NBER WORKING PAPER SERIES

EVIDENCE ON RETRIEVED CONTEXT: HOW HISTORY MATTERS

William N. Goetzmann
Akiko Watanabe
Masahiro Watanabe

Working Paper 29849
http://www.nber.org/papers/w29849

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2022

We thank Jessica Wachter, the memory/beliefs/choice reading group at the University of Pennsylvania, and seminar participants at the University of Amsterdam and the University of Toronto for thoughtful discussions on this research. We thank Milad Nozari and Yi Wang of the International Center for Finance at the Yale School of Management for research support. All errors are the responsibility of the authors. This paper is preliminary - the results may change in later drafts. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Evidence on Retrieved Context: How History Matters
William N. Goetzmann, Akiko Watanabe, and Masahiro Watanabe
NBER Working Paper No. 29849
March 2022
JEL No. G02,G17

ABSTRACT

This paper tests the retrieved context model of Wachter and Kahana (2019) using a long-term panel of economic forecasts by participants in the Livingston Survey. Events in historical time contribute additional explanatory power to a relative time series model. Historical precedents for current macroeconomic conditions appear to be more relevant for extreme quantile forecasts. The results are consistent with the use of the retrieved context mechanism for formulating expectations about asset prices. They also suggest that historical events, not just lagged variables in relative time, matter in economic forecasting.

William N. Goetzmann
School of Management
Yale University
Box 208200
New Haven, CT 06520-8200
and NBER
william.goetzmann@yale.edu

Akiko Watanabe
University of Alberta School of Business
fujimoto@ualberta.ca

Masahiro Watanabe
School of Business - Finance
University of Alberta
Edmonton, Alberta T6G 2R6, Canada
masa@ualberta.ca

1. Introduction

History is a natural reference for economic forecasting, particularly for unusual episodes such as financial crises, inflation shocks and pandemics. Since the 2008 financial crisis, scholars have increasingly sought to extract lessons and make predictions from past financial events. For example, Barro & Ursua (2008) Reinhart & Rogoff (2009), Berkman, Jacobsen and Lee (2011), Krishnamurthy & Muir (2017) and Muir (2017) use the multivariate context of international historical episodes in an event-time framework to test theories about the causes and consequences of financial crises. In these papers and related scholarship, history reveals something that standard macroeconomic time-series models may not. By focusing on events that are economically similar to the subject of study, they can refine predictions about the consequences of unusual or extreme events. In practice, forecasters may also give extra weight to specific historical episodes in predicting future outcomes. This paper tests a proposed mechanism for how recollection of past events may influence predictions about the future.

Events of importance, shocks that change people's lives suddenly and significantly, are almost by definition rare and memorable. Interest in them is also likely to arise when similar contextual conditions occur. For example, Google searches for the terms "stock market crash of 1929" and "stock market crash of 1987" jumped dramatically in the weeks around March, 2020 when the Dow Jones Industrial Average dropped by more than 30%. Similarly, Google search for "Spanish flu" spiked in 2005 (H5N1 outbreak), 2009 (H1N1 SARS outbreak) and March 2020 (SARS COVID 19 epidemic), with almost no interest in the intervening months since 2004 -- the start of the Google trend data. These unsurprising patterns suggest that large, current shocks motivate a search for analogous historical events. At least one well-known investor, Ray Dalio has explicitly cited the relevance of historical macroeconomic context for predicting asset returns. [2]

In fact, it would be surprising if macroeconomic forecasters did not give extra weight to salient historical episodes. In this paper we formulate an econometric framework to test this proposition. Our model is based on the retrieved context theory of Wachter & Kahana (2019), which makes specific predictions about the mechanisms by which the human mind accesses past events. The WK theory is based on experimental evidence that human recall depends on a multi-dimensional set of past conditions. This makes it particularly relevant to the dataset we use in this model: a panel of forecasters who are asked semi-annually to predict a vector of future macroeconomic variables. The WK model is also relevant because the authors draw asset pricing implications from the retrieved context mechanism. If forecasters use the retrieved context mechanism, then further tests of the effects on asset prices may be possible.

We test a more general version of the WK retrieved context model using the semi-annual Livingston Survey which has been conducted from 1948 to the present. We find evidence that professional forecasters participating in the survey placed significant weight on historical macroeconomic events when forecasting the S&P index. These weights and their significance are higher for forecasting tail events: the top and bottom quantiles of stock market growth. Our results

[2] Cf. Kara Chin, Jacqui Frank, and Ali Newhard, "Ray Dalio says the economy looks like 1937 and a downturn is coming in about two years", *Insider* September 18, 2018. https://www.businessinsider.com/ray-dalio-bridgewater-debt-crisis-downturn-coming-about-two-years-2018-9. Accessed 12/1/2021.

suggest that past stock market dynamics are not considered in isolation. Forecasters use a contextual framework that includes factors such as inflation and industrial production.

There is another growing body of evidence about how history matters, in particular individual lived experience. Malmendier and Nagel 2011 & 20166 and Malmendier, Nagel and Yan, 2020 show that agents rely on past personal history to form expectations about asset returns and inflation. This has been termed the experience effect. In related work (Goetzmann, Watanabe and Watanabe, 2021a) we find that the same panel of Livingston Survey economists made significantly different macroeconomic predictions depending upon the date they entered the sample – a rough proxy for the interval of their lived experience over the history of the US economy from the beginning of the 20th century. Those results suggest that forecasters who entered the panel at different dates may not agree on the most relevant historical precedents. If there were such disagreement, it would likely weaken our tests.

The paper is structured as follows. The next section presents our estimation methodology and how it related to the retrieved context theory of Wachter & Kahana (2019). Section 3 presents and discusses the results. Section 4 concludes.

2. Methodology

2.1. Modeling Retrieved Context

In this section we explain our approach to testing Wachter and Kahana (2019). As indicated above, memory is potentially relevant to the formation of expected asset returns. The WK model is developed from psychological models of the mechanism of human memory (Kahana 2012). However, this does not preclude the use of retrieved context to model the accessing of historical rather than personally experienced events. As noted above, Google searches around salient events initiated higher than normal searches for similar events in remote history. We leave unspecified the process by which forecasters are informed about or attend to historical context outside of personal memory.

WK apply the retrieved context literature on human memory to the problem of how the experience effect may operate, and its differential implications for asset prices and investor decision-making. In their model, moments of past personal experience are represented by a time-indexed vector of features which are manifestations of an underlying, economic context. WK propose that recalled inputs to the current decision-making process result from a few, well-established rules of human memory retrieval. These are: (1) *recency*, i.e., human subjects exhibit better memory for recent experiences, (2) *semantic similarity*, i.e., we remember experiences that are most similar in meaning to those we are currently experiencing and (3) *temporal contiguity*, i.e., we remember items that occurred contiguously in time to recently-recalled items.

Livingston forecasts are presumably derived from respondent's hypotheses about the economic context that generates feature vectors. In related work (Goetzmann, Watanabe and Watanabe, 2021a) we found that different entry cohorts of Livingston forecasters predicted significantly different co-movements of S&P index growth and inflation. WK investigate the implications of investor choice conditioned upon forecasts derived from retrieved-context memory. In WK, context will be influenced by recency, semantic similarity and temporal contiguity. For our purposes, semantic similarity is nearly moot because the variables are named.

We construct the feature vector, f, for each semester using the realized six-month growth rates on the S&P index, industrial production, and inflation. These three are selected because they have a history extending into the early 20th century, and thus overlap with the careers of the early Livingston participants. Suppose the economists completing a survey in semester t observe public information released at time $\mathrm{t} - \varepsilon < \mathrm{t}$. We assume that $t - \varepsilon$ is April for a June survey and October for a December survey. We then find the nearest historical match to the realization at $t - \varepsilon$ by choosing a prior semester $\tau < t - 2$ that minimizes the Euclidean distance between the current and past feature vectors,

$$\min_{\tau < t-2} \| f_{t-\varepsilon} - f_\tau \|. \tag{1}$$

We regress economist i's growth forecast of the S&P index over the 6-12 month horizon on context variables with controls over rolling windows:

$$\begin{aligned} \mathrm{fSP}_{i,t}^{6,12} = \mathrm{a} &+ \mathrm{b}_0^{\intercal} \Delta \mathrm{x}_\tau + \mathrm{b}_1^{\intercal} \Delta \mathrm{x}_{\tau+1} + \mathrm{b}_2^{\intercal} \Delta \mathrm{x}_{\tau+2} + \mathrm{c}_\varepsilon^{\intercal} \Delta \mathrm{x}_{t-\varepsilon} + \mathrm{c}_1^{\intercal} \Delta \mathrm{x}_{t-1} + \mathrm{c}_2^{\intercal} \Delta \mathrm{x}_{t-2} \\ &+ h_{1,\mathrm{S}1}^{\intercal} \mathrm{fx}_{\mathrm{i,t-1}}^{0,6} + h_{2,\mathrm{S}1}^{\intercal} \mathrm{fx}_{\mathrm{i,t-2}}^{0,6} + h_{1,\mathrm{S}2}^{\intercal} \mathrm{fx}_{\mathrm{i,t-1}}^{6,12} + h_{2,\mathrm{S}2}^{\intercal} \mathrm{fx}_{\mathrm{i,t-2}}^{6,12} + \mathrm{e}_{i,t}, \end{aligned} \tag{2}$$

where $\Delta x \equiv (\Delta SP, \quad \Delta IP, \quad \Delta CPI)^{\intercal}$ is the vector of growth realizations of the S&P index (*SP*), industrial production (*IP*), and the consumer price index (*CPI*) at the time in subscript; $\mathrm{fx}_{\mathrm{i,t-s}}^{0,6} \equiv (fSP06_{i,t-s} \quad fIP06_{i,t-s} \quad fCPI06_{i,t-s})^{\intercal}$ is the vector of economist i's past growth forecasts over the 0-6 month as of time $t - s$, $s = 1, 2$; $fx_{i,t-s}^{6,12}$ is the corresponding vector over the 6-12 month horizon; and $e_{i,t}$ is the residual. a is the intercept, and vectors b, c, and h are conforming coefficient vectors. Time subscript $t - \varepsilon$ stands for the six-month growth of the variables ending two months prior to survey date t, i.e., April for June surveys and October for December surveys. In particular, $\Delta \mathrm{x}_\tau$ is the growth realization at past time τ that solves the minimization problem in (1), where we take the feature vector f to be Δx.

The panel regression is designed to test whether various mechanisms of the retrieved context model explain individual economist's forecasts. For example, the b_0 coefficients on the $\Delta \mathrm{x}_\tau$ variables capture whether semantic similarity between current and past macroeconomic conditions is relevant to the 6-12 month forecast of S&P index growth. The b_1 and b_1 coefficients on $\Delta \mathrm{x}_{\tau+1}$ and $\Delta \mathrm{x}_{\tau+2}$ allow a test of temporal contiguity, i.e. whether conditions around the most semantically similar past historical context are relevant to the market forecast. The c coefficients on variables $\Delta \mathrm{x}_{t-\varepsilon}$, $\Delta \mathrm{x}_{t-1}$, and $\Delta \mathrm{x}_{t-2}$ capture the effects of recent macroeconomic trends on the forecast. If, for example, economists were using a VAR model to predict S&P growth, these c variables would be significant – this would also be consistent with recency, i.e. better memory for more recent history. Notice that c_ε conforms to the assumption that the information used to make time t forecasts includes data up to two months prior to t.

The set of h coefficients control for the economist's own set of recent macroeconomic forecasts. These might capture anchoring or reliance on prior judgement. We include both the near term and six-month out forecasts over the past two semesters prior to forecast. Since each vector is a triple of three variables, Equation (2) employs 30 explanatory variables and a constant. This corresponds to the following hypothetical logic:

> *"Recent macroeconomic trends look like those at a past time $\tau < t{-}2$. I will use the past realization of S&P growth at $\tau{+}2$ as an input to forecast S&P growth over period t+2."*

Why *t+2*? Why not *t+1*? This is a common question that arises in the use of Livingston Survey data. One might expect near-term predictions to have less noise. In fact, due to the way that Livingston forecasts are collected and calculated, the opposite is true. Expected growth rates are calculated as the log ratio of the predicted S&P index levels. For the first six-month period, this is the ratio of the number specified by the forecaster for a period six months from the end of the current period in which he/she answers the questionnaire, divided by the level of the S&P on the date the questionnaire is filled in. Because the stock market is volatile, and the specific date of completion is unknown, the growth rate implied by the ratio for the near-term growth prediction is observed with noise. Conversely, the ratio of the 12-month prediction to the 6-month prediction has no such noise. Hence, it has become common practice in Livingston research about stock predictions to use the 7-12 month growth forecast as a "cleaner" measure.

2.2. Estimation

We employ three regression methods: ordinary least squares (OLS), partial least squares (PLS), and regularized generalized linear models (RGLM). The last two methods have hyperparameters to be determined. For a given set of hyperparameters, we estimate Equation (2) by each method over expanding rolling windows of at least 80 semiannual periods (40 years) and make out-of-sample predictions in the following periods, explicitly accounting for individual fixed effects. For the PLS and RGLM, we further search over a grid of hyperparameters to optimize a fit metric using the entire set of rolling predictions.[3] Specifically, since the S&P index forecasts start in June 1952 and the other forecasts are already available, requiring two lags of forecasts sets the initial window as June 1953-December 1992. We first demean both the dependent and independent variables for each economist within this window to absorb cross-sectional fixed effects. Following common practice, for methods that compute variances for feature synthesis or distance for regularization purposes—the PLS and RGLM, respectively, in our case—we scale only the independent variables at unit standard deviation. This also allows comparison of the importance of variables by the magnitude of estimated coefficients, as such methods typically do not readily provide standard errors (often for a good reason). We then fit Equation (2) to estimate the coefficients. and predict the forecasts of economists present in June 1993 using the past realizations and lagged forecasts in that period. Finally, the fixed effect coefficients are added back to the model predictions of the economists who are present in both the in-sample and out-of-sample data only, eliminating those who drop in the latter data. The estimation window is increased by one period and the above process is repeated until the last window, June 1953-June 2020, is reached

[3] Because of the use of the entire sample, the optimal choice of hyperparameters is not out-of-sample. Alternatively, we could reserve a validation period after the first rolling window and stick to the hyperparameters determined from the validation period for the rest of the rolling windows to make predictions. Due to the relatively small size of the cross-sectional sample, we opted for using the entire sample to fix the hyperparameters.

and the last prediction is made for December 2020. For the PLS and RGLM, this is done over a grid of hyperparameters, and the optimal hyperparameters are chosen to minimize the root-mean-square error (RMSE) computed from model predictions and actual forecasts during June 1993-December 2020.[4] We report the result using the optimal hyperparameters.

3. Results

3.1. Historical Context Match

The black dots in Figure 1 plot time τ, the historical nearest-neighbor solution to the minimization problem in (1), against current time t. For comparison, the second and third nearest neighbors are shown in blue and red, respectively. While there appear to be horizontal clusters at some historical dates, it is unclear how they add up over individual forecasters. The histogram in Figure 2 shows that some dates tend to be the nearest neighbor date τ compared to others. The five longest bars correspond to June 1925, December 1968, December 1954, June 1947, and June 1968. None of these dates falls on recessions, which is not surprising as expansions have become much longer than recessions since the end of the Great Depression in late 1930s, and normal times would not resemble economic downturns. Four of the above five dates just precede recessions, while one succeeds a recession; according to the NBER, the U.S. economy entered recessions in October 1926, December 1969, and November 1948, and exited a recession in May 1954. Therefore, two semesters ahead ($\tau + 2$) of four nearest-neighbor dates are very close to or already in recessions, and that of December 1954 is well into expansion. If the feature vector at $\tau + 2$ explains individual economists' forecasts, it may be largely picking up those of pessimists at normal times. We will find some support for this conjecture in the analysis below.

3.2. Model Fit

Figure 3 plots the out-of-sample predictions of S&P index growth by Equation (2) against actual forecasts for the three estimation methods. If the model perfectly predicts the economists' forecasts, the dots will fall on the dashed 45-degree line. There are a few outliers that the model cannot predict with reasonable accuracy. The three methods generally produce similar predictions.

3.3. Variable Importance

Table 1 shows the average variable importance over rolling regressions in Equation (2), computed as either the average absolute t-statistic if produced by the OLS method. Otherwise the table reports the average relative magnitude of coefficients as percentages of the largest coefficient (PLS and RGLM). Panel A shows the importance of historical context variables matched by the nearest-neighbor problem in Equation (1), and Panels B and C the importance of recent realizations and lagged individual forecasts, respectively. Figure 4 plots the variable importance in descending order for each method. Not surprisingly, recent realizations tend to be most important in explaining

[4] Other popular fit metrics in the machine-learning literature include R-squared and mean absolute error.

the forecasted S&P growth. For example, the lagged industrial production (ΔIP_{t-1}) is the most important variable regardless of the estimation method, and current inflation in April/October ($\Delta CPI_{t-\varepsilon}$) is always among the top three. Also as expected, the most important forecast is the first lag of the dependent variable, forecasted S&P index growth over the 6-12 month horizon ($fSP612_{i,t-1}$).

In contrast, the importance of historical context variables varies across the methods, but always among the top three in this category is the S&P index growth at τ+2 ($\Delta SP_{\tau+2}$), i.e., in two semesters from the nearest neighbor match. Its average t-statistic from the OLS is 2.51 in Panel A of Table 1, and the average magnitude relative to the largest coefficient is 44.7% and 35.7% by the PLS and RGLM, respectively. Figure 5 plots the importance of $\Delta SP_{\tau+2}$ over time. In all the panels for the three estimation methods, the importance of $\Delta SP_{\tau+2}$ abruptly drops in the middle of the prediction period. It corresponds to June 2009, for which the estimation window first includes the culmination of the world financial crisis in late 2008. Before the crisis, the OLS t-statistic in Panel A hovers around 3 and then dips well below 2, which leads to the aforementioned average of 2.51.

To examine the impact of the financial crisis, we winsorize the two realized S&P index growth series, ΔSP_t and $\Delta SP_{t-\varepsilon}$, at the 1 and 99 percentiles over the rolling estimation period (June 1953 and later). This winsorizes the two lowest and two highest realizations at less extreme values, the lowest of which falls on the crisis. Winsorizing only the left tail or trimming only December 2008 gives a similar result. Table 2 and Figure 6 show the result. In Panel A of Figure 6, the OLS t-statistic now fluctuates around 2.7 after the financial crisis. The post-crisis importance by the PLS and RGLM in Panels B and C is also much higher than the previous figure. This results in the average importance of 2.94, 51.4%, and 39.8% for $\Delta SP_{\tau+2}$ by the three methods in Panel A of Table 2. This confirms the significant impact of the financial crisis. In the following sections, we will also examine the result using the winsorized or pre-crisis sample as warranted.

3.4. Panel Regressions

To mitigate the potentially detrimental effect of an outlier on inference, we focus on the pre-crisis period in this subsection. Table 3 combines the results of two estimation methods: OLS and quantile regressions using the sample ending in December 2007. We leave the treatment of the fixed effect in quantile regressions for future analysis. To allow comparison of the two methods, we report the result without demeaning and scaling the data. There are three panels corresponding to the three main groups of explanatory variables, (A) time τ, the historic nearest-neighbor variable, (B) time *t,* recent lagged macroeconomic realizations, and (C) *h,* economist's individual past forecasts.

3.4.1. Time τ variables

Panel A reports the coefficients and t-statistics for the time τ variables. If the forecaster matched recent trends to an historical context (time τ), and believed that the prior outcome could help predict the current future, we would expect a significant coefficient on $\Delta SP_{\tau+2}$ – the historic market growth aligning with the predicted time period *t*+2. We might also expect other τ period macroeconomic variables to have some significance, although these might be attenuated by the *t* period variables reported in Panel B. First notice that growth in industrial production in the historical period τ is significant. This is consistent with the significant positive covariation of industrial production growth and stock market growth, albeit there is a sign difference for the τ vs. $\tau+1$ realizations. We would expect the coefficient on the time τ variable to be more relevant because this is the inflation that likely matches the most recent realization of inflation prior to time t.[5]

Of relevance to the test of the retrieved context theory and the use of the past, is the positive and significant coefficient on $\Delta SP_{\tau+2}$. The OLS t-statistic of 2.82 corresponds to one of the points just before the dip apparent in Panel A of Figure 5. The quantile regression results estimate coefficients for dependent variable outcomes at the 0.1, 0.25, 0.5, 0.75, and 0.9 quantiles. Notice that the lowest quantile has the highest coefficient. This suggests that the relative importance of the historical context is greater for the most pessimistic forecasts. The second largest coefficient is for the top 10%—the most optimistic forecasts. This pattern suggests a non-linearity: for predictions of big moves in the S&P index, historical macroeconomic context may be more relevant. That said, the difference between either extreme and the median value is statistically marginal. Quantile coefficients on ΔIP_{τ} are also U-shaped, and mostly all significant and positive, consistent with a non-linear relation between S&P forecasts and industrial production at the nearest neighbor date τ. For ΔCPI_{τ}, the coefficient on the lowest quantile is positive and significant, likewise consistent with the hypothesis that extreme market forecasts may incorporate historical precedent more so than median forecasts.

3.4.2. Time *t* variables

Panel B reports the coefficients and t-statistics for t variables: $t–\varepsilon$, $t–1$, and $t–2$, which are current and lagged realized growth rates for the S&P, industrial production and inflation. Note that the $t–\varepsilon$ feature vector f_{t-} is minimally different from f_{τ} by construction. This could introduce multi-collinearity; attenuating or otherwise affecting the coefficient estimates of both sets. With this caveat, we find that that the $t–\varepsilon$ and the $t–1$ lagged value of inflation are both significantly positively related to forecasted S&P growth, and the $t–1$ lagged value for industrial production growth is significantly negatively correlated to forecasted S&P growth. This is difficult to square with prior findings of a strong relationship between forecasted production and stock market growth. Likewise surprising is that ΔS&P(t-2) is significantly negatively related to forecasted S&P growth.

It turns out that these are features of Livingston forecasts rather than multicollinearity. Figure 7 plots the rolling betas of the median forecasted 6-12M S&P growth with respect to lagged realized S&P index growth (ΔSP_{t-1}, Panel A), lagged realized industrial production growth (ΔIP_{t-1}, Panel B), and lagged CPI inflation (ΔCPI_{t-1}, Panel C). The dashed line shows the two standard error bands. The length of the rolling windows is fixed at 40 periods (20 years). We see that the forecasted S&P growth in 12 months loads negatively on lagged S&P and industrial production growth and positively on lagged inflation, although there are periods in which the relation becomes insignificant. Since these are based on simple regressions, multicollinearity is not an issue. Rather, it appears that Livingston economists may have been contrarians in making forecasts. It is also striking that the lagged inflation beta is almost one throughout the 1980's and 1990's and is even larger in earlier years, as if current inflation—as opposed to forecasted future inflation—is

[5] In an unreported robustness check we did not find this sign difference due to near multi-colinearity.

positively related to future stock market returns.[6] The near-unit beta is roughly divided between the coefficients on $\Delta CPI_{t-\varepsilon}$ and ΔCPI_{t-1} in Panel B of Table 3.

3.4.3. Economists' prior forecasts: *h* variables

Panel C reports the coefficients and t-statistics for the *h* variables. The intent of including these variables is to control for the potential desire of the forecaster to align her current forecasts with past predictions. The S&P forecasts are obviously the most relevant for the dependent variable of interest. In the OLS specification, we find this indeed to be the case. Coefficients on the two prior 6 month forecasts, and the t–2 forecast for the 6-12 month S&P forecast are all positive and significant. This is consistent with the persistence of conviction. Indeed, the contrast to the weak to negative relationship of market forecast to past realized S&P growth documented in panel B is particularly interesting. Prior commitment to a market forecast—or at least to a trend—seems stronger than recent market dynamics in influencing economists' predictions about the stock market. It is notable in Panel C that the market forecasts are practically the only significant variables. The only exception is the positive two period lag of the six-to-twelve-month industrial production value. While the findings in Panel C are not directly relevant to the test of the retrieved context theory, they suggest some potential for further investigation of forecaster preference for self-consistency.

3.4.4. Robustness: Winsorized Full Sample

Table 4 shows the result using the winsorized full sample. The coefficients on the historical context variables in Panel A are similar to those in Table 3. Although the magnitude and significance have weakened, the variable of interest, $\Delta SP_{\tau+2}$, remains significant in the OLS estimation (although close to the two-tailed 5% threshold) and exhibits a U shape across the quantiles. The behavior of the recency variables in Panel B and lagged forecasts in Panel C is also similar to Table 3. Finally, to control for career concerns and path-dependent beliefs, we included prior individual macroeconomic variable forecasts. Of these, with one exception, only the coefficients on prior forecasts of S&P growth are significant.

4. Conclusion

The Livingston survey allows us to test predictions of the Wachter and Kahana (2019) retrieved context mechanism for the formation of expectations about stock market returns. The way the test is formulated expands the mechanism to include "collective memory" or history that may not be in the personal life experience of the subject. We use a relatively clean dependent variable, individual forecasts of the S&P index growth over the six months to twelve months period, to eliminate likely errors in the variable. We approximate the contextual mechanism in WK as a matching of current to past feature vectors of macroeconomic variables. This has possible limitations – the closest match may not be the most relevant match, especially given our sole focus on a triplet of indexes, rather than a context that includes larger cultural, economic and historical

[6] In Goetzmann, Watanabe and Watanabe (2021a) we find a near zero beta for the 1980-2000 period on expected future inflation.

trends. We estimated an OLS panel regression that imposes linearity, however we also estimated quantile regressions to test whether historical context variables may be more salient when the market forecast is more extreme. We found evidence that economists used historical precedent as a model for predicting the stock market; the coefficient on the relevant measure of market growth following the best historical match to current macroeconomic conditions is significant. Estimates of other components of the retrieved context model, specifically recency, are also significant. However, some coefficient signs are inconsistent with intuition, suggesting further analysis is needed; it is possible that Livingston economists make contrarian forecasts. In short, history was significant to Livingston Survey participants in their forecasts of stock market growth. Despite limitations to our reduced-form process of identifying relevant context, history clearly mattered, and the way it mattered is consistent with the retrieved context model.

References

Berkman, H., Jacobsen, B. and Lee, J.B., 2011. Time-varying rare disaster risk and stock returns. *Journal of Financial Economics*, *101*(2), pp.313-332.

Goetzmann, W. Watanabe, A. and Watanabe, M., (2021a) Expected Co-Movement, Working paper.

Krishnamurthy, A. and Muir, T., 2017. *How credit cycles across a financial crisis* (No. w23850). National Bureau of Economic Research.

Malmendier, U. and Nagel, S., 2011. Depression babies: do macroeconomic experiences affect risk taking? The Quarterly Journal of Economics, 126(1), pp.373-416.

Malmendier, U. and Nagel, S., 2016. Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), pp.53-87.

Malmendier, U., Nagel, S. and Yan, Z., 2020. The making of hawks and doves. *Journal of Monetary Economics*.

Muir, T., 2017. Financial crises and risk premia. *The Quarterly Journal of Economics*, *132*(2), pp.765-809.

Reinhart, C.M. and Rogoff, K.S., 2009. *This time is different*. princeton university press.

Wachter, J.A. and Kahana, M.J., 2019. *A retrieved-context theory of financial decisions* (No. w26200). National Bureau of Economic Research.

5. Tables

Table 1: Variable Importance. This table shows the average variable importance over rolling regressions in Equation (2) by three estimation methods: ordinary least squares (OLS), partial least squares (PLS), and regularized generalized linear models (RGLM). The variable importance is defined as the average absolute t-statistic for the OLS, and the average relative magnitude of variable coefficients as the percentages of the largest coefficient for the PLS and RGLM. Panel A shows the importance of historical context variables matched by the nearest-neighbor problem in Equation (1), and Panels B and C the importance of recent realizations and lagged individual forecasts, respectively.

(A) Historical Context Variables

Variable	OLS	PLS	RGLM
ΔSP_{τ}	0.28	52.6	33.9
ΔIP_{τ}	2.88	57.0	38.5
ΔCPI_{τ}	0.24	25.4	13.8
$\Delta SP_{\tau+1}$	0.93	14.4	6.3
$\Delta IP_{\tau+1}$	2.03	31.0	24.2
$\Delta CPI_{\tau+1}$	0.65	34.7	26.2
$\Delta SP_{\tau+2}$	2.51	44.7	35.7
$\Delta IP_{\tau+2}$	1.52	12.6	9.3
$\Delta CPI_{\tau+2}$	1.21	15.7	22.4

(B) Recent Realizations

Variable	OLS	PLS	RGLM
$\Delta SP_{t-\varepsilon}$	0.80	42.8	25.6
$\Delta IP_{t-\varepsilon}$	1.58	58.0	11.9
$\Delta CPI_{t-\varepsilon}$	3.24	76.5	70.1
ΔSP_{t-1}	0.19	42.3	5.5
ΔIP_{t-1}	5.48	100.0	100.0
ΔCPI_{t-1}	2.28	79.0	48.1
ΔSP_{t-2}	2.20	50.1	34.5
ΔIP_{t-2}	0.63	0.0	4.3
ΔCPI_{t-2}	0.54	65.3	13.9

(C) Lagged Individual Forecasts

Variable	OLS	PLS	RGLM
$fSP06_{i,t-1}$	2.60	22.9	32.7
$fIP06_{i,t-1}$	1.11	38.5	18.9
$fCPI06_{i,t-1}$	0.76	52.0	6.2
$fSP06_{i,t-2}$	2.21	34.6	29.0
$fIP06_{i,t-2}$	0.91	17.0	12.9
$fCPI06_{i,t-2}$	0.56	53.6	9.8
$fSP612_{i,t-1}$	3.23	60.4	44.6
$fIP612_{i,t-1}$	0.55	21.3	7.9
$fCPI612_{i,t-1}$	0.95	44.1	8.2
$fSP612_{i,t-2}$	1.27	9.3	12.5
$fIP612_{i,t-2}$	2.05	8.2	23.4
$fCPI612_{i,t-2}$	0.70	39.6	3.3

Table 2: Variable Importance: Winsorized sample. This table shows the average variable importance over rolling regressions in Equation (2) using the winsorized sample by three estimation methods: ordinary least squares (OLS), partial least squares (PLS), and regularized generalized linear models (RGLM). The variable importance is defined as the average absolute t-statistic for the OLS, and the average relative magnitude of variable coefficients as the percentages of the largest coefficient for the PLS and RGLM. Panel A shows the importance of historical context variables matched by the nearest-neighbor problem in Equation (1), and Panels B and C the importance of recent realizations and lagged individual forecasts, respectively.

(A) Historical Context Variables

Variable	OLS	PLS	RGLM
ΔSP_{τ}	0.13	52.2	32.7
ΔIP_{τ}	3.45	57.2	44.4
ΔCPI_{τ}	0.10	25.4	11.0
$\Delta SP_{\tau+1}$	1.25	16.3	7.6
$\Delta IP_{\tau+1}$	2.50	31.1	27.6
$\Delta CPI_{\tau+1}$	0.19	32.6	20.5
$\Delta SP_{\tau+2}$	2.94	51.4	39.8
$\Delta IP_{\tau+2}$	1.93	15.5	12.4
$\Delta CPI_{\tau+2}$	0.89	13.6	16.9

(B) Recent Realizations

Variable	OLS	PLS	RGLM
$\Delta SP_{t-\varepsilon}$	0.60	41.8	22.0
$\Delta IP_{t-\varepsilon}$	1.97	58.0	12.3
$\Delta CPI_{t-\varepsilon}$	3.35	76.6	69.9
ΔSP_{t-1}	0.31	43.8	2.5
ΔIP_{t-1}	5.40	100.0	100.0
ΔCPI_{t-1}	2.24	78.9	47.1
ΔSP_{t-2}	2.16	48.7	30.6
ΔIP_{t-2}	0.97	0.0	5.1

(C) Lagged Individual Forecasts

Variable	OLS	PLS	RGLM
$fSP06_{i,t-1}$	2.55	23.3	31.0
$fIP06_{i,t-1}$	1.07	38.5	16.9
$fCPI06_{i,t-1}$	0.78	52.0	4.3
$fSP06_{i,t-2}$	2.27	35.1	27.6
$fIP06_{i,t-2}$	0.84	17.6	11.0
$fCPI06_{i,t-2}$	0.54	53.6	7.4
$fSP612_{i,t-1}$	3.25	60.6	43.3
$fIP612_{i,t-1}$	0.61	21.5	6.1
$fCPI612_{i,t-1}$	0.94	44.1	6.1
$fSP612_{i,t-2}$	1.26	9.5	10.7
$fIP612_{i,t-2}$	2.02	8.5	21.7
$fCPI612_{i,t-2}$	0.69	39.8	1.5

Table 3: Panel Regressions, pre-crisis. This table shows the regression in Equation (2) using the pre-crisis period. The dependent variable is the forecasted S&P return over the 6-12 month horizon from the Livingston Survey. The independent variables include the vectors of past realized (prefix Δ) and forecasted (prefix f) growth rates on the S&P index (SP), industrial production (IP), and the consumer price index (CPI). The numbers in forecasted growths represent the forecast horizon, "06" for the base period-6 month ahead forecasts and "612" for the 6-12 month horizon. τ is the past time that solves the minimization problem in (1). $t-\varepsilon$ stands for two months prior to survey date t, i.e., April for June surveys and October for December surveys. Column (1) employs the ordinary least squares (OLS), and Columns (2) to (5) the quantile regressions for the quantile of the dependent variable shown below the column index. The t-statistics are shown in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Time τ variables (b coefficients)

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
ΔSP_{τ}	-0.018	-0.138	-0.071	0.073	0.080	-0.042
	(-0.184)	(-0.919)	(-0.698)	(1.015)	(0.985)	(-0.245)
ΔIP_{τ}	0.333***	0.585***	0.449***	0.235***	0.232***	0.290*
	(3.171)	(3.878)	(4.057)	(3.567)	(2.627)	(1.895)
ΔCPI_{τ}	0.005	0.487***	0.080	-0.037	-0.177*	-0.241
	(0.036)	(2.867)	(0.646)	(-0.367)	(-1.727)	(-1.073)
$\Delta SP_{\tau+1}$	0.014	0.039*	0.014	0.014	0.009	-0.008
	(1.020)	(1.957)	(0.930)	(1.322)	(0.652)	(-0.366)
$\Delta IP_{\tau+1}$	-0.065**	-0.136***	-0.060**	-0.067***	-0.059**	-0.042
	(-2.234)	(-3.418)	(-2.051)	(-3.099)	(-2.161)	(-0.884)
$\Delta CPI_{\tau+1}$	-0.036	0.133	-0.070	-0.057	-0.093	-0.253**
	(-0.478)	(1.165)	(-0.832)	(-1.145)	(-1.317)	(-1.963)
$\Delta SP_{\tau+2}$	0.038***	0.070***	0.050***	0.021**	0.018	0.064***
	(2.822)	(3.890)	(3.611)	(2.164)	(1.403)	(2.818)
$\Delta IP_{\tau+2}$	0.048*	0.109**	0.063*	0.064***	0.039	0.006
	(1.655)	(2.466)	(1.869)	(2.983)	(1.354)	(0.115)
$\Delta CPI_{\tau+2}$	-0.087	-0.458***	-0.152*	-0.047	-0.040	0.136
	(-0.926)	(-3.086)	(-1.697)	(-0.641)	(-0.430)	(0.831)

Panel B: Time *t* variables (*c* coefficients)

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
$\Delta SP_{t-\varepsilon}$	-0.041	0.080	0.016	-0.133*	-0.136*	-0.031
	(-0.419)	(0.544)	(0.165)	(-1.936)	(-1.728)	(-0.192)
$\Delta IP_{t-\varepsilon}$	-0.186*	-0.340**	-0.249**	-0.117	-0.214**	-0.206
	(-1.710)	(-2.347)	(-2.312)	(-1.503)	(-2.121)	(-1.281)
$\Delta CPI_{t-\varepsilon}$	0.700***	0.807***	0.469**	0.415***	0.658***	0.591*
	(3.547)	(3.051)	(2.239)	(2.901)	(3.679)	(1.779)
ΔSP_{t-1}	0.001	0.022	-0.001	0.006	0.008	0.018
	(0.088)	(1.017)	(-0.041)	(0.579)	(0.642)	(0.820)
ΔIP_{t-1}	-0.280***	-0.304***	-0.184***	-0.208***	-0.208***	-0.272***
	(-5.344)	(-4.290)	(-3.333)	(-5.951)	(-5.791)	(-2.755)
ΔCPI_{t-1}	0.320**	0.160	0.122	0.250**	0.294*	0.718***
	(2.024)	(0.816)	(0.834)	(2.128)	(1.941)	(2.721)
ΔSP_{t-2}	-0.035**	-0.041*	-0.052***	-0.032***	-0.029**	-0.025
	(-2.504)	(-1.943)	(-3.718)	(-3.307)	(-2.488)	(-1.029)
ΔIP_{t-2}	0.036	0.022	0.053	0.005	0.016	0.190**
	(0.788)	(0.310)	(1.046)	(0.156)	(0.429)	(2.455)
ΔCPI_{t-2}	-0.049	0.358	0.065	-0.105	-0.178	-0.249
	(-0.284)	(1.635)	(0.347)	(-0.882)	(-1.200)	(-0.888)

Panel C: Economists' prior forecasts: *h* variables

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
$fSP06_{i,t-1}$	0.052***	0.078***	0.039**	0.029**	0.041**	0.003
	(2.613)	(3.011)	(2.198)	(2.408)	(2.510)	(0.104)
$fIP06_{i,t-1}$	-0.069	-0.098	-0.031	-0.031	-0.042	0.004
	(-1.293)	(-1.496)	(-0.643)	(-0.895)	(-0.933)	(0.045)
$fCPI06_{i,t-1}$	-0.087	-0.203	-0.151	0.086	-0.047	-0.247
	(-0.627)	(-1.241)	(-1.291)	(1.023)	(-0.402)	(-1.394)
$fSP06_{i,t-2}$	0.049**	0.084***	0.039**	0.038***	0.038***	0.028
	(2.472)	(3.232)	(2.024)	(2.688)	(2.873)	(0.880)
$fIP06_{i,t-2}$	0.043	0.052	0.011	0.032	-0.002	-0.062
	(0.843)	(0.859)	(0.247)	(0.961)	(-0.062)	(-0.808)
$fCPI06_{i,t-2}$	0.098	0.034	0.130	0.004	0.081	0.130
	(0.730)	(0.220)	(1.099)	(0.042)	(0.693)	(0.618)
$fSP612_{i,t-1}$	0.073***	0.102***	0.124***	0.094***	0.036**	0.050
	(3.350)	(3.099)	(5.689)	(5.582)	(1.997)	(1.395)
$fIP612_{i,t-1}$	0.040	-0.024	-0.013	-0.001	-0.001	0.245***
	(0.663)	(-0.276)	(-0.204)	(-0.033)	(-0.014)	(2.601)
$fCPI612_{i,t-1}$	-0.140	-0.499***	0.024	-0.146	0.122	-0.013
	(-0.965)	(-3.124)	(0.183)	(-1.357)	(1.278)	(-0.075)
$fSP612_{i,t-2}$	-0.023	-0.018	0.036	-0.002	-0.010	-0.080**
	(-1.066)	(-0.605)	(1.563)	(-0.109)	(-0.510)	(-2.220)
$fIP612_{i,t-2}$	0.124**	0.280***	0.083	0.072*	0.097*	0.130
	(1.961)	(3.365)	(1.358)	(1.651)	(1.860)	(1.288)
$fCPI612_{i,t-2}$	-0.110	-0.218	-0.082	-0.039	-0.088	-0.061
	(-0.778)	(-1.424)	(-0.631)	(-0.446)	(-0.783)	(-0.338)
Constant	-0.0002	-0.060***	-0.025***	0.001	0.025***	0.058***
	(-0.150)	(-24.520)	(-17.220)	(1.092)	(18.863)	(24.403)
Observations	2,355	2,355	2,355	2,355	2,355	2,355
Adjusted R^2	0.117					
F Statistic	11.439*** (df = 30; 2324)					

Table 4: Panel Regressions, winsorized full sample. This table shows the regression in Equation (2) using the winsorized full sample. See the caption of Table 3 for variable descriptions. Column (1) employs the ordinary least squares (OLS), and Columns (2) to (5) the quantile regressions for the quantile of the dependent variable shown below the column index. The t-statistics are shown in parentheses. *, **, and *** represent significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Time τ variables (b coefficients)

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
ΔSP_{τ}	-0.010	-0.164	-0.096	0.041	0.105	0.047
	(-0.113)	(-0.996)	(-1.230)	(0.737)	(1.395)	(0.345)
ΔIP_{τ}	0.379***	0.460***	0.412***	0.277***	0.315***	0.464***
	(3.971)	(2.759)	(4.549)	(4.884)	(3.987)	(3.010)
ΔCPI_{τ}	0.087	0.610***	0.196*	0.103	-0.164*	-0.288
	(0.720)	(3.012)	(1.740)	(1.369)	(-1.916)	(-1.321)
$\Delta SP_{\tau+1}$	0.011	0.028	0.005	0.012	0.016	-0.005
	(0.843)	(1.315)	(0.395)	(1.429)	(1.384)	(-0.245)
$\Delta IP_{\tau+1}$	-0.052**	-0.087*	-0.039**	-0.057***	-0.066***	-0.039
	(-2.116)	(-1.897)	(-1.975)	(-3.620)	(-3.200)	(-1.024)
$\Delta CPI_{\tau+1}$	-0.033	0.051	-0.044	-0.056	-0.026	-0.126
	(-0.488)	(0.364)	(-0.608)	(-1.318)	(-0.567)	(-1.408)
$\Delta SP_{\tau+2}$	0.032***	0.045*	0.042***	0.014*	0.021*	0.043**
	(2.615)	(1.960)	(3.397)	(1.828)	(1.865)	(2.034)
$\Delta IP_{\tau+2}$	0.057**	0.087*	0.063***	0.067***	0.063***	0.064*
	(2.307)	(1.896)	(3.149)	(4.271)	(2.675)	(1.875)
$\Delta CPI_{\tau+2}$	-0.095	-0.481***	-0.164*	-0.064	-0.025	0.059
	(-1.079)	(-3.209)	(-1.824)	(-1.166)	(-0.324)	(0.402)

Panel B: Time *t* variables (*c* coefficients)

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
$\Delta SP_{t-\varepsilon}$	-0.062	0.085	0.043	-0.103*	-0.171**	-0.135
	(-0.742)	(0.548)	(0.571)	(-1.950)	(-2.283)	(-1.012)
$\Delta IP_{t-\varepsilon}$	-0.265***	-0.286*	-0.256***	-0.172***	-0.273***	-0.399***
	(-2.887)	(-1.750)	(-3.249)	(-2.966)	(-3.425)	(-2.706)
$\Delta CPI_{t-\varepsilon}$	0.436***	0.581*	0.238*	0.187*	0.495***	0.089
	(2.685)	(1.764)	(1.812)	(1.737)	(3.402)	(0.352)
ΔSP_{t-1}	0.006	0.031	0.004	0.005	0.007	0.026
	(0.474)	(1.409)	(0.319)	(0.643)	(0.689)	(1.215)
ΔIP_{t-1}	-0.268***	-0.346***	-0.169***	-0.154***	-0.229***	-0.287***
	(-6.106)	(-3.339)	(-4.038)	(-5.198)	(-7.064)	(-3.741)
ΔCPI_{t-1}	0.481***	0.273	0.309***	0.320***	0.444***	1.180***
	(3.693)	(1.269)	(2.737)	(3.779)	(3.457)	(5.226)
ΔSP_{t-2}	-0.028**	-0.038	-0.040***	-0.036***	-0.017	-0.006
	(-2.193)	(-1.562)	(-3.013)	(-4.465)	(-1.553)	(-0.309)
ΔIP_{t-2}	0.035	-0.052	0.047	0.019	0.036	0.173***
	(0.868)	(-0.706)	(1.132)	(0.723)	(1.089)	(2.879)
ΔCPI_{t-2}	0.142	0.185	0.224**	0.106	-0.019	0.098
	(1.113)	(0.689)	(2.024)	(1.279)	(-0.159)	(0.688)

Panel C: Economists' prior forecasts: *h* variables

	Dependent variable: $fSP612_{i,t}$					
	OLS	Quantile Regressions				
	(1)	(2)	(3)	(4)	(5)	(6)
		0.1	0.25	0.5	0.75	0.9
$fSP06_{i,t-1}$	0.053***	0.082***	0.037***	0.034***	0.060***	0.050*
	(2.930)	(3.404)	(2.671)	(3.542)	(4.209)	(1.823)
$fIP06_{i,t-1}$	-0.038	-0.102	-0.006	-0.007	-0.030	0.004
	(-0.804)	(-1.444)	(-0.145)	(-0.236)	(-1.019)	(0.088)
$fCPI06_{i,t-1}$	-0.134	-0.215	-0.239**	-0.003	-0.073	-0.228
	(-1.034)	(-1.077)	(-2.324)	(-0.043)	(-0.695)	(-1.202)
$fSP06_{i,t-2}$	0.049***	0.054**	0.031*	0.032***	0.038***	0.043
	(2.729)	(2.001)	(1.938)	(3.039)	(3.190)	(1.591)
$fIP06_{i,t-2}$	0.049	0.096	0.018	0.033	-0.001	-0.045
	(1.060)	(1.289)	(0.443)	(1.178)	(-0.031)	(-0.698)
$fCPI06_{i,t-2}$	0.032	0.141	0.056	-0.032	-0.043	-0.018
	(0.254)	(0.685)	(0.542)	(-0.460)	(-0.410)	(-0.113)
$fSP612_{i,t-1}$	0.077***	0.126***	0.124***	0.115***	0.046***	0.056**
	(3.729)	(4.552)	(6.627)	(8.373)	(2.754)	(2.008)
$fIP612_{i,t-1}$	0.033	-0.045	-0.010	-0.021	-0.063*	0.144***
	(0.598)	(-0.533)	(-0.211)	(-0.603)	(-1.899)	(2.782)
$fCPI612_{i,t-1}$	-0.150	-0.521**	0.002	-0.072	0.096	-0.046
	(-1.082)	(-2.202)	(0.017)	(-0.821)	(1.404)	(-0.205)
$fSP612_{i,t-2}$	-0.021	-0.012	0.043**	0.010	-0.022	-0.083***
	(-1.029)	(-0.402)	(2.053)	(0.690)	(-1.310)	(-2.719)
$fIP612_{i,t-2}$	0.117**	0.274***	0.084*	0.053	0.117***	0.109
	(2.012)	(2.656)	(1.750)	(1.432)	(2.591)	(1.283)
$fCPI612_{i,t-2}$	-0.134	-0.161	-0.104	-0.086	-0.103	-0.165
	(-0.992)	(-0.702)	(-1.127)	(-1.247)	(-1.085)	(-0.918)
Constant	-0.000	-0.057***	-0.023***	0.0004	0.023***	0.056***
	(-0.000)	(-26.113)	(-18.876)	(0.454)	(20.892)	(25.590)
Observations	2,631	2,631	2,631	2,631	2,631	2,631
Adjusted R^2	0.119					
F Statistic	12.785*** (df = 30; 2600)					

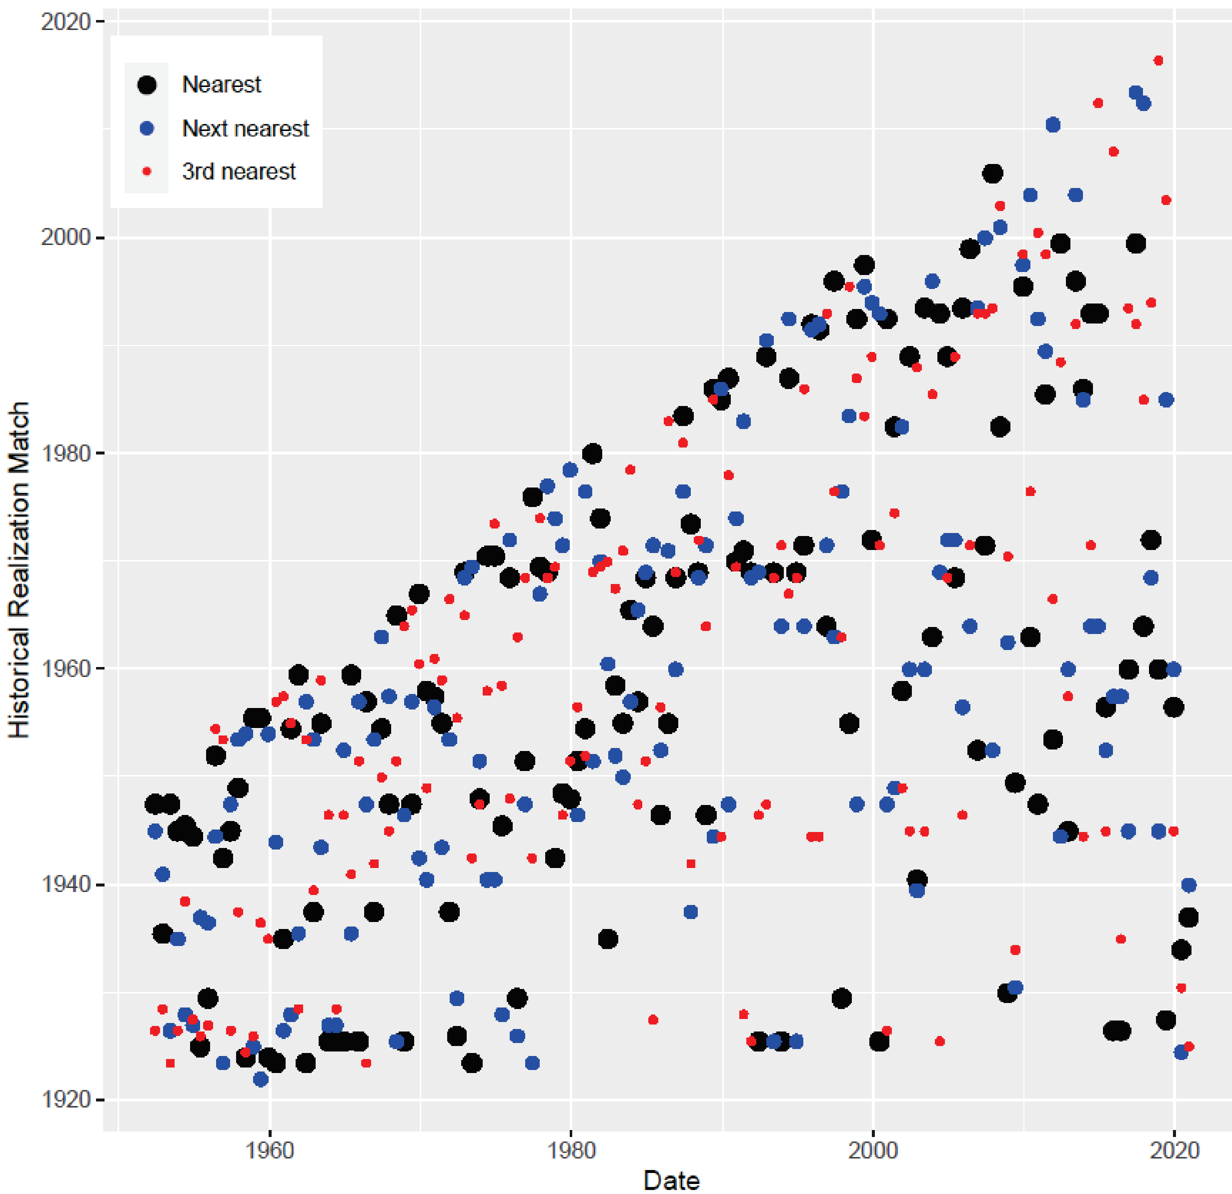

Figure 1: Historical Nearest Neighbors of Realizations. The black dots plot the historical nearest neighbor τ on the vertical axis as the solution to the minimization problem in (1) at time t on the horizontal axis. The blue and red dots are the second and third nearest neighbors defined similarly.

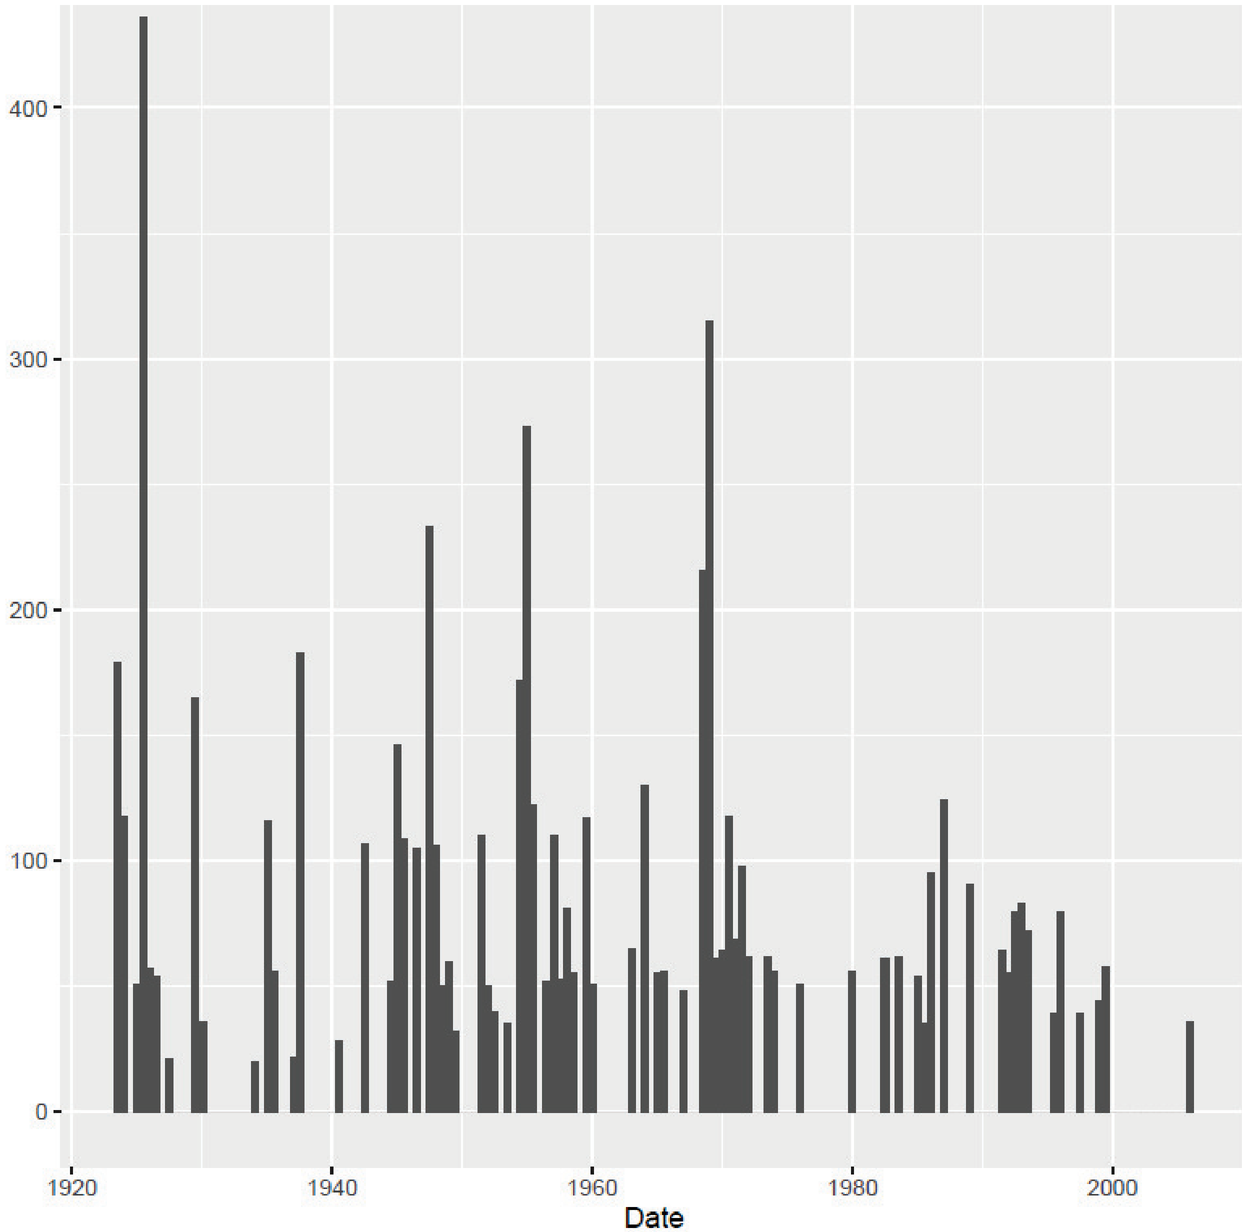

Figure 2: Histogram of the Nearest Neighbor. This figure shows the histogram of the historical nearest neighbor τ on the horizontal axis as the solution to the minimization problem in (1) for all forecasters. Each bar counts the number of times t in (1) that match τ in the panel, i.e., the total number of forecasters present at all t.

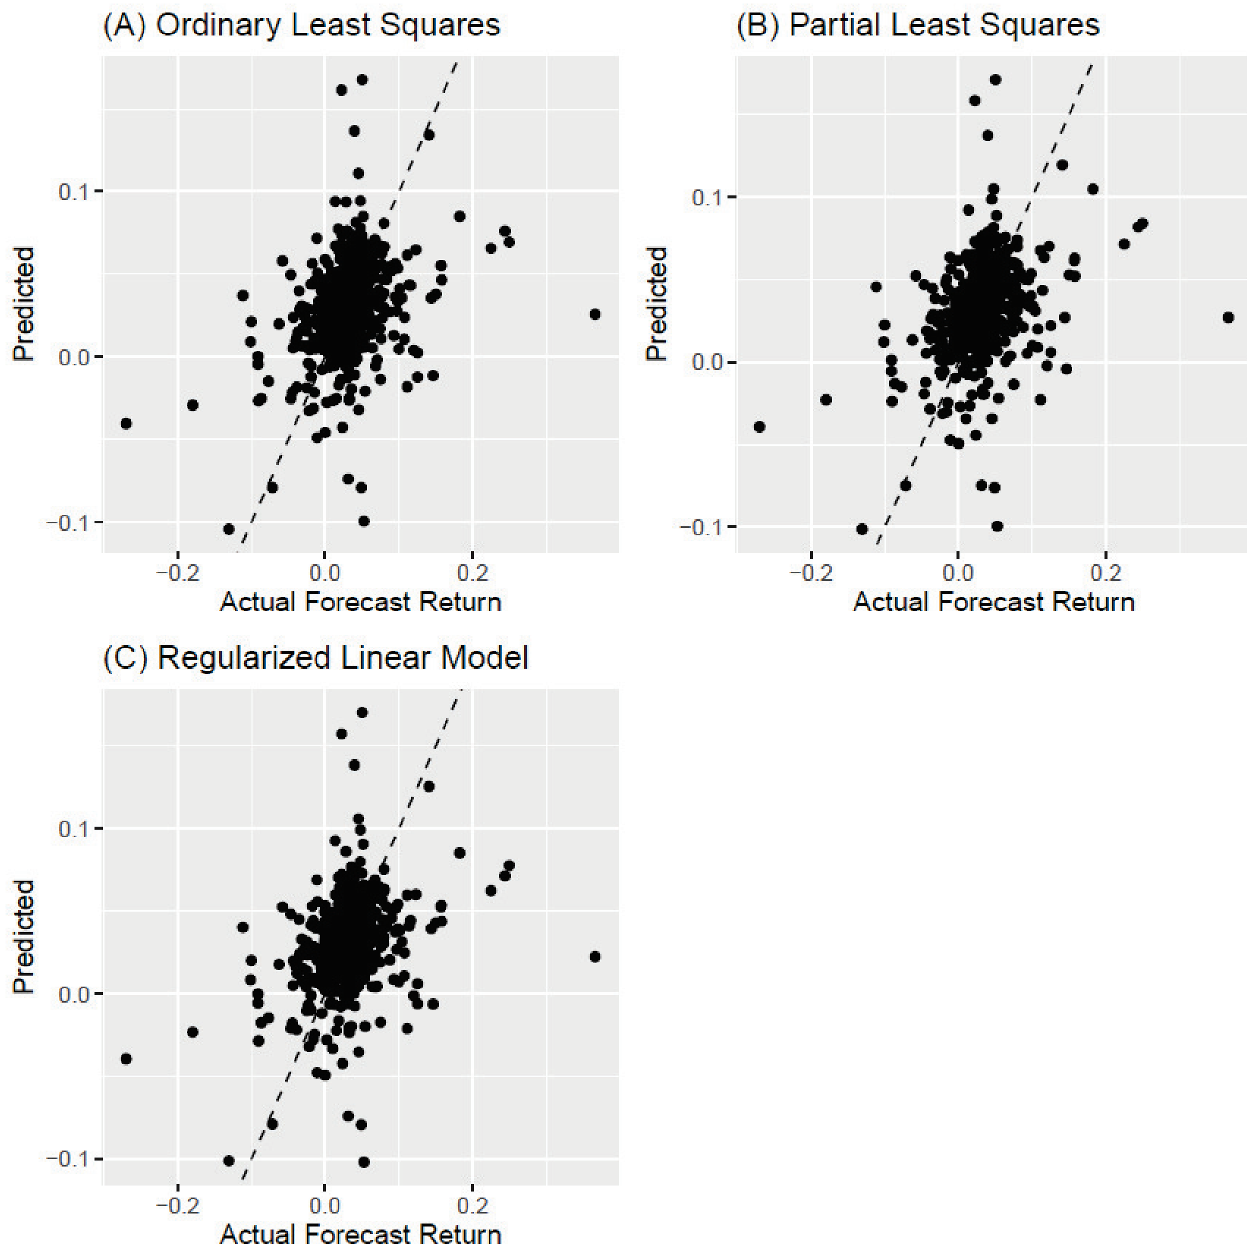

Figure 3: Model Fit. This figure plots the out-of-sample predictions of S&P index growth by Equation (2) against actual forecasts for the three estimation methods. The dashed line is the 45-degree line.

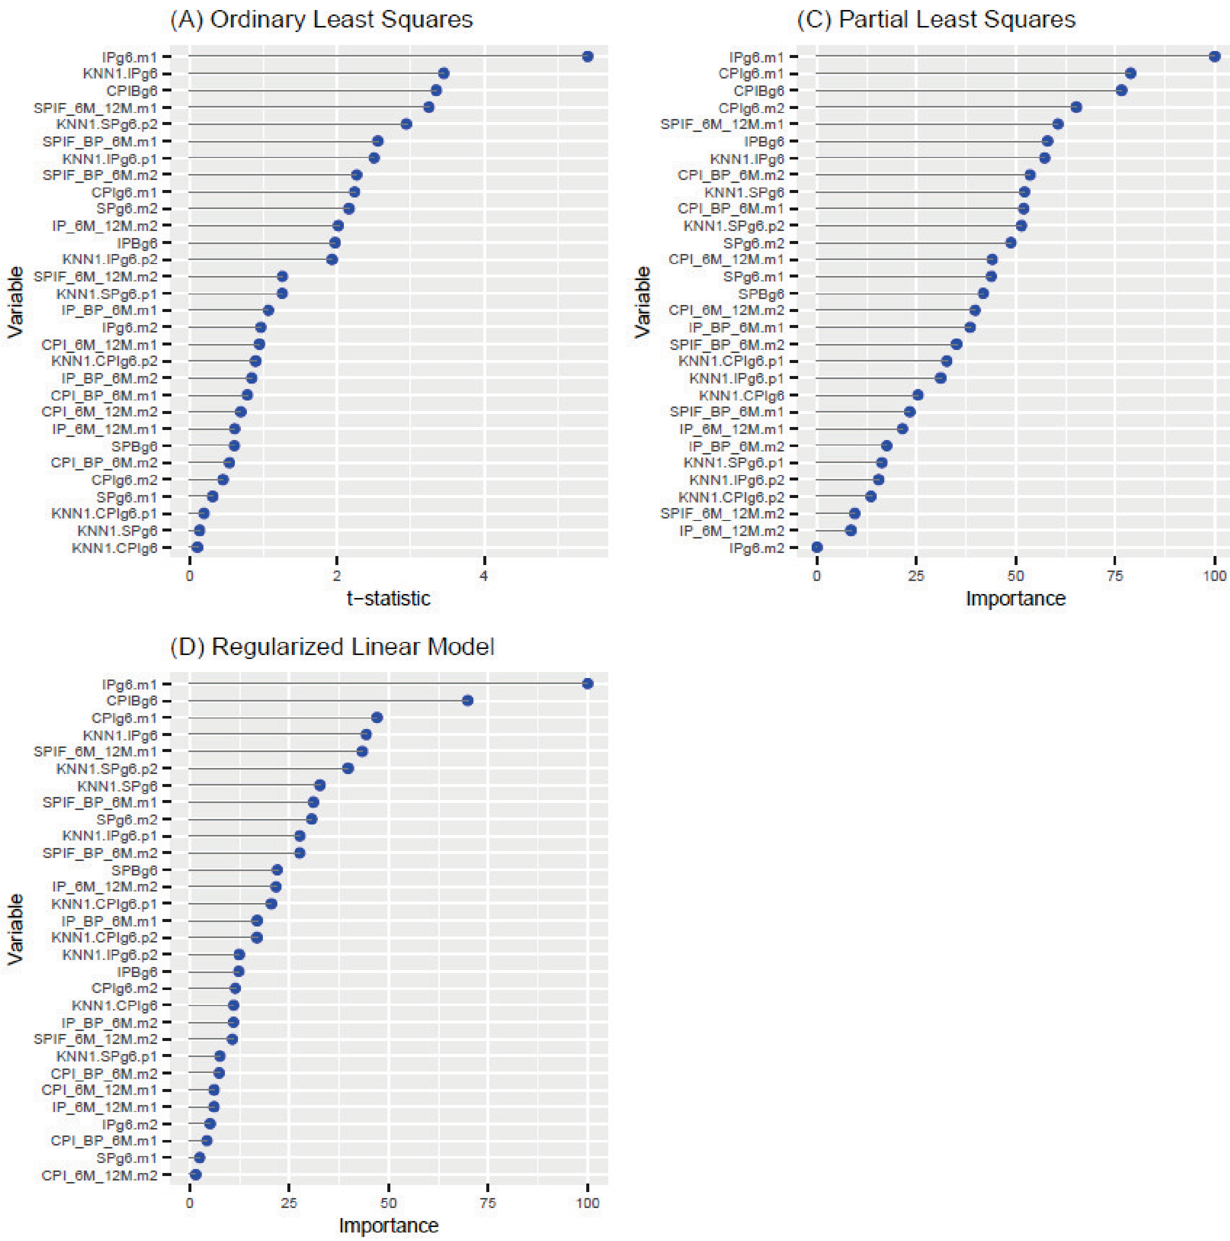

Figure 4: Variable Importance. This figure shows the average variable importance over rolling regressions in Equation (2) by three estimation methods: ordinary least squares (Panel A), partial least squares (Panel B), and regularized generalized linear models (Panel C). See the caption of Table 1 for the definition of variable importance.

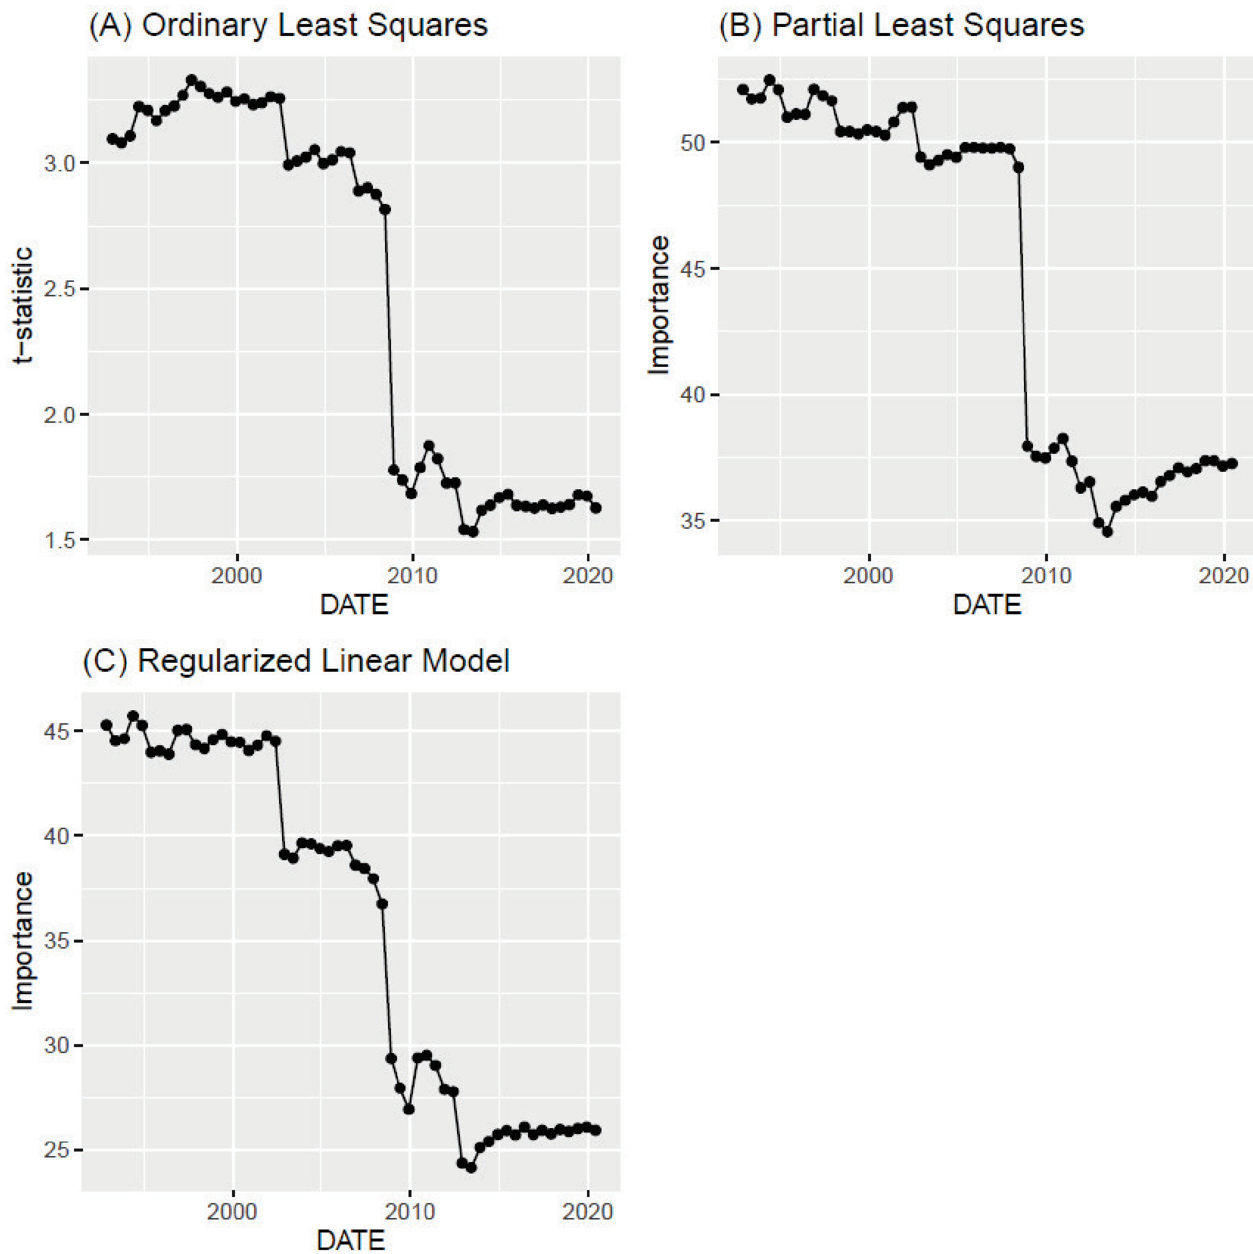

Figure 5: Importance of S&P index growth in two semesters from the nearest neighbor match. This figure plots the importance of $\Delta SP_{\tau+2}$ over time t, where τ minimizes the feature vector distance in Equation (1).

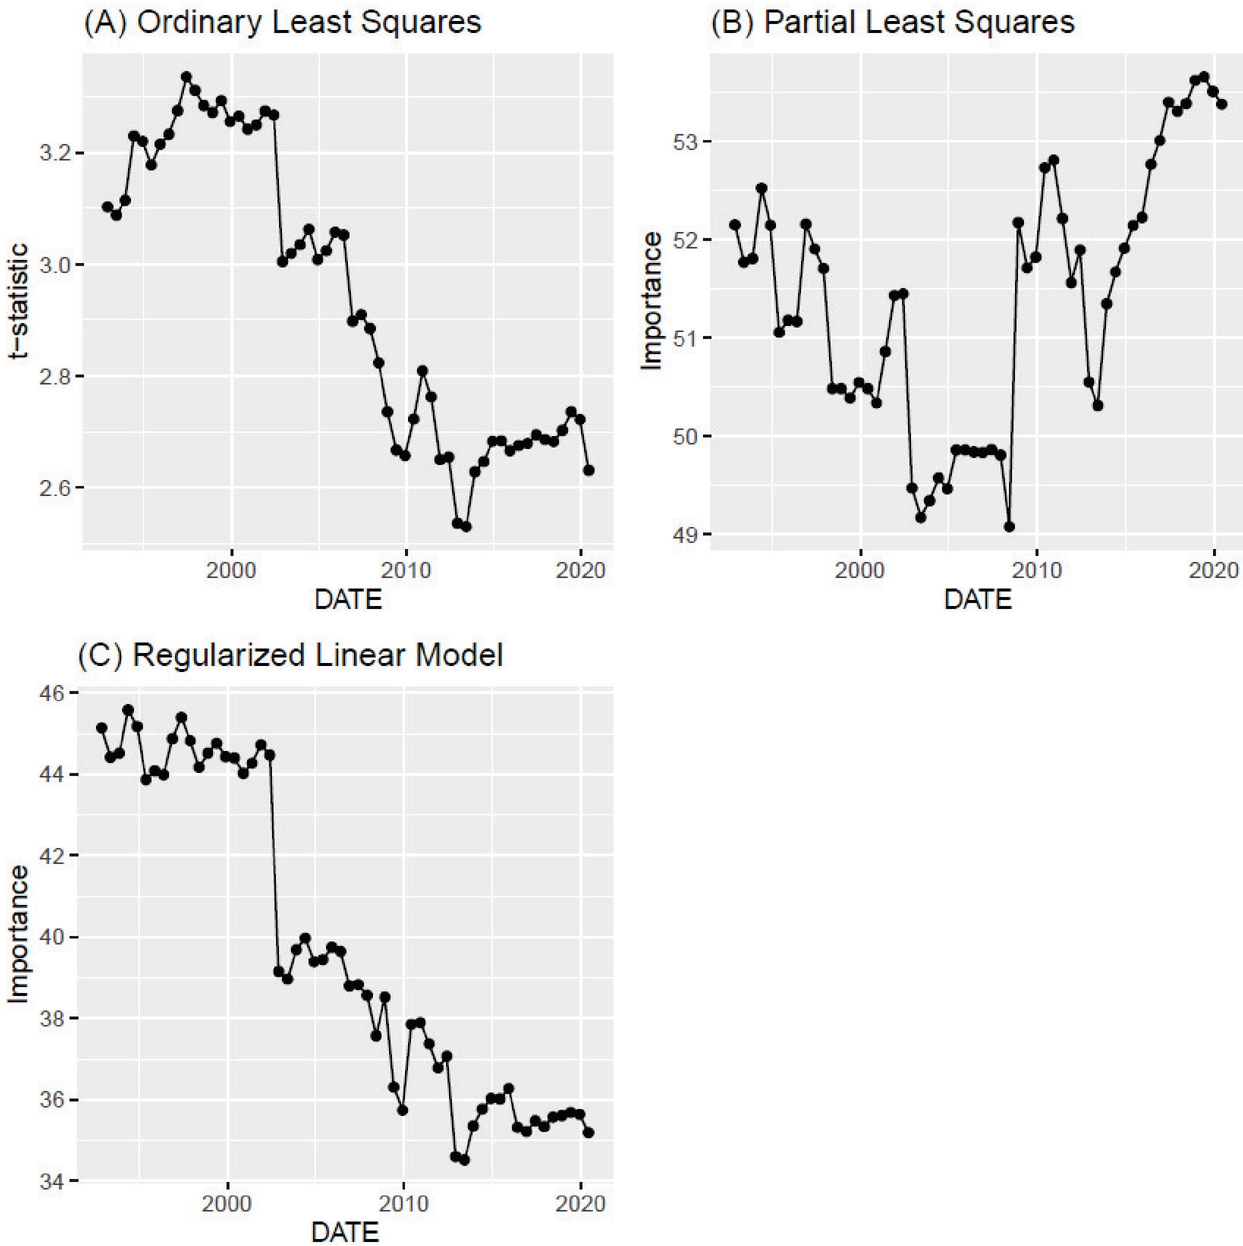

Figure 6: Importance of S&P index growth in two semesters from the nearest neighbor match: Winsorized sample. This figure plots the importance of $\Delta SP_{\tau+2}$ over time t, where τ minimizes the feature vector distance in Equation (1), using the winsorized sample.

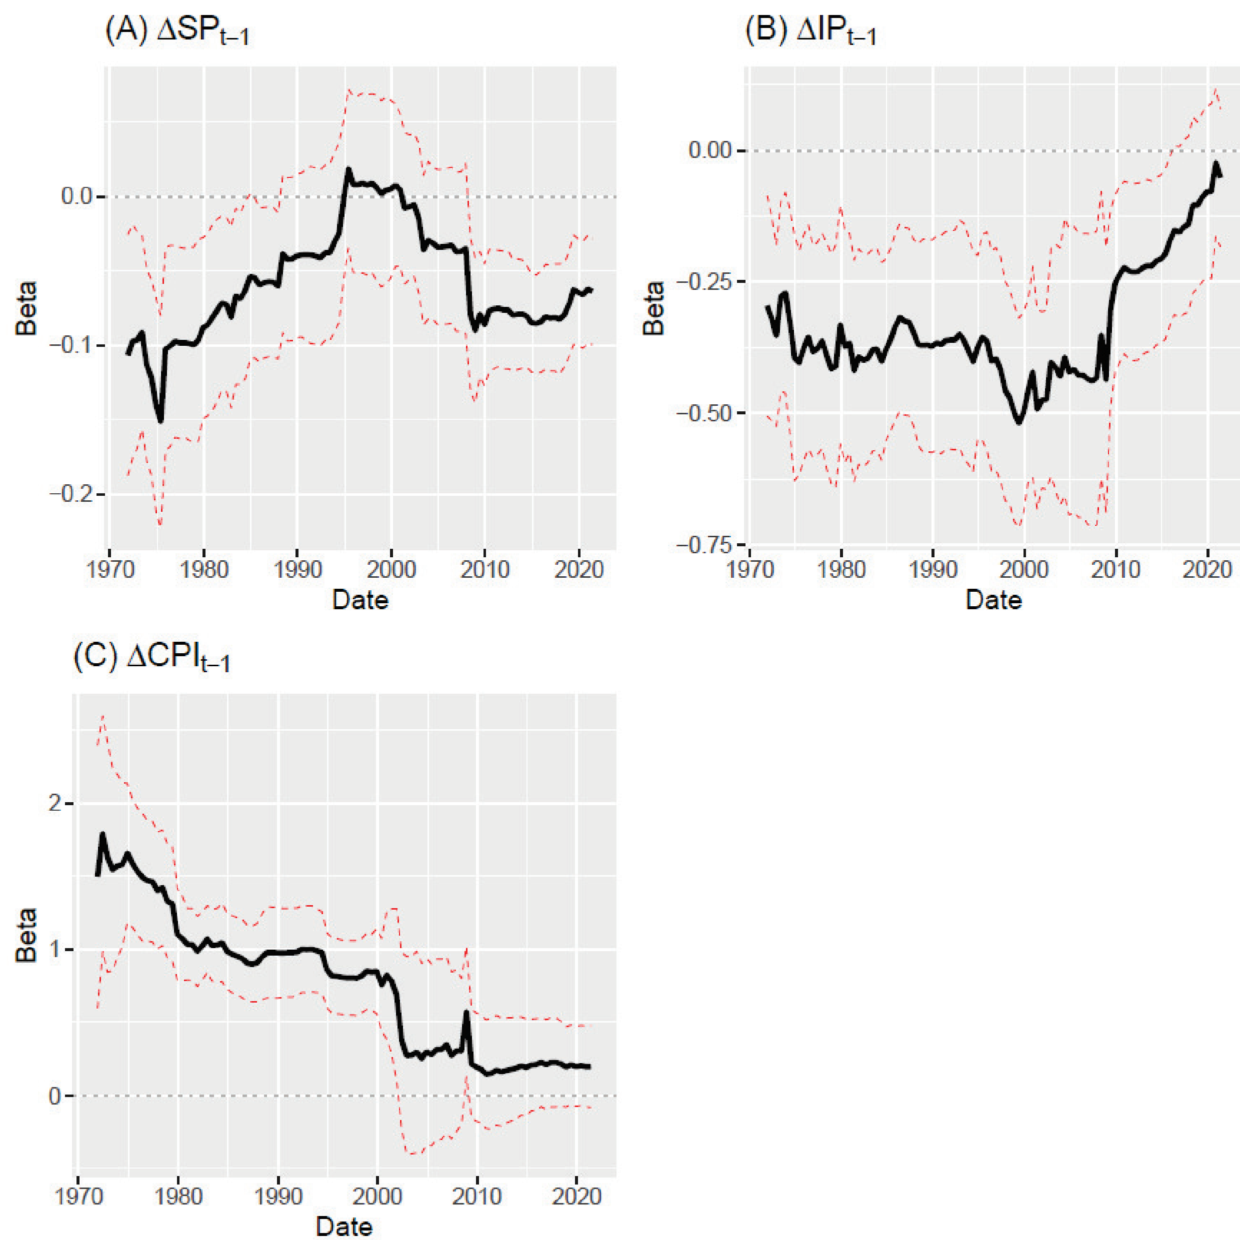

Figure 7: Rolling betas of median forecasted 6-12M S&P growth. This figure plots the rolling betas of the median forecasted 6-12M S&P growth with respect to lagged realized S&P index growth (ΔSP_{t-1}, Panel A), lagged realized industrial production growth (ΔIP_{t-1}, Panel B), and lagged CPI inflation (ΔCPI_{t-1}, Panel C). The length of the rolling windows is fixed at 40 periods