	0.37	0.07	0.10	0.23	2.17
Difference	-0.03	-0.43***	-0.07	0.00	-0.10**	0.03	-0.60***
(<i>p</i> -value)	(0.37)	(0.01)	(0.29)	(0.50)	(0.04)	(0.62)	(0.01)
Panel B: Flexible contract terms for high-growth inventory obsolescence firms							
Obsolescence firms	0.80	0.10	0.30	0.10	0.00	0.30	1.60
Matched firms	0.80	0.40	0.60	0.10	0.00	0.30	2.20
Difference	0.00	-0.30*	-0.30*	0.00	0.00	0.00	-0.60*
(<i>p</i> -value)	(0.50)	(0.06)	(0.09)	(0.50)	(—)	(0.50)	(0.08)
Panel C: Inflexible contract terms for all inventory obsolescence firms							
Obsolescence firms	0.53	0.43	0.40	0.13	0.30	0.10	0.33
Matched firms	0.13	0.13	0.17	0.17	0.33	0.13	-1.10
Difference	0.40***	0.30***	0.23**	-0.03	-0.03	-0.03	0.83***
(<i>p</i> -value)	(0.00)	(0.00)	(0.02)	(0.64)	(0.61)	(0.66)	(0.00)
Panel D: Inflexible contract terms for high-growth inventory obsolescence firms							
Obsolescence firms	0.40	0.60	0.50	0.00	0.20	0.10	0.20
Matched firms	0.30	0.10	0.20	0.20	0.10	0.30	-1.00
Difference	0.10	0.50**	0.30*	-0.20	0.10	-0.20	0.60*
(<i>p</i> -value)	(0.32)	(0.01)	(0.08)	(0.93)	(0.27)	(0.87)	(0.02)

Notes. Sample of all inventory obsolescence firms (panels A and C) consists of 30 matched firm pairs with available supply contracts. Sample of high-growth inventory obsolescence firms (panels B and D) consists of 10 firm pairs with available supply contracts. The following are flexible contract terms: *PROTECTION*—the contract terms include wholesale price caps; *NONBINDING*—the contract requires that the buyer provide the supplier with a *nonbinding* sales forecast, which facilitates supplier’s capacity planning; *TECHNICAL*—the contract guarantees the buyer technical support and advisory services from the supplier; *NONBINDINGINVCOM*—the contract requires that the buyer provide the supplier with a *nonbinding* order forecast, which facilitates supplier’s capacity planning; *RETURNS*—the contract allows for inventory returns or guarantees buybacks; *LEADTIME*—the contract guarantees delivery lead times. *FLEX* is the sum of *PROTECTION*, *NONBINDING*, *TECHNICAL*, *NONBINDINGINVCOM*, *RETURNS* and *LEADTIME*. The following terms are inflexible contract terms: *BINDINGFORECAST*—the contract requires that the buyer provide the supplier with a *binding* sales forecast; *MINORDER*—the contract requires the buyer place a minimum number of orders of specified size in a given period; *MINQUANTITY*—the contract specifies that the buyer place a minimum order quantity in a given period; *BINDINGINVCOM*—the contract requires the buyer maintain a minimum level of inventory; *SOLESOURCING*—the supplier is contractually the only supplier for a particular inventory item; *JOINTVENTURE*—the contract is structured as a joint venture and inhibits supply competition. *INFLEX* is the sum of *BINDINGFORECAST*, *MINORDER*, *MINQUANTITY*, *BINDINGINVCOM*, *SOLESOURCING*, and *JOINTVENTURE*. The *p*-values for the proportions are calculated based one tailed *z*-statistics. The *p*-values for the ordinal variables are calculated base on one tailed *t*-statistics.

*Significant at the 0.10 level; **significant at the 0.05 level; ***significant at the 0.01 level.

supply chain contract flexibility. As with the previous profit margin analysis, we use the same matched firm design in the supply contract tests. Table 8 presents results of the analysis. Panels A and B examine differences in terms of increased flexibility, and panels C and D examine differences in terms of decreased flexibility. Panel A shows that for all inventory obsolescence firms, total contract flexibility is lower than for control firms, with *FLEX* equaling 1.57 for obsolescence firms and 2.17 for matched firms. In panel B, total contract flexibility is also lower for high-growth obsolescence firms. The flexibility terms show that one of the primary drivers of flexibility disparity between the two groups is the nonbinding forecast requirement; obsolescence firms are much less likely to have a nonbinding forecast provision in their contracts.

Panel C shows that for all inventory obsolescence firms, contract provisions that decrease flexibility are also significant, with *INFLEX* equaling 1.90 for obsolescence firms and 1.07 for matched firms. In panel D,

the difference in *INFLEX* is higher for the high-growth obsolescence sample. The contract provisions responsible for driving the difference in *INFLEX* are *BINDINGFORECAST*, *MINORDER*, and *MINQUANTITY*. The significance of these variables is important because it is consistent with the argument that limited supply chain flexibility, particularly in order sizes and quantity, correlates with inventory obsolescence and explains one factor that leads high-growth obsolescence firms to not order differently from moderate-growth firms. The final summary statistic, *FLEX – INFLEX*, for both panels C and D shows a significant difference in flexibility between inventory obsolescence firms and matched firms.

6. Conclusion

Inventory management is a fundamental issue for firms in a broad range of industries. In most practical situations, demand is stochastic and its distribution can change over time. Although it is generally

expected that nonstationary demand processes lead to complex optimal stocking policies, the theoretical inventory literature demonstrates that adaptive base stock policy is an effective substitute for optimal policies at the SKU level, an insight valuable to inventory control analysts who make daily operational decisions. High-level managers and industry analysts, however, are concerned predominantly with behaviors of aggregate inventories, and it is unclear what the adaptive base stock policy implies—if anything—about the behaviors of inventories aggregated at firm or economy levels.

This paper examines SKU-to-firm aggregation by investigating firms' inventory ordering behaviors under nonstationary demand environments. We propose a multi-item model and demonstrate that under reasonable conditions, the optimal inventory policy represents an adaptive base stock policy consisting of two components, the first of which replaces the current period's sales, and the second adjusts the base stock level according to the growth forecast. In other words, under the adaptive base stock policy, firm-level purchases only depend on current sales and change in sales forecast; they should not depend on past sales growth.

Contrary to the theoretical prediction, our results show that inventory purchases are not only a function of current sales and changes in sales forecasts but also past sales growth. High-growth firms purchase less inventory than predicted by the adaptive base stock policy, implying they respond more conservatively to forecasted changes in sales growth than moderate-growth firms do. We further explore this finding by investigating the validity of key assumptions underlying the adaptive base stock policy. We find that demand, inventory holding cost, and supply chain frictions may affect firms' ordering behavior and thus help explain the empirical observations. Specifically, firms' ordering policy is a function of past sales growth because both future demand dynamics and inventory holding costs depend on past sales growth; in addition, firms' inventory holding costs may also be affected by purchasing constraints imposed by supply chain contracts.

Acknowledgments

The authors gratefully acknowledge comments by Serguei Netessine and Matthew J. Sobel; the seminar participants at INSEAD, Washington University in St. Louis, and University of Missouri–St. Louis; the participants of the Supply Chain Management Special Interest Group meeting held in 2011 at the Ross School of Business at the University of Michigan; and the editorial team, which improved the paper. The authors also acknowledge the contribution of Thomson Financial for providing sales forecast data, available through the Institutional Brokers Estimate System.

These data have been provided as part of Thomson's broad academic program to support research.

Appendix

Proof of Proposition 1

Following Graves (1999, Equation (1)), stochastic demand for each item k , $1 \leq k \leq n$ follows an autoregressive integrated moving average (ARIMA) process (there is $n = 1$ in Graves 1999):

$$D_1^k = \mu^k + \epsilon_1^k \quad \text{and} \quad (6a)$$

$$D_t^k = D_{t-1}^k - (1 - \alpha^k)\epsilon_{t-1}^k + \epsilon_t^k \quad \text{for } t = 2, 3, \dots, \quad (6b)$$

where D_t^k is the observed demand for product k during period t and $0 \leq \alpha^k \leq 1$, $0 < \mu^k < \infty$ are known parameters. We assume $\epsilon_t = (\epsilon_t^1, \epsilon_t^2, \dots, \epsilon_t^n)$ are i.i.d. random vectors with $\epsilon_t \sim N(\mathbf{0}, \Sigma)$, where Σ is a $n \times n$ covariance matrix.

We assume the inventory manager first observes the demand vector, \mathbf{D}_t , determines the period t order, receives the order from L periods ago,¹⁶ and then fills demand from inventory. Before ordering during period t , there will, generally, be some stock on hand and some outstanding orders scheduled to arrive during periods $t + 1, \dots, t + L - 1$. Denote by $\mathbf{x}_t = (x_t^1, x_t^2, \dots, x_t^n)$ the inventory position at the beginning of period t (before period t orders are placed). The manager then places an order for the vector

$$\mathbf{q}_t = \mathbf{y}_t - \mathbf{x}_t \quad (7)$$

to be delivered at the beginning of period $t + L$. (We follow Veinott 1965a, p. 209 and Graves 1999, p. 53 to allow $q_t^k < 0$, in which case $c^k q_t^k$ is the rebate received from disposal of q_t^k units of products k during period t .) Under the back-order assumption, if \mathbf{y}_t is the vector of inventory positions after period t orders are placed, then the vector of inventory levels at the end of period $(t + L)$ will be $(\mathbf{y}_t - \mathbf{D}_t - \mathbf{D}_{t+1} - \dots - \mathbf{D}_{t+L})$.

Forecast Model (Graves 1999, p. 52). We define F_{t+1}^k to be the product- k forecast, made after observing demand for item k in time period t , for demand during period $t + 1$:

$$F_1^k = \mu^k \quad \text{and} \quad (8a)$$

$$F_{t+1}^k = \alpha^k D_t^k + (1 - \alpha^k) F_t^k \quad \text{for } t = 1, 2, \dots \quad (8b)$$

By subtracting Equation (8) from (6), one can show by induction that the forecast error is $D_t^k - F_t^k = \epsilon_t^k$, or, equivalently,

$$\mathbf{D}_t - \mathbf{F}_t = \boldsymbol{\epsilon}_t \quad \text{and} \quad \mathbf{D}_t = \mathbf{F}_t + \boldsymbol{\epsilon}_t, \quad (9)$$

where $\mathbf{D}_t = (D_t^1, D_t^2, \dots, D_t^n)$ and $\mathbf{F}_t = (F_t^1, F_t^2, \dots, F_t^n)$. Thus, from (9), we see that the forecast is unbiased because the forecast error is the random noise term for time period t . Following Graves (1999), we note that as of period t , the forecast for demand in any future period equals the forecast for the next period, namely F_{t+1}^k .

¹⁶ We establish results by assuming there is a nonnegative lag L in delivery of each of the n products. Generally, we would allow delivery lag to vary with the product, but the model provides for complexities that have not yet been dealt with where lags are allowed to vary with the product.

Inventory Costs. As is standard, the objective is to choose an ordering policy that minimizes the expected discounted costs. There are three types of costs incurred in a period: ordering, holding, and shortage. Let $\mathbf{c} = (c^1, c^2, \dots, c^n)$, $\mathbf{h} = (h^1, h^2, \dots, h^n)$, and $\mathbf{p} = (p^1, p^2, \dots, p^n)$, respectively, be the purchasing, holding, and back-order cost vectors. Let $\mathbf{a} \cdot \mathbf{b}$ denote a scalar product of two vectors, \mathbf{a} and \mathbf{b} , \mathbb{E} stand for an expectation operator, $\phi \in (0, 1)$ denote a discount factor, and $(\cdot)^+ = \max\{\cdot, 0\}$. We assume for convenience that the cost $\mathbf{c} \cdot (\mathbf{y}_t - \mathbf{x}_t)$ of ordering the vector $(\mathbf{y}_t - \mathbf{x}_t)$ during period t is incurred at the time of delivery during period $t + L$. Then, the holding and shortage penalty costs during period $t + L$ are $\mathbf{h} \cdot (\mathbf{y}_t - \mathbf{D}_t - \mathbf{D}_{t+1} - \dots - \mathbf{D}_{t+L})^+ = \mathbf{h} \cdot (\mathbf{y}_t - \mathbf{D}_t - L\mathbf{F}_{t+1} - \boldsymbol{\epsilon}_{t+1} - \dots - \boldsymbol{\epsilon}_{t+L})^+$ and $\mathbf{p} \cdot (\mathbf{D}_t + \mathbf{D}_{t+1} + \dots + \mathbf{D}_{t+L} - \mathbf{y}_t)^+ = \mathbf{p} \cdot (\mathbf{D}_t + L\mathbf{F}_{t+1} + \boldsymbol{\epsilon}_{t+1} + \dots + \boldsymbol{\epsilon}_{t+L} - \mathbf{y}_t)^+$, respectively. For all \mathbf{y} and t , define

$$G_t(\mathbf{y}) = \mathbf{c} \cdot \mathbf{y} - \phi \mathbf{c} \cdot (\mathbf{y} - \mathbf{D}_t) + \mathbb{E}[\mathbf{h} \cdot (\mathbf{y} - \mathbf{D}_t - L\mathbf{F}_{t+1} - \boldsymbol{\epsilon}_{t+1} - \dots - \boldsymbol{\epsilon}_{t+L})^+] + \mathbb{E}[\mathbf{p} \cdot (L\mathbf{F}_{t+1} + \boldsymbol{\epsilon}_{t+1} + \dots + \boldsymbol{\epsilon}_{t+L} + \mathbf{D}_t - \mathbf{y})^+].$$

Transformation. To solve for the ordering policy, it is convenient to decompose the original problem into two subproblems. In the first subproblem, we set a base stock level to serve the *deterministic* portion of the lead-time demand, $\mathbf{D}_t + L\mathbf{F}_{t+1}$. In the second subproblem, we set a base stock level to serve the *stochastic* portion of demand, namely $\boldsymbol{\epsilon}_{t+1}, \boldsymbol{\epsilon}_{t+2}, \dots$. Define

$$\hat{\mathbf{y}} = \mathbf{y} - \mathbf{D}_t - L\mathbf{F}_{t+1} \tag{10}$$

to be the supply level in excess of the *deterministic* portion of demand. Then the expression for $G_t(\mathbf{y})$ can be written as

$$G_t(\hat{\mathbf{y}}) = \mathbf{c} \cdot \mathbf{D}_t + L(1 - \phi)(\mathbf{c} \cdot \mathbf{F}_{t+1}) + \hat{G}_t(\hat{\mathbf{y}}), \quad \text{where} \\ \hat{G}_t(\hat{\mathbf{y}}) = (1 - \phi)\mathbf{c} \cdot \hat{\mathbf{y}} + \mathbb{E}[\mathbf{h} \cdot (\hat{\mathbf{y}} - \boldsymbol{\epsilon}_{t+1} - \dots - \boldsymbol{\epsilon}_{t+L})^+] + \mathbb{E}[\mathbf{p} \cdot (\boldsymbol{\epsilon}_{t+1} + \dots + \boldsymbol{\epsilon}_{t+L} - \hat{\mathbf{y}})^+].$$

Denote by $f_T(\mathbf{x}_1)$ the expected costs incurred during periods $L + 1, \dots, L + T$ all discounted to the beginning of period $L + 1$ when the following holds (see Veinott 1965a, p. 209): (i) T is sufficiently large; (ii) \mathbf{x}_1 is the initial inventory; (iii) any stock left over after period $T + L$ can be discarded with a return of \mathbf{c} ; (iv) any backlogged demand remaining after period $T + L$ is fulfilled at the cost \mathbf{c} . Then:

$$f_T(\mathbf{x}_1) = \sum_{t=1}^T \phi^{t-1} \hat{G}_t(\hat{\mathbf{y}}_t) + \sum_{t=1}^T \phi^{t-1} (\mathbf{c} \cdot \mathbb{E} \mathbf{D}_t + L(1 - \phi)(\mathbf{c} \cdot \mathbf{F}_{t+1})) - \left(\mathbf{c} \cdot \mathbf{x}_1 - \phi^{T+1} \sum_{t=T+1}^{T+L} \mathbf{c} \cdot \mathbb{E} \mathbf{D}_t \right).$$

Because the term $\sum_{t=1}^T \phi^{t-1} (\mathbf{c} \cdot \mathbb{E} \mathbf{D}_t + L(1 - \phi)(\mathbf{c} \cdot \mathbf{F}_{t+1})) - (\mathbf{c} \cdot \mathbf{x}_1 - \phi^{T+1} \sum_{t=T+1}^{T+L} \mathbf{c} \cdot \mathbb{E} \mathbf{D}_t)$ is not affected by the choice of $\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \dots$, we can drop it from the optimization. The aforementioned objective then becomes:

$$\hat{f}_T(\mathbf{x}_1) = \sum_{t=1}^T \phi^{t-1} \hat{G}_t(\hat{\mathbf{y}}_t).$$

Thus we have reduced our model with ARIMA demand and a delivery lag to an equivalent model studied in Veinott

(1965a). This permits us to apply Theorems 3.1 and 3.2 in Veinott (1965a), which establish the existence of a multi-item base stock policy, say $\hat{\mathbf{y}}^*$ (that does *not* depend on t). Undoing the transformation (10) then implies the following optimal base stock level during period t :

$$\mathbf{y}_t^* = \hat{\mathbf{y}}^* + \mathbf{D}_t + L\mathbf{F}_{t+1}. \tag{11}$$

To derive the period t ordering rule, note that the inventory position at the beginning of period t will be:

$$\mathbf{x}_t = \mathbf{y}_{t-1} - \mathbf{D}_{t-1}. \tag{12}$$

Using (7), (11), and (12), the manager then places an order for the vector

$$\mathbf{q}_t^* = \mathbf{y}_t^* - \mathbf{x}_t = \hat{\mathbf{y}}^* + \mathbf{D}_t + L\mathbf{F}_{t+1} - (\mathbf{y}_{t-1}^* - \mathbf{D}_{t-1}) = \hat{\mathbf{y}}^* + \mathbf{D}_t + L\mathbf{F}_{t+1} - (\hat{\mathbf{y}}^* + \mathbf{D}_{t-1} + L\mathbf{F}_t) + \mathbf{D}_{t-1} = \mathbf{D}_t + L(\mathbf{F}_{t+1} - \mathbf{F}_t).$$

Or, equivalently

$$\mathbf{q}_t^* = \mathbf{D}_t + L(\mathbf{F}_{t+1} - \mathbf{F}_t),$$

which is Equation (1). \square

Derivation of Equation (2)

To derive Equation (2), which we later estimate empirically, we begin with Equation (1). We use the same notation as in the proof of Proposition 1. That is, n is the number of SKUs in a firm’s inventory; $\mathbf{q}_t = (q_t^1, q_t^2, \dots, q_t^n)$ is the vector of period t order quantities (i.e., q_t^k is the number of units of product $1 \leq k \leq n$ ordered during period t); \mathbf{D}_t and \mathbf{F}_t are the vectors of period t demands and period t demand forecasts (see Equations (6) and (8)); \mathbf{S}_t is the vector of sales quantities; \mathbf{c} is the vector of procurement costs, and \mathbf{r} is the vector of selling prices. Also, we assume that there is a common percentage markup, m , for each of the n products.¹⁷ As before, $\mathbf{a} \cdot \mathbf{b}$ denotes a scalar product of two vectors, \mathbf{a} and \mathbf{b} .

From Equation (1), we readily see that for each product $1 \leq k \leq n$ it must be true that:

$$q_t^k = D_t^k + L(F_{t+1}^k - F_t^k), \quad 1 \leq k \leq n.$$

Multiply both sides of the previous equation by the procurement cost of product k , c^k , and divide by $\mathbf{c} \cdot \mathbf{S}_t$ to obtain:

$$\frac{c^k q_t^k}{\mathbf{c} \cdot \mathbf{S}_t} = \frac{c^k D_t^k}{\mathbf{c} \cdot \mathbf{S}_t} + L \frac{c^k (F_{t+1}^k - F_t^k)}{\mathbf{c} \cdot \mathbf{S}_t}, \quad 1 \leq k \leq n. \tag{13}$$

Then in Equation (13), multiply both the numerator and the denominator of $(c^k (F_{t+1}^k - F_t^k)) / (\mathbf{c} \cdot \mathbf{S}_t)$ by $(1 + m)$ and sum over k . These steps yield

$$\frac{\mathbf{c} \cdot \mathbf{q}_t}{\mathbf{c} \cdot \mathbf{S}_t} = \frac{\mathbf{c} \cdot \mathbf{D}_t}{\mathbf{c} \cdot \mathbf{S}_t} + L \frac{\mathbf{r} \cdot (\mathbf{F}_{t+1} - \mathbf{F}_t)}{\mathbf{r} \cdot \mathbf{S}_t}. \tag{14}$$

In Equation (14), $\mathbf{c} \cdot \mathbf{q}_t$ denotes inventory purchases in dollars, $\mathbf{c} \cdot \mathbf{S}_t$ is cost of goods sold in dollars, $\mathbf{r} \cdot \mathbf{S}_t$ represents revenue in dollars, and $\mathbf{r} \cdot (\mathbf{F}_{t+1} - \mathbf{F}_t)$ is the change in demand

¹⁷ Note that under this assumption, selling prices can still vary across products.

forecast in dollars. Next, we show how Equation (14) can be simplified further to yield Equation (2) for the three forecast models tested in §5.1—(i) the perfect foresight, (ii) the analyst forecast, and (iii) the naive forecast.

For forecast model (i), there is $S_t = F_t$ by definition. Plugging it into Equation (14) gives

$$\frac{c \cdot q_t}{c \cdot S_t} = \frac{c \cdot D_t}{c \cdot S_t} + L \frac{r \cdot (F_{t+1} - F_t)}{r \cdot F_t}. \quad (15)$$

If we set $b_0 = (c \cdot D_t)/(c \cdot S_t)$ (which should be equal to 1 by the back-order assumption), $b_1 = L$, and $\Delta FSALES_{t+1} = (r \cdot (F_{t+1} - F_t))/(r \cdot F_t)$, then the previous equation reduces to Equation (2).

For forecast model (ii), $F_t = S_t$ if the mean analyst forecast error is zero. For many firms in our analysis, F_t is obtained by averaging sales forecasts of a large number of analysts and the average forecast error should be negligible. Thus Equation (14) also becomes (15). The forecast model (iii) is an exponential smoothing model (8) with $\alpha^k = 1$ for all k . Under this model, $r \cdot S_t = r \cdot F_t + r \cdot \epsilon_t$. Term $r \cdot \epsilon_t$ is negligible if a firm's product line is sufficiently diversified (for example, it carries multiple product categories that have independent market demands; or, within each category, the product variants are substitutes with negatively correlated demands). Thus, with (iii), Equation (15) is a reasonable approximation of Equation (14) unless all products in a firm's product line are perfect complements. □

References

Allen EJ, Larson CR, Sloan RG (2013) Accrual reversals, earnings and stock returns. *J. Accounting Econom.* 56(1):113–129.

Bollerslev T (1986) Generalized autoregressive conditional heteroskedasticity. *J. Econometrics* 31(3):307–327.

Cachon G, Olivares M (2010) Drivers of finished-goods inventory in the U.S. automobile industry. *Management Sci.* 56(1):202–216.

Cachon GP, Randall T, Schmidt GM (2007) In search of the bullwhip effect. *Manufacturing Service Oper. Management* 9(4):457–479.

Caliendo M, Kopeinig S (2008) Some practical guidance for the implementation of propensity score matching. *J. Econom. Surveys* 22(1):31–72.

Caplin AS (1985) The variability of aggregate demand with (S, s) inventory policies. *Econometrica* 53(6):1395–1409.

Cermeno R, Grier KB (2001) Modeling GARCH processes in panel data: Theory, simulations and examples. Working paper, University of Oklahoma, Norman.

Cermeno R, Grier KB (2006) Conditional heteroskedasticity and cross-sectional dependence in panel data: An empirical study of inflation uncertainty in the G7 countries. *Contributions Econom. Anal.* 274:259–277.

Chen H, Frank MZ, Wu O (2005) What actually happened to the inventories of American companies between 1981 and 2000? *Management Sci.* 51(7):1015–1031.

Chen H, Frank MZ, Wu O (2007) U.S. retail and wholesale inventory performance from 1981 to 2004. *Manufacturing Service Oper. Management* 9(4):430–456.

Cotter J, Tuna I, Wysocki PD (2010) Expectations management and beatable targets: How do analysts react to explicit earnings guidance? *Contemporary Accounting Res.* 23(3):593–624.

Fama EF, MacBeth JD (1973) Risk, return, and equilibrium: Empirical tests. *J. Political Econom.* 81(3):607–636.

Francis J, Hanna JD, Vincent L (1996) Causes and effects of discretionary asset write-offs. *J. Accounting Res.* 34:117–134.

Gaur V, Kesavan S (2009) The effects of firm size and sales growth rate on inventory turnover performance in the U.S. retail sector. Argawal N, Smith SA, eds. *Retail Supply Chain Management*, Chap. 7, Vol. 122 (Springer-Verlag, Berlin), 25–52.

Gaur V, Fisher ML, Raman A (2005) An econometric analysis of inventory turnover performance in retail services. *Management Sci.* 51(2):181–194.

Graves SC (1999) A single-item inventory model for a nonstationary demand process. *Manufacturing Service Oper. Management* 1(1):50–61.

Hendricks KB, Singhal VR (2005a) Association between supply chain glitches and operating performance. *Management Sci.* 51(5):695–711.

Hendricks KB, Singhal VR (2005b) An empirical analysis of the effect of supply chain disruptions on long-run stock price performance and equity risk of the firm. *Production Oper. Management* 14(1):35–52.

Hendricks KB, Singhal VR (2009) Demand-supply mismatches and stock market reaction: Evidence from excess inventory announcements. *Manufacturing Service Oper. Management* 11(3):509–524.

Johnson GD, Thompson HE (1975) Optimality of myopic inventory policies for certain dependent demand processes. *Management Sci.* 21(11):1303–1307.

Kahn JA (1987) Inventories and the volatility of production. *Amer. Econom. Rev.* 77(4):667–679.

Kesavan S, Kushwaha T (2011) Do inventory holding costs and stockout costs explain inventory investment behavior of U.S. public retailers during economic shocks? Working paper, Kenan-Flagler Business School, Chapel Hill, NC.

Kesavan S, Gaur V, Raman A (2010) Do inventory and gross margin data improve sales forecasts for U.S. public retailers? *Management Sci.* 56(9):1519–1533.

Kessinger C, Pieper H (2005) Managing risk with structured supply agreements. Chae A, Fromm H, eds. *Supply Chain Management on Demand*, Chap. 7 (Springer-Verlag, Berlin), 143–164.

Khan A, Thomas JK (2007) Inventories and the business cycle: An equilibrium analysis of (S, s) policies. *Amer. Econom. Rev.* 97(4):1165–1188.

Lai R (2006) Inventory and stock market. Working paper, Harvard University, Cambridge, MA.

Lamond BF, Sobel MJ (1995) Exact and approximate solutions of affine reservoir models. *Oper. Res.* 43(5):771–780.

Lovejoy WS (1992) Stopped myopic policies in some inventory models with generalized demand process. *Management Sci.* 38(5):688–707.

Matsumoto DA (2002) Management's incentives to avoid negative earnings surprises. *Accounting Rev.* 77(3):483–514.

Montgomery DC, Johnson LA (1976) *Forecasting and Time Series Analysis*, 1st ed. (McGraw Hill, New York).

Morton TE, Pentico DW (1995) The finite horizon nonstationary stochastic inventory problem: Near-myopic bounds, heuristics, testing. *Management Sci.* 41(2):334–343.

Mosser PC (1991) Trade inventories and (S, s) . *Quart. J. Econom.* 106(4):1267–1286.

Nissim D, Penman SH (2001) Ratio analysis and equity valuation: From research to practice. *Rev. Accounting Stud.* 6(1):109–154.

Porteus EL (2002) *Foundations of Stochastic Inventory Theory*, 1st ed. (Stanford University Press, Palo Alto, CA).

Rajagopalan S, Malhotra A (2001) Have U.S. manufacturing inventories really decreased? An empirical study. *Manufacturing Service Oper. Management* 3(1):14–24.

Rumyantsev S, Netessine S (2007) What can be learned from classical inventory models? A cross-industry exploratory investigation. *Manufacturing Service Oper. Management* 9(4):409–429.

Song J, Zipkin PH (1996) Managing inventory with the prospect of obsolescence. *Oper. Res.* 44(1):215–222.

- Stoll HR (2000) Friction. *J. Finance* 55(4):1479–1514.
- Veinott AF Jr (1965a) Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.* 12(3): 206–222.
- Veinott AF Jr (1965b) Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes. *Oper. Res.* 13(5):761–778.
- Vyas D (2011) The timeliness of accounting write-downs by U.S. financial institutions during the financial crisis of 2007–2008. *J. Accounting Res.* 49(3):823–860.
- Wagner HM (2002) And then there were none. *Oper. Res.* 50(1): 217–226.
- Zipkin PH (2000) *Foundations of Inventory Management*, 1st ed. (McGraw Hill, New York).