# Necessary Evidence For A Risk Factor's Relevance*

Alex Chinco[†], Samuel M. Hartzmark[‡], and Abigail B. Sussman[§]

May 15, 2020

**Abstract**

Textbook finance theory assumes that investors strategically try to insure themselves against bad future states of the world when forming portfolios. This is a testable assumption, surveys are ideally suited to test it, and we develop a framework for doing so. Our framework combines survey experiments with field data to test this assumption as it pertains to any candidate risk factor. We study consumption growth to demonstrate the approach. While participants strategically respond to changes in the mean and volatility of stock returns when forming their portfolios, there is no evidence that investors view this canonical risk factor as relevant.

**Keywords:** Risk Factors, Expected Returns, Correlation Neglect, Asset Pricing

*We would like to thank Elena Asparohova, Nicholas Barberis, Justin Birru, Ing-Haw Cheng, James Choi, Tony Cookson, William Diamond, Xavier Gabaix, Stefano Giglio, Niels Gormsen, Robin Greenwood, Arpit Gupta, Lars Hansen, Xing Huang, Alex Imas, Brian Kelly, Ralph Koijen, Ben Matthies, Marina Niessner, Lubos Pastor, Kelly Shue, Amir Sufi, David Solomon, Johannes Stroebel, Jessica Wachter, George Wu, Jeff Wurgler, and Anthony Zhang as well as seminar participants at Michigan State, Michigan, Chicago Booth, TCU, Utah, Georgia State, Yale, Rochester, and Colorado for extremely helpful comments and suggestions. We would also like to thank Abigail Bergman, Nicholas Herzog, and Kari Greenswag for excellent research assistance. Becky White and Jasmine Kwong provided invaluable help with data collection. A prior version was circulated under the title *Risk-Factor Irrelevance*.

[†]University of Illinois, Gies College of Business. chinco@illinois.edu

[‡]University of Chicago, Booth School of Business and NBER. samuel.hartzmark@chicagobooth.edu

[§]University of Chicago, Booth School of Business. abigail.sussman@chicagobooth.edu

# 1 Introduction

The textbook explanation for why some assets have higher expected returns than others involves "extending the principles behind fire and casualty insurance to investment portfolios" (Cochrane, 1999). Investors in standard asset-pricing models worry about not having enough money during bad future states of the world. So, they pay more for assets that tend to realize high returns in these bad future states, giving such assets high current prices and low expected returns.

Different models argue that investors worry about different kinds of bad future states. When a model argues that $X$ is a priced risk factor, it is saying that asset prices are moving because investors are trying to hedge the specific kind of bad future state associated with drops in $X$ (whatever that may be). The textbook approach to testing such a model involves looking for a purely econometric relationship between average returns and correlations with $X$. While such evidence is consistent with $X$ being priced, it is not sufficient to prove the case (Lewellen et al., 2010).

This paper proposes a new and complementary approach for testing a separate necessary condition of these models. Textbook models assert that asset markets are "in reality big insurance markets" (Cochrane, 1999). Thus, equilibrium asset prices must reflect the "strategic behavior of investors who wish to hedge against future adverse changes" (Campbell and Viceira, 1999) in particular risk factors. Such investors must worry about and value insurance against these specific risk factors. This represents a testable hypothesis that has been largely unexplored. We develop a survey-based framework for testing it.

Survey experiments are the ideal tool for this purpose. While an asset can provide insurance even if investors are not initially aware of this fact, the equilibrium price of that asset can only contain an insurance premium if its insurance value is commonly understood ahead of time. Home owners can generally explain what they are paying for when they buy fire insurance. Drivers can typically explain what they are paying for when they buy car insurance. If investors are pricing assets by extending these same principles to financial markets, then they should also be able to explain what they are paying to insure themselves against when forming their portfolios.

It is possible for asset prices to move for reasons that investors do not fully understand (Friedman, 1953). However, a key insight in this paper is that a priced risk factor cannot be one of these hidden reasons. For $X$ to be a relevant risk factor, asset prices must be moving because of something that investors do fully understand: each asset's value as insurance against $X$. It should be straightforward to find evidence that investors think about $X$ as something to be hedged using a well-designed survey. Yet, a purely econometric approach cannot provide such evidence. The only way to know what investors are thinking is to ask them.

The framework in this paper can be used to evaluate the relevance of nearly any proposed risk factor. We demonstrate it using a case study that focuses on the most influential risk factor in the academic literature: consumption growth (CCAPM; Lucas, 1978). We choose consumption growth both because it is the most-studied risk factor and because "*all factor models are derived as specializations of the consumption-based model* (Cochrane, 2001, p. 151; emphasis in original)."[1] While habit (Campbell and Cochrane, 1999) and long-run-risk (Bansal and Yaron, 2004) models add state variables to the CCAPM, investors still value hedging drops in future consumption (see Section 4.4). Similarly, in heterogeneous-agent models (e.g., Constantinides and Duffie, 1996), investors still try to hedge consumption shocks of various kinds, and we can use our framework to directly investigate these permutations.

The first step in our framework examines revealed preferences. We ask each participant multiple questions about how they would allocate an endowment between a portfolio of stocks and a riskless bond based on data describing stock returns and consumption growth. We show each participant the actual historic time series of consumption growth. Each participant sees the same consumption-growth time series in every question they answer. However, we present different simulated stock returns in each question. We vary the mean, volatility, and correlation between stock returns and consumption growth across simulations. Participants see numeric values for the mean, volatility, and correlation as well as cumulative graphs of both time series.

We survey a wide range of investor types: finance professionals (including some professional traders), Mechanical Turkers, Booth MBA students, and clients at a meeting of a large asset manager. We designed our survey experiment to make it as easy as possible for all participants to follow the textbook asset-pricing logic. We gave them intuitive instructions, removed all superfluous information, and presented difficult-to-find data in a straightforward manner. In online versions, all participants included in our analysis also had to pass a comprehension test demonstrating that they understood the definitions of all statistical and economic concepts.

The resulting data strongly suggests that participants adjust their demand for stocks based on their perceptions of future market conditions. Participants invest more in stocks when average stock returns are higher ($p$-value < 0.01%), and they invest less when stock returns are more volatile ($p$-value < 0.01%). Our survey participants demonstrate that they are thinking about risk and return as captured by the mean and volatility when investing. Moreover, their reactions to changes in these two parameters are consistent with textbook logic. In short, participants understand their investment task and respond to some parameters in our experimental setup exactly as expected.

However, participants do not respond to changes in the correlation between stock returns and consumption growth when forming their portfolios ($p$-value = 99.9%). Textbook theory assumes that investors view this correlation as having primary importance. Yet, in our setting, there is no change in their demand when stock returns move from being uncorrelated with consumption growth,

[1]The quote continues: "Many authors of factor-model papers disparage the consumption-based model [i.e., the consumption capital asset-pricing model, CCAPM], forgetting that their factor model *is* the consumption-based model plus extra assumptions that allow one to proxy for marginal-utility growth from some other variables."

$\rho = 0.00$, to having a correlation of $\rho = 0.45$. Standard calibrations suggest a CCAPM investor would demand an 11% risk premium to accommodate this increase.

A key benefit of running survey experiments is that it is straightforward to conduct robustness checks. There could be confusion as to what words a lay person would use to describe consumption growth. We use the term *'economic growth'* in our baseline version, but we find the same null response to changes in correlations when presenting participants with variables called *'personal income'*, *'personal wealth'*, *'house prices'*, or *'industrial production'*. Participants could be confused about how to interpret the numeric value of a correlation (even though a definition is provided, participants passed a comprehension check, and some of them are professional traders). Yet, we find similar results when we include additional details on the scale of correlation, when we show participants a scatterplot rather than time-series charts, and when we use words (*'high'*, *'medium'*, *'low'*, or *'none'*) rather than numbers. The graph of cumulative stock returns and consumption may be confusing rather than helpful, but we find similar results in a treatment where no graph is provided. If anything, participants' investment decisions are even harder to square with textbook finance theory when we do things, such as using a scatterplot or only providing participants with the numbers, to further emphasize and clarify consumption-growth correlations.

While these revealed-preference results suggest that participants ignore an asset's correlation with consumption growth when investing, it is possible they are thinking about this parameter in a manner not captured by our various empirical setups. To address this concern, the second component of our framework directly asks participants how they made their investment decisions using a two-stage procedure. First, we asked participants whether they consider mean stock returns, stock-return volatility, and/or the correlation between stock returns and consumption growth when investing. Then, whenever a participant said that they did consider a given parameter, we followed up by asking them how they used this parameter when forming their portfolio.

Participants reported caring about the mean and volatility of stock returns, but most participants (57%) stated that the correlation between stock returns and consumption growth did not play any role in their decision making. Furthermore, among the 43% of participants who did report thinking about correlations, roughly 3 out of 4 reported increasing their demand for stocks when stock returns were *more* correlated with consumption growth—the opposite of what a textbook investor is assumed to do. Across all participants, only 11% reported thinking about consumption-growth correlations in a manner consistent with textbook theory.

The fact that the *Wall Street Journal* does not routinely announce or discuss an asset's correlation with consumption growth should already give us pause. If investors were interested in insuring themselves against drops in consumption, it would be trivial for financial news outlets to report this statistic. It is possible that investors may already know these numbers, leaving no need to report them. To explore this possibility, we examine mutual-fund prospectuses to see what these funds list as their key risks and objectives. Even if everyone is aware of a fund's risk-factor correlations, a fund must still include them in its prospectus for legal reasons. Presumably most investors are aware that the market can go up or down, yet every prospectus discusses this risk. By contrast, not a single fund lists a correlation with any aggregate risk factor in its prospectus.

The aim of this paper is to change the standard of evidence that economists demand of factor models. To explain why one restaurant is more popular than another, you could model diners who prefer nutrients that are more prevalent in the first restaurant's dishes (Ang, 2014, §6.2). But, if you asked people, you would realize that nutrient content is not a good basis for understanding why diners choose. In essence, we are asking people about their investment decisions and realizing that consumption-growth correlations are not a good basis for understanding why markets move.

Our results underscore the importance of understanding not only whether a model is empirically successful but also why. When an asset's correlation with $X$ does not explain average returns, there are two possibilities to consider. It could be that $X$ is not a relevant risk factor, or it could be that $X$ is a relevant risk factor in spite of its empirical shortcomings. Researchers can use our framework to figure out which possibility is correct. If investors are not actively considering $X$, a researcher should discard it. Whereas, if investors are actively considering $X$, a researcher should add features to his model to accommodate $X$ as a priced risk factor.

A researcher proposing that $X$ is a priced risk factor can use our framework to strengthen their claim by showing that investors think about correlations with $X$ as dictated by their model and make investment decisions based on this reasoning. Further, researchers can use our survey-based framework to generate and test new models for which empirical datasets may not be readily available. Researchers promoting a behavioral bias or trading friction can apply our framework to address concerns about competing risk-based explanations. Seminar audiences, conference discussants, and journal editors can and should be asking for this necessary evidence of risk-factor relevance. Without it, they should be much more skeptical that any proposed $X$ is a priced risk factor.

## 1.1 Contribution to the Literature

This paper adds to a large literature testing factor models of asset prices.[2] Many empirical patterns in returns have been documented in both the time series (see Cochrane, 2017, for a recent review) and the cross-section (see McLean and Pontiff, 2016, for a recent review). Often multiple risk-factor-based explanations are proposed for the same pattern in returns, even though the explanations "have little in common economically with each other" (Lewellen et al., 2010). For example, there are more than ten different macro-finance modeling paradigms that attempt to explain the same aggregate market moments (Cochrane, 2017). The standard approach to

[2]When we refer to $X$ as a risk factor, we mean that asset prices reflect investors' "intertemporal hedging demand" for $X$ (Merton, 1969, 1971, 1973). While the term risk factor is sometimes colloquially used to mean any form of return predictability, this is not the intended meaning of in this paper.

disentangling these models is to use advanced econometrics (Chinco et al., 2020; Feng et al., 2020; Freyberger et al., 2020; Harvey and Liu, 2020; Bryzgalova, 2017; Kozak et al., 2018). This paper takes a different tack. Instead of re-examining market data in a new way, we test whether investors are following the core economic logic underpinning a proposed model.

Our paper contributes to the literature that examines how people use and perceive correlations. A number of papers examine biases in how correlations are perceived (Jennings et al., 1982; Matthies, 2018; Ungeheuer and Weber, 2019; Laudenbach et al., 2019). Our setting removes the influence of such a channel by providing the specific value of a correlation. Researchers have examined how misunderstanding correlations can lead to biased decision making in non-financial settings (Enke and Zimmermann, 2017; Levy and Razin, 2015) and have also demonstrated that when combining risky assets into a portfolio, participants do not appropriately account for the correlation structure between the assets (Eyster and Weizsacker, 2016; Kallir and Sonsino, 2009; Matthies, 2018). Our paper adds to this literature by demonstrating that examining how investors utilize risk-factor correlations can be used to test a proposed factor model.

This paper complements a literature examining decision making in experimental settings and how this decision making relates to rational expectations and market outcomes (e.g., Plott and Sunder, 1988; Smith et al., 1988). A number of papers examine whether people behave like classic mean-variance investors (Bossaerts and Plott, 2004; Huber et al., 2019; Kroll et al., 1988; Kroll and Levy, 1992) and recent work studies the Lucas (1978) model in laboratory settings (Asparouhova et al., 2016; Crockett et al., 2019). While some of these experiments find support for the predictions of the CCAPM in lab settings, they focus on empirically analyzing outcomes. While these papers represent important contributions to the literature, similar to the literature that tests asset pricing models, they examine a separate necessary condition of this model's relevance.[3] We contribute to this literature by presenting a framework to directly evaluate the economic reasoning that investors use and whether this reasoning is consistent with a proposed model.

Prior work has looked at how people frame financial decisions—e.g., over individual positions (Odean, 1998), across positions (Frydman et al., 2017), and over portfolios (Hartzmark, 2014)—and how this framing influences portfolio decisions. Other research has examined which attributes investors are drawn to—e.g., saliency (Barber and Odean, 2007), sustainability (Hartzmark and Sussman, 2019), and dividend payments (Harris et al., 2015; Hartzmark and Solomon, 2017). A separate line of survey papers examines more macro questions about consumption (Di Maggio et al., 2020) and investors' agreement with particular lines of economic reasoning (Choi and Robertson, 2019; Giglio et al., 2019; Liu et al., 2020). There is also a literature exploring investor inattention to key asset-pricing quantities (Gabaix, 2014, 2016). This paper contributes by showing that investors

[3]For example, Gode and Sunder (1993) provides an experimental setting where the "allocative efficiency of a double auction derives largely from its structure, independent of traders' motivation."

do not frame investment decisions in a manner that utilizes correlations with consumption growth as an input.

# 2 Survey Design

Textbook asset-pricing models make a core claim about investors' economic reasoning. Investors in these models are worried about not having enough money during particular kinds of bad future states of the world. So, they are willing to pay more for assets that tend to realize high returns in these bad future states, giving such assets low expected future returns going forward. These models are specifically arguing that asset prices are moving because of something that investors understand: each asset's value as insurance against key risk factors. Because the insurance value of an asset must be generally understood by investors for it to be priced, survey experiments are the ideal tool for evaluating the necessary conditions for risk-factor relevance. This section describes the survey experiment we developed to do so.

## 2.1 Instructions

We designed the experiment to use the bare minimum of mathematical and statistical concepts, and we provided definitions for every concept that we referenced. These definitions were not meant to be technical. We used plain language and intuitive concepts so participants could easily understand them. The following describes our baseline online treatment. We discuss the minor modifications we made for different participant populations and robustness treatments below.

Upon entering the survey and completing a consent form, participants were given instructions explaining that they will be making investment decisions based on economic growth and stock market performance. The instructions included intuitive and straightforward definitions of key terms:

> *Economic growth refers to how well the economy as a whole is doing. It is commonly reported as Gross Domestic Product (GDP) which is a measure of the goods and services produced in the US economy. The information about the stock market is for a mutual fund that passively invests in a broad blue-chip stock market index, such as the S&P 500 or the Dow Jones Index. The value of the mutual fund reflects the value of its investments, so when the stocks it invests in have a higher price, the value of the mutual fund will be higher.*

Each participant was informed that they would be seeing annualized numeric values for the mean, volatility, and correlation between the two time series as well as a graph showing the cumulative performance of both stock returns and the economy as a whole. To make sure that it was clear what this meant, we provided each participant with the definition of mean, volatility, and correlation:

> *When the average per year is higher you should expect greater increases in value in a given year, corresponding to steeper increases in the line displayed. When uncertainty is*

*higher, you should expect greater swings, for example higher highs and lower lows are more likely than if uncertainty is low. When a correlation is higher, this means that if one series goes up, the other is more likely to go up too, and if it goes down, the other is also more likely to go down.*

To ensure that participants understood the concepts, they answered several multiple-choice comprehension questions about the definitions of economic growth, average growth, uncertainty, and correlation. We only included data from the subset of people who correctly answered these questions on their first try in our empirical analysis.[4]

The first portion of our survey experiment asks participants to allocate an $1,000 endowment between stocks and bonds based on information about both economic growth and stock returns. Specifically, we ask them to allocate money between a "mutual fund that passively invests in a broad blue-chip stock market index, such as the S&P 500 or the Dow Jones Index," and "a bond earning 2%." See Figure 1a for a sample question. Participants in the online versions see 10 questions of this type. Questions are randomly selected from a larger set of 36 possible questions.

For each question, a participant observes a time series of cumulative stock returns and a time series of aggregate consumption as well as summary statistics describing the mean, volatility, and correlation between these two time series. Participants are told that the numeric values provided to them are stable predictors of future returns for that particular fund. We also tell participants that each round of questioning is unrelated to the last. After participants finished this first set of questions, we asked them a set of follow-up questions about the economic reasoning behind their portfolio decisions using a two-stage procedure.

## 2.2 Design Choices

Our goal is to develop the best way to ask investors whether they view a proposed risk factor as relevant. While seemingly simple, there are multiple ways that one could pose such a question. As with any tool, there are better and worse ways to implement a survey design. Poor econometric technique can lead to erroneous conclusions from data; whereas, a well-defined analysis can lead to meaningful insights. The same is true of surveys. In this section, we discuss why we designed the survey as we did and why we think this design provides the most accurate and interpretable responses.

We examine a simplified and stylized setting so that we can get the most information out of our participants while avoiding a number of potential pitfalls. In our stylized setting, we can define all relevant terms to make sure the concepts are understood. Further, in this setting we can capture

[4] Finance professionals continued the experiment only if they correctly answered all of the comprehension questions correctly on the first try (493 participants). MTurkers completed the full survey irrespective of their responses to the comprehension check, but we only included data from the 322 participants who passed the comprehension test in our empirical results.

behavior that an economist would describe as consumption hedging irrespective of how the participant would describe it. An alternative method would be to directly ask participants about the strategy they use in their own personal portfolio. While it may sound simple to an economist, personal-portfolio strategy is an intricate problem involving many variables (e.g., intra-household decision-making, preferences over non-financial aspects of assets, biases favoring certain stocks over others). It would likely be difficult for most people to answer this question in practice.

A common concern in survey design is attribute substitution (Kahneman and Frederick, 2002) whereby participants do not respond to the complicated question that they are asked, but rather answer a related question that is easier to answer. Participants faced with a difficult question, such as a question about how important a particular strategy is for their own personal investment decisions, may instead respond to a simpler question, such as a question about whether that strategy sounds reasonable. By designing a simple, concrete task, our setting leaves little ambiguity about what question the respondents are answering.

We were careful not to ask leading questions, which can produce experimenter-demand effects (Schwarz, 1999). Experimenter demand can bias survey results when a researcher's choice of survey questions influences participants' answers to these questions. Suppose a participant does not actually try to hedge consumption and has never thought about doing so. When asked about such a strategy in a survey experiment, a participant might anticipate that a researcher would not be asking this question unless it was widely used or recommended. As a result, he might falsely claim that this is how he invests.

This is why, when we follow up and ask participants about their economic reasoning, we use a two-stage question. In the first stage, we ask whether participants considered a parameter—e.g., "Did you think about mean stock returns when investing?" Then, in the second stage, we ask about how they invested based on this parameter—e.g., "When mean stock returns were higher, did you try to hold more or less stock?" This design choice sidesteps experimenter-demand effects because, even if participants believe the researcher would like them to report caring about a given parameter, it is not clear which direction they should report desiring. Thus, we jointly examine consideration and direction of each parameter as our main outcome variables. If participants say they consider a given variable, but utilize it in inconsistent ways, it is suggestive that participants' responses are the result of experimenter demand.

Our survey experiment allows us to cross validate participants' answers across multiple question types to provide converging evidence towards the answer. By using multiple methods, the intention is for the cross-validation to ameliorate concerns about using any single technique. Consistent evidence across question types and within participants suggests that our survey experiment produces valid responses to the questions being asked.

We use consumption growth as a case study in this paper to demonstrate our framework, but our survey design can evaluate the relevance of almost any candidate risk factor. For most candidate risk factors, a researcher should simply be able to take our framework, replace consumption growth with a new risk factor, and run the same survey experiments. In supplemental material, we provide a step-by-step guide that researchers can use to run their own online survey experiments (Bergman et al., 2020). However, we recognize that some risk factors might have idiosyncrasies requiring further tailoring of our approach. There is an active literature studying how to design surveys which we cannot do justice to here. As such, we emphasize that researchers should abide by current best practices in survey design when deviating from our benchmark survey.[5]

## 2.3 Participant Populations

We surveyed four participant populations. The goal was to solicit responses from a broad swath of the investing population and minimize the possibility that there was an important investor type we did not reach. As such, we examined a range of investors in terms of sophistication and wealth. Specific investor types might play a special role in markets, but they do not play such a role in textbook asset-pricing models. The CCAPM (Lucas, 1978) studies a representative investor as do popular models built on top of the CCAPM framework (e.g., Bansal and Yaron, 2004; Campbell and Cochrane, 1999). These models do not give any guidance as to which specific investor populations we should survey.

The first participant pool is comprised of people who work in the finance or banking industries. We used CloudResearch, a service that specializes in connecting researchers with unique and hard-to-reach sample populations,[6] to recruit these participants. CloudResearch has access to more than 50 million online panelists worldwide. We present summary statistics for the 493 finance professionals who completed our survey in panel (a) of Table 1. This participant pool was fairly wealthy, with 56% earning more than \$100$k$ per year. 45% of the finance professionals we surveyed were under the age of 40. Furthermore, these participants tended to invest their own money, with 90% reporting that they owned either individual stocks or mutual funds. We also asked participants about their job function, and 28% of this population stated that their job involved investing in financial securities.

The second population we examined is drawn from Amazon's Mechanical Turk (MTurk) marketplace. Research examining this platform finds that participants recruited through MTurk, who are commonly referred to as *'MTurkers'*, tend to perform similarly on tasks (Casler et al., 2013) and perform better on attention checks (Hauser and Schwarz, 2016) than traditional participant pools recruited in labs. MTurkers also represent a more diverse set of participants (Paolacci and Chandler, 2014). We present summary statistics for the 322 MTurkers who completed our survey in panel (b) of Table 1. MTurkers have lower incomes when compared to the finance professionals, with

[5] We strongly encourage economists to consult colleagues who are well versed in survey design.
[6] See https://www.cloudresearch.com/products/prime-panels/.

only 13% of MTurkers earning more than \$100$k$ per year. MTurkers also tended to be younger than the finance professionals, with 72% being under 40 years of age. A lower fraction of MTurkers owned financial securities, but even in this sample 65% reported owning either stocks or a mutual fund.

The third population we examined is drawn from MBA students at the University of Chicago Booth School of Business. Our sample consisted of 308 participants, of which 38% reported having previously worked in the finance industry. We gave our survey to students enrolled in various MBA courses at the business school, including sections of a core *'investments'* class. Panel (c) of Table 1 contains summary statistics for this participant pool.

The fourth population we examined represents clients attending a conference of a large asset manager. Conference attendees were mostly wealthy investors and portfolio managers. This sample consists of 93 participants who completed our survey. A condition of the survey was that we cannot disclose the name of the asset manager or their clients, but we can safely assume that conference attendees fit the textbook description of a sophisticated investor and generally manage large sums of money on their clients behalf. Panel (d) of Table 1 contains summary statistics for this final participant pool.

## 2.4 Survey Variations

The above description applies to the online version of our survey experiment we gave to the finance-professionals and MTurker participant pools. We administered the MBA-student survey using pen and paper during a class break. We surveyed clients at the asset-manager conference using tablets at a designated booth.

We presented the MBA-student and asset-manager samples with abbreviated instructions and definitions. We did not ask comprehension checks due to time constraints and because both groups are likely to be familiar with the basic concepts. The goal of the design was to parsimoniously present the same information as in the online survey experiment to a group of people that have more knowledge of financial markets. The internet appendix includes examples of instructions from each of these versions of the survey experiment.

For the MBA-student and asset-manager participant pools, we also reduced the number of questions to fit page limits and time constraints. The MBA students saw 5 investment-decision questions. The asset-manager sample saw 4 investment-decision questions. Half of the investors at the asset managers' conference were asked a percent-allocation question, dividing a 100% allocation across the two options. The results were materially similar.

## 2.5 Data Simulation

Our experiment involves time-series data on economic growth and stock returns.

The economic-growth time series each participant sees is always the same across all questions. This time series represents seasonally-adjusted quarterly US GDP, $\Delta \log C_t \overset{\text{def}}{=} \log(GDP_t) - \log(GDP_{t-1})$, from 1980Q1 to 2018Q4 (i.e., $T = 156$ observations in total). Average annualized GDP growth during this sample period is $\mu_{\Delta \log C} = 5.2\%$ while the annualized volatility of GDP growth is $\sigma_{\Delta \log C} = 1.6\%$.

We simulate stock returns using some combination of the following parameter values:

$$\begin{aligned} \mu_R &\in \{4\%, 6\%, 8\%\} \\ \sigma_R &\in \{10\%, 15\%, 20\%\} \\ \rho &\in \{0.00, 0.15, 0.30, 0.45\} \end{aligned}$$

where $\mu_R$ denotes the mean annualized stock return, $\sigma_R$ denotes annualized stock-return volatility, and $\rho$ denotes the correlation between stock returns and consumption growth.

The typical stock-market return is 6% per year, and annualized return volatility is roughly 15% (Cochrane, 2001). The CCAPM says an asset's expected excess return should be

$$\mu_R = \gamma \times (\, \rho \cdot \sigma_R \cdot \sigma_{\Delta \log C} \,)$$

where $\gamma$ denotes investors' coefficient of relative risk aversion. To match the equity premium, we need to assume $\gamma \approx 100$ (Campbell, 2003), and the annualized volatility of consumption is about $\sigma_{\Delta \log C} = 1.6\%$ in our sample period. So, according to the CCAPM, investors should view a mutual fund with a 6% per year average returns as underpriced and have high demand for its shares when $\rho = 0.00$; whereas, they should see this same 6%-per-year fund as overpriced when $\rho = 0.45$:

$$\underbrace{100 \times \big(0.00 \cdot 15\% \cdot 1.6\%\big) = 0\%}_{\substack{\text{CCAPM-implied expected excess return} \\ \text{in low correlation setting (underpriced)}}} < \quad 6\% \text{ per year} \quad < \underbrace{10.8\% = 100 \times \big(0.45 \cdot 15\% \cdot 1.6\%\big)}_{\substack{\text{CCAPM-implied expected excess return} \\ \text{in high correlation setting (overpriced)}}}$$

In other words, moving from $\rho = 0.00$ to $\rho = 0.45$ should cause a CCAPM investor to increase the expected excess return he would demand for holding an asset from zero to roughly double the sample average that we observe.

For each set of parameter values, we first draw $T = 156$ iid realizations $\Delta \log Z_t \overset{\text{iid}}{\sim} \mathrm{N}(0,1)$. Then, we orthogonalize these 156 random draws with respect to the realized set of 156 consumption growth shocks, $\widetilde{\Delta \log Z}_t \overset{\text{def}}{=} \big(\Delta \log Z_t - \widehat{\mathrm{E}}[\Delta \log Z_t | \widetilde{\Delta \log C}_t]\big) \big/ \widehat{\mathrm{Sd}}[\Delta \log Z_t | \widetilde{\Delta \log C}_t]$, where $\widetilde{\Delta \log C}_t \overset{\text{def}}{=} \big(\Delta \log C_t - \widehat{\mathrm{E}}[\Delta \log C_t]\big) \big/ \widehat{\mathrm{Sd}}[\Delta \log C_t]$ denotes the consumption-growth shocks. We do this to avoid the error-in-variables problem caused by our finite sample period, which contains 156 quarters. If we skip this step, the resulting error-in-variables problem is not large, but since we can control everything about our experimental setting, we try to remove all avoidable sources of error from our results.

We simulate stock returns using the formula below:

$$R_t = \mu_R + \sigma_R \times \left( \rho \cdot \widetilde{\Delta \log C}_t + \sqrt{1-\rho^2} \cdot \widetilde{\Delta \log Z}_t \right)$$

Because of the orthogonalization step, the resulting stock-return time series has a mean of exactly $\mu_R$, a volatility of exactly $\sigma_R$, and a correlation with consumption growth of exactly $\rho$. We show participants a line labeled *'stock market'* representing the cumulative returns to investing $1 in this portfolio in 1980 on a log scale. See Figure 1a. For each set of parameter values, we run the simulation using 5 different random-number seeds.

Finally, the realized consumption-growth time series is quite smooth. So, for a subset of participants, we added orthogonalized noise to the consumption to make it easier to see comovement between the time series and stock returns. See Figure 1b. Our results for all tests are nearly identical when we add noise, so we combine both samples in our analysis.[7]

# 3 Main Results

According to textbook finance theory, markets fluctuate because intelligent forward-looking investors are worried about not having enough money during certain kinds of bad future states of the world. If an asset's returns are less correlated with the associated risk factors, then the asset is more likely to have positive returns during these specific bad future states. As a result, investors should be willing to pay more for the asset today, and the asset should have lower expected returns going forward because investors recognize the risk-factor correlations and adjust their demand accordingly.

For such a model to explain why markets actually fluctuate, investors must be following this logic. People must prefer assets that have lower risk-factor correlations and adjust their demand based on this information. This is not what we find for consumption growth. We find no evidence that people trade the way textbook investors are assumed to trade or think the way textbook investors are assumed to think.

## 3.1 Investment Decisions

We begin by exploring how participants' investment decisions change as we vary mean stock returns, stock-return volatility, and the correlation between stock returns and consumption growth. We estimate regressions of the following form:

$$\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{mean}_{i,q} + \hat{\gamma} \cdot \texttt{volatility}_{i,q} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$$

The dependent variable, $\texttt{stockFrac}_{i,q}$, is the fraction of the endowment that the $i$th participant invests in stocks when answering the question in round $q$. The variables $\texttt{mean}_{i,q}$, $\texttt{volatility}_{i,q}$,

[7]See Internet Appendix for the results presented separately for each sample.

and $\texttt{correlation}_{i,q}$ represent the mean, volatility, and correlation with consumption growth used to simulate the stock returns for that question. We estimate all $t$-statistics and $p$-values for the investment-decision regressions using standard errors clustered by participant.

Table 2 reports the results of these regressions. Column (1) regresses the fraction invested in the stock market on only the mean stock return. We estimate a slope coefficient of 3.24, which is statistically significant at the 1% level. This coefficient implies that participants increased the fraction invested in stocks by $12.96\% = (8\% - 4\%) \times 3.24$ in response to a move from the lowest mean return, 4%, to the highest mean return, 8%. Across our four participant pools, people invested 60% of their endowment in stocks on average. Thus, a 4% change in expected returns increased the proportion allocated to stocks by about 20%.

Column (2) repeats the regression using stock-return volatility rather than mean stock returns as the sole right-hand-side variable. We estimate a slope coefficient of −0.61, which is again statistically significant at the 1% level. This coefficient implies that a 10% drop in stock-return volatility—i.e., a move from the highest volatility, 20%, to the lowest, 10%—results in a 6.1% increase in stock investment. The results in columns (1) and (2) of Table 2 suggest that participants respond to changes in the mean and volatility of stock returns. Moreover, they do so consistent with a textbook investor who likes higher means and dislikes higher volatilities.

The results are quite different when examining correlations in column (3). This column repeats the regression using correlation with consumption growth as the sole right-hand-side variable. We find no measurable change in participants' behavior in response to a change in the correlation between stock returns and consumption growth. The estimated coefficient is −0.0000235 (rounded to 0.00) with a $t$-statistic of 0.00 and a $p$-value of 99.9%. In addition to being statistically insignificant, this point estimate is also economically small. In response to an increase in correlation of 0.45, the regression implies that participants would only decrease their allocation to stocks by $0.001\% = (0.45 - 0) \times 0.0000235$. This is one tenth of one basis point. Our participants simply do not adjust their demand in response to changes in the correlation between stock returns and consumption growth, the canonical risk factor in textbook finance theory.

In column (4) of Table 2, we include all three of these right-hand-side variables in the same regression and find nearly identical results. Participants adjust their demand in response to changes in the mean and volatility of stock returns but not in response to changes in the correlation between stock returns and consumption growth. This is expected since we randomly chose the mean, volatility, and correlation parameters used to simulate the stock returns for each question.

To examine whether these results are driven by participant-specific attributes, we add participant fixed effects in column (5). The coefficient on the correlation parameter hardly changes from column (3). Another concern is that participants might change their behavior over the course of the experiment. To account for this, we introduce question-order fixed effects in column (6) and find similar results. Finally, column (7) adds both participant and question-order fixed effects which again results in unchanged point estimates.

The results in Table 2 show that, on average, participants strongly respond to changes in the mean and volatility of stock returns but ignore changes in the correlation between stock returns and consumption growth. It remains possible, however, that these pooled results hide the behavior of a subset of participants who act differently. To address this concern, we re-estimate the coefficients in column (3) of Table 2 on each participant pool. We report these investment-decision results by participant pool in Table 3. Column (2) in this table adds participant-pool fixed effects to the specification in column (3) of Table 2 to capture that fact that there are differences in the average fraction invested in the stock market across participant pools. The estimated slope coefficient of 0.00, which is identical to column (3) of Table 2, reveals that these differences are not driving our results.

In columns (3)-(6) of Table 3, we re-estimate the regression separately for each participant pool. We find a coefficient on the correlation parameter of 0.00 for the finance professionals, 0.00 for the MTurkers, 0.07 for MBA students, and −0.03 for the asset-manager sample. The MBA students have the only point estimate that is marginally statistically significant, but it has a positive sign rather than the negative sign predicted by theory. In addition to being generally statistically insignificant, these point estimates are all at least two orders-of-magnitude smaller than the point estimate on mean returns in column (1) of Table 2.

We also examine our results based on participant characteristics. Figure 2 graphs regression coefficients estimated over various subgroups of our participant pools. Every subgroup of participants we examine exhibits similar behavior. Old and young participants; male and female participants; participants with incomes greater than \$100$k$ and those with incomes less than \$100$k$; participants who think they invest wisely and those who do not think they invest wisely all ignore changes in correlation. We repeat the analysis for all participants who state that they own either stocks or mutual funds. We repeat the analysis for the subset of 136 finance professionals in our sample who stated that their job involved trading financial securities. Financial professionals who trade securities for a living do not adjust their demand in response to changes in correlations. In every sub-sample, participants react to changes in means; they react to changes in volatilities; they ignore changes in correlations.

We would like to emphasize that the correlation changes we explore in this paper, from $\rho = 0.00$ to 0.45, should be sufficient to generate large differences in investor behavior. While it would be possible to further extend the range of correlations in our study, to our knowledge, there is not a general belief that stock returns and consumption growth have a correlation close to one or that they are negatively correlated. As discussed earlier, investors should view a mutual fund with a 6% per year average return as underpriced and have high demand for its shares when $\rho = 0.00$; whereas, they should see this same 6%-per-year fund as overpriced and have low demand when $\rho = 0.45$.

Moving from $\rho = 0.00$ to 0.45 should cause a CCAPM investor to demand an additional 11% per year risk premium—a change that is roughly double the average excess return on the market, $\mu_R = 6\%$ per year. Our experiment suggests that, in response to an 11% increase in mean return, a participant would increase their stock allocation by $35.64\% = 11\% \times 3.24$, but they do not respond to a shift in consumption-growth correlations as our models would suggest.

Another benchmark worth considering is the magnitude of the shift in correlations required to solve the well-known equity-premium puzzle (Mehra and Prescott, 1985). Standard measures of consumption growth based on nondurables and services have a correlation with the excess market returns of about 0.38. Savov (2011) measures consumption growth using garbage and finds a correlation of 0.54. This shift in correlations of 0.16 is large enough "to formally resolve the joint equity premium-risk-free rate puzzle." We show that participants are not adjusting their demand in response to a change in correlations that is roughly 3 times as large. Thus, investors in a textbook asset-pricing model should view these changes in correlations as economically significant.

We describe the variable of interest as *'economic growth'* in our baseline experiment. We do so because we think it is an intuitive concept that most participants should understand. Further, it closely maps to the relevant variable described by textbook theory. With that said, it could be that participants are trying to hedge something similar to economic growth but think of it in different terms. If such a channel was at work, we would be wrongly rejecting consumption growth as a priced risk factor due to labeling.

Table 4 shows that this is not the case. The results in each column are analogous to column (4) in Table 2, but *'economic growth'* is relabeled as *'personal income'*, *'personal wealth'*, *'house prices'*, or *'industrial production'*. We use the same underlying data and the same simulated series for each of these. The only change is to relabel the variable from *'economic growth'* to the alternative term and to change the definition in the instructions. The coefficient on $\texttt{correlation}_{i,q}$ is statistically indistinguishable from zero in every column. The wording we use to describe the risk factor does not affect our results.

In Table 5, we explore concerns related to the display format of correlations and show that these concerns are unlikely to explain our findings. Maybe participants are unable to interpret a numeric value for a correlation? There are *a priori* reasons to doubt that a lack of numerical savvy explains our results. Participants passed a comprehension test and reacted to numerical changes in the mean and volatility of stock returns. Nevertheless, column (1) in Table 5 shows that our results do not change if we replace the numerical values of $\rho \in \{0.00, 0.15, 0.30, 0.45\}$ with the text $\rho \in \{$*'none'*, *'low'*, *'medium'*, *'high'*$\}$.

Another concern could be that the scale of the correlation statistic is confusing. If participants are unaware that correlations fall in a range between negative one and one, $\rho \in [-1, 1]$, perhaps they are unaware that the changes in correlation in our experiment are quite large. In column (4) of

Table 5, we show results for our investment-decision questions when we include additional information about the range of the correlation and how to interpret it. Again, nothing changes about our results. The coefficient on $\texttt{correlation}_{i,q}$ is still almost zero ($p$-value = 85.35%). Our findings are not driven by participants failing to understand the scale of correlations.

But, maybe this result is specific to time-series plots? Scatterplots do more to visually highlight a correlation, so we run the experiment displaying such a plot instead of a time series graph. The fact that media outlets do not generally choose this graphical format, which accentuates correlations, already suggests that this statistic is less likely to be relevant to investors. Column (3) in Table 5 shows that, if we present participants with questions containing scatterplots rather than time-series graphs, they are even more likely to behave in the opposite way from a textbook investor. The coefficient on $\texttt{correlation}_{i,q}$ is positive (0.15) and statistically significant at the 1% level.

While the time-series graph was included as it seemed the closest to standard displays of information in the financial press, it could be that participants were mislead by the graph altogether. Column (2) in Table 5 shows that, if we ask participants the same investment-decision questions as before but remove the time-series graph, participants are again even less likely to follow textbook asset-pricing logic. The coefficient on $\texttt{correlation}_{i,q}$ in column (2) of Table 5 is 0.18 and statistically significant at the 1% level. So, instead of investing less in stocks when stock returns are more correlated with consumption growth as a textbook investor would do, participants invest more in stocks when stock returns are more correlated. If anything, participants' investment decisions are even harder to square with textbook theory when we remove the graph.

We go to great lengths to clearly communicate the correlation between stock returns and consumption growth to our participants. Every survey variant we study delivers the same result: participants do not respond to changes in consumption-growth correlations the way a textbook investor should. The textbook explanation for why asset prices move starts with the assumption that investors are trying to insure themselves against drops in future consumption (Merton, 1973). If investors are actually trying to do so, it should be fairly straightforward to find evidence that investors are trying to accomplish this goal. For example, participants in every pool strongly respond to the mean and volatility of stock returns in all survey variations. If hedging consumption growth is a core determinant of their investment strategy, it seems unlikely that subtle differences in the framing of exactly what is being hedged should lead to such dramatically different results.

We think our robustness checks span most reasonable variations in terms describing both aggregate and idiosyncratic consumption growth. It is possible, though, that a future researcher might use our survey experiment to show that participants respond to changes in the correlation between stock returns and some other variable $X$, which is a slight variant of consumption growth empirically. However, finding that participants try to insure themselves against negative shocks to consumption growth when it is described as $X$ but not when it is described as economic growth, income growth, wealth growth, house-price appreciation, or changes in industrial production would only underscore the importance of using a well-designed survey experiment to test investors' economic reasoning. It would suggest that, even though consumption growth and $X$ look econometrically similar, investors view these two variables quite differently. Such a result would highlight the value of our framework. It would point to an important distinction that investors are drawing when pricing assets that existing econometric methods are blind to.

## 3.2 Economic Reasoning

The results above suggest that participants do not adjust their demand for stocks in response to changes in the correlation between stock returns and consumption growth, the canonical risk factor in textbook asset-pricing models. With that said, perhaps participants are considering consumption-growth correlations in a manner that is not captured by our regressions. To address this concern, we follow-up the investment-decisions questions by asking participants about the economic reasoning for their choices using a two-stage procedure.

We first ask participants whether they considered mean stock returns, stock-return volatility, and the correlation between consumption growth and stock returns when making their investment decisions. Let $\texttt{consider}_i$ be an indicator variable which is equal to one if the $i$th participant reported thinking about a parameter when making their investment decisions. The first row of Table 6 shows that 77% of participants considered average stock returns when making their investment decisions, 59% of participants considered stock-return volatility, but only 43% of participants considered the correlation between stock returns and consumption growth. Correlations are the least likely to be considered.

In the second stage, we ask participants who said they considered a given parameter about the direction in which they used this information. Let $\texttt{textbookLogic}_i$ be an indicator variable that is equal to one if the $i$th participant reported thinking about a parameter using textbook asset-pricing logic. This variable equals zero if a participant does not consider the variable at all. The third row of Table 6 examines whether the participants that considered a variable did so in a manner consistent with textbook theory. It shows that 76% of the participants who considered average stock returns when making their investment decisions tried to buy more stocks when average stock returns were higher. Likewise, 72% of the participants who considered stock-return volatility tried to buy more stocks when this parameter was lower. In short, 3 out of 4 participants who considered the mean and volatility of stock returns when investing did so in the textbook direction.

The opposite is true for consumption-growth correlations. The right-most number in the third row of Table 6 shows that only 24% of the participants who considered consumption-growth correlations tried to buy more stocks when this correlation dropped. Even the participants who considered consumption-growth correlations stated that they did so in the opposite direction posited by textbook explanations.

The difference between these findings and the intuition of textbook asset-pricing theory is stark. Most participants did not consider the correlation between stock returns and consumption growth, even after being given its numeric value and being asked directly about it. Theory suggests this should be a central parameter of interest, yet it was not considered as important by 57% of participants. Moreover, of the participants who did consider this correlation, most did so in the opposite direction of what standard models would suggest—3 out of 4 participants who considered consumption-growth correlations tried to buy *more* stocks when stock returns were *more* correlated with consumption growth. This means that 3 out of 4 participants who claimed they were thinking about the correlation between stock returns and consumption growth were trying to hold more stocks when stock returns were a worse hedge against bad economic times. The results suggest the non-responsiveness of participants' demand to changes in risk-factor correlations in the first portion of our experiment simply reflects how participants think about their investment decisions.

Table 7 repeats the analysis in Table 6 separately for each participant pool to show that the results are not driven by the less financially sophisticated participants in our sample. Column (3) shows that only 35% of the investors at the asset manager's conference (33 out of 93 participants) said that they considered the stock market's correlation with consumption growth when making their investment decisions. Column (9) shows that less than half of those 33 investors (48% or 16 participants) said that they tried to invest more when holding stocks was a better hedge against drops in consumption. Across all participant pools, investors tended to not think about correlations or to do so in a manner inconsistent with textbook theory, mirroring the results from the first portion of the experiment.

One of the benefits of using converging evidence across multiple question types is that we can use this structure to cross-validate question responses. We find that 11% of participants report thinking about consumption-growth correlations like a textbook investor should. One possibility is that these participants strongly follow the textbook logic and perhaps are particularly important for driving asset prices. Another possibility is that this 11% is largely noise.

Table 8 examines participants' investment decisions based on their answers to the economic-reasoning questions in the second part of the experiment. The findings indicate that the 11% of participants who followed textbook logic are likely explained by noise. Each entry in Table 8 is a regression of the fraction invested in stocks on the indicated parameter using data on a specific subpopulation. The first and second rows report results only for those participants who said they either did not consider or did consider a particular parameter. The first row of columns (1) and (2) shows that participants adjusted their demand for stocks in response to changes in the mean and volatility of stock returns even when they said they were not explicitly considering these parameters. When participants did report thinking about either the mean or volatility of stock returns, their portfolio response was even stronger as shown in the second row of those two columns. In contrast, the first two rows of column (3) in Table 8 show that there was no change in participants' demand in response to changes in consumption-growth correlations regardless of whether they reported thinking about the parameter.

The third and fourth rows of Table 8 report results only for participants who said they considered a given parameter. The third row reports results for participants who stated they did not think about the parameter using textbook logic. The fourth row reports results for those who told us they thought about the parameter like a textbook investor would. In columns (1) and (2), the demand responsiveness to the mean and volatility of stock returns is stronger when participants told us that they were using textbook logic. Participants who said they were thinking about mean stock returns and trying to buy more stocks when this mean was higher did so, $4.06\% \gg 1.18\%$. Likewise, participants who said they were thinking about stock-return volatility and trying to buy less stock when this volatility was higher again followed through on their reported aims, $-1.17\% \ll 0.05\%$. By contrast, the demand responsiveness to correlation changes is nearly zero for all participants who told us they considered the parameter. The participants who told us they were trying to use textbook logic have demand that is indistinguishable from those who told us they were trying to take the opposite approach.

Table 9 further shows that the results are the same when we relabel *'economic growth'* as *'personal income'*, *'personal wealth'*, *'house prices'*, and *'industrial production'*. In every case, participants were more likely to report thinking about the mean and volatility of stock returns when forming their portfolio. Conditional on doing so, they were more likely to think about the mean and volatility of stock returns like a textbook investor should. Just like before, among the roughly 1/3 of participants who reported thinking about the risk-factor correlation at all, most did so using the opposite of textbook logic. Of the participants who thought about a risk-factor correlation at all, it was never the case that more than 20% thought about it using textbook logic. Put differently, 4 out of 5 participants who reported thinking about the correlation between stock returns and their income, wealth, house prices, or production when investing tried to buy more stocks when they were more correlated with this risk factor.

## 3.3 Reporting Correlations

If an investor from your favorite asset-pricing model were to peruse a popular financial-news source, he would be puzzled by the lack of information related to correlations with macroeconomic variables. An investor who cared about risk-factor correlations would presumably like this information to be presented as clearly as possible. Suppose that investors did not care only about correlations with consumption growth, but instead cared about both correlations with consumption growth as well as correlations with slow moving habit or with estimates of long-run risks. There is no reason why such correlations could not also be widely reported. From a revealed-preference perspective, it seems likely that real-world investors do not demand this information.

While puzzling, it could be that real-world investors do care about risk-factor correlations in spite of the fact that they are not widely reported. To examine whether this is plausible, we examined a variety of news sources to search for evidence that investors view risk-factor correlations as an important component of their investment decisions. Financial authorities provide educational documents to investors describing the risks that investors might consider relevant for their investments. We examined such documents from FINRA, the SEC, the Financial Conduct Authority, and the Ontario Security Commision.[8] While these documents list many specific risks, we were not able to identify any discussion about how correlations with potential risk factors were relevant. For example, in *'Investment Risk, Explained'* provided by FINRA, there are nine specific sources of risks listed for investors to consider. The discussions surrounding each of these risks pertain to uncertainty and volatility. There is no discussion of hedging, correlations, or risk factors.

If the more relevant group is professional investors, it is possible that such documents simply reflect the viewpoints of uneducated investors. Professional investors use a variety of risk assessment tools, such as those included in Bloomberg. While these products contain a number of options to assess various aspects of portfolio risk, none calculates correlations with macro risk factors as a default input.[9] If professional investors viewed such correlations as an important aspect of how they form their own portfolios or their clients' portfolios, it would be quite surprising that the tools they used lacked this basic feature.

To analyze a more systematic source of information, we examined mutual-fund prospectuses. A fund is required to report its investment objectives and risks regardless of how newsworthy these objectives and risks are. Funds also have discretion to highlight a variety of other potential aspects of the fund. For example, each Vanguard fund includes a *'plain talk'* section in its prospectus that attempts to explain investing concepts or strategies using straightforward language. Thus, if a fund thought that its correlation with an aggregate risk factor was an important component of investors' decision making, it could and should present information about this statistic in its prospectus. If a fund had correlations with macroeconomic risks that would make investors want to buy more of the fund, then it would likely say so in its prospectus to drive flows. If a fund's legal department believed that there was some possibility of being sued by an investor who viewed a correlation with a risk factor as a relevant risk, then including it in the list of potential risks would be an obvious

[8]For an example from each agency, see *Investment Risk, Explained* provided by FINRA (`https://www.finra.org/investors/insights/investment-risk`), *What is Risk* provided by the SEC (`www.investor.gov/introduction-investing/investing-basics/what-risk`), *Assessing suitability: Establishing the risk a customer is willing and able to take and making a suitable investment selection* provided by the FCA (`www.fca.org.uk/publication/finalised-guidance/fsa-fg11-05.pdf`), and *9 types of investment risk* from the Ontario Security Commission (`www.getsmarteraboutmoney.ca/invest/investing-basics/understanding-risk/types-of-investment-risk/`).

[9]We examined documentation from Bloomberg and Factset. We also discussed the details of these tools with a number of industry participants at a variety of different large financial firms. None reported any standard option in any of the available tools which would be relevant for hedging macro risk factors, such as aggregate consumption growth.

step to minimize liabilities from such a suit.

Mutual funds do not report these numbers or discuss their correlations with macro risk factors in their prospectuses. For example, the Vanguard 500 (VFIAX), which has nearly \$500$b$ in assets under management, stated its investment objective as tracking a benchmark index. It did not discuss any aggregate risk factors or its correlation with such variables. Under principle risks, the fund lists stock-market risk, "which is the chance that stock prices overall will decline," as well as investment-style risk, "which is the chance that returns from large-capitalization stocks will trail returns from the overall stock market," but it never talks about exposure to aggregate risk factors. There is a further discussion of risks later in the prospectus, but this is largely related to volatility: "stock markets tend to move in cycles, with periods of rising prices and periods of falling prices." Funds report a variety of statistics about their past performance such as fees, taxes, distributions, and performance. The Vanguard 500 fund's prospectus reports 171 numeric values in tables and figures, with even more values in the text. None of these numbers corresponds to the correlation between the fund's returns and a macro risk factor.

We systematically reviewed the mutual-fund prospectuses of the largest 25 US mutual funds, which jointly held about \$3.7$t$ at the time. Table 10 summarizes the results. We ranked US open-ended funds listed on Morningstar Direct as of July 30, 2019 based on their share-class asset value. We examined each fund's prospectus for five characteristics related to risk-factor correlations. Did a fund report a numerical value for the correlation between the fund's performance and any macroeconomic variable? Did a fund graph its performance with any macroeconomic variable? In the section on risks, did a fund list its return correlation with any macroeconomic variable other than the stock market itself? In the section on objectives, did a fund list an objective related to its correlation with a macroeconomic variable, such as exposure to an aggregate risk? Finally, we searched the text of each prospectus for the words *'covary'*, *'covariance'*, *'correlate'*, and *'correlation'*, counting the number of times these words appear in the document.

Mutual-fund prospectuses lack information about how a fund's returns covary with aggregate risk factors. None of the funds report numeric or graphical information relating their performance to macroeconomic fundamentals. No fund lists its correlation with macroeconomic outcomes as an investment risk, and no fund lists hedging an aggregate risk factor as an investment objective. Perhaps the closest thing we find is that some funds list reasons why the market may be volatile. For example, the Fidelity Contra Fund (FCNTX) warns its investors that "stock markets are volatile and can decline significantly in response to adverse issuer, political, regulatory, market, or economic developments." Notably, these are descriptions of why there may be volatility in returns not of how the fund's returns might be correlated with such variables. Our word search reveals that 22 of 25 funds fail to use any word related to correlation or covariance in their prospectus. The 3 prospectuses that do contain one of these words use them in a way that is unrelated to macroeconomic risk, as they reference the fund's tracking error or its relationship to derivative securities.

The underlying assumption in textbook finance theory is that investor are using assets' correlations with various risk factors to construct insurance against bad future states of the world. However, risk-factor correlations almost never get reported in the financial news, are not discussed in important fund documents, and are not used as the default settings in professional risk-analysis tools. It would be trivial for any of these sources to include such a factor, so their absence suggests that investors do not demand this information. It suggests that correlations are unlikely to represent a central explanation for asset-pricing movements as real world investors do not view them as relevant to their investment decisions.

# 4 Discussion

Textbook asset-pricing models make a core claim about investors' economic reasoning, and surveys are the ideal tool for evaluating this claim. Our results suggest that the canonical risk factor, consumption growth, is not relevant to investors' portfolio decisions. In this section, we aim to put this finding into the appropriate context. We address several potential concerns that naturally arise when using survey experiments to study financial markets. We outline the connection between our results and the existing finance survey literature. We describe how our results connect to the existing theoretical literature, and we describe what they imply for model development going forward.

## 4.1 Potential Concerns with Survey Experiments

There are several common concerns about using survey evidence to study financial markets. We have carefully designed our survey experiment with these concerns in mind. We outline and address these concerns in this section.

Are participants behaving differently in the experiment than they would in the real world? Yes. But, this concern is unlikely to explain our findings. First, we made it as easy as possible for participants to follow textbook asset-pricing logic. We removed any superfluous information. We presented difficult-to-find information about correlations in a straightforward manner. When using an online platform to survey finance professionals and MTurkers, we only examined participants who passed a multi-question comprehension test on the first try. Second, our participants consistently followed the textbook asset-pricing logic when responding to changes in the mean and volatility of stock returns. They only deviated from this logic for questions about the correlation between stock returns and consumption growth. While our experimental setting is different from the real world, for this to explain our findings it would have to impact only participants' responses to questions about correlations and not their responses to questions about means and volatilities. If anything, our straightforward provision of information about correlations should lead participants to be more likely to incorporate this parameter in their decision making.

Were participants simply confused? No. This explanation is unlikely to explain our results given that we provided participants with intuitive definitions of the mean stock return, stock-return volatility, and the correlation between stock returns and consumption growth. We then followed up on these definitions with a test for comprehension, eliminating any participants who did not get all questions correct on the first try. While we designed our experiment to be as straightforward as possible, there are clearly alternative ways one could present information that might lead to different results. The results in Tables 4, 5, and 9 suggest that slight changes to the question wording or the information display do not qualitatively affect our results.

Don't people form expectations in a biased way? Yes. But, our experiment purposefully removes this channel. There is a growing literature that uses surveys to examine how investors form beliefs about various economic parameters (see Section 4.2). Our survey design abstracts from such issues because we directly give participants the exact consumption-growth-correlation parameter value, leaving no need for them to form expectations. While the question of how people form expectations is important, the results of our survey experiment are not influenced by the answer to this question.

Do participants have the right incentives? Yes. Participants consistently act on and think about changes in the mean and volatility of stock returns in every part of our survey experiment. It is certainly true that the stakes are much higher in real-world markets than in our survey experiment. With that said, if the lack of financial incentives was distorting participant responses, then why is there no distortion in their answers to the questions about the mean and volatility of stock returns?

Are we surveying the right people? Probably. While our participant pools could be missing some important investors, they do capture a broad swath of the investing public. It is true that specific investor types might play a special role in financial markets. But, it is also true that these investors play no such role in textbook asset-pricing models. The CCAPM (Lucas, 1978) studies a representative investor as do popular models built on top of the CCAPM framework (Bansal and Yaron, 2004; Campbell and Cochrane, 1999). Every one of our participant pools yields identical results, and standard finance theory offers no explanation for why each of these participant pools should be ignored. If we are not surveying the right subset of investors, then economists are not modeling the right subset of investors, either.

Can we ignore the findings due to the Friedman critique? No. In its original form, the Friedman critique was not a critique at all. It was an explanation for why markets might move in ways that are more complicated than the ones captured by researchers' current models.[10] However, this logic has

[10]He states: "Consider the problem of predicting the shots made by an expert billiard player. It seems not at all unreasonable that excellent predictions would be yielded by the hypothesis that the billiard player made his shots *as if* he knew the complicated mathematical formulas. . . It is only a short step from these examples to the economic hypothesis that under a wide range of circumstances individual firm behave as if they were seeking rationally to maximize their expected returns." That being said, when you actually ask an advanced billiard player about what they are doing, you tend to get very complicated mathematical responses. e.g., there is a popular 1995 book on the subject called *The Science of Pocket Billiards* (Koehler, 1995). There is simply no substitute for asking

been broadly interpreted by some to imply that researchers can ignore investors' economic reasoning.

This reinterpretation goes too far. The central purpose of writing down an asset-pricing model is to explain why investors think asset $A$'s price should be low while asset $B$'s price should be high. Textbook models make a core claim about what investors are trying to do. It is the economic restrictions imposed by a model's preference specification, information structure, and budget constraint that make a model useful. These are the things that allow researchers to perform counterfactual analyses—i.e., to know what the prices of each asset would have been had market conditions been different. It is essential to accurately model the problem investors are trying to solve.

We fully recognize that all models require simplifying assumptions. A model as detailed as the thing being modeled would be useless.[11] With that said, if a model states that equilibrium asset prices are moving today because investors are trying to hedge their future exposure to $X$, the answer to the question *'Are investors trying to hedge their exposure to $X$?'* is not a nuance that can be ignored. It represents the core strategy that investors are supposed to be following. It is the very thing we are modeling; it is the thing we need a map to navigate in the first place.

Can we ignore the results because you can get survey respondents to say anything? No. While it is possible to generate spurious results using surveys, this is also true of other forms of empirical analysis and of theoretical models. You can support any result by pointing to bad research. You can lie with data, but that does not mean there are no reliable econometric results. It is possible to make vacuous theoretical arguments, but that does not mean all of theory is without merit. Surveys are no different. They are one more approach to learning about financial markets. There is an active literature studying how to design surveys which we cannot do justice to here. We have discussed at length the reasoning behind the survey design choices we make (see Section 2.2). A reader who disagrees with the design can and should use our framework to implement a different design (within the best practices of survey design) to demonstrate the merit of the concern.

Are surveys the right tool to examine risk-factor relevance? Yes. Survey experiments are an ideal tool for asking questions about risk-factor relevance. This is because, according to the logic of textbook asset-pricing models, a priced risk factor must be generally understood by investors. Priced risk factors cannot be a surprise to everyone in the market. An asset's equilibrium price can only contain an insurance premium if investors generally agree that it provides an insurance value worth paying for. If asset prices are moving because investors are trying to hedge their future exposure to $X$, then investors must be aware of each asset's exposure to $X$, and they should be able to explain this reasoning to researchers.

people what they are doing; this is true of both investors and pool sharks.

[11] Lewis Carroll highlighted the problem in his 1894 novel, *Sylvie And Bruno Concluded*, when he described the obvious problem with using an ultimate map, representing the world on a unitary scale, a mile to a mile: "It has never been spread out, yet. The farmers objected. They said it would cover the whole country, and shut out the sunlight."

There are clearly cases in financial markets where this is exactly what is going on. For example, "commodity futures markets have had a long history of assisting commodity producers to hedge their commodity-price risks (Cheng and Xiong, 2014)." Further, "exchange rates are a major source of uncertainty for multinationals (Jorion, 1990)" and FX forward markets exist so firms can hedge this risk. Likewise, "sovereign CDS contracts function as insurance contracts that allow investors to buy protection against the event that a sovereign defaults (Longstaff et al., 2011)." In all these examples, changes in risk-factor exposure affect the equilibrium price of a contract.

If someone were to ask a commodity-futures trader or FX-market investor about his decision, the investor would likely be able to state the relevant risk in a way that was consistent with textbook asset-pricing logic. For example, investors trading oil-futures are aware that today's price is likely determined by exposure to future oil shocks. When the CEO of Southwest airlines was asked about why they were active in the oil-futures market, he stated that the company "loaded up years ago on hedges against higher fuel prices."[12] This risk is commonly understood, which is why it is priced.

## 4.2 Survey Evidence

This paper was inspired by a burgeoning literature that uses survey evidence to study financial markets. We contribute to this literature by showing how to use such evidence in a new way. Rather than using surveys to estimate investors' beliefs about a particular parameter that researchers already know to be important, we use surveys to test whether investors view a parameter as relevant to their portfolio decisions in the first place.

As pointed out by Adam et al. (2020), survey evidence "allows researchers to consider alternatives to the rational-expectations assumption in an empirically disciplined way." Amromin and Sharpe (2013) explores how investors use current economic conditions to inform their stock-market expectations. Armona et al. (2018) looks at how home buyers form expectations about future house-price appreciation. Greenwood and Shleifer (2014) give evidence that investors form expectations by over-extrapolating. Giglio et al. (2019) studies how the trading behavior of institutional investors is related to their return expectations and past market experience. And, Liu et al. (2020) uses survey evidence to study the reasons for excess trading in the stock market.

Although clearly related, we are using surveys with a slightly different goal in mind. We are not trying to estimate investors' beliefs about a particular parameter, such as expected returns. In fact, we are providing our survey participants with all relevant information so there is no need for them to construct such expectations. Instead, we are looking for evidence that a particular parameter—an asset's correlation with consumption growth—plays any role in investors' decisions.

Choi and Robertson (2019) examine a related, though distinct, question. The authors asked investors whether they considered a wide range of theories to be important for their own financial

[12] Bailey, J. "Southwest Airlines gains advantage by hedging long-term oil contracts." *New York Times.* Nov 28, 2007.

decision making. This is an important and generative idea that has garnered substantial interest in the field. However, their survey design does not aim to estimate the zero effect needed to test a specific risk factor's relevance. Our approach does, using multiple methods to provide converging evidence for a specific risk-factor's relevance.

To illustrate, when we asked participants about their economic reasoning, we found that 43% reported thinking about the correlation between stock returns and consumption growth. This number is qualitatively similar to the 29% of participants in the Choi and Robertson (2019) study who stated that the return covariance with consumption was *'very important'* or *'extremely important'*. However, when we asked our participants a follow-up question about how they used this information, we found that roughly 3 out of 4 people who considered consumption-growth correlations were trying to buy *more* stocks when this correlation was higher, inconsistent with the logic of the CCAPM. Furthermore, when we examined the 11% of participants who reported using textbook logic, we found no evidence that their portfolio decisions reflected this economic reasoning.

## 4.3 Centrality of Consumption Hedging

The framework we develop in this paper can be used to evaluate the relevance of any proposed risk factor. To demonstrate the framework, however, we had to select a specific risk factor for our case study. We chose consumption growth as our main variable of interest. In this section, we discuss why our results pertaining to this specific variable also have implications for most other textbook models.

Consumption growth is the sole state variable of the CCAPM. The current approach to addressing the empirical failings of the CCAPM is to introduce new risk factors to this model. The idea is that, if exposure to consumption growth cannot entirely explain why markets fluctuate, then exposure to consumption growth and additional state variables can. These theories work by amplifying the importance of differences in risk-factor correlations. As discussed below, even though these models introduce additional state variables to downplay the importance of consumption-hedging motives, it is still crucial to these models that investors are trying to insure themselves against future negative shocks to consumption. Contrary to what we find, an agent in these models should still respond to our framework in a manner consistent with textbook logic.

### 4.3.1 External Habit

To illustrate, consider how the results in this paper apply to the Campbell and Cochrane (1999) external-habit model. This model studies a representative investor trying to maximize a modified utility function, $U_t \overset{\text{def}}{=} \frac{(C_t - H_t)^{1-\gamma} - 1}{1-\gamma}$. In this model, investors have power utility over consumption in excess of a slow-moving benchmark, $H_t$, which represents the level of consumption that investors have become accustomed to—i.e., habit. Formally, $\log H_t \overset{\text{def}}{=} \lambda \cdot \sum_{\ell=0}^{\infty} \phi^{\ell} \cdot \log C_{t-\ell}$ where $\lambda > 0, \phi \in (0,1)$. Thus, drops in consumption following booms are more painful for investors in this model.

The key new state variable in this model is an investors' surplus consumption ratio given by $X_t \overset{\text{def}}{=} \frac{C_t - H_t}{C_t}$. Growth in the model-implied stochastic discount factor (SDF) is then given by $\Delta \log M_{t+1} = \log \beta - \gamma \cdot \Delta \log C_{t+1} - \gamma \cdot \Delta \log X_{t+1}$. In what initially looks like a nod to Merton (1973) ICAPM logic, the model seems to say that average stock returns could be high either because they covary with consumption growth or because they covary with growth in the surplus-consumption ratio:

$$\mathrm{E}[R_{t+1}] - R_f \approx \gamma \times \mathrm{Cov}[\Delta \log C_{t+1}, R_{t+1}] + \gamma \times \mathrm{Cov}[\Delta \log X_{t+1}, R_{t+1}]$$

But, notice that *'either/or'* is not actually the right conjunction to use here though. The second term on the right hand side is not independent of the first. If investors are not trying to insure themselves against drops in consumption (if the first term is zero), they cannot be trying to insure themselves against drops in the surplus-consumption ratio either (the second term must be zero).[13]

This means that surplus consumption is not a separate risk factor. It is a way of amplifying investors' desire to hedge consumption-growth shocks. Note expected returns in Campbell and Cochrane (1999) can be rewritten as:

$$\mathrm{E}[R_{t+1}] - R_f \approx \tfrac{\gamma}{X_t} \times \mathrm{Cov}[\Delta \log C_{t+1}, R_{t+1}]$$

In other words, habit is a model where investors have time-varying risk aversion with respect to consumption shocks. Since the average surplus-consumption ratio is $\mathrm{E}[X_t] \approx 0.5$, this means that variation in risk aversion due to habit roughly doubles the risk premium investors demand for their exposure to consumption shocks. Investors in a habit model must still want to insure themselves against shocks to consumption otherwise the model will not work.

### 4.3.2 Long-Run Risk

The Bansal and Yaron (2004) long-run-risk model uses a different preference specification and state variable, but the result is the same. The effects of shocks to the new state variable on asset prices cannot exist if investors do not try to hedge their exposure to consumption-growth shocks.

The long-run-risk model assumes that investors have Epstein and Zin (1989) recursive preferences over current and future consumption. These preferences introduce an additional slow-moving state variable just like in the habit model. It is again the case that, if investors are not trying to insure themselves against future drops in consumption, then they cannot be trying to insure themselves against the new risk factor either. If there is no desire to hedge

[13] To see why, consider rewriting the first term as $\gamma \times \mathrm{Cov}[\Delta \log C_{t+1}, R_{t+1}] = \gamma \times (\rho \cdot \sigma_{\Delta \log C} \cdot \sigma_R)$. For changes in $\rho$ to have no effect on expected returns via this first term, $\partial_\rho\big(\gamma \times \rho \cdot \sigma_{\Delta \log C} \cdot \sigma_R\big) = 0$, it must be that at least one of the following conditions holds: $\gamma = 0$, $\sigma_{\Delta \log C} = 0$, or $\sigma_R = 0$. But, when you inspect the definition of habit, $H_t$, it is clear that $\sigma_{\Delta \log C} = 0$ implies that $\sigma_{\Delta \log X} = 0$. So, if any of these three conditions holds, the second term must be zero as well, $\gamma \times \mathrm{Cov}[\Delta \log X_{t+1}, R_{t+1}] = \gamma \times (\varrho \cdot \sigma_{\Delta \log X} \cdot \sigma_R) = 0$.

consumption-growth shocks, then there is no long-run risk either (Epstein et al., 2014). Another way of making this same point is to notice that the long-run-risk model is formally equivalent to a model where investors are ambiguity averse with respect to parameters of the consumption-growth process (Hansen and Sargent, 2008; Epstein and Schneider, 2010; Bidder and Dew-Becker, 2016).

#### 4.3.3 Rare Disasters

There is a class of models built on top of the CCAPM framework that our results do not directly speak to: rare-disaster models à la Rietz (1988), Barro (2006), and Gabaix (2012). We find that investors are not trying to insure themselves against normal-times variation in consumption growth, but this is not a direct test of whether asset prices are moving because investors want to insure themselves against extreme shocks to consumption.

That being said, even if a rare-disasters model fits the data, it is not obvious that real-world investors think along these lines. For example, Giglio et al. (2019) finds that investors who believe a future disaster is more likely do not believe that average returns will be higher to compensate for this additional risk—i.e., they do not believe in a disaster risk premium. The assumption that investors are trying to insure themselves against rare disasters is testable, and researchers should use the framework we develop in this paper to test it.

#### 4.3.4 Heterogeneous Agents

Most of the discussion in our paper (and in the literature) relates to representative-agent models, but our paper also has implications for heterogeneous-agent models. For example, in the Constantinides and Duffie (1996) model, heterogeneous investors try to hedge shocks to their own personal income. When this hedging demand is aggregated, the model predicts that aggregate consumption growth looks like a priced risk factor. We directly show in Tables 4 and 9 that participants are not trying to hedge shocks to their own personal income.

But, suppose we had found the opposite results. Would such a finding suggest that consumption growth is a relevant risk factor? No. It would suggest that shocks to personal income are relevant. Shocks to aggregate consumption and shocks to personal income are two different things. Even if these two kinds of shocks sometimes coincide, researchers can still imagine future scenarios where they do not. Economists should care about economic mechanism behind a link between average returns and correlations with $X$ for the same reasons that labor economists care about which economic mechanism explains the link between going to college and earning more money.

Further, investors likely differ along a variety of dimensions, whether it be preferences, wealth, expectations, or any number of alternatives. While models that focus on these various dimensions have been proposed, most heterogeneous agent models retain risk-factor hedging, be it aggregate or personal, as the core strategy that investors pursue. Our results test this mechanism and suggest it is unlikely to accurately represent how investors approach portfolio decisions.

## 4.4 Model Development

The results in this paper shed light on how researchers can both test new models and respond to the empirical shortcomings of existing models. First, researchers can use our survey-based framework to generate and test new models for which empirical datasets may not be readily available. e.g., researchers promoting a behavioral bias or trading friction can apply our framework to address concerns about competing risk-based explanations.

In addition, the uniform results across the various survey questions and participant pools suggest that investors' desire to insure themselves against consumption shocks is unlikely to explain why asset prices move. This suggests subsequent models that add complexity to this foundation are also unlikely to do so. Suppose our results had suggested otherwise. Suppose they indicated that participants strongly responded to the correlation between stock returns and consumption growth consistent with textbook models. Further suppose they showed that participants reported thinking about an asset's correlation with consumption growth as suggested by the CCAPM. Such results would not solve the CCAPM's empirical shortcomings, but they would support the current approach to solving them. The results would have implied that investors cared about consumption-growth correlations in a more complicated way than is captured by the CCAPM. Thus, models introducing new risk-factor correlations to the CCAPM, such as habit (Campbell and Cochrane, 1999) and long-run risk (Bansal and Yaron, 2004), would have a strong foundation to build on. This is not what we find, though.

This does not mean that the CCAPM is a bad model. Even if a factor model does not explain how the world actually works, it can still make important normative prescriptions. For example, it could be that more people should be hedging shocks to consumption growth even if they currently are not doing so. Demonstrating to investors a behavior that they should adopt could have important welfare implications.

Finally, our results do not suggest that investors are irrational. They suggest that investors are not trying to hedge shocks to consumption growth. Even though the CCAPM is the starting point for textbook asset-pricing theory, real-world investors are not irrational for not following the CCAPM's logic. If investors at a conference run by a large asset manager are not trying to hedge consumption growth, it does not mean that these investors are doing the wrong thing. It means that economists are using the wrong model.

# 5 Conclusion

If $X$ is a priced risk factor—if the hypothesis that asset prices are moving because investors are trying to insure themselves against bad future shocks to $X$ is true—then investors must be <u>*trying*</u> to insure themselves against bad future shocks to $X$. This goal must be something investors are generally aware of. This is a necessary condition that must hold for any $X$ to be a relevant risk factor. We develop a framework for testing it using both survey experiments and field data.

We apply this framework to evaluate the relevance of aggregate consumption growth, the canonical risk factor in the academic literature. We find no evidence that investors prefer assets with lower consumption-growth correlations or adjust their portfolio holdings based on this reasoning.

Going forward, when a researcher proposes a new risk factor, $X$, we should ask for evidence that investors actually think and trade based on the logic of their model. Economists typically apply a strict hierarchy of explanations. If there is a risk-based explanation for an asset-pricing phenomenon, then this explanation is viewed as correct even if there are other more parsimonious models. Our results cast doubt on this hierarchy. They suggest that economists should be more skeptical of risk-based explanations in the absence of supporting evidence that the model actually captures how investors price assets.

# References

Adam, K., D. Matveev, and S. Nagel (2020). Do survey expectations of stock returns reflect risk-adjustments? *Journal of Monetary Economics*.

Amromin, G. and S. Sharpe (2013). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science*.

Ang, A. (2014). *Asset management: A systematic approach to factor investing*. Oxford University Press.

Armona, L., A. Fuster, and B. Zafar (2018). Home-price expectations and behavior: Evidence from a randomized-information experiment. *Review of Economic Studies*.

Asparouhova, E., P. Bossaerts, N. Roy, and W. Zame (2016). "lucas" in the laboratory. *Journal of Finance*.

Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset-pricing puzzles. *Journal of Finance*.

Barber, B. and T. Odean (2007). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*.

Barro, R. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*.

Bergman, A., A. Chinco, S. Hartzmark, and A. Sussman (2020). Survey curious? start-up guide and best practices for online survey experiments. *Working Paper*.

Bidder, R. and I. Dew-Becker (2016). Long-run risk is the worst-case scenario. *American Economic Review*.

Bossaerts, P. and C. Plott (2004). Basic principles of asset-pricing theory: Evidence from large-scale experimental financial markets. *Review of Finance*.

Bryzgalova, S. (2017). Spurious factors in linear asset-pricing models. *Working Paper*.

Campbell, J. (2003). Consumption-based asset pricing. *Handbook of the Economics of Finance*.

Campbell, J. and J. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock-market behavior. *Journal of Political Economy*.

Campbell, J. and L. Viceira (1999). Consumption and portfolio decisions when expected returns are time varying. *Quarterly Journal of Economics*.

Carroll, L. (1894). *Sylvie and Bruno concluded.* New York; London: Macmillan.

Casler, K., L. Bickel, and E. Hackett (2013). Separate but equal? A comparison of participants and data gathered via Amazon's MTurk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*.

Cheng, I.-H. and W. Xiong (2014). Financialization of commodity markets. *Annual Review of Financial Economics*.

Chinco, A., A. Neuhierl, and M. Weber (2020). Estimating the anomaly base rate. *Journal of Financial Economics*.

Choi, J. and A. Robertson (2019). What matters to individual investors? Evidence from the horse's mouth. *Journal of Finance*.

Cochrane, J. (1999). Portfolio advice for a multi-factor world. *Chicago Federal Reserve, Economic Perspectives*.

Cochrane, J. (2001). *Asset pricing.* Princeton University Press.

Cochrane, J. (2017). Macro-finance. *Review of Finance*.

Constantinides, G. and D. Duffie (1996). Asset pricing with heterogeneous consumers. *Journal of Political Economy*.

Crockett, S., J. Duffy, and Y. Izhakian (2019). An experimental test of the Lucas asset-pricing model. *Review of Economic Studies*.

Di Maggio, M., A. Kermani, and K. Majlesi (2020). Stock-market returns and consumption. *Journal of Finance*.

Enke, B. and F. Zimmermann (2017). Correlation neglect in belief formation. *Review of Economic Studies*.

Epstein, L., E. Farhi, and T. Strzalecki (2014). How much would you pay to resolve long-run risk? *American Economic Review*.

Epstein, L. and M. Schneider (2010). Ambiguity and asset markets. *Annual Review of Financial Economics*.

Epstein, L. and S. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica*.

Eyster, E. and G. Weizsacker (2016). Correlation neglect in portfolio choice: Lab evidence. *Working Paper*.

Feng, G., S. Giglio, and D. Xiu (2020). Taming the factor zoo. *Journal of Finance*.

Freyberger, J., A. Neuhierl, and M. Weber (2020). Dissecting characteristics non-parametrically. *Review of Financial Studies*.

Friedman, M. (1953). *Essays in positive economics*. University of Chicago Press.

Frydman, C., S. Hartzmark, and D. Solomon (2017). Rolling mental accounts. *Review of Financial Studies*.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics*.

Gabaix, X. (2016). Behavioral macroeconomics via sparse dynamic programming. *Working Paper*.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2019). Five facts about beliefs and portfolios. *Working Paper*.

Gode, D. and S. Sunder (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*.

Greenwood, R. and A. Shleifer (2014). Expectations of returns and expected returns. *Review of Financial Studies*.

Hansen, L. and T. Sargent (2008). *Robustness*. Princeton University Press.

Harris, L., S. Hartzmark, and D. Solomon (2015). Juicing the dividend yield: Mutual funds and the demand for dividends. *Journal of Financial Economics*.

Hartzmark, S. (2014). The worst, the best, ignoring all the rest: The rank effect and trading behavior. *Review of Financial Studies*.

Hartzmark, S. and D. Solomon (2017). The dividend disconnect. *Journal of Finance*.

Hartzmark, S. and A. Sussman (2019). Do investors value sustainability? A natural experiment examining ranking and fund flows. *Journal of Finance*.

Harvey, C. and Y. Liu (2020). False (and missed) discoveries in financial economics. *Journal of Finance*.

Hauser, D. and N. Schwarz (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject-pool participants. *Behavior Research Methods*.

Huber, J., S. Palan, and S. Zeisberger (2019). Does investor risk perception drive asset prices in markets? Experimental evidence. *Journal of Banking & Finance*.

Jennings, D., T. Amabile, and L. Ross (1982). Informal covariation assessment: Data-based vs. theory-based judgements. *Judgment Under Uncertainty: Heuristics and Biases*.

Jorion, P. (1990). The exchange-rate exposure of US multinationals. *Journal of Business*.

Kahneman, D. and S. Frederick (2002). Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*.

Kallir, I. and D. Sonsino (2009). The neglect of correlation in allocation decisions. *Southern Economic Journal*.

Koehler, J. (1995). *The science of pocket billiards*. Sportology Publications.

Kozak, S., S. Nagel, and S. Santosh (2018). Shrinking the cross-section. *Journal of Financial Economics*.

Kroll, Y. and H. Levy (1992). Further tests of the separation theorem and the capital asset-pricing model. *American Economic Review*.

Kroll, Y., H. Levy, and A. Rapoport (1988). Experimental tests of the separation theorem and the capital asset-pricing model. *American Economic Review*.

Laudenbach, C., M. Ungeheuer, and M. Weber (2019). How to alleviate correlation neglect. *Working Paper*.

Levy, G. and R. Razin (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review*.

Lewellen, J., S. Nagel, and J. Shanken (2010). A skeptical appraisal of asset-pricing tests. *Journal of Financial Economics*.

Liu, H., C. Peng, W. Xiong, and W. Xiong (2020). Resolving the excessive trading puzzle: An integrated approach based on surveys and transactions. *Working Paper*.

Longstaff, F., J. Pan, L. Pedersen, and K. Singleton (2011). How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*.

Lucas, R. (1978). Asset prices in an exchange economy. *Econometrica*.

Matthies, B. (2018). Biases in the perception of covariance. *Working Paper*.

McLean, D. and J. Pontiff (2016). Does academic research destroy stock-return predictability? *Journal of Finance*.

Mehra, R. and E. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics*.

Merton, R. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*.

Merton, R. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*.

Merton, R. (1973). An intertemporal capital asset-pricing model. *Econometrica*.

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*.

Paolacci, G. and J. Chandler (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science*.

Plott, C. and S. Sunder (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica*.

Rietz, T. (1988). The equity risk-premium: A solution. *Journal of Monetary Economics*.

Savov, A. (2011). Asset pricing with garbage. *Journal of Finance*.

Schwarz, N. (1999). Self-reports: How the questions shape the answers. *American Psychologist*.

Smith, V., G. Suchanek, and A. Williams (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*.

Ungeheuer, M. and M. Weber (2019). The perception of dependence, investment decisions, and stock prices. *Journal of Finance*.

# A Figures

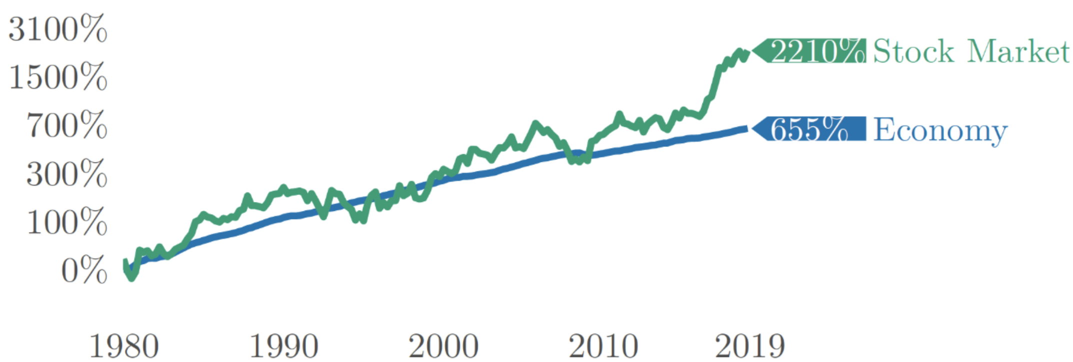


| | Economy | Stock Market |
|---|---|---|
| Average annual growth rate: | 5.2% | 8.0% |
| Average level of uncertainty: | 1.6% | 20.0% |
| Correlation between economic growth and the stock market: | 0.15 | |

Given the information presented above, how would you allocate your $1,000 in savings between a mutual fund that invests in a broad blue-chip stock market index and a bond earning 2%? (*must total $1,000)*

| | |
|---|---|
| Mutual Fund | 0 |
| Bond | 0 |
| Total | 0 |

**Figure 1a. Sample Investment-Decision Question without noise.** This figure shows a sample question about investment decisions from the first part of our survey experiment. The economic growth line is GDP without noise added.

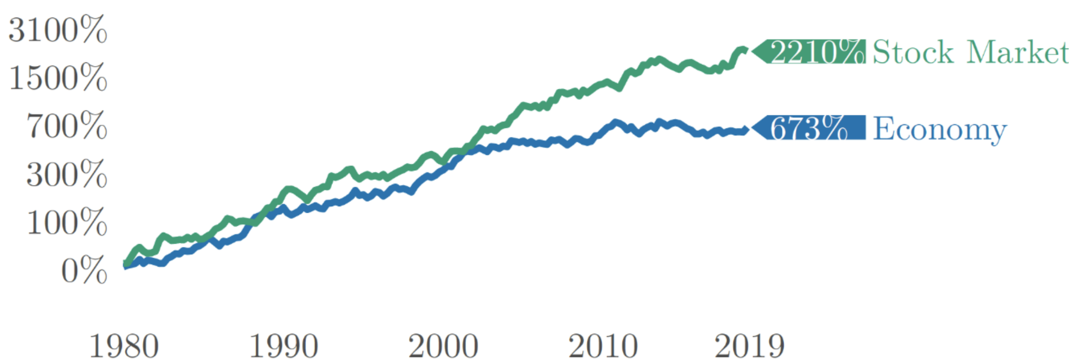


| | Economy | Stock Market |
|---|---|---|
| Average annual growth rate: | 5.2% | 8.0% |
| Average level of uncertainty: | 8.1% | 10.0% |
| Correlation between economic growth and the stock market: | 0.00 | |

Given the information presented above, how would you allocate your $1,000 in savings between a mutual fund that invests in a broad blue-chip stock market index and a bond earning 2%? (*must total $1,000)*

| | |
|---|---|
| Mutual Fund | 0 |
| Bond | 0 |
| Total | 0 |

**Figure 1b. Sample Investment-Decision Question with Noise.** This figure shows a sample question about investment decisions from the first part of our survey experiment. The economic growth line is GDP with noise added.

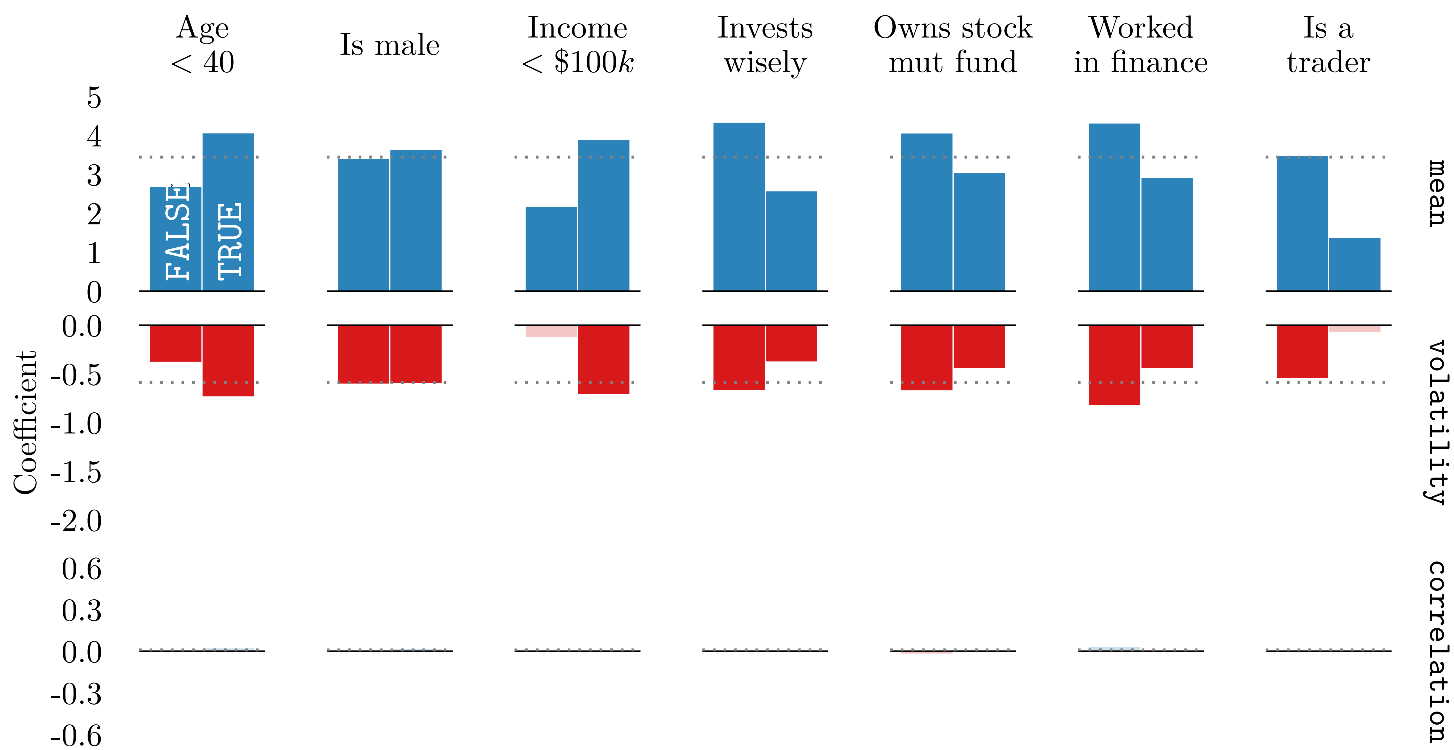


**Figure 2. Investment Decisions by Participant Characteristics.** This figure reports regression results corresponding to column (4) in Table 2 for different subsets of our participant pools. From top to bottom, each set of three bars represents the slope coefficients from the regression $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{mean}_{i,q} + \hat{\gamma} \cdot \texttt{volatility}_{i,q} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. The right-hand-side variables correspond to the average stock returns, $\texttt{mean}_{i,q}$, stock-return volatility, $\texttt{volatility}_{i,q}$, and the correlation between stock returns and consumption growth, $\texttt{correlation}_{i,q}$, used to simulate the data for each question. The y-axis is scaled so that a change in the given parameter of low to high would match the scale of the y-axis for the mean graphs. Opaque bars are significant at the 5% level using standard errors clustered by participant. Transparent bars are insignificant. Blue bars denote positive values. Red bars denote negative values. The horizontal dotted gray lines correspond to coefficient values from Table 2 column (4).

# B Tables

| a) Finance Professionals ($N$ = 493) | # | Avg | Sd | Min | Med | Max |
|---|---|---|---|---|---|---|
| Stock Fraction | | 0.56 | 0.24 | 0.00 | 0.50 | 1.00 |
| Age < 40 | 221 | 0.45 | | | | |
| Is male | 215 | 0.44 | | | | |
| Income < $\$100k$ | 217 | 0.44 | | | | |
| Owns stock or mutual funds | 442 | 0.90 | | | | |
| Is a trader | 136 | 0.28 | | | | |

| b) MTurkers ($N$ = 322) | | | | | | |
|---|---|---|---|---|---|---|
| Stock Fraction | | 0.58 | 0.27 | 0.00 | 0.60 | 1.00 |
| Age < 40 | 232 | 0.72 | | | | |
| Is male | 210 | 0.65 | | | | |
| Income < $\$100k$ | 281 | 0.87 | | | | |
| Own stocks or mutual funds | 209 | 0.65 | | | | |
| Worked in finance | 29 | 0.09 | | | | |

| c) MBA Students ($N$ = 308) | | | | | | |
|---|---|---|---|---|---|---|
| Stock Fraction | | 0.67 | 0.30 | 0.00 | 0.70 | 1.00 |
| Age < 40 | 304 | 0.99 | | | | |
| Is male | 183 | 0.59 | | | | |
| Worked in finance | 118 | 0.38 | | | | |

| d) Asset Managers ($N$ = 93) | | | | | | |
|---|---|---|---|---|---|---|
| Stock Fraction | | 0.67 | 0.25 | 0.00 | 0.70 | 1.00 |

**Table 1. Summary Statistics.** This table presents summary statistics describing the four participant pools in our survey experiment. Stock Fraction: average fraction of their endowment that each participant invests in stocks; computed using data at the participant×question level. All remaining rows are computed at the participant level. #: represents the number of participants who answered *'Yes'*. Asset managers were not asked background questions prior to taking our survey due to time constraints.

Dependent Variable: $\texttt{stockFrac}_{i,q}$

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| $\texttt{mean}_{i,q}$ | 3.24*** | | | 3.22*** | | | |
| | (14.52) | | | (14.51) | | | |
| $\texttt{volatility}_{i,q}$ | | −0.61*** | | −0.60*** | | | |
| | | (10.07) | | (9.97) | | | |
| $\texttt{correlation}_{i,q}$ | | | 0.00 | 0.00 | 0.01 | 0.00 | 0.01 |
| | | | (0.00) | (0.16) | (0.69) | (0.06) | (0.68) |
| Participant FE | | | | | ✓ | | ✓ |
| Question FE | | | | | | ✓ | ✓ |
| # Obs | 10,062 | 10,062 | 10,062 | 10,062 | 10,062 | 10,062 | 10,062 |
| Adj. $R^2$ | 0.04 | 0.01 | 0.00 | 0.05 | 0.43 | 0.01 | 0.44 |

**Table 2. Investment Decisions.** This table shows how participants' investment decisions vary with average stock returns, $\texttt{mean}_{i,q}$, stock-return volatility, $\texttt{volatility}_{i,q}$, and the correlation between stock returns and consumption growth, $\texttt{correlation}_{i,q}$. This table uses observations on all participant pools. Each column reports the results of a different regression of the form $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{mean}_{i,q} + \hat{\gamma} \cdot \texttt{volatility}_{i,q} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. Columns (5) and (7) include participant fixed effects. Columns (6) and (7) include question-order fixed effects. The numbers in parentheses are $t$-stats computed using standard errors clustered by participant. $^{\star}$, $^{\star\star}$, and $^{\star\star\star}$ indicate statistically significant estimates at the 10%, 5%, and 1% levels.

Dependent Variable: $\texttt{stockFrac}_{i,q}$

| | All Participants | | Finance Professionals | MTurkers | MBA Students | Asset Managers |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $\texttt{correlation}_{i,q}$ | 0.01 | 0.00 | 0.00 | 0.00 | 0.07$^{\star}$ | −0.03 |
| | (0.69) | (0.33) | (0.11) | (0.12) | (1.88) | (0.38) |
| Participant FE | ✓ | | ✓ | ✓ | ✓ | ✓ |
| Pool FE | | ✓ | | | | |
| # Obs | 10,062 | 10,062 | 4,930 | 3,220 | 1,540 | 372 |
| Adj. $R^2$ | 0.43 | 0.02 | 0.37 | 0.43 | 0.51 | 0.38 |

**Table 3. Investment Decisions by Participant Pool.** This table shows how the investment decisions of different types of participants change in response to an asset's return correlation with consumption growth, $\texttt{correlation}_{i,q}$. Each column reports the results of a different regression of the form $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. Columns (1) and (2) report results using all participant pools. Column (3) reports results only for finance professionals. Column (4) reports results only for MTurkers. Column (5) reports results only for MBA students. And, column (6) reports results only for asset managers attending the investor conference. Columns (1), (3), (4), (5), and (6) include participant fixed effects. Column (2) includes pool fixed effects. The numbers in parentheses are $t$-stats computed using standard errors clustered by participant. $^{\star}$, $^{\star\star}$, and $^{\star\star\star}$ indicate statistically significant coefficient estimates at the 10%, 5%, and 1% levels.

Dependent Variable: $\texttt{stockFrac}_{i,q}$

| | Personal Income (1) | Personal Wealth (2) | House Prices (3) | Industrial Production (4) |
|---|---|---|---|---|
| $\texttt{mean}_{i,q}$ | $6.34^{\star\star\star}$ | $4.35^{\star\star\star}$ | $2.82^{\star\star\star}$ | $2.71^{\star\star\star}$ |
| | (10.15) | (8.16) | (4.92) | (5.13) |
| $\texttt{volatility}_{i,q}$ | $-0.71^{\star\star\star}$ | $-0.51^{\star\star\star}$ | $-0.75^{\star\star\star}$ | $-0.15$ |
| | (4.40) | (3.37) | (4.55) | (0.82) |
| $\texttt{correlation}_{i,q}$ | $-0.02$ | 0.04 | 0.04 | 0.06 |
| | (0.35) | (1.15) | (0.97) | (1.45) |
| # Obs | 1,110 | 1,860 | 1,340 | 1,700 |
| Adj. $R^2$ | 0.16 | 0.08 | 0.05 | 0.03 |

**Table 4. Investment Decisions Using Different Terms.** This table shows how participants' investment decisions vary with average stock returns, $\texttt{mean}_{i,q}$, stock-return volatility, $\texttt{volatility}_{i,q}$, and the correlation between stock returns and a risk factor, $\texttt{correlation}_{i,q}$, when that variable is labeled as *'personal income'*, *'personal wealth'*, *'house prices'* or *'industrial production'*. Each column reports the results of a different regression of the form $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{mean}_{i,q} + \hat{\gamma} \cdot \texttt{volatility}_{i,q} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. Each column uses data on a separate population of MTurkers. Numbers in parentheses are $t$-stats computed using standard errors clustered by participant. $^\star$, $^{\star\star}$, and $^{\star\star\star}$ indicate statistically significant estimates at the 10%, 5%, and 1% levels.

Dependent Variable: $\texttt{stockFrac}_{i,q}$

| | Text Only | Additional $\rho$ Instructions | Scatterplots | No Graphs |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| $\texttt{mean}_{i,q}$ | $4.11^{\star\star\star}$ | $3.52^{\star\star\star}$ | $1.31^{\star\star\star}$ | $1.93^{\star\star\star}$ |
| | (6.52) | (6.82) | (2.89) | (4.64) |
| $\texttt{volatility}_{i,q}$ | $-0.58^{\star\star\star}$ | $-0.71^{\star\star\star}$ | $-0.74^{\star\star\star}$ | $-0.37^{\star\star}$ |
| | (3.35) | (4.51) | (3.99) | (2.26) |
| $\texttt{correlation}_{i,q}$ | 0.05 | 0.01 | $0.15^{\star\star\star}$ | $0.18^{\star\star\star}$ |
| | (1.29) | (0.18) | (3.85) | (4.61) |
| # Obs | 1,290 | 1,800 | 1,520 | 1,870 |
| Adj. $R^2$ | 0.07 | 0.06 | 0.03 | 0.03 |

**Table 5. Investment Decisions Using Treatment Variations.** This table shows how participants' investment decisions vary when we provided them with information about the correlation between stock returns and consumption growth in a different format. Each column reports the results of a different regression of the form $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{mean}_{i,q} + \hat{\gamma} \cdot \texttt{volatility}_{i,q} + \hat{\delta} \cdot \texttt{correlation}_{i,q} + \hat{\varepsilon}_{i,q}$. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. In column (1), participants saw $\rho \in \{$*'none'*, *'low'*, *'medium'*, *'high'*$\}$ reported as text rather than $\rho \in \{0.00, 0.15, 0.30, 0.45\}$. In column (2), participants received additional instruction on how to interpret a correlation coefficient. In column (3), participants saw a scatterplot rather than a time-series plot. In column (4), participants saw the numeric values of the mean, volatility, and correlation but no time-series graphs. Each column uses data on a separate population of MTurkers. Numbers in parentheses are $t$-stats computed using standard errors clustered by participant. $^\star$, $^{\star\star}$, and $^{\star\star\star}$ indicate statistically significant estimates at the 10%, 5%, and 1% levels.

| | Mean $\mu_R$ (1) | Volatility $\sigma_R$ (2) | Correlation $\rho$ (3) |
|---|---|---|---|
| $\Pr[\texttt{consider}_i]$ | 0.77 [0.04] | 0.59 [0.07] | 0.43 [0.03] |
| $\Pr[\texttt{textbookLogic}_i]$ | 0.58 [0.05] | 0.44 [0.08] | 0.11 [0.01] |
| $\Pr[\,\texttt{textbookLogic}_i \,\vert\, \texttt{consider}_i = \mathit{true}\,]$ | $0.76^{+++}$ [0.02] | $0.72^{+++}$ [0.03] | $0.24^{---}$ [0.03] |

**Table 6. Economic Reasoning.** This table depicts the rate at which participants reported thinking about average stock returns ($\mu_R$), stock-return volatility ($\sigma_R$), and the correlation between stock returns and consumption growth ($\rho$) when making their investment decisions. $\texttt{consider}_i$ is an indicator variable for whether the $i$th participant thought about a parameter at all. $\texttt{textbookLogic}_i$ is an indicator variable for whether the $i$th participant thought about this parameter using textbook asset-pricing logic. If $\texttt{consider}_i = \mathit{false}$ then $\texttt{textbookLogic}_i = \mathit{false}$ as well. This table uses observations on all participant pools. Numbers in square brackets are standard errors clustered by participant pool. In the bottom row, we use $^{+}$, $^{++}$, and $^{+++}$ to indicate probabilities greater than 0.50 with statistical significance at the 10%, 5%, and 1% levels; whereas, we use $^{-}$, $^{--}$, and $^{---}$ to indicate probabilities less than 0.50 at the same significance levels.

| Dependent Variable: | $\texttt{consider}_i$ | | | $\texttt{textbookLogic}_i$ | | | $\texttt{textbookLogic}_i \vert \texttt{consider}_i = \textit{true}$ | | |
|---|---|---|---|---|---|---|---|---|---|
| | $\mu_R$ | $\sigma_R$ | $\rho$ | $\mu_R$ | $\sigma_R$ | $\rho$ | $\mu_R$ | $\sigma_R$ | $\rho$ |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| $\texttt{isFinancePro}_i$ | 0.72 | 0.52 | 0.38 | 0.56 | 0.36 | 0.10 | $0.78^{+++}$ | $0.69^{+++}$ | $0.27^{---}$ |
| | [0.02] | [0.02] | [0.02] | [0.02] | [0.02] | [0.01] | [0.02] | [0.03] | [0.03] |
| $\texttt{isMTurker}_i$ | 0.77 | 0.53 | 0.48 | | | 0.08 | | | $0.17^{---}$ |
| | [0.02] | [0.03] | [0.03] | | | [0.02] | | | [0.04] |
| $\texttt{isMBAstudent}_i$ | 0.88 | 0.77 | 0.48 | 0.69 | 0.58 | 0.11 | $0.78^{+++}$ | $0.76^{+++}$ | $0.23^{---}$ |
| | [0.02] | [0.02] | [0.03] | [0.03] | [0.03] | [0.02] | [0.03] | [0.03] | [0.04] |
| $\texttt{isAssetManager}_i$ | 0.62 | 0.54 | 0.35 | 0.37 | 0.40 | 0.17 | $0.59^{+}$ | $0.74^{+++}$ | 0.48 |
| | [0.05] | [0.05] | [0.05] | [0.05] | [0.05] | [0.03] | [0.05] | [0.05] | [0.06] |
| # Obs | 1,216 | 1,216 | 1,216 | 894 | 894 | 1,216 | 683 | 543 | 525 |
| Adj. $R^2$ | 0.77 | 0.61 | 0.44 | 0.60 | 0.46 | 0.11 | 0.77 | 0.72 | 0.26 |



**Table 7. Economic Reasoning by Participant Pool.** This table depicts the rate at which different participant pools reported thinking about average stock returns ($\mu_R$), stock-return volatility ($\sigma_R$), and the correlation between stock returns and consumption growth ($\rho$) when making their investment decisions. Each column reports results for a separate regression of the form: $\texttt{y}_i = \hat{\alpha} \cdot \texttt{isFinancePro}_i + \hat{\beta} \cdot \texttt{isMTurker}_i + \hat{\gamma} \cdot \texttt{isMBAstudent}_i + \hat{\delta} \cdot \texttt{isAssetManager}_i + \hat{\varepsilon}_{i,q}$. The dependent variable is an indicator variable capturing whether/how the $i$th participant thought about a parameter. $\texttt{consider}_i$ is an indicator variable for whether the $i$th participant thought about a parameter at all. $\texttt{textbookLogic}_i$ is an indicator variable for whether the $i$th participant thought about this parameter using textbook asset-pricing logic. If $\texttt{consider}_i = \textit{false}$, then $\texttt{textbookLogic}_i = \textit{false}$ as well. The right-hand-side variables are indicators for which population the $i$th participant belongs to. Numbers in square brackets are standard errors. In columns (7), (8), and (9), we use $^{+}$, $^{++}$, and $^{+++}$ to indicate probabilities greater than 0.50 with statistical significance at the 10%, 5%, and 1% levels; whereas, we use $^{-}$, $^{--}$, and $^{---}$ to indicate probabilities less than 0.50 at the same significance levels.

Dependent Variable: $\texttt{stockFrac}_{i,q}$

| | $\texttt{mean}_{i,q}$ (1) | $\texttt{volatility}_{i,q}$ (2) | $\texttt{correlation}_{i,q}$ (3) |
|---|---|---|---|
| $\texttt{consider}_i = \mathit{false}$ | $2.03^{\star\star\star}$ | $-0.37^{\star\star\star}$ | $-0.01$ |
| | (4.30) | (4.34) | (0.59) |
| $\texttt{consider}_i = \mathit{true}$ | $3.62^{\star\star\star}$ | $-0.77^{\star\star\star}$ | $0.01$ |
| | (14.36) | (9.16) | (0.67) |
| $(\texttt{consider}_i = \mathit{true})\,\&\,(\texttt{textbookLogic}_i = \mathit{false})$ | $1.18^{\star\star}$ | $0.05$ | $0.03$ |
| | (2.22) | (0.26) | (1.12) |
| $(\texttt{consider}_i = \mathit{true})\,\&\,(\texttt{textbookLogic}_i = \mathit{true})$ | $4.06^{\star\star\star}$ | $-1.17^{\star\star\star}$ | $-0.03$ |
| | (11.68) | (10.38) | (0.71) |

**Table 8. Investment Decisions by Economic Reasoning.** This table presents regression results showing how participants' investment decisions vary with average stock returns, $\texttt{mean}_{i,q}$, stock-return volatility, $\texttt{volatility}_{i,q}$, and the correlation between stock returns and consumption growth, $\texttt{correlation}_{i,q}$. This table uses data on all participant pools. Each entry in the table represents the estimated slope coefficient, $\hat{\beta}$, of a separate regression of the form $\texttt{stockFrac}_{i,q} = \hat{\alpha} + \hat{\beta} \cdot \texttt{x}_{i,q} + \hat{\varepsilon}_{i,q}$ where $\texttt{x}_{i,q} \in \{\texttt{mean}_{i,q}, \texttt{volatility}_{i,q}, \texttt{correlation}_{i,q}\}$ using the specified sub-population for a given row. The dependent variable is the fraction invested in stocks, $\texttt{stockFrac}_{i,q}$. $\texttt{consider}_i$ is an indicator variable for whether the $i$th participant thought about a parameter at all. $\texttt{textbookLogic}_i$ is an indicator variable for whether the $i$th participant thought about this parameter using textbook asset-pricing logic. If $\texttt{consider}_i = \mathit{false}$, then $\texttt{textbookLogic}_i = \mathit{false}$ as well. Numbers in parentheses are $t$-stats computed using standard errors clustered by participant. $^\star$, $^{\star\star}$, and $^{\star\star\star}$ indicate statistically significant estimates at the 10%, 5%, and 1% levels.

| a) Personal Income | Mean $\mu_R$ (1) | Volatility $\sigma_R$ (2) | Correlation $\rho$ (3) |
|---|---|---|---|
| $\Pr[\texttt{consider}_i]$ | 0.65 [0.04] | 0.48 [0.05] | 0.28 [0.04] |
| $\Pr[\texttt{textbookLogic}_i \mid \texttt{consider}_i = \mathit{true}]$ | $0.99^{+++}$ [0.01] | $0.85^{+++}$ [0.05] | $0.19^{---}$ [0.09] |

| b) Personal Wealth | | | |
|---|---|---|---|
| $\Pr[\texttt{consider}_i]$ | 0.69 [0.04] | 0.53 [0.03] | 0.37 [0.04] |
| $\Pr[\texttt{textbookLogic}_i \mid \texttt{consider}_i = \mathit{true}]$ | $0.87^{+++}$ [0.03] | $0.83^{+++}$ [0.04] | $0.12^{---}$ [0.05] |

| c) House Prices | | | |
|---|---|---|---|
| $\Pr[\texttt{consider}_i]$ | 0.68 [0.05] | 0.54 [0.05] | 0.31 [0.04] |
| $\Pr[\texttt{textbookLogic}_i \mid \texttt{consider}_i = \mathit{true}]$ | $0.92^{+++}$ [0.02] | $0.96^{+++}$ [0.02] | $0.20^{---}$ [0.06] |

| d) Industrial Production | | | |
|---|---|---|---|
| $\Pr[\texttt{consider}_i]$ | 0.62 [0.04] | 0.48 [0.04] | 0.38 [0.04] |
| $\Pr[\texttt{textbookLogic}_i \mid \texttt{consider}_i = \mathit{true}]$ | $0.87^{+++}$ [0.03] | $0.85^{+++}$ [0.04] | $0.06^{---}$ [0.03] |

**Table 9. Economic Reasoning Using Different Terms.** This table depicts the rate at which participants reported thinking about average stock returns ($\mu_R$), stock-return volatility ($\sigma_R$), and the correlation between stock returns and four different risk factors ($\rho$) when making their investment decisions. Each panel uses data on a separate population of MTurkers and reports results when the risk factor is labeled as *'personal income'*, *'personal wealth'*, *'house prices'*, or *'industrial production'*. $\texttt{consider}_i$ is an indicator variable for whether the $i$th participant thought about a parameter at all. $\texttt{textbookLogic}_i$ is an indicator variable for whether the $i$th participant thought about this parameter using textbook asset-pricing logic. If $\texttt{consider}_i = \mathit{false}$ then $\texttt{textbookLogic}_i = \mathit{false}$ as well. Numbers in square brackets are standard errors clustered by participant pool. In the bottom row we use $^{+}$, $^{++}$, and $^{+++}$ to indicate probabilities greater than 0.50 with statistical significance at the 10%, 5%, and 1% levels; whereas, we use $^{-}$, $^{--}$, and $^{---}$ to indicate probabilities less than 0.50 at the same significance levels.

| | | TNA ($\$bil$) | | Mentions of | Mentions related to other macro variables | | Other macro correlation info? | |
|---|---|---|---|---|---|---|---|---|
| | | Share Class | Fund | correlat(ion\|e) covar(iance\|y) | §Investment Risks | §Investment Objectives | Num. | Graph |
| Vguard 500 Idx, Adm | VFIAX | 276 | 483 | 0 | 0 | 0 | No | No |
| Vguard Tot Stock Mkt Idx, Adm | VTSAX | 225 | 814 | 0 | 0 | 0 | No | No |
| Fidelity 500 Idx | FXAIX | 198 | 198 | 4 | 0 | 0 | No | No |
| Vguard Tot Stock Mkt Idx, Instl Pl | VSMPX | 170 | 814 | 0 | 0 | 0 | No | No |
| Vguard Tot Intl Stock Idx, Inv | VGTSX | 146 | 382 | 0 | 0 | 0 | No | No |
| Vguard Tot Stock Mkt Idx, I | VITSX | 140 | 814 | 0 | 0 | 0 | No | No |
| Vguard Tot Stock Mkt Idx, Inv | VTSMX | 139 | 814 | 0 | 0 | 0 | No | No |
| Vguard Instl Idx, Instl Pl | VIIIX | 115 | 229 | 0 | 0 | 0 | No | No |
| Vguard Institutional Idx, I | VINIX | 114 | 229 | 0 | 0 | 0 | No | No |
| Vguard Tot Intl Stock Idx, Instl Pl | VTPSX | 112 | 382 | 0 | 0 | 0 | No | No |
| Vguard Tot Bond Mkt II Idx, Inv | VTBIX | 107 | 182 | 0 | 0 | 0 | No | No |
| Vguard Tot Bond Mkt Idx, Adm | VBTLX | 100 | 229 | 0 | 0 | 0 | No | No |
| Fidelity Contrafund | FCNTX | 95 | 122 | 0 | 0 | 0 | No | No |
| Amer Funds Gr Fund of Amer, A | AGTHX | 91 | 196 | 0 | 0 | 0 | No | No |
| Vguard Wellington, Adm | VWENX | 89 | 105 | 0 | 0 | 0 | No | No |
| Vguard Tot Bond Mkt II Idx, I | VTBNX | 75 | 182 | 0 | 0 | 0 | No | No |
| Vguard Tot Intl Stock Idx, Adm | VTIAX | 75 | 382 | 0 | 0 | 0 | No | No |
| Amer Funds Incm Fund of Amer, A | AMECX | 74 | 111 | 0 | 0 | 0 | No | No |
| Amer Funds American Bal, A | ABALX | 72 | 150 | 0 | 0 | 0 | No | No |
| Amer Funds Europa Gr, R6 | RERGX | 71 | 162 | 0 | 0 | 0 | No | No |
| Dodge & Cox Stock | DODGX | 71 | 71 | 1 | 0 | 0 | No | No |
| Vguard 500 Idx, Instl Select | VFFSX | 70 | 483 | 0 | 0 | 0 | No | No |
| PIMCO Income Instl | PIMIX | 67 | 128 | 3 | 0 | 0 | No | No |
| Amer Funds Cap Income Bldr, A | CAIBX | 65 | 105 | 0 | 0 | 0 | No | No |
| Vguard Interm-Term Tx-Ex, Adm | VWIUX | 65 | 68 | 0 | 0 | 0 | No | No |
| | | 2821 | 3734 | 8 | 0 | 0 | | |

**Table 10. Mutual-Fund Prospectuses.** This table describes how the 25 largest US mutual funds talk about risk-factor correlations in their prospectuses. "Mentions of correlat(ion|e) covar(iance|y)" counts the number of times *'correlation'*, *'correlate'*, *'covariance'*, or *'covary'* appear in a fund's prospectus. "Mentions related to macro variables" counts the number of times a fund mentions its exposure to a macroeconomic variable in either the Investment Risks or Investment Objectives sections of its prospectus. "Other macro correlation info?" is an indicator variable for whether a prospectus addressed the fund's correlation with macroeconomic variables in some other way.