Seasonal Reversals in Expected Stock Returns*

Matti Keloharju   Juhani T. Linnainmaa   Peter Nyberg

October 2018

Abstract

Stocks tend to earn high or low returns relative to other stocks every year in the same calendar month (Heston and Sadka 2008; Keloharju, Linnainmaa, and Nyberg 2016). In this paper, we show that these seasonalities are balanced by seasonal reversals: a stock that has a high expected return relative to other stocks in one month has a low expected return relative to other stocks in the other months. The seasonalities and seasonal reversals add up to zero over the calendar year. Our evidence suggests that return seasonalities are likely due to temporary mispricing. Seasonal reversals are economically large, statistically highly significant, and they resemble, but are distinct from, long-term reversals. A factor that estimates expected returns from both average same- and other-month returns has a t-value of 9.93, and it is robust throughout the 1963–2016 sample period.

*Keloharju is with the Aalto University School of Business, CEPR, and IFN. Linnainmaa is with the University of Southern California and NBER. Nyberg is with the Aalto University School of Business. We thank Chris Hrdlicka, Mark Kamstra, Owen Lamont, and Jon Lewellen for insightful comments. Financial support from the Academy of Finland, Inquire Europe, and Nasdaq Nordic Foundation is gratefully acknowledged.

1 Introduction

Stocks that are winners in a given month tend to continue to outperform stocks that are losers in that same calendar month, for up to 20 years (Heston and Sadka 2008; Keloharju, Linnainmaa, and Nyberg 2016). For example, if a stock has performed well (poorly) relative to other stocks in March in the past, we can expect it to offer a high (low) return relative to other stocks also next March. At the same time, there is little evidence of persistent differences in expected returns between stocks (Keloharju, Linnainmaa, and Nyberg 2018). These two findings—long-term predictability in the form of seasonalities and the lack of long-term differences in expected returns—are seemingly at odds with each other.

Figure 1 illustrates this seeming contradiction. The red line is about seasonalities. We assign stocks into deciles each month based on their average same-month return over the prior 20-year period; for example, at the end of February 2009 we form the portfolios by stocks' average returns in March 1989, March 1990,..., and March 2008. We compute value-weighted returns for the resulting portfolios over the following ten years and report t-values associated with the difference between the top and bottom deciles. Return seasonalities are the spikes in the figure: a high March return in the past predicts a high March return not just this year (horizon = month 1 in the figure) but also far out into the future. The black line is about unconditional differences in average returns. Following Keloharju, Linnainmaa, and Nyberg (2018), we assign stocks into portfolios based on a combination of 34 return predictors; we use the set of accounting-based predictors that show the most persistence in the original study. Unlike seasonalities, the differences between the top and bottom deciles vanish in a few years. Although stock returns are significantly predictable far out into the future, we are unable to identify persistent differences in stocks' expected returns.

To reconcile these two facts, we hypothesize that seasonalities must be offset by seasonal reversals. For example, if a stock's expected return in March exceeds that of the market average, its total expected return in the other months must fall below the market average by the same amount. Seasonalities must

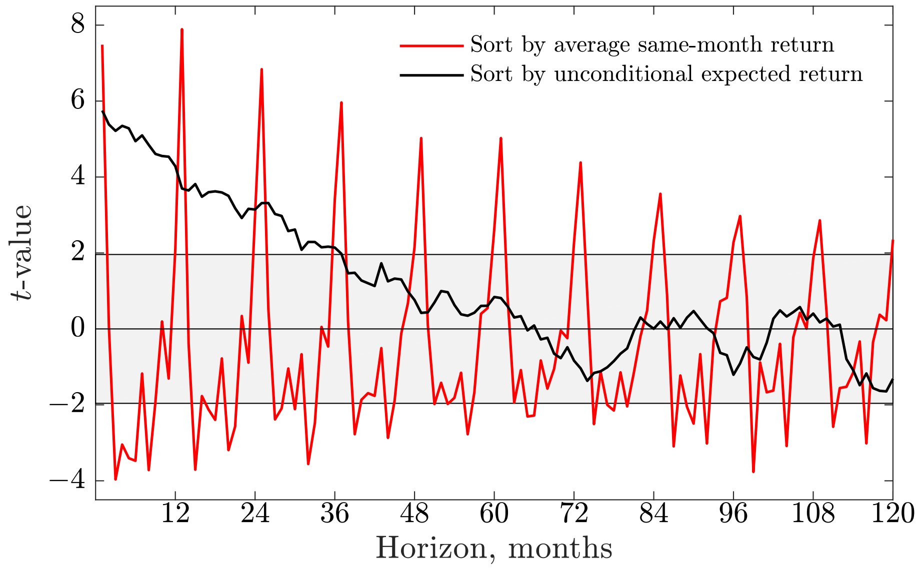

Figure 1: **Long-lasting seasonalities and the lack of long-term differences in expected returns.** We assign stocks into deciles each month either by the average same-month return over the prior 20-year period (red line) or by a combination of 34 accounting-based predictors (black line). The historical same-month return is computed from the viewpoint of horizon = month 1 returns. We compute value-weighted returns for the resulting portfolios over the next 10 years. This figure reports t-values associated with the differences between the top and bottom deciles. The shaded area indicates estimates that are significant at the 5% level.

add up to zero over the calendar year for them not to leave a trace in unconditional expected returns. The adding-up constraint we hypothesize is not a tautology. If a stock's expected return *relative* to the other stocks is high in one month, it does not have to have a low expected return relative to other stocks in the other months. It would be tautological to state that a stock with a high expected return in one month relative to its own time-series mean must earn a low expected return in the other months relative to this mean. The adding-up constraint is a statement about *cross-sectional* differences in expected returns, not about time-series differences in expected returns.

We first test whether the adding-up constraint holds in the data. We can gauge this by computing the correlation between a stock's expected return in one month, proxied by its historical average returns in that month, and the sum of its expected returns in the other months. If the adding-up constraint holds perfectly and if expected returns can be observed without noise, the correlation is −1. In reality, the noise in the expected return estimates biases the correlation towards zero. We therefore assess the extent to which this constraint holds in the data using simulations. We simulate data from a model in which the adding-up constraint holds perfectly and the simulated returns are as noisy as true returns. Both the simulations and the data give the same correlation estimate of -0.06. The data are thus consistent with a model in which cross-sectional differences in expected returns cancel out over the calendar year.

Seasonalities are not confined to monthly equity returns in the U.S. If seasonalities are balanced by seasonal reversals, we would expect to find seasonal reversals wherever seasonalities are found. Following Keloharju, Linnainmaa, and Nyberg (2016), we measure seasonalities and seasonal reversals in daily stock returns, country equity indexes, and commodity returns. We find seasonal reversals in all of them. In addition, we also document seasonalities and seasonal reversals in international stock returns and country government bond indexes.

Our insights on seasonal reversals improve the predictive power of seasonal trading strategies. Given that realized returns are noisy, same- and other-calendar-month returns both contain independent information about future expected returns.[1] A factor that sorts stocks based on the same-minus-other-calendar-month difference earns an average return of 67 basis points per month with a t-value of 9.93, a notable increase from the seasonality factor's average return of 61 basis points per month (t-value = 8.37). Neither seasonalities nor seasonal reversals subsume each other, consistent with them containing independent information about expected returns.

Seasonalities and seasonal reversals are unrelated to short-term reversals, momentum, and long-term reversals. Although seasonal reversals resemble long-term reversals, different mechanisms drive them. The average return on the long-term reversal factor is 29 basis points per month (t-value = 2.95), but its correlations with size and value render its three-factor model alpha statistically insignificant (Fama and French 1996; Asness, Moskowitz, and Pedersen 2013). The seasonal reversal factor's three-factor model

[1] If the expected returns could be measured without noise, a perfect adding-up relationship between seasonalities and seasonal reversals would make one of them redundant in a predictive regression.

alpha, by contrast, is significant with a t-value of 6.17. The addition of the momentum and long-term reversal factors lower this t-value, but only to 5.33. That is, the seasonal reversal factor is more than just another version of the long-term reversal factor.

Return seasonalities must be due to time-variation in the price of risk, the quantity of risk, or mispricing. We view our evidence about seasonal reversals as suggesting that seasonalities are unlikely due to seasonal variation in the price or quantity of risk. A risk factor's premium may be higher in one month because the underlying risk matters more, or is perceived as being more costly to bear, in that month than others. However, if so, why is the risk premium lower in all other months by an exactly offsetting amount? That is, the risk-based explanation posits no reason for the seasonalities to add up to zero.[2]

Although seasonal reversals could balance out seasonalities by luck, the fact that we observe reversals not only in U.S. equities at the monthly frequency, but also at the daily frequency and in other asset classes, casts doubt on this explanation. A more plausible explanation is that seasonalities emanate from temporary mispricing. Heston, Korajczyk, and Sadka (2010) find seasonalities in *intraday returns*: a stock's return over a 30-minute interval predicts its return over the same interval for up to 40 trading days. They attribute this seasonality to traders consistently trading in the same direction at the same time of the day. This consistency in supply or demand represents a shock that is imperfectly absorbed by the rest of the market, thereby generating a price effect. These intraday seasonalities cancel out because they represent just temporary deviations from fundamental prices. Our results suggest that most return seasonalities across multiple asset classes, and even at the monthly frequency, may also be due to temporary mispricing induced by the predictable trading of investors.

The rest of the paper is organized as follows. Section 2 describes how different assumptions about

[2] Suppose that all information—not just about assets but also about human capital, economy, and so forth—is only released once a year in December, and that there is no leakage of information and no asymmetric information. In this world, risky assets earn all of their risk premiums in December; outside December, risky assets are effectively riskless, as all market participants know that no other market participant receives any information pertaining to asset valuations outside December. Here, risk premiums do not reverse: the high December risk premium is not offset by a negative non-December risk premium.

the nature of the cross-sectional variation in expected returns alter the predictive relationship between returns and lagged returns. Section 3 measures seasonalities and seasonal reversals using Fama and MacBeth (1973) regressions. Section 4 calibrates a model with seasonalities and long-term reversals to quantify the extent to which seasonalities add up to zero. Section 5 constructs seasonality and seasonal reversal factors, and examines their relation to short-term reversals, momentum, and long-term reversals. Section 6 measures seasonal reversals in daily stock returns, international stock returns, country equity indexes, country government bond indexes, and commodity returns. Section 7 concludes.

2 Seasonalities and seasonal reversals in expected stock returns

In this section we analyze how different assumptions about the nature of the cross-sectional variation in expected returns alter the predictive relation between past and contemporaneous returns. Panel A of Figure 2 plots the average Fama and MacBeth (1973) coefficients from univariate regressions of month t returns against month $t-k$ returns,

$$r_{it} = a + b \times r_{i,t-k} + \varepsilon_{it}, \tag{1}$$

where r_{it} is stock i's return in month t. Panel B is similar to Panel A except that it predicts returns with past annual returns. This distinction between monthly and annual returns is important for seasonal reversals. We estimate all regressions in Figure 2 using lags up to 10 years.

The first subpanel in the upper left reports the estimates that use the monthly CRSP return data from January 1963 through December 2016 on stocks listed on NYSE, Amex, and NASDAQ.[3] The negative coefficient at the first lag is about short-term reversals; the positive coefficients up to the year mark are about momentum; and the spikes are about the seasonalities in stock returns (Heston and

[3] We exclude securities other than ordinary common shares. We use CRSP delisting returns; if a delisting return is missing and the delisting is performance-related, we impute a return of −30% (Shumway 1997). Later in this study, we include book-to-market as a control variable. We use the book values of equity from the annual Compustat files, supplemented with the Davis, Fama, and French (2000) data, and follow the standard conventions to time this information.

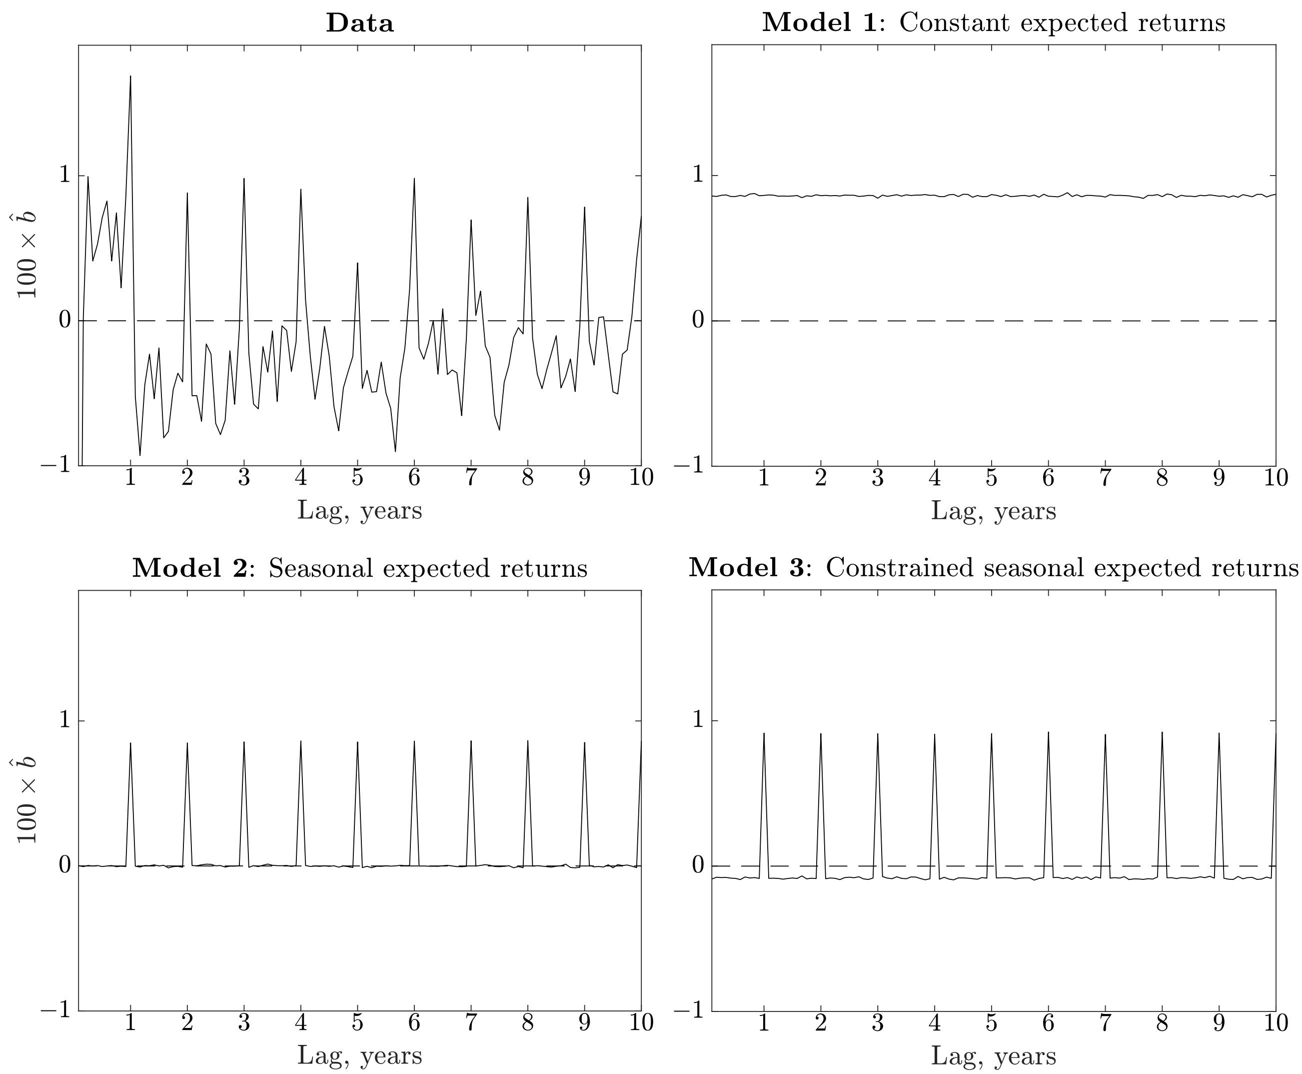

Panel A: Regressions against past monthly returns

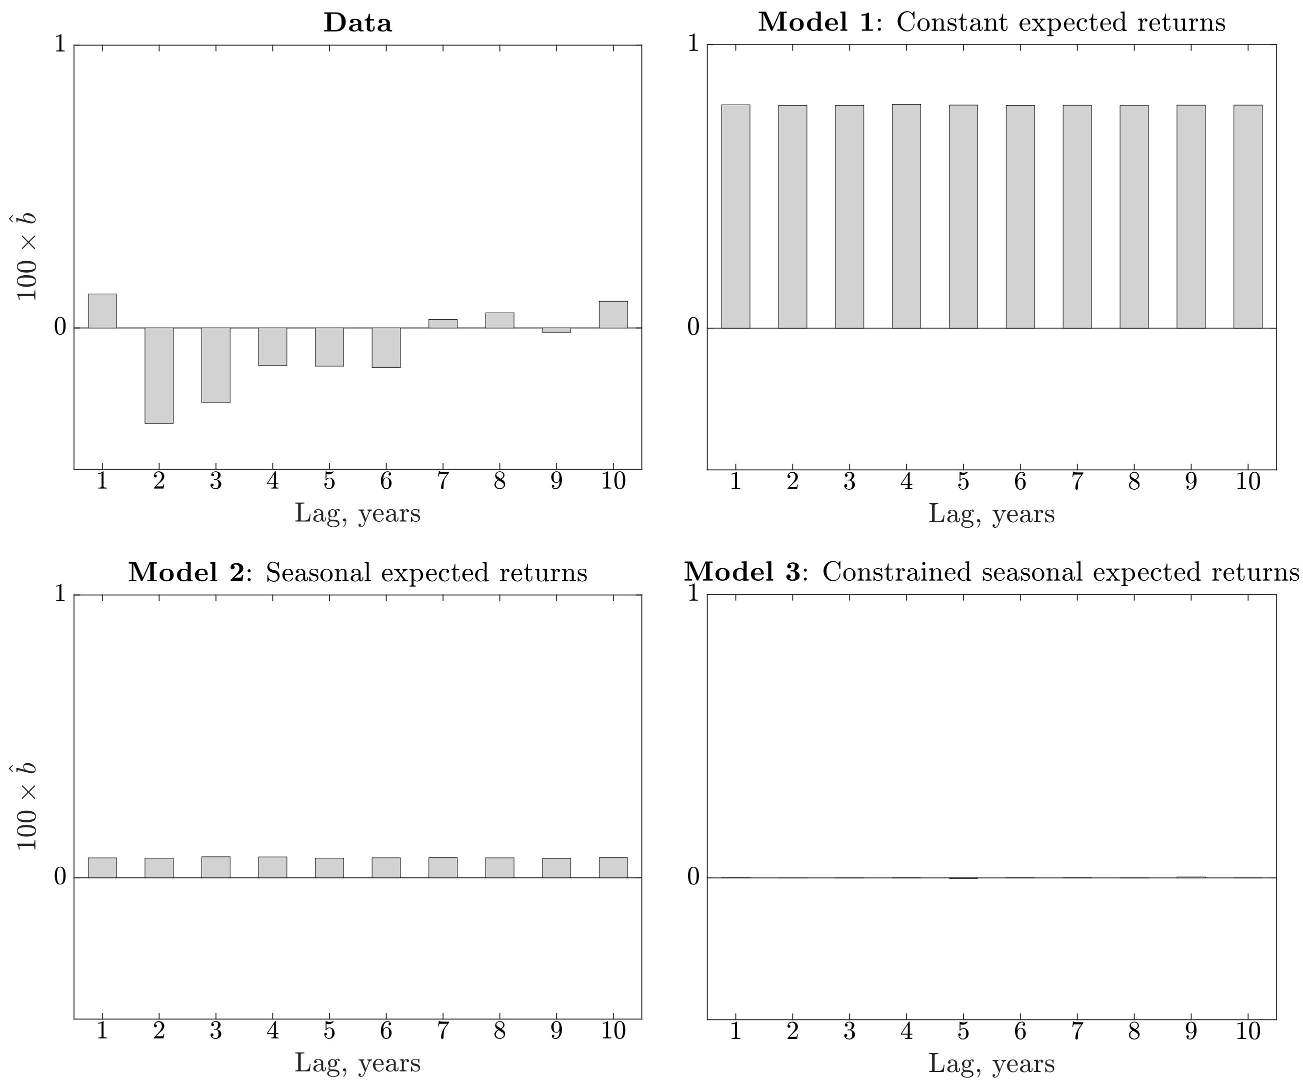

Panel B: Regressions against past annual returns

Figure 2: **Fama-MacBeth regressions: Data versus theory.** This figure reports estimates from univariate Fama-MacBeth regressions that predict the cross section of monthly returns with past monthly (Panel A) and past annual (Panel B) returns using lags up to ten years. The first subpanel uses return data on NYSE, Amex, and Nasdaq stocks from January 1963 through December 2016. The other subpanels simulate data under different assumptions about expected returns. In Model 1, expected stock returns are constant. In Model 2, expected stock returns vary by calendar month. In Model 3, expected stock returns vary by calendar month and satisfy the adding-up constraint. This adding-up constraint restricts the sum of each stock's expected returns to zero.

Sadka 2008).

The first model, illustrated in the upper right subpanel, is one in which stock returns contain a persistent expected return component but no seasonal component. We draw stock returns from the process

$$r_{it} = \mu_i + \epsilon_{it}, \tag{2}$$

where ϵ_{it} is I.I.D. In this model, a stock's expected return could be 5% per year every month of the year; for another stock it could be 8% per year. With persistent differences in expected returns, realized returns exhibit "poor man's momentum:" month $t-k$ return predicts the cross section of month t returns because we explain $\mu_i + \epsilon_{it}$ with $\mu_i + \epsilon_{i,t-k}$—that is, the same expected return appears both on the left- and right-hand sides of the regression.[4] The theoretical regression coefficient at any lag k therefore equals

$$\hat{b}_k = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2}. \tag{3}$$

This positive predictive relationship holds both when we explain today's returns with monthly (Panel A) and annual (Panel B) returns. The model's predictions therefore profoundly contradict the data. Setting aside the short-term reversals and momentum, the main difference between the actual data and the simulated data is that, in the model, past monthly returns predict today's returns at all lags. In the data, this positive predictive relationship holds only at annual lags in the monthly regressions. In Panel B, past returns after year 1 are typically negatively correlated with the cross section of monthly returns.

In the second model, depicted in the lower left subpanel, stocks' expected returns display seasonal variation. The stock return process is

$$r_{it} = \mu_{i,m(t)} + \epsilon_{it}, \tag{4}$$

where ϵ_{it} is I.I.D. and $\mu_{i,m(t)}$ is stock i's expected return in calendar month $m(t) = 1, \ldots, 12$. This

[4] See, for example, Lo and MacKinlay (1990), Conrad and Kaul (1998), and Berk, Green, and Naik (1999) for discussions of this mechanism.

model says that a stock's expected return in October could differ from its expected return in November. Because we predict the *cross section* of returns, we do not specify the level of expected returns; it washes out from the regression estimates. In this model, we assume that a stock's expected return in one month is independent of its expected returns in the other months. That is, we do not impose any constraint on the sum of the expected returns, $\mu_{i,1} + \cdots + \mu_{i,12}$. Panel A's Fama-MacBeth regression coefficient then equals

$$\hat{b}_k = \begin{cases} \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} & \text{at annual lags,} \\ 0 & \text{at non-annual lags.} \end{cases} \tag{5}$$

This model is consistent with the data with respect to the seasonal spikes. Past same-month returns positively predict today's return because both contain the same expected return component, $\mu_{i,m}$. At non-annual lags, past returns have no prediction power on expected returns today.

Panel B of Figure 2 shows that, in this model, historical annual returns still positively predict today's returns. The reason is that some stocks have predominantly positive seasonalities—$\mu_{i,1} + \cdots + \mu_{i,12} \gg 0$—while others have predominantly negative seasonalities—$\mu_{i,1} + \cdots + \mu_{i,12} \ll 0$. A stock might, for example, have high expected returns for six months of the year, and expected returns close to zero for the rest of the year. Annual returns inform about these sums of seasonalities. A stock with a high realized annual return is more likely a stock with more positive than negative seasonalities. Therefore, without a constraint on the seasonalities in expected returns, past annual returns positively predict today's returns, just as they would in a model with constant expected returns. This positive correlation between today's returns and past annual returns contradicts the data.

The third model, illustrated in the lower right subpanel, imposes the following adding-up constraint on the seasonalities,

$$\mu_{i,1} + \mu_{i,2} + \cdots + \mu_{i,12} = 0. \tag{6}$$

In Panel A's monthly regressions, the regression coefficient at annual lags is the same in this model as in the model without the constraint. However, because of the "adding-up constraint" in equation (6), a stock's realized return in, say, January is informative about its expected returns both in January and in all other months. A stock with an unusually high January expected return must have unusually low expected returns throughout the rest of the year. A stock's expected return in January, therefore, negatively relates to its expected return in, for example, February:

$$\mu_i^{\text{Jan}} = -\left(\mu_i^{\text{Feb}} + \cdots + \mu_i^{\text{Dec}}\right) = -\mu_i^{\text{Feb}} + \text{“noise”}. \tag{7}$$

The Fama-MacBeth regression coefficient under Model 3 equals

$$\hat{b}_k = \begin{cases} \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} & \text{at annual lags, and} \\ \frac{-\frac{1}{11}\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} & \text{at non-annual lags.} \end{cases} \tag{8}$$

This model is consistent with several features of the data. First, similar to the model without the constraint, the seasonalities in expected returns generate annual spikes in the monthly regression coefficients. Second, because the seasonalities add up to zero, the non-annual regression coefficient are pushed downwards. These negative troughs are the seasonal reversals. Third, because every stock's annual expected return equals zero, annual realized returns do not predict differences in future returns.

However, the model also is inconsistent with some aspects of the data. The short-term reversals and momentum are short-run, autocorrelation-like effects, and a model with only persistent variation in expected returns cannot match these features. Similarly, the long-term reversals of De Bondt and Thaler (1985) cannot only be about seasonal reversals. Panel B of Figure 2 shows that these negative coefficients are present also in annual regressions. These negative coefficients cannot emanate from seasonalities alone. In a model with just seasonal variation in expected returns, the coefficients are either positive, as in the second model, or zero, as in the third model. They cannot be negative. Negative correlations must emanate either from negative serial correlations or from positive cross-covariances across assets (Lo and MacKinlay 1990).

Table 1: Average annual and non-annual returns in Fama-MacBeth regressions

This table presents average Fama and MacBeth (1973) regression slopes and their t-values from cross-sectional regressions that predict monthly returns. The regressions use data from January 1963 through December 2016 for all stocks (regressions 1–3), all-but-microcaps (regressions 4–6), and 48 value-weighted Fama-French industries (regressions 7–9). Microcaps are stocks with market values of equity below the 20th percentile of the NYSE market capitalization distribution. We cross-sectionally demean the data before computing the same- and other-month returns, $\bar{r}_{\text{same-month}}$ and $\bar{r}_{\text{other-month}}$. Both averages use up to 20 years of historical data. The average non-annual return skips a year; that is, to predict the cross section month t return, the first term in the other-month average is the month $t-13$ return. Regression estimates are multiplied by the factor of 100.

Explanatory	All stocks			All-but-microcaps			Industries		
variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
log(ME)	−0.07	−0.07	−0.06	−0.06	−0.07	−0.07	−0.04	−0.03	−0.03
	(−2.25)	(−2.01)	(−1.95)	(−1.97)	(−2.41)	(−2.42)	(−1.61)	(−1.16)	(−1.09)
log(BE/ME)	0.30	0.20	0.18	0.22	0.12	0.10	0.15	0.02	0.01
	(5.84)	(4.33)	(4.00)	(3.91)	(2.45)	(2.06)	(1.78)	(0.26)	(0.06)
r_1	−5.54	−5.58	−5.62	−3.39	−3.52	−3.56	4.74	4.80	4.24
	(−15.54)	(−15.82)	(−16.00)	(−8.07)	(−8.56)	(−8.76)	(4.48)	(4.49)	(4.07)
$r_{12,2}$	0.46	0.44	0.43	0.42	0.41	0.40	1.24	1.21	1.17
	(2.98)	(2.92)	(2.86)	(2.21)	(2.16)	(2.13)	(3.49)	(3.39)	(3.35)
$r_{60,13}$			−0.06			−0.05			−0.05
			(−2.10)			(−2.10)			(−0.60)
$\bar{r}_{\text{same-month}}$	5.47	4.67	4.93	6.52	6.14	6.56	19.10	17.98	19.22
	(9.88)	(8.17)	(8.56)	(8.70)	(8.10)	(8.66)	(5.92)	(5.46)	(5.87)
$\bar{r}_{\text{other-month}}$		−18.51	−16.05		−16.36	−12.84		−42.94	−35.25
		(−6.57)	(−5.71)		(−4.50)	(−3.46)		(−3.33)	(−2.64)

3 Seasonalities and seasonal reversals in Fama-MacBeth regressions

Table 1 reports estimates from Fama-MacBeth regressions that predict the cross section of monthly stock returns. We estimate these regressions for all stocks, all-but-microcaps, and for 48 value-weighted Fama-French industry portfolios. Microcaps are stocks with market values of equity below the 20th percentile of the NYSE market capitalization distribution as of the end of month $t-1$.

Regressions (1), (4), and (7) predict returns using log-size, log-book-to-market, past-month return, the prior one-year return skipping a month, and the average same-calendar-month return. We compute this average return from cross-sectionally demeaned returns using up to 20 years of historical data.[5] Average returns decrease in size and increase in both book-to-market and momentum. These three effects are statistically significant both for all stocks and for the all-but-microcaps sample. The estimated slope on the average same-calendar-month return is positive and statistically significant; its t-value is 9.88 in the regressions that include all stocks, and 8.70 in the all-but-microcaps sample. This effect, which is consistent with the estimates in Heston and Sadka (2008) and Keloharju, Linnainmaa, and Nyberg (2016), is economically large. The coefficient estimate of 5.5 in the full sample, for example, implies that a 1% difference in past average same-calendar-month returns between two stocks predicts a 0.055% difference in these stocks' returns this month.

Regressions (2), (5), and (8) add the average other-month return to the model. The estimated slope on this variable is negative and statistically significant. Its t-value is -6.57 in the full sample and -4.50 in the all-but-microcaps sample. This effect is economically even larger: a 1% difference in the average other-month returns in regression (2) translates into a -0.19% difference in monthly returns today.

The fact that both the same- and other-month returns remain significant is consistent with seasonalities being balanced by seasonal reversals. Although both variables measure the same underlying quantity—the stock's expected return $\mu_{i,m(t)}$—they are incrementally informative because returns are noisy signals of expected returns.

Because the average same- and other-month returns are closely related to long-term reversals, we add these reversals as an additional control in regressions (3), (6), and (9). We use the usual definition of long-term reversals, measuring stock returns over the prior five-year period and skipping a year. The addition of these long-term reversals has only a modest effect on the slope estimates for the average same- and other-month returns. The coefficients and t-values on the average same-month return increase slightly, and those on the average other-month return decrease slightly. The long-term reversal variable

[5]If all stocks have the same amount of historical data, the cross-sectional demeaning does not change the estimates because the demeaning shifts all averages up or down by the same amount. Demeaning ensures that the average same-month returns of stocks with different amounts of historical data are comparable (Keloharju, Linnainmaa, and Nyberg 2016).

itself is significant with a t-value of -2.10 in both the full sample and in the all-but-microcaps sample.

The industry estimates in regressions (7) through (9) show that seasonal reversals are also present in the returns of value-weighted industry portfolios. Because these portfolios are well-diversified, this significance suggests that seasonal reversals are unlikely to emanate from any stock-specific effects. Although some patterns in industry returns differ from those in stock returns—most importantly, industries display significant momentum already in month 1 (Moskowitz and Grinblatt 1999) and the cross-industry value effect is, at best, weak (Cohen and Polk 1996; Novy-Marx 2013)—the patterns related to seasonalities and seasonal reversals are quite similar. The industry results also further highlight the difference between seasonal and long-term reversals. While the long-term reversals estimate is within just one standard error from zero in regression (9), the t-value associated with seasonal reversals is -2.64.

4 Measuring seasonal reversals

4.1 Model

In this section we calibrate a model to the data to assess the extent to which return seasonalities cancel out over the calendar year. We simulate data from a model and choose the parameters to fit the annual spikes and non-annual troughs of the data subpanel of Figure 2 Panel A. In this model, we continue with the assumption that a stock's realized return equals its seasonal expected return plus noise,

$$r_{it} = \mu_{i,m(t)} + \epsilon_{it}. \tag{9}$$

We generate the $\mu_{i,m(t)}$s as follows. First, we generate 12 draws from a normal distribution,

$$\mu_{i,m}^{e} \sim N(0, \sigma_{\mu}^{2}) \text{ for } m = 1, \ldots, 12, \tag{10}$$

and then demean the resulting draws,

$$\mu_{i,m} = \mu^e_{i,m} - \frac{1}{12}\sum_{m=1}^{12}\mu^e_{i,m}. \tag{11}$$

These expected returns $\mu_{i,m}$s thus perfectly satisfy the adding-up constraint: a high expected return in one month is offset by correspondingly lower expected returns in the other months.

Because stock returns exhibit long-term reversals, and because long-term reversal also induce a negative cross-sectional correlation between stock returns and lagged stocks returns, we let the return innovations ϵ_{it} exhibit such reversals after the one-year mark. Specifically, we assume that this return innovation equals

$$\epsilon_{it} = \sum_{k=13}^{120}\delta_k \xi_{i,t-k} + \xi_{i,t}, \tag{12}$$

where $\xi_{it} \sim N(0, \sigma^2_\xi)$. With $\sum_{k=13}^{120}\delta_k < 0$, this assumption builds in long-term reversals. Because our interest is calibrating the model to the seasonalities and reversals in the data, we do not model the short-term reversals and momentum. For simplicity, we assume that δ_k is non-positive and that it changes linearly in k, that is, $\delta_k = \min(\bar{\delta} + k \times \gamma, 0)$. This assumption permits the possibility that the reversals strengthen or weaken over time.

4.2 Calibration

The model is characterized by four parameters: σ^2_μ generates the seasonalities in expected returns; $\bar{\delta}$ and γ generate the long-term reversals; and σ^2_ξ determines the amount of noise in individual stock returns.

We simulate data by taking the full CRSP database as the starting point. We then replace the true returns with returns simulated from the model. This procedure ensures that the number of firms each month matches the number of firms in the actual data. We choose the parameters to match the autocorrelation patterns shown in Panel A of Figure 2. We match, between the simulations and the

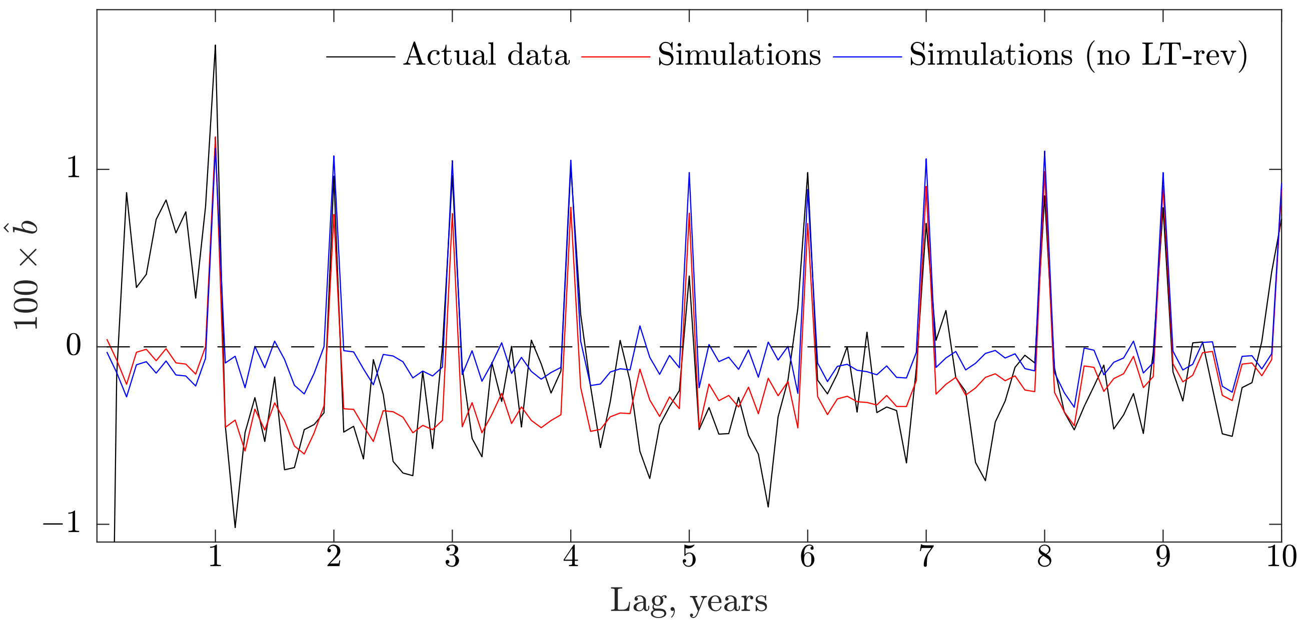

Figure 3: **Seasonalities, seasonal reversals, and long-term reversals.** The black line represents the estimates from univariate Fama-MacBeth regressions that predict the cross section of monthly returns with past monthly returns using lags up to ten years. The red and blue lines represent the same coefficients computed using simulated data with the same dimensions as in the actual data. In the model, stock returns' expected returns vary by calendar month and add up to zero, and return innovations display long-term reversals from month $t-120$ to $t-13$. These long-term reversals dampen over time. The model is calibrated to match the cross-sectional variance of stock returns and the annual spikes and non-annual troughs in the Fama-MacBeth regressions. The red line simulates data from the full model; the blue line shuts down long-term reversals, leaving only seasonalities and seasonal reversals in the model.

data, the cross-sectional variance of stocks returns and the coefficients from regressions that predict the cross section of monthly returns with past returns. The explanatory returns consist of the annual returns in months $t-12$, $t-24$, $\ldots$, $t-120$; in addition, we include the average non-annual returns over the prior ten years, skipping a year. That is, the first of these non-annual regressions uses the average return from month $t-23$ to month $t-13$; the second uses the average return from month $t-35$ to month $t-25$; and so forth. We find the parameters with the Simulated Method of Moments, using the identity matrix as the weighting matrix to match these 20 moments—one cross-sectional variance, 10 annual regression coefficients, and 9 non-annual regression coefficients—between the data and the model.

Figure 3 shows the average Fama-MacBeth coefficients from regressions that predict the cross section of monthly returns with lagged returns. The black line represents the estimates that use the actual data. These estimates are the same as those reported in Panel A of Figure 2. The red and blue lines report estimates that are based on one simulation each. The red line simulates one set of data from the full model with both seasonalities and long-term reversals. The blue line simulates data with otherwise the same parameters except that it shuts down long-term reversals by setting $\bar{\delta} = \gamma = 0$.

The model is not designed to match short-term reversals and momentum, so the red line (simulation) differs substantially from the black line (data) up to the one-year mark. However, after this point, the model matches the key features of the return data. Both the seasonal spikes and the non-seasonal troughs are of the same magnitude. This similarity indicates that real data are consistent with a model in which seasonal reversals completely balance out seasonalities.

4.3 Correlations in average calendar-month returns: Data versus simulations

A correlation between a stock's expected return in one month and the sum of its expected returns in the other months is a measure of the extent to which the seasonalities satisfy the adding-up constraint. This correlation is -1 if this constraint holds perfectly. We measure the correlation from average returns. We first take all stocks with at least 10 years of data over the entire sample period. We then cross-sectionally demean the data and compute, for each stock, the average return in each calendar month. We reorganize the data so that we have 12 observations for each stock: a stock's average January return aligned with the sum of its average February-through-December returns, and similarly for the other months. We then estimate the regression

$$\bar{r}_{i,m} = a + b \times \sum_{m' \neq m} \bar{r}_{i,m'} + e_{i,m}. \tag{13}$$

The slope coefficient estimate, $\hat{b}$, from this regression equals -0.057, and it is statistically significant with a t-value of -33.1.[6] It is important to emphasize that this regression is not a predictive regression. We estimate this regression to measure the degree to which average returns in one month are related to the sum of the average returns in the other months.

The negative slope coefficient indicates that a stock that earned, on average, high returns in one month earned, on average, lower returns in the other months. Because we demean the data, this negative correlation is not due to the seasonal patterns in market-wide returns (Kamstra, Kramer, and Levi 2003). The fact that $\hat{b}$ is statistically significantly negative—with a t-value of -33.1!—alone indicates that expected returns exhibit at least *some* amount of seasonal reversal.

How closely do seasonalities in expected returns satisfy the adding-up constraint? The estimate of -0.057 is substantially higher than -1, but this estimate is biased towards zero because of an errors-in-variables problem. The explanatory variable $\sum_{m' \neq m} \bar{r}_{i,m'}$ is noisy, so the -0.057 estimate only indicates that the adding-up constraint holds in the data to some extent; it does not quantify the degree to which this constraint holds. To get a sense of how noisy signals realized returns are of expected returns, consider the Fama-MacBeth slope coefficients from annual lags. The estimates in Figure 3 are, on average, just below 0.01. This estimate indicates, by equation (8), that just under 1% of the cross-sectional variance of stock returns emanates from differences in expected returns. The downward bias in b in regression equation (13) is therefore substantial.

To assess the magnitude of the bias, we use the same simulated data that underlie Figure 3 and estimate the regression in equation (13). The slope coefficient from this regression is -0.058 with a t-value of -46.6. That is, the negative slope estimate of -0.057 in the data is consistent with a model in which the seasonalities in expected returns perfectly cancel out.

[6] We compute the standard error by block bootstrapping the data by calendar month. We draw calendar months in blocks with replacement, recompute stocks' average returns, and repeatedly re-estimate the regression in equation (13). This bootstrapping procedure uses only time-series variation in returns to quantify the amount of estimation uncertainty about b in regression equation (13).

Table 2: Average annual and non-annual returns in Fama-MacBeth regressions: Alternative formation periods

This table presents average Fama and MacBeth (1973) regression slopes and their t-values from cross-sectional regressions that predict monthly returns. The regressions predict returns using the average of all past returns, average of same-month returns, average of other-month returns, and the average difference between the same- and other-month returns. Each of these specifications is estimated as a separate univariate regression. Row "Year 1" uses data from $t-12$ through $t-1$; row "Year 2–5" uses returns from $t-60$ through $t-13$; and so forth. The regressions use data from January 1963 through December 2016 for all stocks, all-but-microcaps, and 48 value-weighted Fama-French industries. Microcaps are stocks with market values of equity below the 20th percentile of the NYSE market capitalization distribution. We cross-sectionally demean the data before computing the averages of past returns.

	Construction of historical average return							
	All		Same-month return		Other-month return		Same-month − Other-month	
Years	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$
				All stocks				
1	2.10	1.14	1.70	6.88	0.20	0.11	1.46	5.65
2–5	−13.63	−4.17	3.07	5.83	−17.45	−5.73	3.91	7.83
6–10	−7.58	−2.78	4.19	7.23	−12.27	−5.02	4.60	8.64
11–15	−0.06	−0.02	4.41	6.93	−4.97	−2.06	4.27	7.10
16–20	−3.94	−1.57	3.73	5.30	−8.36	−3.35	3.79	5.71
				All-but-Microcaps				
1	7.59	3.26	1.65	4.47	6.25	2.76	1.01	2.74
2–5	−10.93	−3.21	2.67	4.31	−14.18	−4.36	3.63	6.16
6–10	−5.21	−1.68	3.97	6.24	−9.07	−3.19	4.35	7.40
11–15	−0.89	−0.35	3.40	5.20	−4.00	−1.63	3.40	5.42
16–20	−3.66	−1.44	3.31	4.45	−7.05	−2.84	3.50	4.97
				Industries				
1	24.93	5.54	4.94	4.47	22.54	5.22	2.06	1.89
2–5	−2.20	−0.33	1.10	0.55	−3.58	−0.55	2.93	1.59
6–10	−19.04	−2.67	6.83	3.37	−25.13	−3.64	8.52	4.30
11–15	−6.05	−0.98	6.04	3.21	−9.68	−1.64	6.91	3.70
16–20	−14.14	−2.21	5.01	2.30	−16.13	−2.68	6.00	2.94

4.4 Comparing same-month and other-month regressions

We measure seasonalities and seasonal reversals in Table 2 by comparing slope coefficients from regressions that predict the cross section of monthly returns with all, same-month, and other-month returns. For example, when we predict returns this month with average returns from five years prior to two years prior (line 2–5 in the table), the "All" regression uses the average of all returns from month $t-60$ to $t-13$; the "Same-month" regression uses the average of returns in months $t-60$, $t-48$, $t-36$, and $t-24$; and the "Other-month" regression uses the average of all other returns.

Differences between "all," "same-month," and "other-month" coefficients measure the extent to which seasonalities reverse. To see why, assume, as in equation (4), that stock returns take the form of

$$r_{it} = \mu_{i,m(t)} + \epsilon_{it}.$$

The regression coefficient that explains the cross section of monthly returns with lagged same-month returns is then $\hat{b}_{\text{same-month}} = \frac{\text{cov}(r_{i,t}, r_{i,t-k})}{\text{var}(r_{i,t-k})} = \frac{\sigma_\mu^2}{\sigma_\mu^2+\sigma_\epsilon^2}$ for lags k that are multiples of 12. If seasonalities in expected returns do not reverse, the coefficients from the "all" and "same-month" regressions are equal. If there are seasonal reversals, the value of the "same-month" coefficient exceeds that of the "all" coefficient.

If seasonalities in expected returns do not reverse—that is, if a stock's expected return in one month is unrelated to that in the other months, $\text{cov}(\mu_{i,m}, \mu_{i,m'}) = 0$ for all $m \neq m'$—then a stock's average return over a year contains the same information as the lagged same-month return. Moreover, averaging leaves the signal-to-noise ratio unchanged. Suppose, for example, that we estimate a cross-sectional regression of month t returns against the average return from month $t-24$ to $t-13$. Assuming the serial independence of both $\mu_{i,m(t)}$ and ϵ_{it}, the regression slope is

$$\hat{b}_{\text{all}}^{\text{no reversals}} = \frac{\text{cov}(r_{i,t}, \frac{1}{12}\sum_{k=13}^{24} r_{i,t-k})}{\text{var}\left(\frac{1}{12}\sum_{k=13}^{24} r_{i,t-k}\right)} = \frac{\frac{1}{12}\sigma_\mu^2}{\frac{1}{(12)^2}\left(12\sigma_\mu^2 + 12\sigma_\epsilon^2\right)} = \frac{\sigma_\mu^2}{\sigma_\mu^2 + \sigma_\epsilon^2} = \hat{b}_{\text{same-month}}.$$

If seasonalities in expected returns completely reverse, then the coefficient from the "all" regression will be zero. In the example above, the covariance $\text{cov}(r_{i,t}, \frac{1}{12}\sum_{k=13}^{24} r_{i,t-k})$ decomposes into two parts: the covariance with return in month $t-24$ is $\frac{1}{12}\sigma_\mu^2$, but, by the adding-up constraint, the covariance with the returns in the other months is $-\frac{1}{12}\sigma_\mu^2$. That is, if $\mu_{i,1} + \cdots + \mu_{i,12} = 0$, the average return over a year does not contain any information about expected returns. If seasonalities reverse at least partially, then $\hat{b}_{\text{all}}^{\text{reversals}} < \hat{b}_{\text{same-month}}$. We can measure seasonalities and seasonal reversals by comparing the "all" regression coefficient to the same-month coefficient or, alternatively, we can make comparisons between the same-month and other-month coefficients.

In Table 2 we measure the strength of seasonal reversals by comparing the same- and other-month regression coefficients. If seasonalities in stock returns reverse, then the coefficient from the same-month regression will be positive (and significant), that from the other-month regression will be negative (and significant), and that from the "all" regression coefficient will be zero (and insignificant). The issue from the statistical testing perspective is the inconvenient prediction that the all-coefficient is zero: an estimate that is *not* statistically significantly different from zero cannot be construed as evidence for *accepting* the null hypothesis. A comparison of the same- and other-month regression coefficients circumvents this issue. In the presence of seasonalities and seasonal reversals, these two coefficients are predicted to differ from zero with opposite signs.

Table 2 reports these coefficients from regressions that predict the cross sections of monthly stock and industry returns. We explain returns with average returns in years 1, 2–5, 6–10, 11–15, and 16–20. We use alternative lags because of long-term reversals. The negative "all" coefficient of -13.6 (t-value $= -4.17$) for years 2 through 5 in the full sample is consistent with these long-term reversals. However, when we regress today's stock returns against the five-year average return from year 11 to 15, the average "all" coefficient is close to zero. Therefore, by skipping ten years, we appear to skip over most—or all—of long-term reversals. Over the same 11 to 15-year period, the average same-month return coefficient is significant with t-values of 6.93; the other-month return is significant with a t-value of -2.06, and the difference between the two has a t-value of 7.10. Stepping back even further in time, the t-values associated with all returns, same-month returns, and other-month returns are -1.57, 5.30, and -3.35 when we skip 15 years before we begin measuring average returns.

These estimates suggest that the seasonalities in individual stock returns reverse completely. First, the significantly positive same-month coefficient, as before, indicates that there are seasonalities in expected stock returns. Second, the fact that the average same-month coefficient exceeds the annual coefficient indicates that some of these seasonalities reverse. Third, the finding that the annual coefficient is close to zero is consistent with the perfect reversal of the seasonalities in individual stock returns.

The estimates for the value-weighted industry portfolios suggest that the seasonalities in expected industry returns reverse perfectly as well. In the regression with the 11 to 15 years formation period, for example, the t-values associated with the same-month and other-month returns are 3.21 and -1.64, and the difference between the two has a t-value of 3.70. Seasonal reversals in industry portfolios relate to reversals in individual stock returns. Keloharju, Linnainmaa, and Nyberg (2016) show that a substantial part of seasonalities in individual stock returns stems from seasonalities in industry returns. At the same time, expected returns do not appear to vary significantly across industries (Moskowitz and Grinblatt 1999). The existence of seasonal reversals reconciles these two sets of findings.

If seasonalities reverse, both the same- and other-month average returns predict the cross section of stock returns through the same mechanism: a high average December return, for example, predicts high December returns, but so must a low average *non*-December return. The regressions in the last column of Table 2 predict returns using the difference between the same- and other-month average returns. If both the same- and other-month average returns are noisy versions of the same economic signal—the seasonal return component—then this combination should better predict returns than either of the two proxies in isolation. Consistent with this prediction, outside the one-year momentum effect, the t-values associated with the estimates in the last column are always higher than those associated with the estimates in the other columns.

5 Seasonality, seasonal reversal, and long-term reversal factors

5.1 Average monthly factor returns and correlations

The Fama-MacBeth regressions of Table 1 suggest that average same- and other-calendar-month returns are informative about future returns. We measure the usefulness of these signals from the investment perspective by constructing HML-like factors that select stocks by their average past returns. We construct a seasonality factor (ANN) by sorting stocks into six portfolios by size and the average same-calendar-month return using monthly rebalancing. The return on this factor is the difference between the value-weighted returns on the two high-average portfolios and the two low-average portfolios. Similarly, we construct a seasonal reversals factor (NANN) using the same methodology, except that we sort stocks by their average other-calendar-month returns. Because this is a reversal factor, we compute the return on the factor as the difference between the two low-average and the two high-average portfolios. Finally, we construct an annual-minus-non-annual factor by sorting stocks by the difference between the average same- and other-calendar-month returns.

Table 3 reports the monthly percent returns for these factors and their correlations. We also report, for comparison, the same statistics and correlations for the market, size, value, momentum, and long-term reversals factors. The long-term reversals factor is another HML-like factor that chooses stocks by their five-year-skip-a-year return (Fama and French 1996).

The average returns on the seasonality and seasonal reversal factors are economically large and statistically significant. The seasonality factor earns an average return of 61 basis points per month (t-value of 8.37); the seasonal reversal factor earns 45 basis points (t-value of 4.89); and the combination of the two—the annual-minus-non-annual factor—earns 67 basis points (t-value of 9.93).

The difference in the predictive powers of the seasonality and annual-minus-non-annual factors is large. The moderate difference in the levels of t-values (8.37 versus 9.93) is not the right comparison because these estimates are so far out in the tails of the distributions. In terms of likelihoods, a

Table 3: Monthly percent returns on factors and correlations

This table reports average monthly percent returns, standard deviations, and t-values for various factors (Panel A) and the monthly return correlations (Panel B). The first four factors are those of the Carhart (1997) four-factor model: market, size, value, and momentum. The other factors are HML-like factors that first sort stocks into six portfolios by market capitalization and the sorting variable. The long-term reversal factor (LTREV) sorts stocks by their five-year return skipping a year; the seasonality factor (ANN) sorts stocks by their average same-calendar-month return; the seasonal reversal factor (NANN) sorts stocks by their average other-calendar-month return; and the annual-minus-non-annual factor (AMN) sorts stocks by the difference between the average same-calendar-month and other-calendar-month return. The data are demeaned in each cross section before computing the average same- and other-calendar-month returns. Both averages use up to 20 years of historical data. The factor return data are from January 1963 through December 2016.

Panel A: Monthly percent returns

Factor	Name	Average return	Standard deviation	t-value
MKTRF	Market	0.52	4.41	3.00
SMB	Size	0.23	3.08	1.86
HML	Value	0.38	2.82	3.46
UMD	Momentum	0.66	4.21	4.02
LTREV	Long-term reversal	0.29	2.49	2.95
ANN	Seasonality	0.61	1.85	8.37
NANN	Seasonal reversal	0.45	2.33	4.89
AMN	Annual − Non-annual	0.67	1.71	9.93

Panel B: Monthly return correlations

Factor	MKTRF	SMB	HML	UMD	LTREV	ANN	NANN	AMN
MKTRF	1							
SMB	0.29	1						
HML	−0.26	−0.21	1					
UMD	−0.13	0.00	−0.19	1				
LTREV	−0.02	0.26	0.45	−0.07	1			
ANN	0.18	0.03	−0.23	−0.05	−0.13	1		
NANN	−0.51	−0.25	0.72	0.00	0.45	−0.15	1	
AMN	−0.02	−0.07	0.06	−0.03	0.03	0.89	0.20	1

move from a t-value of 8.37 to 9.93 is equivalent to a move from a t-value of 1.96 to 5.56; that is, the difference between an estimate being statistically significant at the 5% level versus there being overwhelming evidence against the null hypothesis.[7]

[7]This comparison is based on the following computation. The t-values of 8.37 and 9.93 correspond to p-values of 5.762×10^{-17} and 3.083×10^{-23}. The latter event is therefore less probable than the first event by a factor of 1.9 million. If we start from a p-value of 0.05 (t-value = 1.96), an event that is proportionally as improbable has a p-value of 0.000000027—or a t-value of 5.56.

The seasonality and seasonal reversal factors correlate differently with the market, value, and long-term reversals factors. Whereas the seasonality factor correlates positively with the market and negatively with both the value and long-term reversal factors, these correlations have the opposite signs for the seasonal reversal factor. The seasonal reversal factor's correlation with the market is -0.51; that with the value factor is 0.72; and that with the long-term reversal factor is 0.45. As a consequence of these offsetting correlations, the annual-minus-non-annual factor is nearly uncorrelated with these other factors, with correlations ranging from -0.07 to 0.06.

5.2 Incremental information

In Table 4 we examine the incremental information of the seasonality, seasonal reversal, and long-term reversal factors. In the first column of Panel A, for example, we regress the monthly returns on the seasonality factor on the returns on the market, size, and value factors. These spanning regressions assess the extent to which the left-hand side factor (here, the seasonality factor) contains information that is not present in the set of the right-hand side factors (here, the market, size, and value factors).

The alphas from these spanning regressions have two interpretations. The first interpretation pertains to the investment problem. A statistically significant alpha indicates that an investor, who currently trades the right-hand side factors, can increase his portfolio's Sharpe ratio significantly by also trading the left-hand side factor. The second interpretation relates to comparing different asset pricing models. A statistically significant alpha indicates that an asset pricing model that adds the left-hand side factor to the set of right-hand factors is statistically superior to a model that contains only the right-hand side factors (Barillas and Shanken 2017).

In the first column's regression, which explains seasonalities with the three-factor model, the alpha is 64 basis points per month with a t-value of 8.79. As suggested by Table 1's Fama-MacBeth regressions, the seasonality factor thus contains a substantial amount of information about expected returns. The second and third columns add to the right-hand side the long-term reversal and seasonal reversal factors, but the addition of these factors does not materially lower the alpha. In column (3)'s model with both

Table 4: Incremental information of the seasonality, seasonal reversal, and long-term reversal factors

This table reports estimates from spanning regressions. In Panel A, the left-hand side variable is the monthly return on the seasonality factor or the seasonal reversal factor. In Panel B, it is the return on the annual-minus-non-annual factor or the long-term reversal factor. The explanatory variables are the four factors of Carhart's (1997) four-factor model—market, size, value, and momentum—and, depending on the specification, the seasonality or the seasonal reversal factor. t-values of the alpha are reported in parentheses. The factor return data are from January 1963 through December 2016.

Panel A: Seasonality factor and seasonal reversal factor

	Dependent variable					
Explanatory	Seasonality factor			Seasonal reversal factor		
variables	(1)	(2)	(3)	(4)	(5)	(6)
			Monthly percent alpha			
$\hat{\alpha}$	0.64	0.67	0.62	0.35	0.30	0.25
	(8.79)	(8.10)	(7.32)	(6.17)	(5.33)	(3.91)
			Factor loadings			
MKTRF	0.06	0.05	0.08	−0.18	−0.17	−0.18
SMB	−0.03	−0.03	−0.01	−0.02	−0.08	−0.08
HML	−0.14	−0.14	−0.21	0.52	0.43	0.44
UMD		−0.03	−0.04		0.04	0.04
LTREV		−0.02	−0.05		0.23	0.23
ANN						0.09
NANN			0.15			
			Factor loadings, t-values			
MKTRF	2.34	2.08	2.94	−11.13	−11.61	−11.94
SMB	−0.99	−0.86	−0.44	−0.81	−4.10	−4.02
HML	−3.35	−3.08	−3.99	18.23	14.87	15.15
UMD		−0.90	−1.06		2.27	2.43
LTREV		−0.37	−1.18		7.94	8.07
ANN						2.62
NANN			2.51			

of these additional factors, the alpha is 62 basis points per month with a t-value of 7.32.

The other columns of Panel A show that the seasonal reversal factor also contains information that is not present in the Carhart (1997) four-factor model, the long-term reversal factor, and the seasonality factor. Seasonal reversal factor's monthly three-factor model alpha is 35 basis points (t-value = 6.17).

Panel B: Annual-minus-non-annual factor and long-term reversal factor

	Dependent variable			
	Annual-minus-non-annual factor		Long-term reversals factor	
Explanatory variables	(1)	(2)	(3)	(4)
		Monthly percent alpha		
$\hat{\alpha}$	0.66	0.67	0.04	−0.08
	(9.70)	(8.53)	(0.50)	(−0.96)
		Factor loadings		
MKTRF	0.01	0.00	0.01	0.10
SMB	−0.04	−0.05	0.29	0.30
HML	0.03	0.01	0.46	0.20
UMD		−0.01		0.00
LTREV		0.03		
ANN				−0.06
NANN				0.49
		Factor loadings, *t*-values		
MKTRF	0.29	0.21	0.22	3.54
SMB	−1.27	−1.58	7.18	8.10
HML	0.72	0.24	10.39	3.88
UMD		−0.21		−0.14
LTREV		0.76		
ANN				−1.19
NANN				7.74

This alpha falls only slightly to 30 basis points (t-value 5.33) when we also control for long-term reversals. Long-term reversals are therefore largely unrelated to seasonal reversals, and the seasonal reversal factor contains information that is not present in the long-term reversal factor. Regression (6) includes the seasonality factor on the right-hand side. The alpha is now 25 basis points per month with a t-value of 3.91. This estimate indicates that seasonalities and seasonal reversals contain independent information about the cross section of returns. This estimate is consistent with both variables providing noisy (but independent) estimates of the seasonal component of expected returns.

Panel B of Table 4 reports estimates from spanning regressions with either the annual-minus-non-annual factor or the long-term reversal factor as the dependent variable. The annual-minus-non-annual factor, which estimates expected returns from both the average same- and other-calendar-month returns, earns a monthly three-factor model alpha of 66 basis points (t-value of 9.70). Because this factor does not correlate substantially with momentum or long-term reversals, its alpha in a full model including these factors remains at 67 basis points (t-value of 8.53).

The long-term reversal factor, by contrast, does not contain information beyond that contained in the three-factor model. This factor's three-factor model alpha is 4 basis points (t-value = 0.50). The culprits for this loss of significance are the size and value factors. Long-term losers are, on average, small value stocks, so the long-term reversal factor loads positively on both the size and value factors (Fama and French 1996). The insignificant alpha indicates that an investor who already trades the market, size, and value factors would not benefit from trading also the long-term reversal factor. The finding that the seasonal reversal factor earns a statistically significant alpha in a regression that includes, among other things, size, value, and the long-term reversal factors is therefore important. The significance of this alpha indicates that the seasonal reversal factor is not just a different measure of long-term reversals, but a distinct effect.

Table 5: Maximum Sharpe Ratios

This table reports weights and ex-post maximum Sharpe ratios for strategies formed from market, value, size, momentum, seasonality, seasonal reversal, and long-term reversal factors. The seasonality, seasonal reversal, and long-term reversal factors are HML-like factors that first sort stocks into six portfolios by market capitalization and the sorting variable. The long-term reversal factor sorts stocks by their five-year return skipping a year; the seasonality factor sorts stocks by their average same-calendar-month return; the seasonal reversal factor sorts stocks by their average other-calendar-month return; and the annual-minus-non-annual factor sorts stocks by the difference between the average same-calendar-month and other-calendar-month return. The data are demeaned in each cross section before computing the average same- and other-calendar-month returns. Both averages use up to 20 years of historical data. The factor return data are from January 1963 through December 2016.

	Portfolio weights (%)								
						Seasonality factors			Sharpe
#	MKTRF	SMB	HML	UMD	LTREV	ANN	NANN	AMN	ratio
1	100								0.41
2	21	11	42	27					1.07
3	21	10	40	26	3				1.08
4	7	6	25	14	2	47			1.69
5	24	11	3	16	−11		57		1.34
6	11	7	10	11	−4	38	28		1.81
7	10	7	18	13	0			53	1.73

5.3 Maximum Sharpe ratios

The spanning regressions of Table 4 indicate that investors would have earned higher Sharpe ratios by trading the seasonalities in expected returns in addition to the other factors. In Table 5, we quantify these results by constructing ex-post maximum Sharpe ratio portfolios from various combinations of factors listed in Table 3. We use monthly factors returns from January 1963 through December 2016 to find the tangency portfolio for different combinations of factors and then compute and report the Sharpe ratio associated with this portfolio.

The first line indicates that an investor would have earned an annualized Sharpe ratio of 0.41 by investing in the market portfolio between January 1963 and December 2016. The second row shows that the size, value, and momentum factors were very valuable. The optimal combination of these factors would have earned a Sharpe ratio of 1.08, and the investor would have obtained this Sharpe ratio by investing two-thirds of his portfolio in the value and momentum factors.

The remaining rows of Table 5 add to this set of factors different combinations of the long-term reversal, seasonality, and seasonal reversal factors. Consistent with its insignificant alpha in Table 4's regressions, the addition of the long-term reversal factor does not change the maximum Sharpe ratio. Both the seasonality and seasonal reversal factors, however, substantially increase the Sharpe ratio. The seasonal reversal factor alone increases it from 1.08 to 1.34; the seasonality factor alone increases it to 1.69; and the addition of both of them increases it to 1.81. This last increase is consistent with both factors containing information about expected returns that is not present in the other. The row with both the seasonality and seasonal reversal factors shows that the optimal portfolio is now two-thirds invested in these two factors alone; the remaining one-third is approximately evenly divided between the market, size, value, and momentum factors. The last row shows that the maximum Sharpe ratio is nearly as high, 1.73, when we replace the two factors with the annual-minus-non-annual factor.

The shifts in portfolio weights in Table 5 that result from adding the seasonal factors are substantial. The shifts indicate that the seasonal effects—seasonalities and seasonal reversals—are very large in magnitude relative to the non-seasonal return predictors such as value and momentum.

5.4 Persistence

Seasonalities and seasonal reversals are persistent throughout the sample period. Figure 4 illustrates this persistence by reporting the t-values associated with the average factor returns for ten-year rolling windows. The first data points in the figure, for example, correspond to a ten-year window from January 1963 to December 1972. The seasonality, seasonal reversal, and annual-minus-non-annual factors earn average returns of 58 basis points (t-value = 4.30), 28 basis points (t-value = 1.55), and 66 basis points (t-value = 5.14) over this ten-year period, and it is these t-values that we report in Figure 4.

None of the factors averages a negative return over any of the ten-year periods. The t-value for the combination factor, which sorts stocks into portfolios by the difference between the average same- and other-calendar-month returns, is above 2.5 for every ten-year period between January 1963 and

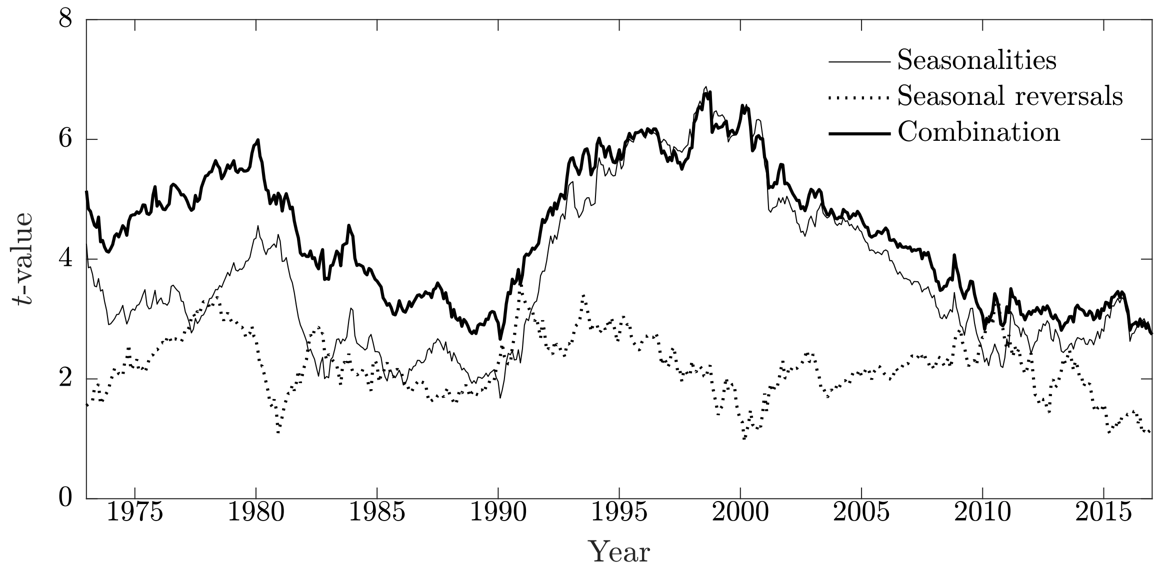

Figure 4: **Seasonality and seasonal reversal factors, 1963–2016.** This figure plots the t-values associated with the risk premiums on three factors. The seasonality factor sorts stocks by their average same-calendar-month return; the seasonal reversal factor sorts stocks by their average other-calendar-month return; and the annual-minus-non-annual factor sorts stocks by the difference between the average same-calendar-month and other-calendar-month return. We compute average factor returns using rolling ten-year windows, and report the t-values for these rolling windows. The 1975 points, for example, are the t-values of the factors for the 10-year period up to 1975.

December 2016.

The persistence of these factor premiums is important. Heston and Sadka (2008) document the seasonality effect in stock returns using data up December 2002. The last few years of Figure 4 are therefore entirely out of sample for the seasonality factor. This out-of-sample persistence of the effect alone suggests that return seasonalities and seasonal reversals are most likely not an artifact of data-mining (McLean and Pontiff 2016; Linnainmaa and Roberts 2018). At the same time, this degree of persistence is puzzling, casting doubt on the risk-based explanations for the cross-sectional variation in expected returns.

If the seasonality and seasonal reversal factors earn premiums to compensate for some combination of risks, these risks appear not to have materialized over the past 50 years. An alternative explanation, which would reconcile both the persistence of the seasonalities and seasonal reversals, and the fact that seasonal reversals seem to balance out seasonalities, is that seasonalities are due not to risk-

Table 6: Daily Fama-MacBeth regressions

This table presents average Fama and MacBeth (1973) regression slopes and their t-values from cross-sectional regressions that predict daily returns. The regressions use data from January 1963 through December 2016 for all stocks. The same-weekday average is, e.g., the average Monday return during the period reported in the years column when the regression predicts the cross section of Monday returns; in regressions that predict Tuesday returns, it is the average Tuesday return; and so forth. "Same − Other" in the last column is the average same-weekday return minus the average other-weekday return.

	All days		Same weekday		Other weekday		Same − Other	
Years	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$
1	12.73	9.75	7.39	22.05	3.02	3.02	4.99	19.57
2–5	−3.31	−1.65	10.08	16.23	−15.28	−9.60	9.33	18.31
6–10	0.54	0.26	9.91	13.55	−13.11	−7.47	8.77	13.91
11–15	3.91	1.73	8.70	10.40	−9.25	−4.79	7.47	10.30
16–20	−1.91	−0.72	8.16	8.93	−13.42	−6.05	7.48	9.46

based mechanisms but to temporary mispricing. Seasonalities may be induced by investors trading consistently in the same direction in the same periods. These effects can persist in the data if the resulting seasonalities are not large enough to be exploited as stand-alone anomalies because of their high turnover. At the same time, even if the associated round-trip transaction costs are prohibitively large, investors could benefit from these seasonalities by using them to time their trades (Heston, Korajczyk, and Sadka 2010; Novy-Marx and Velikov 2016).

6 Other seasonalities

6.1 Daily data

Keloharju, Linnainmaa, and Nyberg (2016) show that stocks also have strong seasonalities at the day-of-the-week level. For example, a stock that has historically done well on Mondays has higher expected returns on Mondays. We now show that, similar to monthly seasonalities, also daily seasonalities are balanced by seasonal reversals.

Table 6 reports estimates from Fama-MacBeth regressions that predict the cross section of daily returns. We predict returns using the average return computed over all days, the same weekdays, the other weekdays, and the difference between the two. We estimate these averages using historical data in years 1, 2 through 5, 6 through 10, 11 through 15, and 16 through 20.

The estimates in the all-column are similar to the regressions reported in Table 2 except that the dependent variable is a daily instead of a monthly return. Consistent with Keloharju, Linnainmaa, and Nyberg (2016), daily returns are highly seasonal: same-weekday returns have significantly positive coefficients and the other-weekday returns have significantly negative coefficients beyond the first year. Because these two effects cancel each other out—that is, seasonal reversals balance out seasonalities—the estimates in the first column are close to zero beyond the first year.

In the context of daily returns, the idea that seasonalities emanate from temporary mispricing appears the most plausible. Chan, Leung, and Wang (2004), for example, find that stocks with high institutional holdings earn higher average returns on Mondays than those with low institutional holdings. These higher-than-average Monday returns may stem from institutions' excess demand on Mondays, a demand that is not absorbed by the rest of the market without a price impact. If so, we would expect non-Monday returns to be lower by an offsetting amount; Monday returns are higher only because excess institutional demand temporarily inflates share prices.

6.2 Countries, bonds, and commodities

Heston and Sadka (2010) show that return seasonalities exist not only in the cross section of U.S. stock returns (Heston and Sadka 2008) but also within international stock markets. Keloharju, Linnainmaa, and Nyberg (2016) show that seasonalities are not confined to equity returns; they document seasonalities also in stock market indexes and commodity returns. Using a time series of over 200 years, Baltussen, Swinkels, and van Vliet (2018) find evidence of seasonalities also in government bond indexes and foreign exchange. If seasonalities in different markets and asset classes emanate from the same mechanism, we would expect to find seasonal reversals in other markets and asset classes as well. Following the studies listed above, we measure seasonal *reversals* in the cross sections of international

Table 7: Seasonalities and seasonal reversals in international stock returns, country-level stock and bond indexes, and commodities

This table reports estimates from cross-sectional regressions that predict the cross section of international stock returns (Panel A) and average returns for strategies that trade seasonalities in country equity indexes, country government bond indexes, and commodities (Panel B). International stock market data in Panel A are from Datastream and include all developed countries excluding the United States. We remove the 20 percent of the smallest firms from each country each month. Stocks from country c are included in the sample in month t if the month-t cross section in country c has at least 50 stocks. We follow the methodology in Table 2 and estimate Fama and MacBeth (1973) regressions that predict the cross section of international stock returns using prior stock returns over different horizons. The regressions include country fixed effects. We compute average stock returns using all past returns, same-month returns, other-month returns, and the difference between the same- and other-month returns. In Panel B we construct portfolios from country-level stock and bond index and commodity portfolios each month based on average same- or other-month returns, or the difference between the same- and other-month averages. The number of commodities in the high and low portfolios is two until December 1993 and three after this point as the number of commodities with sufficient historical return data increases to 15. The country index strategy has three countries in the long and short legs throughout the sample. In Panel A, the first cross-sectional regression has January 1987 returns on the left-hand side; the last regression uses December 2016 return data. In Panel B we begin forming commodity and country index portfolios at the end of December 1974 and end the sample in December 2011. We report t-values in parentheses.

Panel A: International stock returns

	Construction of historical average return							
	All		Same-month return		Other-month return		Same-month − Other-month	
Years	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$	$\hat{b}$	$t(\hat{b})$
1	3.39	3.51	1.16	5.08	2.96	2.89	0.59	3.17
2–5	−7.90	−5.58	1.39	4.64	−9.02	−6.43	2.09	7.39
6–10	−1.67	−1.32	2.76	7.64	−4.08	−3.51	2.87	8.91
11–15	−4.24	−2.44	2.25	5.03	−7.22	−4.16	2.49	5.77
16–20	−0.11	−0.05	2.44	4.86	−3.42	−1.72	2.32	4.70

Panel B: Country-level stock and bond indexes and commodities

	Sort by:		
Asset	Same-month return	Other-month return	Same − Other
Countries: Equity	0.45	−0.29	0.50
	(2.29)	(−1.45)	(2.84)
Countries: Government bonds	0.11	−0.13	0.19
	(1.04)	(−0.92)	(1.88)
Commodities	0.90	−0.19	0.92
	(2.19)	(−0.48)	(2.31)

stock returns, stock market indexes, government bond indexes, and commodities.[8]

Panel A of Table 7 reports estimates from cross-sectional regressions that predict the cross section of international stock returns. These regressions are the same as those reported in Table 2 for the U.S. stock market. We include in the sample stocks from all developed countries except the United States. We remove the 20 percent of the smallest firms from each country each month. Stocks from country c are included in the sample in month t if the month-t cross section in country c has at least 50 stocks. The left-hand side return data in the cross-sectional regressions start in January 1987 and end in December 2016. The international results in Panel B are similar to the U.S. results. The same-month regressions show that the same-month stock return significantly predicts future stock returns up to 20 years into the future; this seasonality result is consistent with the results in Heston and Sadka (2010). These seasonalities, however, are again offset by seasonal reversals: the other-month regressions are negative and statistically significant at the 5% level from years 2 to 15, and at the 10% in the 16-to-20 year specification. Because of these reversals, the average of "all" prior returns in the first column either do not predict the cross section of stock returns (years 6–10 and 16–20) or significantly *negatively* predict these returns.

In Panel B of Table 7 we form portfolios using same- and other-month returns using three sets of assets: country equity indexes, country government bond indexes, and commodities. The country-level analyses use data from 15 countries: Austria, Belgium, Canada, Denmark, France, Germany, Italy, Japan, Netherlands, Norway, Spain, Sweden, Switzerland, UK, and the U.S. The government bond data are from Bloomberg; we use the total return indexes on the 10-year government bonds. The commodity return data are returns on futures contracts on Aluminium, Copper, Nickel, Lead, Zinc, Brent, Gas Oil, Crude Oil, Gasoline, Heating Oil, Natural Gas, Cotton, Coffee, Cocoa, Sugar, Soybean, Kansas Wheat, Corn, Wheat, Lean Hogs, Feeder Cattle, Live Cattle, Gold, and Silver.[9]

[8] We do not study foreign exchange as Baltussen, Swinkels, and van Vliet (2018) find that a seasonal strategy has generated statistically insignificant returns in this asset class during the past 50 years.

[9] We thank Ralph Koijen, Toby Moskowitz, Lasse Pedersen, and Evert Vrugt for providing these data. See Koijen et al. (2013) for a description of these data.

Seasonalities are balanced out by seasonal reversals in each of these asset classes. Seasonalities and seasonal reversals are the weakest for bonds but, even there, the sort based on the difference between the average same- and other-month returns generates a return spread of 19 basis points per month (t-value = 1.88). In the equity indexes and commodities, the long-short portfolios based on the same-minus-other difference have t-values of 2.84 and 2.31.

The results on seasonalities and seasonal reversals in monthly U.S. equity returns could be dismissed as a chance finding. Despite the high t-values, ones that well exceed the 3.0 bound suggested by, for example, Harvey, Liu, and Zhu (2016), seasonal reversals might offset seasonalities in one asset class (U.S. equities) and at one frequency (monthly) just by luck. The results in this section suggest that seasonalities are almost certainly intertwined without seasonal reversals. It seems that wherever we see return seasonalities, we also see offsetting seasonal reversals.

7 Conclusions

While seasonalities in stock returns last for up to 20 years (Heston and Sadka 2008), the differences in unconditional expected returns vanish within five years after portfolio formation (Keloharju, Linnainmaa, and Nyberg 2018). These two findings represent a puzzle: how can we reliably predict how well a stock will perform in a particular month 20 years later, but we cannot say with any degree of confidence which of two stocks has a higher expected returns after just five years?

These two findings can be reconciled only if return seasonalities are balanced out by seasonal reversals. If an asset's expected return exceeds that of the other assets in March by 5%, its total return in the remaining months—that is, the combined January, February, April, ..., December return—must be below that of the other assets by 5%. This adding-up constraint must hold for return seasonalities not to leave a trace in long-term expected returns.

A thought experiment best illustrates how monthly returns can be predictable up to 20 years without there being persistent differences in expected returns. Suppose we have identified a stock that earns reliably higher returns in December; we can therefore predict today that this stock's return in December in 20 years will be high as well. However, from the viewpoint of investing—or, equivalently, discounting cash flows—seasonal reversals render this predictability inconsequential. A stock's cumulative return over a calendar year equals $(1 + \mathrm{E}(\tilde{r}_{i,\mathrm{jan}})) \times (1 + \mathrm{E}(\tilde{r}_{i,\mathrm{feb}})) \times \cdots \times (1 + \mathrm{E}(\tilde{r}_{i,\mathrm{dec}})) - 1 \approx \mathrm{E}(\tilde{r}_{i,\mathrm{jan}}) + \mathrm{E}(\tilde{r}_{i,\mathrm{feb}}) + \cdots + \mathrm{E}(\tilde{r}_{i,\mathrm{dec}})$. We cannot earn the December return 20 years later without also earning the January-through-November returns 20 years later; and the point of the adding-up constraint is that these *other* returns perfectly offset the high December return.

We show that seasonal reversals indeed balance out seasonalities. These reversals are economically large: a factor that estimates stocks' expected returns from both same-month and other-month returns earns an average return of 67 basis points per month, which is significant with a t-value of 9.93. This factor contains information other factors do not include. In a spanning regression against the market, size, value, momentum, and long-term reversal factors, the monthly alpha remains at 67 basis points, now with a t-value of 8.53.

Our results illustrate how profoundly the data are inconsistent with many other assumptions about the nature of the cross-sectional variation in expected returns. One a priori plausible assumption about cross-sectional differences in expected returns would be that some of these differences are persistent and non-seasonal. Under this assumption, stock returns would be expected to positively correlate with all past returns. This is not what we observe in the data: stock returns positively predict the cross section of stock returns only at annual lags; at non-annual lags, the predictive relationship is negative because of long-term reversals and seasonal reversals.

In addition to monthly and daily U.S. stock returns, we find seasonal reversals also in international stocks, country equity and government bond indexes, as well as commodity returns. The fact that return seasonalities "add up to zero" over the calendar year suggests that seasonalities may be due to temporary mispricing rather than seasonal variation in either the quantity or price of risk. The risk-based explanations do not call for an adding-up constraint: a risk factor may be more "important" in one month than other months, but there is no reason for these variations in risk premiums to add up to zero; yet, in both equities and in the other asset classes they appear to do so. A mispricing story, where an asset's price is temporary pushed above or below its fundamental value by traders who consistently trade assets in the same direction in the same time periods, calls for such a constraint: seasonal reversals occur when the temporary price deviations subside. If so, models and analyses of intermediation and institutional trading, such as those in Lou (2012), Vayanos and Woolley (2013), and Lou, Polk, and Skouras (2016), appear to hold the key to understanding the source of seasonalities—and seasonal reversals—in asset returns.

REFERENCES

Asness, C., T. J. Moskowitz, and L. H. Pedersen (2013). Value and momentum everywhere. *Journal of Finance 68*(3), 929–985.

Baltussen, G., L. Swinkels, and P. van Vliet (2018). P-hacking: Trend, momentum, value, carry, seasonality and BAB tested since 1800. Working paper, Erasmus School of Economics.

Barillas, F. and J. Shanken (2017). Which alpha? *Review of Financial Studies 30*(4), 1316–1338.

Berk, J. B., R. C. Green, and V. Naik (1999). Optimal investment, growth options, and security returns. *Journal of Finance 54*(5), 1553–1607.

Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance 52*(1), 57–82.

Chan, S. H., W.-K. Leung, and K. Wang (2004). The impact of institutional investors on the Monday seasonal. *Journal of Business 77*(4), 967–986.

Cohen, R. B. and C. Polk (1996). An investigation of the impact of industry factors in asset-pricing tests. Working Paper, London School of Economics.

Conrad, J. and G. Kaul (1998). An anatomy of trading strategies. *Review of Financial Studies 11*, 489–519.

Davis, J. L., E. F. Fama, and K. R. French (2000). Characteristics, covariances, and average returns: 1929 to 1997. *Journal of Finance 55*(1), 389–406.

De Bondt, W. F. M. and R. H. Thaler (1985). Does the stock market overreact? *Journal of Finance 40*(3), 379–395.

Fama, E. F. and K. R. French (1996). Multifactor explanations of asset pricing anomalies. *Journal of Finance 51*(1), 55–84.

Fama, E. F. and J. D. MacBeth (1973). Risk, return and equilibrium: Empirical tests. *Journal of Political Economy 71*(3), 607–636.

Harvey, C. R., Y. Liu, and H. Zhu (2016). ...And the cross-section of expected returns. *Review of Financial Studies 29*(1), 5–68.

Heston, S. L., R. A. Korajczyk, and R. Sadka (2010). Intraday patterns in the cross-section of stock returns. *Journal of Finance 65*(4), 1369–1407.

Heston, S. L. and R. Sadka (2008). Seasonality in the cross-section of stock returns. *Journal of Financial Economics 87*(2), 418–445.

Heston, S. L. and R. Sadka (2010). Seasonality in the cross-section of stock returns: The international evidence. *Journal of Financial and Quantitative Analysis 45*(5), 1133–1160.

Kamstra, M. J., L. A. Kramer, and M. D. Levi (2003). Winter blues: A SAD stock market cycle. *American Economic Review 93*(1), 324–343.

Keloharju, M., J. T. Linnainmaa, and P. Nyberg (2016). Return seasonalities. *Journal of Finance 71*(4), 1557–1590.

Keloharju, M., J. T. Linnainmaa, and P. Nyberg (2018). Long-term discount rates do not vary across firms. University of Southern California working paper.

Koijen, R. S. J., T. J. Moskowitz, L. H. Pedersen, and E. B. Vrugt (2013). Carry. NBER Working Paper No. 19325.

Linnainmaa, J. T. and M. R. Roberts (2018). The history of the cross section of stock returns. *Review of Financial Studies 31*(7), 2606–2649.

Lo, A. W. and A. C. MacKinlay (1990). When are contrarian profits due to stock market overreaction? *Review of Financial Studies 3*(2), 175–205.

Lou, D. (2012). A flow-based explanation for return predictability. *Review of Financial Studies 25*(12), 3457–3489.

Lou, D., C. Polk, and S. Skouras (2016). A tug of war: Overnight versus intraday expected returns. London School of Economics working paper.

McLean, R. D. and J. Pontiff (2016). Does academic research destroy stock return predictability? *Journal of Finance 71*(1), 5–32.

Moskowitz, T. J. and M. Grinblatt (1999). Do industries explain momentum? *Journal of Finance 54*(4), 1249–1290.

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics 108*(1), 1–28.

Novy-Marx, R. and M. Velikov (2016). A taxonomy of anomalies and their trading costs. *Review of Financial Studies 29*(1), 104–147.

Shumway, T. (1997). The delisting bias in CRSP data. *Journal of Finance 52*(1), 327–340.

Vayanos, D. and P. Woolley (2013). An institutional theory of momentum and reversal. *Review of Financial Studies 26*(5), 1087–1145.