PRICE AGNOSTIC DEMAND *

SAMUEL M. HARTZMARK

ABIGAIL B. SUSSMAN

November 2023

Abstract

We document that stock-market return expectations are largely unrelated to the level of the stock market. This disconnect emerges because investors lack strong views as to what the market price should be. Investors are unable to identify extreme deviations from actual market prices and are thus willing to accept a wide range of prices as the actual market price, even holding information about fundamentals constant. This is at odds with standard asset pricing models which assume that investors are uncertain about the future, but understand what the current market price should be and whether the actual price is different from this. This price agnostic demand can help explain a number of puzzling empirical patterns in asset prices.

*Hartzmark is with Boston College Carroll School of Management and the NBER, Sussman is with the Chicago Booth School of Business. We would like to thank Francesca Bastianello, Alex Chinco, Niels Gormsen, Eric So and David Solomon for helpful comments and seminar participants at the Briq Beliefs Workshop, Aalto and Boston College.

We might define an efficient market as one in which price is within a factor of 2 of value, i.e., the price is more than half of value and less than twice value... By this definition, I think almost all markets are efficient almost all of the time. -Fisher Black, AFA Presidential Address (1986)

I. Introduction

In almost any asset pricing model of the stock market, agents view current information, form expectations of the relevant state variables, and use these to predict the future market price level. The current market price equals the model's prediction based on the current state variables and is assumed to be well understood. Whether an investor focuses directly on their prediction for the future market price, or subtracts and divides that expectation by the current price to examine returns has no bearing on how they value owning a fraction of the stock market, or how the market is priced.

While the *actual* current market price is clearly known, the seemingly innocuous assumptions above mask a strong assumption about the perception of what the market value *should* be. For this equivalence of thinking in terms of prices and returns to hold, there must be a strong conviction about the appropriate level of the stock market and whether the actual level reflects this. For example, assume that the S&P 500 is currently trading at 4,000. Further assume that given the current information on fundamentals, an investor from your favorite asset pricing model thinks the market price one year later will be 4,200, or equivalently, they expect a return of 5%. What would happen to such an investor if they were in the exact same situation, but the price of the S&P 500 was 3,000?[1] With identical fundamental information, such an investor would still predict that next year's price would be 4,200. In this case, the investor would predict a 40% return over the coming year, drastically higher than the 5% return when the price was 4,000.

This paper proposes a different framework of investor behavior in which investors are agnostic to what the current market price should be. Consequently, expectations of *changes* in the market price (i.e., return expectations) are similar irrespective of the current market price. Such expectations largely adapt to the current level of the market price, taking that price as given and adjusting from that point. We note that investors behaving in this manner need not be irrational. Indeed, if the above Fisher Black (1986) quote describes how markets function (in a manner very much at odds with asset pricing models), using fundamentals to predict next period's price without regard to the current price is not sensible.

[1] We note in most standard rational models prices do not deviate from fundamentals, so such a price movement would be an off equilibrium outcome. The models do make predictions that investors would notice and respond to such an outcome if it ever did occur.

Returning to our example, let's say an investor who is largely agnostic about what the current level of the market price should be expects a return of 5% when the S&P 500 is at 4,000. In the setting where the market was at 3,000, but fundamental information was otherwise identical, they would also predict returns of about 5%. The return expectation remains the same because such an investor does not have a strong opinion about what the level of the stock market should be and thus views the market being at 3,000 or 4,000 as roughly the same rather than being dramatically mispriced. Thus, a key testable prediction for price agnostic demand versus almost every other asset pricing model is whether or not investors have a strong beliefs about the *current* correct level of the stock market.

In this paper, we provide evidence that investors exhibit considerable uncertainty about what the level of the stock market should be. In a variety of settings, we find that investors largely take the market price as given. When faced with extreme market mispricing, agents are unable to identify it meaningfully and do not form return expectations based on it. When asked, sophisticated financial professionals do not think it is possible to use economic information to determine the market price within 10% of its true value without seeing it. This finding can help explain puzzles such as excess volatility and why markets are inelastic. It suggests a reinterpretation of the finding that cash flow news is largely unimportant for explaining market variation. Our results militate against rational discount rate variation and suggest instead that a key driver is uncertainty as to what the appropriate level of the market should be.

We provide evidence consistent with price agnostic demand representing the majority of return expectations over annual horizons. We begin with evidence from laboratory settings using participants from Prolific (an online pool similar to Mturk), MBA students at the University of Chicago Booth School of Business, and participants at a meeting of one of the world's largest investment managers. Students at Booth and the asset manager sample are meant to represent a more sophisticated class of investors. Many Booth students have worked in the finance industry prior to their MBA, have received significant financial education, and go on to manage investments after their time at Booth. The asset manager, who unfortunately must remain anonymous, manages large sums of money on behalf of large institutional clients and represents an ideal example of a sophisticated and deep-pocketed market participant.

In our baseline experiment, we provide participants with a variety of data from a specific point in time. Specifically, we provide the interest rate, inflation, prior economic growth and Fed forecasts of future economic growth reflecting a particular day in the past. Our goal was to provide information

that is readily available in the real world and also represents inputs that could plausibly account for the level of asset prices.

We also provide participants with the level of the S&P 500 and the price to earnings ratio, but we vary whether this level reflects the true value corresponding to the time the other fundamental information is drawn from. We randomly manipulate the level of the S&P 500 to be anywhere from -50% to 50% of the actual value of the corresponding day in the past. The price to earnings ratio incorporates this manipulated S&P 500 value as well. We explain this structure to participants, though we do not state the amount that prices are manipulated by. Participants online see 10 such scenarios from different dates in the past.

Participants' first task is to try and identify how much the displayed level of the stock market deviated from the actual value on the date for which the fundamental information was provided. If participants can infer roughly the actual level of either the stock market or price to earnings from the data, it should be straightforward to identify large deviations. For example, participants should be able to report that the market level was manipulated to be higher or lower in scenarios where the market was, for example, 40% higher or 40% lower respectively. Further, participants should be able to give a rough idea of the magnitude and correctly identify that the market was 40% higher than it actually was, rather than 10% higher. On the other hand, if agents exhibit agnostic demand and largely take the market price as given, agents will likely have little ability to identify such shifts to the market price.

While participants are able to identify some mispricing, we find that their ability to do so is quite limited. When faced with overpricing scenarios, with an average of 20% overpricing, the average perceived mispricing reported is 2.8%. When presented with underpricing scenarios, with an average of -20% mispricing, participants report an average of -1.2% perceived mispricing. Including participant or date fixed effects yields similar results. Overall, participants struggle to identify economically meaningful deviations from actual market prices.

One potential concern is that sophisticated market participants have skill at such a task, while our online panel does not. Indeed, our sophisticated asset manager sample does appear better at detecting mispricing than the rest of our sample, though they still miss most of the mispricing. Thus, for this sample, we directly ask whether they think it is possible to know the level of the market within 10% based on economic fundamentals without seeing it. The majority of our sample of asset managers, 68%, says this task is not possible. We followed up with the subset of respondents who said this *was*

possible and asked whether the experiment provided the data they would want to understand the level of the market. All but one said that it did, suggesting that there is not an obvious missing variable that would increase the precision of this task. Together, the results suggest that even sophisticated market participants exhibit behavior consistent with price agnostic demand.

Consistent with our proposed framework, much of the limited ability to perceive mispricing seems to be due to a reliance on the displayed market price. While participants do not respond strongly to prices being shifted from the actual price, they do respond strongly to the displayed information. When P/E ratios are particularly high (above 40), respondents perceive the market to be overpriced by 12%. When it is particularly low (below 15), respondents perceive -5% mispricing. Given that the displayed P/E ratios are correlated with the level of mispricing induced, much of the limited success that participants have in predicting mispricing is explained by variation in the displayed P/E ratios. This result suggests that participants are understanding and engaging with the task. In addition, it is broadly consistent with research in psychology finding that people have a hard time understanding numbers (e.g., the market price) in isolation, but have an easier time understanding them in context (e.g., relative to earnings; Kahneman and Frederick (2005); Kahneman (2003)).

We also ask participants to estimate what the level of the S&P 500 actually was 12 months after the snapshot of time presented to them (i.e., predict the future level of the S&P 500 relative to the date they are considering). If agents could perceive mispricing that they believed would correct over the next 12 months, there should be a strong negative impact on perceived future returns based on the level of actual mispricing. Participants in settings with induced overpricing expect returns of -0.5% over the following year while participants with induced underpricing expect returns of 6.2%. We again see some directional ability to identify mispricing, but of much smaller magnitude than we would expect in any standard model.

Once again, this ability is dwarfed, and largely explained, by the displayed P/E ratio. Participants seeing a P/E ratio above 40 expect the market to return -11% returns over the following year. Participants seeing a P/E below 15 expect 12% returns. Participants use fundamentals to predict future returns, but this is largely about forming expectations of the future without regard to whether the current price of the market is different from what it actually was.

While it appears that investors are unable to identify economically large shifts in the level of the market price, there are a number of potential concerns, as with any lab experiment. One such concern is that variation in the level of mispricing presented to participants may have led to confusion or

difficulty in reporting their perceived mispricing. Thus, we run a simpler version of the experiment where investors are presented with a scenario where the market is underpriced by 25%, overpriced by 25% or the actual market level. Each scenario has an equal probability of occurring. Participants are asked simply whether the market is overpriced, underpriced, or the actual level. Participants identify the correct scenario 37% of the time, better than pure chance of 33%. With that said, 25% mispricing is extreme, so the fact that agents are not easily able to identify the direction of these large swings makes it likely that such investors do not form an independent view of the current market price based on fundamentals but instead largely take it as given.

While we provide commonly reported data to investors, there is substantially more real world data for investors to digest in actual market outcomes. This gap may lead to reasonable concerns about the external validity of the result. We aim to mitigate this concern by focusing on first order shifts in mispricing rather than subtle nuance. While subtle details will of course be important, first order issues, such as the market being over or under priced by 40%, should not depend on such subtlety. We also ask the subset of our asset manager sample who thought the task was possible whether we provided the relevant data to achieve this task, and their response was yes. As long as investors have a rough idea of the appropriate level of the market price based on the most important metrics, it should be straightforward to identify mispricing of this magnitude.

Finally, we examine a variety of real world examples to show that investors generally seem to form return expectations largely based on price agnostic demand. We first note that in the financial press, while there is a large amount of financial data reported, there is minimal information that could be used to pin down the market price. This is unsurprising as market timing based on perceptions of deviations from the proper price level is not a particularly popular strategy. Consistent with expectations being largely price agnostic, but with a small positive weight on a fundamental component, Asness et al. (2015) argues that market timing should comprise a small amount of a trading strategy at most. Similar arguments can be found by academics (e.g., Samuelson), or finance professionals (e.g., Jack Bogle) who argue against market timing generally, but in extreme settings (e.g., the tech bubble) decreased their equity exposure. Again, this is consistent with a small positive weight in that the fundamental component of demand which only becomes relevant once the fundamental based mispricing becomes extreme enough to overcome the small weight.

We also examine real world forecasts of expected returns. Outside of the P/E ratio, most fundamental variables in the models we study are not related to the current level of the market price,

consistent with price agnostic demand. Further, there is minimal conviction in the model's predictions, or the mechanism driving those predictions. In a recent prediction, Vanguard suggested the market was 30% overpriced which led it to shade down return expectations slightly and suggest no changes to investment behavior.[2] Similarly, both Goldman Sachs and Bank of America report overvalued markets expecting annual returns over the next 10 years of 3% and 0%-1% respectively. How much conviction was behind these estimates for the next year? The Wall Street Journal summarised the conclusion as "Buy stocks anyway, because the next year looks great."[3]

Our paper adds to the literature using experimental methods to test the predictions of asset pricing models. For example, Chinco et al. (2022) show that while nearly all asset pricing models assume investors are attempting to hedge risk factors, this is not something investors are trying to do. Andre et al. (2023) shows investors respond to news in ways difficult to explain with standard theory and Choi and Robertson (2020) and Bender et al. (2022) show how investors say they view asset prices. Our paper contributes to this literature by showing that investors lack conviction about fundamental based estimates of the market price and thus largely exhibit price agnostic demand.

This paper contributes to the literature on understanding price volatility in the stock market. Prior research has found that the market is excessively volatile when compared to fundamentals (e.g. Shiller et al. (1981)) and relatedly, that most market price movements can be attributed to discount rate variation rather than cash flow variation (e.g. Cochrane (1992)). When faced with these facts, the literature focused on explanations where agents were rationally aware of time-varying discount rates. Under such explanations, investors understand the current market price and also could ask for relevant information to understand a current discount rate. Our results do not support either of these predictions but instead support an alternative interpretation of the empirical facts. Namely, given the extreme uncertainty as to what the level of the market should be, there is a wide bound of values that investors will treat as roughly equivalent.

Our paper therefore also contributes to the literature on market efficiency and arbitrage (e.g. Shleifer and Vishny (1997)). A lack of conviction as to what the appropriate level of the market price will be suggests an unwillingness to trade based on perceptions of the current market price.[4] If market participants are unable to identify large shifts in market prices as influencing future returns, they

[2] https://corporate.vanguard.com/content/corporatesite/us/en/corp/articles/investment-economic-outlook-feb-2024.html

[3] https://www.wsj.com/finance/stocks/stock-market-overvalued-forecasts-2025-e073e1d4?

[4] While our focus is the overall market, Ben-David and Chinco (2024) examine equity analysts of individual firms and finds their forecasts are based on trailing P/E. Of course expectations should be formed based on the future, not the past, suggesting equity analysts similarly lack conviction on what the appropriate level of firm prices should be.

would have no reason to trade based on their perceived level of mispricing. Of course, the assumption that prices are equal to the efficient price is driven by the arbitrage mechanism of investors being willing to make such trades. This suggests the level of the aggregate stock market is much less likely to be close to an efficient fundamental based price and can help to account for the excess volatility and discount rate variation discussed above.

Finally, our paper also adds to the recent debate on whether financial markets are inelastic (e.g. Gabaix and Koijen (2021), Hartzmark and Solomon (2021)). Textbook models largely leave trade unmodelled given an assumption that deep pocketed arbitrageurs will actively trade to correct any perceived mispricing. To do so, an arbitrageur must have a strong conviction that there is mispricing and that there is a profitable trade to be made in attempting to correct it. Our results suggest this is unlikely to be the case. If there is a wide range of prices that arbitrageurs are willing to be agnostic about, then when trade shifts prices they will be unlikely to strongly trade to correct it. Thus price agnostic demand represents one potential mechanism that leads to inelastic markets.

II. Price Agnostic Demand Framework

In this section, we present a simple framework to illustrate the intuition of our argument.

In standard models, agents do not need to see the current price to know what it is. Once they understand the relevant state variables, they would know what the price should be. As Cochrane (2009) writes, you can "summarize asset pricing by two equations," the first of which is $p_t = E(m_{t+1}x_{t+1})$ and the second is $m_{t+1} = f(data, parameters)$. Of course agents in an asset pricing model know the form of f() along with the data and the parameters. Thus, while state variables and f() are different across models, the fact that agents inherently know what the price should be, even when it is not presented to them, is a nearly universal assumption in asset pricing theories.

Assume there is an f() from a given asset pricing model that transforms the relevant fundamentals into expectations of future prices F_{t+1}. The input to this function, f(), will be based on all relevant state variables and their expectations. Thus, for a rational agent in a standard asset pricing model we know that:

$$E_t[P_{t+1}] = F_{t+1}$$

If this is the expectation of future prices, then expected returns would simply be:

$$E_t[R_{t+1}] = \frac{E_t[P_{t+1}] - P_t}{P_t} = \frac{F_{t+1} - P_t}{P_t}$$

Further, if this is how investors form expectations and respond to fundamentals, P_t would reflect those fundamentals and represent the rational, agreed-upon price based on the information at time t. If for some reason that was not the case, agents would realize this. That is because expected returns are perceived relative to this price, so if the price were far lower than that from their model, they would view expected returns as significantly higher.

The key assumption our paper adds is that people's opinion of the correct current price is closely linked to the price that they observe, irrespective of whether the prevailing market price reflects fundamentals. In our example, fundamental state variables are not ignored, but they receive significantly less weight than in standard models. More formally, we assume that when agents form expectations of future prices, they put weight θ on fundamental information and weight $(1-\theta)$ on the price agnostic component of demand. The price agnostic component takes the current price P_t and multiplies it by an expected return (unrelated to the current price level) Z_t. Thus,

$$E_t[P_{t+1}] = \theta F_{t+1} + (1-\theta)P_t(1+Z_t)$$

In its simplest form, Z would be a constant. This would be equivalent to the random walk hypothesis. Given the current price level, an investor would always expect the next price to be a constant return amount above the current one. Of course, the return multiplier, Z, could be more complex and could interact with other state variables and be time varying, e.g. Z_t. The important thing for our setting is that such complexities could lead to time variation in Z, but remain a multiplicative function of the current price.

Given Z_t, it is straightforward to write expected returns as:

$$E_t[R_{t+1}] = \frac{\theta F_{t+1} + (1-\theta)P_t(1+Z_t) - P_t}{P_t}$$

This illustrates two extreme cases. When $\theta = 1$, we have the textbook case. The current price has no impact on future price expectations because fundamental state variables are the only input to expectations. This is akin to standard rational models. When $\theta = 1$, the numerator of the formula is the gap between the expected future price based on fundamentals and whatever the current market price is. If the current price were below fundamental values (meaning $P_t < F_t$), investors would expect very high returns. And vice-versa with the opposite.

Further, it is important to note that such a setting in the real world implies that investors generally

think the price *today* is a function of fundamentals F_t. Otherwise, there would be no reason to expect tomorrow's price to be a function of fundamentals. Thus, for an agent who believes that returns follow a process where $\theta = 1$, it must be that they also generally believe the current price reflects fundamentals in a similar manner.

When $\theta = 0$, the fundamentals that describe the rational value of the stock market do not matter for forecasting future prices. Investors simply look at the current price and use this as the only input to predict the future return. In practice, this means that expected returns are not influenced by whether the current price is above or below what fundamental values imply.

If theta equals 0, the first term cancels out. Further, let us assume the case with constant returns of Z=8%. Subbing this in would yield:

$$E_t[R_{t+1}] = \frac{1.08 * P_t - P_t}{P_t} = 0.08$$

In many ways, this paper is a study of θ. While asset pricing models suggest it should be 1, the evidence that we find is that it is closer to 0 (though still positive). This means that agents only minimally weight fundamental value when forming expectations of the next period's price. As long as the price is not extremely different from what fundamental value would suggest, expectations of the next period's price are largely a transformation of the current price. Similarly, expected returns are not greatly influenced by the level of the market price relative to its fundamental value. Said differently, in most cases, prices are not drastically different from fundamentals and most demand is thus largely agnostic to what the price should be. Only when the price is extremely different than what fundamentals imply will F_t have a meaningful effect on expectations of future prices and future returns.

This model is meant to display the simple intuition and implications of being agnostic about the level of the market price, but of course the influence of such a channel need not take this specific functional form. One likely deviation, consistent with the results in Section IV, comes from how fundamentals influence return expectations. When we document how such information is actually used, we find that it is not to predict fundamental values. Rather, it is as an input to a regression based on historical data. Such a channel may lead to fundamentals mattering, but in a way quite different from what standard theory suggests. While we think this framework is illustrative of the intuition of the price agnostic demand we study, we leave how best to realistically model our results

for future research.

III. Identifying Mispricing

Our baseline experiment evaluated whether people can identify deviations from the market price and whether this influences their perception of returns. We find that participants' personal perceptions of mispricing influence their perceptions of future returns. However, actual mispricing has little effect on the perception of mispricing.

A. Participants

We recruited participants for our baseline study (N=289) from three separate populations. First, we recruited 198 participants through Prolific Academic, a platform similar to Amazon's Mechanical Turk. In addition, we recruited 60 MBA students and 31 visitors to an event hosted by a deep pocketed asset manager held in May of 2023.

B. Experimental Design

Our goal was to design a setting in which participants have enough information to understand the current level of the stock market according to standard models and intuition. We then wanted to see if people were able to identify when a displayed market level was significantly different from the value that actually occurred at a given point in time.

We begin by randomly choosing days between 1971 and 2021. We wanted to use actual dates with actual data to avoid having to specify a model to generate data. If participants were unable to detect mispricing relative to model generated data, this could simply reflect that they had a different model of how the world worked than the one we chose. By choosing actual dates we can be agnostic as to how asset markets actually work or should work and just take these observations of examples of how they have worked in the past.

We displayed the interest rate, inflation, prior year's economic growth, the Fed economic growth forecast, the level of the S&P 500 index, and the market price-to-earnings ratio. Participants were informed, accurately, that these numbers were the actual values from the selected day. Our goal in selecting these variables was to provide commonly discussed and presented statistics that could plausibly be used for our task. Some of these variables are also directly relevant in many asset pricing models, or at least tangentially related to certain state variables. We think these variables represent the most commonly discussed variables for such a task, are widely used, and are often a component of professional projections of future returns.

The values for the S&P 500 index and the market price to earnings ratio could be different from those that actually occurred. These values could be the level that actually occurred on that day, or they could be higher or lower by up to 50%. For each day, we randomly select a number X% between 10% and 50%. One-third of participants viewed the actual value, one-third viewed values higher by X% and one-third viewed values lower by X%. Participants were told that the market value and associated price-to-earnings ratio could be the same, higher, or lower than what actually occurred.

We provided intuitive definitions for all terms along with historical ranges to help participants interpret these values. We required all participants to pass comprehension checks on economic concepts and on the experimental structure in order for them to be included in the experiment.

Once participants finished the instructions and passed comprehension checks, they were shown the first scenario. They were asked "Considering the metrics above, do you believe that the S&P 500 Index Level and PE Ratio reflect levels that are higher, lower, or the same as the actual index level and PE Ratios from the date displayed in the table?" A slider was displayed below where subjects could select values from -50% to 50%. After this, participants were asked what they thought the value of the S&P 500 actually was 12 months later.

C. Design Choices

We aimed to use research methods that would most accurately assess whether investors understand what the current market price of an asset should be based on relevant information about underlying fundamentals. In an ideal experiment, we would identify a setting where all economic variables were identical, but the market price varied, and we would test whether investors would notice the price variation. For example, if we assumed that the S&P 500 were trading at 4,000 today, we could randomize investors into worlds in which the S&P 500 was at 3,000 or 5,000 while holding all other variables constant and test whether investors would recognize mispricing in the latter settings.

Since this thought experiment is not possible, we turn to experimental scenarios that aim to mimic the variation as closely as possible. In this section, we discuss why we designed the experiment as we did and why we believe this design provides meaningful and interpretable responses.

A primary goal of the experimental design is to make it as easy as possible for participants to identify mispricing when it occurs. In other words, we aim to provide a conservative test. Therefore, in all of our settings, we examine mispricing of at least 10%. Thus, the main conclusions from this experiment will depend on a meaningful ability of investors to determine the true price based on

major shifts from this price rather than subtle variation. We ask participants to identify the presence of mispricing, rather than the correct price, because we believe this is an easier task.

In addition, we aim to provide participants with only the most relevant information rather than distracting them with additional statistics that are less likely to be used to determine price. In our baseline experiment, we provide information about the interest rate, inflation, prior year's economic growth and the Fed economic growth forecast. While additional information is available in the world, we restricted the amount of information to encourage participants to focus on the most relevant metrics and maximize the chance that they would be able to identify mispricing when it occurred.

One main design choice was to rely on fundamental information and prices from historical dates. This stands in contrast to an artificial environment in which we create fundamental information and map that onto prices based on standard asset pricing models and examine participants' ability to identify whether the price maps onto what the asset pricing model would have predicted. An artificial environment would lead to a joint hypothesis problem in which we are simultaneously testing investor's beliefs about the model and their ability to identify the true price based on fundamental information. In the case of historical data, we are not testing a theoretical mapping between fundamental information and price, but instead testing a relationship between fundamentals and price that has actually occurred.

One concern about this experiment, particularly with an unsophisticated population, is that we are asking participants questions that they may be unable to answer due to confusion around the task itself (Bergman et al., 2020). To help address this concern, we provide participants with intuitive definitions of each of the key terms along with reading checks to ensure that participants understand the key concepts and task. Further, we ask participants to rate their confidence in their responses and find that results are largely robust to examining only participants with high confidence in their response. Importantly, we also test whether the same patterns emerge across investors with varying degrees of sophistication and find that they do.

In addition, we analyze patterns in participants' responses that could be used to indicate that they are internally coherent and indicative of an understanding of the task. For example, we find that the displayed price to earnings ratio serves as an input into evaluations of mispricing. We view weighting of this metric as a signal of general task understanding and attention to the metrics displayed.

Another common concern in the case of direct questioning is that of experimenter demand (Bergman et al., 2020). In other words, participants may be providing responses that they think the person de-

signing the survey wants them to give rather than reporting their true beliefs. In this case, we largely circumvent this issue by being upfront in the survey instructions, making participants aware of the set of possible responses, and aligning the experimenter's goals with that of the survey-taker. In each case, the clear goal is to try to identify whether mispricing occurs, and if so how much.

Finally, as with any experimental design, the task may be flawed for some other unknown reason. Thus, for our sophisticated asset manager sample, we directly ask them a much more general question about their ability to determine market price as a function of fundamentals, outside of this experimental context. Specifically we ask them:

> *Do you believe that it is possible for investors to use information on economic fundamentals from a given day to determine the market price on that day within 10% of the true value?*

If the design of the experiment is fundamentally flawed, but in the real world the task is straightforward, these respondents should be able to let us know that this is the case by answering yes. This would tell us that in the real world these sophisticated investors think they can easily accomplish the task we are attempting to emulate in the lab.

D. Experimental Results

D.1 Baseline Results

We begin by examining whether or not people were able to use the information to understand the prevailing market price. We find that people were only able to do so to an extremely limited extent.

Figure 1 shows the main result across all three populations. The left of the figure shows the actual average level of mispricing. The maroon bar shows that when the market was underpriced, on average it was underpriced by about -20%. The blue bar shows when it was overpriced it was overpriced by about 20%.

The right three bars show that perceptions were not particularly close to this reality. The perceived level in the low state was -1.1%. This is about 19 percentage points higher than the actual, with a t-statistic on the difference of -22.59. The perceived level of mispricing when there was no mispricing was slightly above zero at 1.37%, with a t-statistic on the difference of 1.74. When there was positive mispricing the perceived level was 2.68%, with a t-statistic on the difference of 22.01. The order of these three bars is monotonic, which is consistent with some skill at this task and also consistent with responding logically to the experimental environment. Overall, participants miss the majority of the mispricing that they are exposed to.

Table 1 explores this more formally using a variety of measures of accuracy. We first ask whether people can accurately detect direction. We count an answer as directionally correct if they write down a positive number for positive mispricing, a negative number for negative mispricing or zero for no mispricing. The first row shows that the correct direction is guessed 36% of the time. This is better than random guesses which would yield 33% correct answers (with a t-stat relative to random of 2.49), but not particularly close to being perfectly accurate (with a t-statistic on the difference relative to being accurate of 68). The next row shows perceived mispricing, which is the same as in Figure 1. Participants do increase their perceived prices monotonically, moving from low to none to high, but not nearly to the extent that they should.

Next we examine the forecast error. The row shows that when mispricing is too low (with an actual average mispricing of -20%), perceived mispricing is too high by about 18 percentage points. When mispricing is positive, perceived mispricing is too low by about 17 percentage points. The none column is not different from zero.

We finally turn to the absolute value of errors. The average absolute value is about 23 and is fairly consistent across the three columns. This shows that guesses are not particularly responsive to the actual level of mispricing, regardless of the scenario.

Next, we explore how perceptions of mispricing varied with levels of mispricing. Figure 2 shows a scatterplot of the perceived mispricing relative to the actual level of mispricing. The perceived mispricing is the average level of mispricing for a given day with a given level of mispricing. There is indeed a slight positive trend in these dots, but it is fairly moderate compared to accurately identifying mispricing (shown by the gray line). Thus, the graph again illustrates that people simply are not particularly sensitive to the induced mispricing.

Table 2 explores this sensitivity more formally by regressing perceived mispricing on actual mispricing. The first column runs this regression without controls and shows a coefficient of 0.077. This is positive and statistically significant, consistent with some responsiveness to mispricing. With that said, this number is much closer to zero than it is to 1 (if mispricing was perfectly reported). The next column adds a person fixed effect. There is a slight increase in the coefficient, to 0.088, but it remains similar. Next, we put in a fixed effect for a given date. This holds fixed the real information known on that date. The coefficient remains similar at 0.071. The final column includes both and again finds a similar result. This suggests that while there is some sensitivity to induced mispricing, the ability to perceive this is fairly small.

The results suggest that while participants can detect some mispricing, they miss the majority of it. One potential explanation is that this is driven by the less sophisticated respondents on Prolific which obfuscate the skill in this task from the Booth MBA students or the sophisticated asset managers.

Table 3 repeats the summary stats analysis separately by participant type. The MBA students do not appear notably better than Prolific participants. They guess correctly 36% of the time compared to 35% for Prolific and if anything have higher absolute errors of 25.6, compared to 23.3 for Prolific.

While the MBAs display no evidence of skill, the deep pocketed asset managers seem to do better at this task than either of the other groups. They are correct 36% of the time, and they have absolute errors of 20.8. Thus it appears that they are able to do this task with more precision than either of the other subject groups.

Table 4 repeats this analysis and finds a similar pattern. Prolific and MBA students have a coefficient on actual mispricing of 0.08, but the asset managers are about twice as responsive with a coefficient of 0.17.

These results suggest that more sophisticated participants are able to more effectively use the data in our task and respond to it. With that said, even these experts struggle to respond meaningfully to large swings in mispricing. While they do better, they remain much closer to missing all mispricing than they do to being able to identify it all. There is no evidence that these experts would meaningfully understand the market price without seeing it first.

D.2 Expected Returns

While our experiment induces mispricing and asks participants to identify it, this is a somewhat non-standard task. In the real world, it is much more likely that a participant focuses on questions more closely related to what they expect market performance to be. Thus, in this section we ask about expected returns.

To do so, we ask what people expect the level of the market actually was 12 months in the future relative to the data shown, and we use this to calculate expected returns. As discussed, this is the most natural question to ask an agent from most asset pricing models and perhaps is the most relevant question of interest to investors in the real world. Further, in asking about the future actual price we are directly testing the mechanism assumed to keep markets efficient, namely that if mispricing were ever to exist, investors would recognize this.[5] The answer to this question should include both

[5]Otherwise they could not engage in trade to correct it.

the expected returns relative to what economic fundamentals suggest, but also the correction of the induced mispricing.

Figure 2 shows how expected returns vary with actual mispricing. The pattern is similar to that in terms of induced mispricing. There is a subtle negative slope to the dots, but of a degree far too small relative to what was induced. If you knew that the displayed market price was manipulated to be 40% lower than the true one and would correct in the next 12 months, you should expect returns of nearly 100% the following year relative to the displayed price. Actual expectations with such extreme mispricing are closer to 0%.

Table 5 shows how expected returns vary with displayed mispricing. Consistent with the figure, the coefficients are negative, but quite small in magnitude at roughly -0.05. This means that if the market were overpriced by 25% (based on which return expectations should be greater than 20%), expected returns are about 1.25 percentage points lower than they would be otherwise. Consistent with a small positive theta, mispricing influences return expectations, but only very slightly.

D.3 Responsiveness to other variables

The results are consistent with respondents having only a small amount of skill at detecting mispricing. One potential concern is that respondents simply were not paying attention to the experiment and ignoring or failing to understand the information presented to them. Thus, we next turn to see whether there were meaningful differences in responding to the variables that were displayed.

One plausible concern would be that participants were confused and unable to respond to the information, or that they were not paying attention to what was being displayed. If this were the case, our setting would not provide meaningful information towards our question of interest. Examining responses to the variables displayed allows us to test whether participants appeared to be paying attention to this information. In addition, to the extent that participants do respond to the information provided, this analysis helps us understand what variables participants consider important inputs into their decisions.

Before turning to the specific variables, we examine variation in mispricing by day in Figure 4 Panel B. Recall that each day had a specific X drawn where some subjects saw the price increased by X (the green dots), some decreased by X (the orange dots) and some saw the actual value (the blue dots). Each X is unique, so each column of 3 dots shows perceived mispricing where all of the economic fundamentals were the same (other than price and P/E). What is clear from this figure is

that there was clustering on a given day in that the dots all tend to be higher or lower in terms of perceived mispricing. This is consistent with investors putting some weight on fundamentals, but lacking significant ability to identify actual mispricing.

Table 6 regresses perceptions of mispricing on each of the variables displayed. The first 6 columns examine each variable in a univariate regression while the seventh column includes them all. All regressions include a participant fixed effect and standard errors clustered by participant. We find strong positive relationships between perceived mispricing and the displayed price to earnings level, the level of the S&P 500 index and an inverse relation with the level of the risk free rate.

We explore the price to earnings ratio further in Figure 4. The table graphs the average response for the perception of mispricing for a given level of P/E ratio displayed. The chart shows a strong positive relationship. As the P/E ratio increased, so did perceived mispricing. When the P/E ratio displayed was in the teens, perceived mispricing was around -5%. When the P/E ratio was above 40, perceived mispricing was above 10%.

This suggests that respondents were paying attention to the experiment and responding to the P/E ratio. Further, looking at the spread of dots it is clear that they were responding to the displayed P/E ratio, but unable to detect most of the induced mispricing. Altogether, the results suggest that participants were paying attention and responding to the numbers that were being displayed to them. While they were responding to the prompts of the experiment, they were unable to detect the majority of the induced mispricing.

It is worth noting that this result is consistent with broader literature in psychology on how people process numbers. People have a hard time identifying what value means in isolation and tend to evaluate magnitudes in the context of relevant benchmarks (e.g., Kahneman (2003); Kahneman and Frederick (2005). In one canonical example, when participants were asked about their willingness to travel 20 minutes to save $5, they reported much greater willingness to travel when the total purchase cost $15 (a 33% savings) versus $125 (only a 4% savings) (Tversky and Kahneman, 1981).

As numbers get larger, the challenge of identifying and interpreting magnitude becomes more extreme and people become less sensitive to differences Dehaene (2003); Kahneman et al. (2000). For example, when participants reported their willingness to fund protective measures to prevent 2,000, 20,000, or 200,000 bird deaths the reported amounts remained around $80 in all cases Desvousges et al. (1993). The challenges of assessing magnitude can be mitigated when people have a basis of comparison Hsee et al. (2005). Against this research backdrop, it seems reasonable that our participants

would have a hard time identifying the true value of the stock market in isolation, but that they would be better at identifying situations where the price is high or low relative in relative terms (i.e., the P/E ratio).

Next we ask how future perceived returns were influenced by the displayed variables. We find that displayed P/E continues to strongly negatively predict future returns as does the level of the market. Inflation and the risk free rate are also positively associated with expected returns.

Figure 4 Panel B shows expected returns relative to the displayed price to earnings ratio. There is a strong pattern from the top left to bottom right. When P/E is quite low, expected returns average in the teens. When they are quite high expected returns move to slightly negative.

These base results are consistent with participants paying attention to the information being displayed and responding to it. The information influences both perceptions of mispricing and expectations of future returns.

Another concern is that agents are able to identify mispricing, but that some of the mispricing in our setting is too subtle to observe. Before empirically addressing this concern, we highlight that the minimum level of mispricing in our setting is 10%. This level is economically very significant. If investors can perceive extreme mispricing, but levels such as 10% are not sufficiently extreme as to be detectable, this would have significant ramifications for asset pricing models. Repeating the analysis only for observations where mispricing is greater than 20% in absolute value yields similar results.

These results suggest that participants are unable to assess mispricing. As an additional check, we next examine how expected returns vary with perceptions of mispricing. If a participant understands our experiment, then expected returns should be lower when beliefs about mispricing are higher.

Table 8 regresses expected returns on perceived mispricing. We find a statistically significant negative coefficient. When participants perceived mispricing to be lower, they expected future returns to be lower. Column 2 includes the actual level of mispricing. We find no meaningful impact for this variable and the insignificant point estimate is positive, not negative. This result is consistent with agents meaningfully responding to questions consistently across both returns and expected mispricing. However, agents have minimal ability to detect economically sizable mispricing.

D.4 Alternative Experiments

Perhaps agents struggle to understand the magnitude of mispricing, but are able to perceive its direction. For example, agents may miss nuance, but correctly identify that the setting with 20% mispricing

is overpriced and -20% mispricing is underpriced. To examine whether this is the case, we ask whether participants can correctly identify the direction of mispricing. We examine whether participants provided a positive number in a scenario with positive mispricing, a negative number with negative mispricing, and 0 in a setting with the correct price.

While the ability to predict mispricing is not zero, neither is it particularly large. In our setting, someone choosing at random would get the direction correct 1/3 of the time. Participants do better than random, correctly guessing 36% of the time (t-statistic of 1.88). Participants correctly identify cases where there is no mispricing at a lower rate (8%) than overpricing (50%) or underpricing (42%). Thus it could be that this number is so low largely because of difficulty identifying cases where there is no mispricing. To account for such an issue we next limit ourselves only to settings with under or over pricing. In these cases, participants correctly identify mispricing 49% of the time, insignificantly different from the benchmark of random.

To further examine whether a simpler setting may improve accuracy, we ran a separate experiment where the price could be the actual price, 25% higher or 25% lower. Participants were instructed this was the case. After seeing a given scenario subjects were simply asked if the displayed price was the true price, manipulated too high or manipulated too low.

We find generally similar results to the baseline. Guessing at random would yield correct guesses 1/3 of the time, or about 33%. In this simple setting subjects improve upon this accuracy by being correct 36% of the time (t-stat that it's equal to 1/3 of 3.67).

D.5 Direct Reporting

We designed our experiments with the goal of giving participants the best chance at meaningfully responding to induced mispricing, if it were possible for them to do so. While we find minimal skill at actually completing this task, it is possible that this reflects poor design of the task itself.

Thus for the sophisticated asset manager sample, we asked:

> *Do you believe that it is possible for investors to use information on economic fundamentals from a given day to determine the market price on that day within 10% of the true value?*

Most of the asset managers (68%) responded that this was not possible. This suggests that the responses reflect the fact that investors cannot know the level of the market without seeing it, rather than that the task was poorly designed.

Finally, for those who thought it was possible, we wanted to know whether the information we provided was sufficient. Indeed, if we designed the task well, but did not give the requisite statistics, this could explain our results as well. Thus, those who said the task was possible we asked:

> *Do you believe that the information on economic fundamentals provided in this survey (risk-free rate, inflation, prior GDP growth, expected GDP growth and volatility) were sufficient to determine the market price within 10% of the true value?*

90% of these respondents said yes. Thus, it is not obvious that there was more data that we could have provided for more accurate responses.

This sample is comprised of deep pocketed, sophisticated asset managers familiar with financial data and financial markets. The fact that most of these respondents do not think it is possible to use fundamentals to understand the current level of the market suggests that our asset pricing models are missing a fundamental aspect of how financial markets function.

IV. Real World Evidence

The information displayed to investors typically reflects what those investors consider important. Financial information tools, the financial press, and industry participants that market and sell financial products are each in the business of providing investors the information they would like to see. These are competitive businesses where success is dictated by successfully providing such information to investors.

This section explores the common sources of information that are provided to investors in an attempt to find examples of tools that are useful for identifying market mispricing. We examine the financial press, financial information tools (e.g. Bloomberg), investor educational tools, financial products, the statements of practitioners and the academic literature on market timing. We find that while there are a myriad of topics that are emphasized and an avalanche of data accompanying it, almost none of this focuses on identifying market mispricing.

While many investors seem to be largely unresponsive to the current level of the stock market, this is not universally the case. Of course, a number of people do attempt to understand what a reasonable current price for the market is, usually with a goal of forecasting future market returns. This is consistent with our theory framework and a positive theta. A variety of techniques are used, including some consisting of seemingly unsophisticated appeals to one's intuition. However, many such forecasts are based largely on the current level of the market price relative to some measure of

fundamentals. We ask whether such forecasts suggest that people, particularly sophisticated market participants, have strong feelings on precisely what the level of the market should be today and, further, whether they are able to use this information to meaningfully forecast returns.

Before returning to this question, we need guidance as to what such a valuation model can imply. Finance theory suggest there are three broad concepts that must be understood to know the current level of the stock market without seeing it and likely all three are relevant to some degree in the real world. The first two, cash flow and (rational) discount rate information, are consistent with standard efficient markets. If markets are not efficient, an investor would also need to understand the behavioral bias or market frictions inducing this deviation as a third component of the price level.

Under classic efficient market theory with constant discount rates, cash flow information is the only information that an investor needs to know to understand the current market price. If markets are efficient, then the price of the market simply reflects these fundamentals and future expected returns are always the same. Thus, to understand what the current market price should be, of course investors need to understand the current market fundamentals as they relate to cash flows.

The second is discount rate news. Perhaps people are currently very risk averse and demand a high future return. Thus, the market has a lower price all else equal. A market participant would need to understand these dynamics to understand the current level of the market price since the price would be lower than if cash flow news was the same, but investors were relatively more risk tolerant.

Finally, the market could be mispriced relative to these efficient market forces. Perhaps investors are overly optimistic and the market price is too high relative to efficient benchmarks. Perhaps institutions are stretched and forced to sell off assets at a discount, resulting in a lower price than these benchmarks. Regardless of the specific mechanism, this mechanism must be understood to know the current market price.

Thus, to understand the current level of the market without seeing it, someone needs to understand each of these in turn. Further, someone predicting future market returns would be quite interested in understanding which mechanism accounts for the current level of the market price. A high price based on cash flow news should not lead to different forecasts of future returns. Slow moving discount rate movements likely predict future returns with different dynamics and time horizons than behavioral biases, which in turn likely has different predictability than when institutional frictions will be relaxed.

With that in mind, we now turn to the actual forecasts that are made and note that such a certainty

of explaining current valuation ratios is absent. To a first approximation, these forecasts compare current valuation ratios to historic valuation ratios without regard for what mechanism caused them to achieve those levels.

For one somewhat representative example, we take Vanguard's Capital Markets model, a proprietary model that makes forecasts of future returns for a wide array of asset classes. One particular focus is on the level of the US stock market. The major discussion for the 2024 forecasts was the CAPE ratio, which was at the 99th percentile of historical value. The forecast notes how extremely high this is and states that it is about 30% higher than what their model suggests the price of the market should be.

What does it make of this extreme though? It shades down expected returns slightly, to 4.3% per year as compared to 5.7% the prior year, but higher than the 3.3% from two years prior.[6] And how confident are they that they really understand what's going on? Not very. They state that "while valuations are undoubtedly high right now, that doesn't mean they can't go higher in the near term. They are not a market-timing tool. And even over extended periods, valuations are not infallible predictors of outperformance or underperformance." So what is their advice? They "would not encourage investors to make drastic changes to their asset allocation." This example is consistent with a small weight on market fundamentals and a small positive theta. Even in an extreme setting such as the 99th percentile of valuation ratio, there is little confidence that their model of the right value of the stock market is the correct one.

Another recent example comes from Goldman Sachs.[7] In a report from October 2024, Goldman predicted that annual returns over the next ten years should be 3%. This was meaningfully below the expectations reported by other firms and received significant media attention. While the Goldman Sachs model was largely similar to other models, such as including P/E, it differentiated itself by also including a variable for market concentration.[8] The report discusses how excluding this variable would result in a forecast of 7% returns per year, roughly the average of the other forecasters.

So, how confident should an investor be that Goldman's model is correct and the other models are wrong? Empirically, the relationship between concentration and market returns is marginal. Includ-

[6] 2024: https://corporate.vanguard.com/content/corporatesite/us/en/corp/articles/investment-economic-outlook-feb-2024.html; 2023: https://institutional.vanguard.com/content/dam/inst/iig-transformation/insights/pdf/2022/venmo-outlook-2023-beating-back-inflation.pdf and 2022: https://www.ch.vanguard/content/dam/intl/europe/documents/en/usd-vanguard-economic-and-market-outlook-2022.pdf

[7] https://www.gspublishing.com/content/research/en/reports/2024/10/18/29e68989-0d2c-4960-bd4b-010a101f711e.html

[8] In addition, their model includes variables for recession frequency, 10-year treasury yield, and the S&P 500 ROE.

ing the variable in the model from 1930-2014 yields a t-stat of -1.8 (compared to -6.1 for P/E). Thus, its inclusion is difficult to motivate empirically. In practice, other financial firms could choose to use such a variable in their model, but they had not yet done so. The report also notes that in the past Goldman Sachs had not thought to include the variable. It was not present in their original 2012 model. In 2012, Goldman predicted returns of 8% per year which ex-post was too low as realized returns were 13.3%. Including concentration in the prior model would have predicted returns of 14% over this period. Given the relatively weak predictability of the variable over the whole sample, it seems likely that the choice to include concentration was driven at least in part by trying to do something with respect to the large recent realized forecast error, regardless of whether such a change was warranted to improve future forecasts.

A noteworthy aspect of these models is they are not trying to solve the problem that asset pricing models say they should be focusing on. In other words, they are not trying to directly predict the future level of the market based on fundamentals. For example, the Goldman model puts returns on the left hand side and regresses these on a number of variables, most of which are not influenced by the current level of the market price (e.g., market concentration or profitability). This seems much closer to the agnostic component of demand as opposed to fundamental based. The model does include a P/E ratio which is related to the level of the market price. With that said, P/E can predict returns in different ways (e.g., irrational bubble versus rational discount rate variation), which would likely lead to different predictions of next period's market level. There is no attempt to differentiate what P/E is capturing. Overall, the models seem to have some variables capturing market level, that subtly influence return expectations. This is consistent with demand largely being price agnostic, but with a theta that is non-zero, but small.

The theme that guessing the level of the market price should be a component, but only a minor component, of investing decisions is commonly echoed. For example, (Asness et al., 2015), discuss market timing based on fundamentals as a "sin" and that you should "sin a little." In our agnostic demand framework, we would translate this as a small positive theta. Of course when things get more extreme, even a small positive theta should have an impact. So when the dot com bubble increased, Paul Samuelson suggested "market timing is an investing sin, and for once I recommend that you sin a little" (as cited in (Asness et al., 2015)). Jack Bogle, a famous proponent of buy and hold without regard to the level of the market also decreased his exposure to the stock market at this time.[9]

[9] https://www.youtube.com/watch?v=k6ra5POdsYg

We note that what makes the dot com bubble "different" is that these investors had a clear sense of why valuation ratios were so high, which was behavioral mispricing. They did not think it was plausible that the prices could be justified based on fundamentals or discount rates and this lead the period to be a notable exception. Most situations are much less extreme. Even in the current period where some metrics look very high relative to historical standards, many are cautious as to whether this truly represents deviations from what the market price should be.

Finally we note that while forecasts of the level of the market exist, they are not particularly common or attention grabbing. If a core component of investing decisions was driven by things like the CAPE ratio, or any number of other metrics that could be used to forecast market level prices, it would be likely that such statistics and tools to forecast them would be more commonly displayed to investors.

In general, such statistics or tools are almost never shown. The Wall Street Journal displays thousands of statistics for many securities every day. The vast majority of statistics are related to past price changes, volatilities, volume, relationship to 52 week high and such. Clicking on a market index leads to more statistics about past prices. The print version highlights these measures as well as a P/E measure, though the P/E measure is not highlighted or included in any graphs.

One interesting aspect of a world of price agnostic demand is that it is similar in many ways to the classic concept of market efficiency without time varying returns. Of course, under this theory, market returns were constant and prices were efficient in that they reflected all cash flow news, so differences in realized returns were totally unpredictable. Thus, if markets functioned in this way, there would be no need to ever try and predict future returns based on the level of the market. In price agnostic demand, the lack of return predictability is basically the same, but the mechanism is quite different.

V. Discussion

The evidence in this paper suggests that investors do not have a strong opinion on what the level of the market price should be. While theoretically inconsistent with standard models, does this actually matter for our understanding of financial markets? Are there important implications to be gleaned from price agnostic demand, or can we continue to use models "as if" investors understand the level of the market, even though they do not.

Perhaps the most obvious implication is that if a large component of demand is price agnostic, arbitrage forces must be much weaker than is commonly assumed. If traders lack conviction on what

the level of the stock market should be, then investors will not actively trade to correct perceived mispricing in the stock market. Without such investors, there is no reason to expect that the market needs to be particularly close to its correct level in the short run. Recall, Vanguard claims the market is 30% higher than what it thinks it should be. Under standard models, this would represent a compelling opportunity to actively trade. Vanguard suggests traders take no action based on this (in the context of the model) extreme mispricing.

Loosening this arbitrage constraint can help to explain a number of other findings in finance. For example, Shiller et al. (1981) finds that market prices are far too volatile to be explained by fundamentals. This was a major challenge to asset pricing models based on fundamentals based prices. If there is a wide range of prices that a given investor will accept as normal though, than it is likely that prices can move much more than underlying fundamentals.

This is closely related to the recent literature that highlights the importance of trading in markets, the influence of price pressure and the inelastic market hypothesis (e.g. Gabaix and Koijen (2021), Hartzmark and Solomon (2021)). Gabaix and Koijen (2021) note that there is a relatively small amount of arbitrage capital dedicated to trading against mispricing. Agnostic demand presents an explanation, namely that investors generally don't hae the knowledge to implement such strategies.

An open question in Gabaix and Koijen (2021) is why investors allocate money the way that they do. Even if markets are inelastic, why wouldn't households allocate money in a manner consistent with fundamentals, such that markets behaved as if they were elastic and efficient. Our paper suggests a solution. Investors simply aren't allocating money based on the current level of the market because they lack conviction on whether it's correct.

This also helps to explain why market timing strategies are not particularly profitable. While over long enough horizons, fundamentals are indeed important, over shorter term horizons (over which investors typically evaluate their performance), investors are not allocating money based on these fundamentals. They do not drive performance because people are not trading strongly based on them.

Combining price agnostic demand with evidence of the lack of profitability of market timing strategies and evidence of inelastic markets suggests a reevaluation of the horizon over which classic fundamental based asset pricing models can explain. Historically, the field has been comfortable thinking such models are relevant over periods such as a month, a quarter, or a year. These results suggest this is not the case though. While the constraints of fundamentals do indeed seem somewhat

relevant, the horizon of this relevance is likely closer to the 5-10 year horizons over which fundamentals do exhibit strong predictability. A large component of this explanation lies in the ubiquity of price agnostic demand.

VI. Conclusion

Our paper presents evidence that investors have very little conviction of what the level of the stock market should be. This suggests that their demand for assets and their expected returns largely adapts to whatever the current level of the market may be.

In our experimental setting we find that participants are quite poor at identifying large deviations from actual market prices. This leads to an interesting question of whether it would be possible (absent knowledge of what the variables actually were on those pass dates) to mathematically back out the correct answer given the information provided. The answer is no. This is basically a restatement of the fact that most market movements are explained by discount rate news.

With that said, we think our results suggest the need to reevaluate the finding that most market movements cannot be explained by fundamentals. The conclusion drawn from this literature is that price movements were actually rationally driven by discount variation. Under these models market participants understand these variations which explain the current market price. If such mechanisms explained our results, then we would not be providing the relevant state variables to market participants. However, they could have requested them from us and did not. If you thought the market level was largely pinned down by the perceived risk of a rare disaster, or time variation in risk aversion, such variables could have been requested. As an aside, proxies for such variables could be easily discussed in the financial media as well. No one suggested we should provide such a variable though, making it unlikely that such a mechanism explains our results.

Our explanation is much simpler. Pinning down the level of the market price is perceived to be very difficult. This means that there are a large range of prices that people are willing to accept as normal. When things get extreme relative to these expectations, such as during the tech bubble, fundamental variables do begin to have relevance for expected returns. For a wide range of values though, demand largely adapts to the prevailing market price. This fundamental aspect of financial markets is largely absent from modern asset pricing models, but has wide ranging implications for how markets function.

References

Andre, Peter, Philipp Schirmer, and Johannes Wohlfart, "Mental models of the stock market," 2023.

Asness, Clifford, Antti Ilmanen, and Thomas Maloney, "Back in the Hunt," *Institutional Investor*, 2015.

Ben-David, Itzhak and Alex Chinco, "Expected EPS$\times$ Trailing P/E," Technical Report, National Bureau of Economic Research 2024.

Bender, Svetlana, James J Choi, Danielle Dyson, and Adriana Z Robertson, "Millionaires speak: What drives their personal investment decisions?," *Journal of Financial Economics*, 2022, *146* (1), 305–330.

Bergman, Abigail, Alex Chinco, Samuel M Hartzmark, and Abigail B Sussman, "Survey curious? start-up guide and best practices for running surveys and experiments online," *Start-Up Guide and Best Practices For Running Surveys and Experiments Online (October 05, 2020)*, 2020.

Black, Fischer, "Noise," *The journal of finance*, 1986, *41* (3), 528–543.

Chinco, Alex, Samuel M Hartzmark, and Abigail B Sussman, "A new test of risk factor relevance," *The Journal of Finance*, 2022, *77* (4), 2183–2238.

Choi, James J and Adriana Z Robertson, "What matters to individual investors? Evidence from the horse's mouth," *The Journal of Finance*, 2020, *75* (4), 1965–2020.

Cochrane, John H, "Explaining the variance of price–dividend ratios," *The Review of Financial Studies*, 1992, *5* (2), 243–280.

_ , *Asset pricing: Revised edition*, Princeton university press, 2009.

Dehaene, Stanislas, "The neural basis of the Weber–Fechner law: a logarithmic mental number line," *Trends in cognitive sciences*, 2003, *7* (4), 145–147.

Desvousges, William H, F Reed Johnson, Richard W Dunford, Sara P Hudson, K Nicole Wilson, and Kevin J Boyle, "Measuring natural resource damages with contingent valuation: tests of validity and reliability," in "Contributions to economic analysis," Vol. 220, Elsevier, 1993, pp. 91–164.

Gabaix, Xavier and Ralph SJ Koijen, "In search of the origins of financial fluctuations: The inelastic markets hypothesis," Technical Report, National Bureau of Economic Research 2021.

Hartzmark, Samuel M and David H Solomon, "Marketwide predictable price pressure," *Available at SSRN 3853096*, 2021.

Hsee, Christopher K, Yuval Rottenstreich, and Zhixing Xiao, "When is more better? On the relationship between magnitude and subjective value," *Current Directions in Psychological Science*, 2005, *14* (5), 234–237.

Kahneman, Daniel, "Maps of bounded rationality: Psychology for behavioral economics," *American economic review*, 2003, *93* (5), 1449–1475.

_ **and Shane Frederick**, "A model of heuristic judgment," *The Cambridge handbook of thinking and reasoning*, 2005, *267*, 293.

_ **, Ilana Ritov, David Schkade, Steven J Sherman, and Hal R Varian**, "Economic preferences or attitude expressions?: an analysis of dollar responses to public issues," *Elicitation of preferences*, 2000, pp. 203–242.

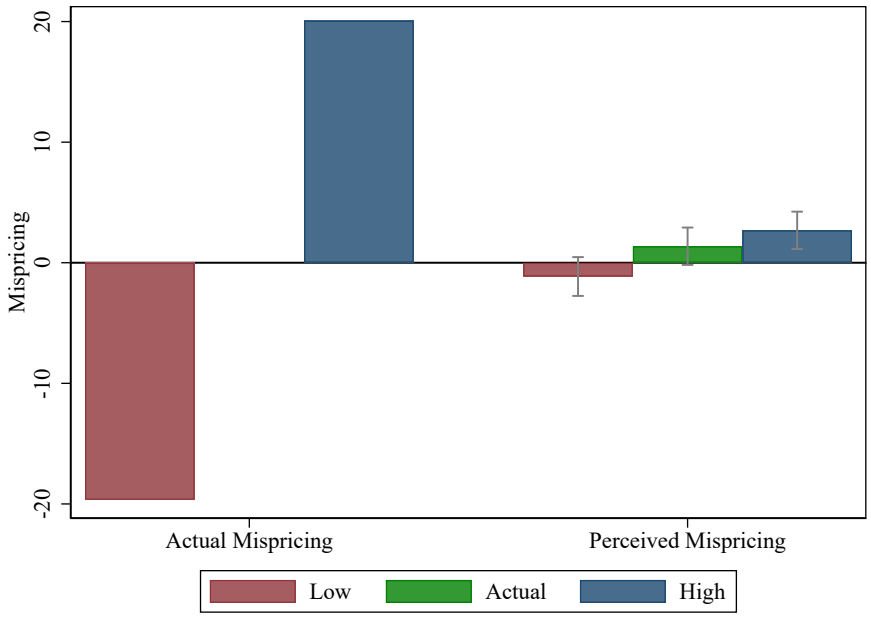
Shiller, Robert J et al., "Do stock prices move too much to be justified by subsequent changes in dividends?," 1981.

Shleifer, Andrei and Robert W Vishny, "The limits of arbitrage," *The Journal of finance*, 1997, *52* (1), 35–55.

Tversky, Amos and Daniel Kahneman, "The framing of decisions and the psychology of choice," *science*, 1981, *211* (4481), 453–458.

Figure 1: Actual vs. perceived Mispricing by direction

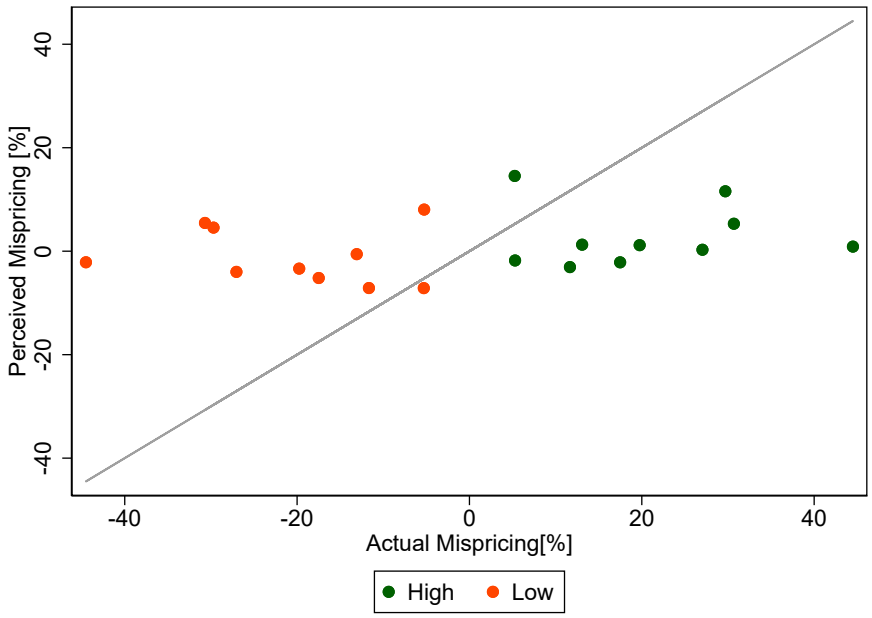
(a) Panel A: Actual vs. perceived mispricing



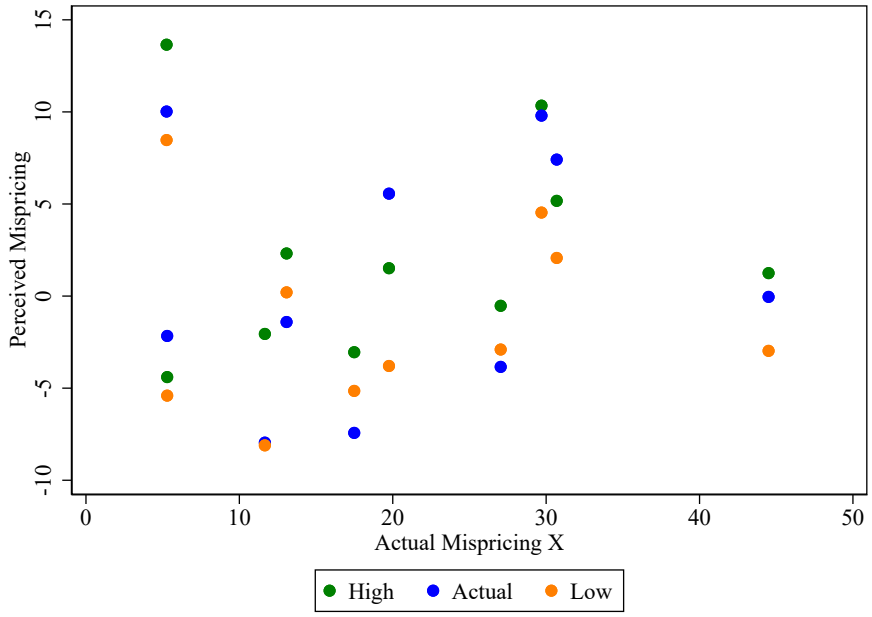
These figures shows the average mispricing by scenarios with negative mispricing (maroon), no mispricing (green) and positive mispricing (blue). The left three bars show the actual average amount of mispricing in each scenario. The right three bars show the perceived level. Grey bars indicate the 95% confidence interval.

Figure 2: Perceived Mispricing on actual Mispricing

(a) Panel A: Perceived Mispricing by actual Mispriciving



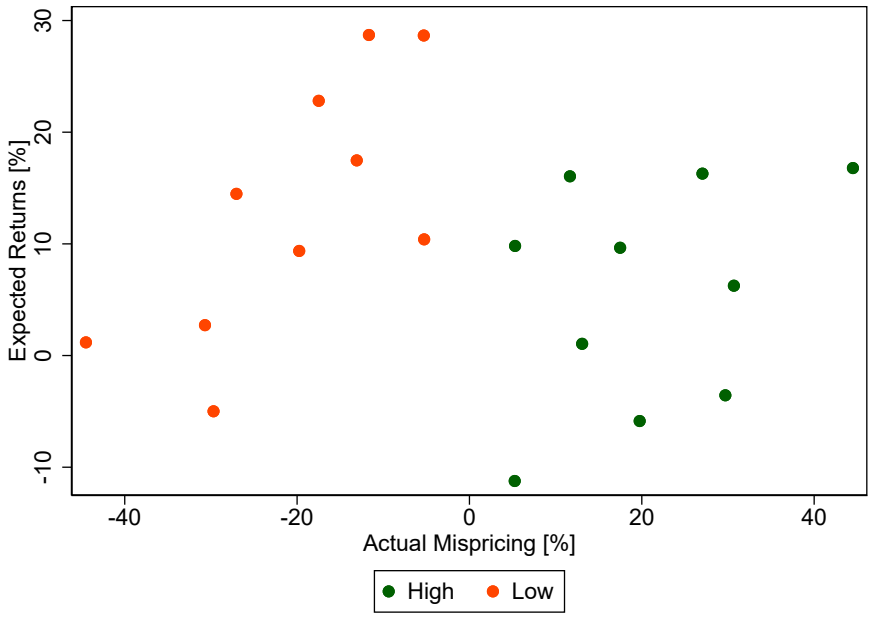
(b) Panel B: Perceived Mispricing by Date



These figures show how mispricing and expected returns vary with the actual level of mispricing. Each dot is the average answer for the given level of mispricing. Panel A shows the answer given for the level of perceived mispricing and the grey line shows the actual level of mispricing. Panel B shows perceived mispricing based on the level of X chosen for a given date, where for that date participants saw prices X% higher, X% lower or the actual level. Green dots display answers from scenarios where the displayed P/E ratio was above the actual, orange dots show answers from the actual scenario, and blue dots show answers from the low scenario.

Figure 3: Expected Returns on actual Mispricing

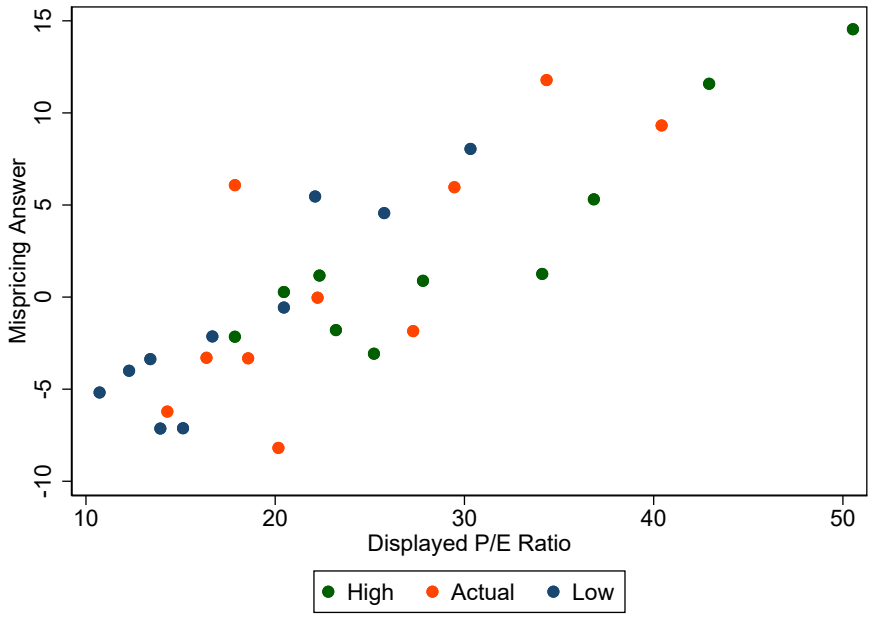
(a) Panel A: Expected Returns



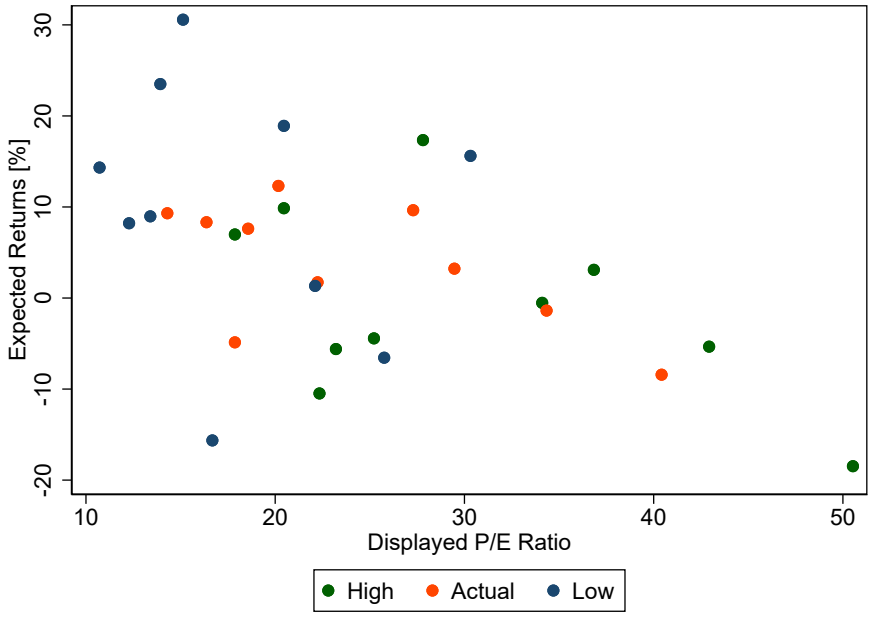
This figure show how expected returns vary with the actual level of mispricing. Each dot is the average answer for the given level of mispricing. Green dots display answers from scenarios where the displayed P/E ratio was above the actual, orange dots show answers from the actual scenario, and blue dots show answers from the low scenario.

Figure 4: Mispricing and Return variation on displayed P/E Ratio

(a) Panel A: Perceived Mispricing



(b) Panel B: Expected Returns



These figures show how mispricing and expected returns vary with the price to divided ratio that is displayed. Each dot is the average answer for the given level of P/E ratio. Panel A shows the answer given for the level of mispricing. Panel B shows the expected returns. Green dots display answers from scenarios where the displayed P/E ratio was above the actual, orange dots show answers from the actual scenario, and blue dots show answers from the low scenario.

Table 1: Perceived Mispricing

Actual Mispricing	All	Low	None	High
	(1)	(2)	(3)	(4)
Correct Direction	0.36	0.46	0.08	0.53
	(-67.79)	(-32.19)	(-99.23)	(-27.70)
Perceived Mispricing	0.95	-1.14	1.37	2.69
	(2.09)	(22.59)	(1.74)	(-22.01)
Error	0.97	18.51	1.37	-17.40
	(1.68)	(20.36)	(1.74)	(-19.76)
Absolute Error	23.40	26.93	18.15	24.87
	(67.66)	(43.24)	(38.17)	(39.12)
Observations	2552	873	827	852

This table shows summary stats on perceived mispricing. Column 1 shows all observations, column 2 shows scenarios with negative mispricing, column 3 shows scenarios with no mispricing and column 4 shows scenarios with positive mispricing. The first measure is whether the correct direction for mispricing was chosen (e.g. positive mispricing guessed in a high mispricing setting). The next measure is the mean of mispricing. Following that is the error (perceived mispricing minus actual) and the last is the absolute value of the error. t-statistics for the test of whether the measure is equal to the correct response are below each value.

Table 2: Perceived Mispricing varying with actual mispricing

	(1)	(2)	(3)	(4)
Mispricing	0.0774***	0.0877***	0.0706***	0.0804***
	(2.95)	(3.53)	(2.70)	(3.28)
YM FE	No	No	Yes	Yes
ID FE	No	Yes	No	Yes
R^2	0.00394	0.203	0.0589	0.257
Observations	2552	2552	2552	2552

This table shows regressions of perceived mispricing on actual mispricing. Columns 3 and 4 include date of observation fixed effects and columns 2 and 4 include ID fixed effects. Standard errors are clustered by ID with t-statistics in parenthesis.

Table 3: Perceived Mispricing by Subject Type

Panel A: Prolific

Actual Mispricing	All	Low	None	High
	(1)	(2)	(3)	(4)
Correct Direction	0.35	0.45	0.08	0.52
	(-60.16)	(-28.25)	(-89.08)	(-24.58)
Perceived Mispricing	0.89	-1.15	1.02	2.80
	(1.63)	(20.95)	(1.17)	(-20.77)
Error	0.84	19.27	1.02	-17.75
	(1.28)	(18.49)	(1.17)	(-18.29)
Absolute Error	23.29	27.36	17.75	24.62
	(59.45)	(38.02)	(33.60)	(34.69)
Observations	1980	665	649	666

Panel B: MBA Students

Actual Mispricing	All	Low	None	High
	(1)	(2)	(3)	(4)
Correct Direction	0.36	0.46	0.11	0.46
	(-25.15)	(-12.35)	(-28.39)	(-11.72)
Perceived Mispricing	-0.05	-0.56	2.10	-1.39
	(0.01)	(6.02)	(0.82)	(-7.06)
Error	0.14	15.15	2.10	-18.62
	(0.01)	(6.05)	(0.82)	(-6.97)
Absolute Error	25.57	26.56	21.30	28.13
	(26.23)	(16.56)	(14.01)	(15.34)
Observations	356	132	105	119

Panel C: Asset Managers

Actual Mispricing	All	Low	None	High
	(1)	(2)	(3)	(4)
Correct Direction	0.38	0.47	0.04	0.66
	(-18.56)	(-9.13)	(-40.99)	(-5.87)
Perceived Mispricing	3.15	-2.08	3.42	8.79
	(2.69)	(8.00)	(1.40)	(-4.72)
Error	3.76	17.75	3.42	-11.73
	(2.13)	(6.73)	(1.40)	(-3.78)
Absolute Error	20.82	23.68	17.15	21.57
	(19.31)	(12.35)	(12.04)	(9.97)
Observations	216	76	73	67

This table shows summary stats on perceived mispricing by subject type. Panel A shows prolific subjects, B shows MBA and C shows asset managers. Column 1 shows all observations, column 2 shows scenarios with negative mispricing, column 3 shows scenarios with no mispricing and column 4 shows scenarios with positive mispricing. The first measure is whether the correct direction for mispricing was chosen (e.g. positive mispricing guessed in a high mispricing setting). The next measure is the mean of mispricing. Following that is the error (perceived mispricing minus actual) and the last is the absolute value of the error. t-statistics for the test of whether the measure is equal to the correct response are below each value.

Table 4: Perceived Mispricing varying with actual mispricing

	Mturk		MBA		Asset Managers	
	(1)	(2)	(3)	(4)	(5)	(6)
Mispricing	0.0653**	0.0781***	0.103	0.0768	0.165**	0.167**
	(2.28)	(3.06)	(1.29)	(0.91)	(2.54)	(2.36)
YM FE	No	Yes	No	Yes	No	Yes
ID FE	No	Yes	No	Yes	No	Yes
R^2	0.00311	0.338	0.00347	0.232	0.0239	0.276
Observations	1980	1980	356	356	216	216

This table shows regressions of perceived mispricing on actual mispricing. Columns 1 and 2 show MTurkers, columns 3 and 4 show MBA students and columns 5 and 6 show Asset Managers. Even columns include date of observation fixed effects and ID fixed effects. Standard errors are clustered by ID with t-statistics in parenthesis.

Table 5: Expected Returns varying with actual mispricing

	(1)	(2)	(3)	(4)
Mispricing	-0.0533**	-0.0636***	-0.0411**	-0.0499***
	(-2.44)	(-3.05)	(-2.04)	(-2.75)
YM FE	No	No	Yes	Yes
ID FE	No	Yes	No	Yes
R^2	0.00232	0.224	0.192	0.413
Observations	2550	2550	2550	2550

This table shows regressions of expected returns on actual mispricing. Columns 3 and 4 include date of observation fixed effects and columns 2 and 4 include ID fixed effects. Standard errors are clustered by ID with t-statistics in parenthesis.

Table 6: Perceived Mispricing based on displayed data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
P/E	0.526***						0.264***
	(7.19)						(2.94)
SP Level		0.00381***					0.00189*
		(7.03)					(1.97)
Inflation			-0.442**				-0.305
			(-2.46)				(-1.25)
Risk Free				-1.510***			-0.260
				(-6.17)			(-0.91)
GDP					0.420***		-0.0721
					(4.27)		(-0.33)
Exp GDP						1.432**	-0.162
						(2.45)	(-0.21)
R2	0.246	0.249	0.200	0.236	0.203	0.199	0.254
Observations	2552	2552	2552	2552	2552	2552	2552

This table shows how perceived mispricing varies with the information displayed. In columns 1 through 6 the answer given for percentage mispricing is regressed on each of the variables indicated in the left most column. All regressions include subject fixed effects. Standard errors are clustered by subject, with t-statistics in parentheses.

Table 7: Expected Returns based on displayed data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
P/E	-0.580***						-0.213***
	(-15.30)						(-3.32)
SP Level		-0.00445***					-0.00426***
		(-15.04)					(-5.64)
Inflation			1.118***				0.290
			(7.46)				(1.33)
Risk Free Rate				2.034***			0.170
				(14.57)			(0.78)
GDP Growth					-0.0684		0.932***
					(-0.75)		(5.04)
Expected GDP Growth						-1.945***	-1.461**
						(-3.66)	(-2.33)
R2	0.348	0.364	0.258	0.374	0.234	0.238	0.418
Observations	1782	1782	1782	1782	1782	1782	1782

This table shows how expected returns vary with the information displayed. In columns 1 through 6 the answer given for percentage mispricing is regressed on each of the variables indicated in the left most column. All regressions include subject fixed effects. Standard errors are clustered by subject, with t-statistics in parentheses.

Table 8: Perceived Mispricing and Expected Returns on Actual Mispricing

Panel A: Mispricing and Returns on Actual mispricing

	Perceived Mispricing			Returns		
	(1)	(2)	(3)	(4)	(5)	(6)
Mispricing	7.850***	1.036	0.107	-0.145	10.64***	13.63***
	(3.04)	(0.37)	(0.04)	(-0.08)	(4.98)	(6.38)
P/E		0.204**	0.253**		-0.421***	-0.505***
		(2.03)	(2.46)		(-5.49)	(-6.64)
SP Level		0.00271***	0.00256**		-0.00204***	-0.00344***
		(3.38)	(2.56)		(-3.61)	(-4.59)
Inflation			-0.245			0.378*
			(-0.91)			(1.74)
GDP Growth			-0.240			1.097***
			(-1.02)			(5.81)
Risk Free Rate			-0.113			-0.134
			(-0.36)			(-0.60)
Expected GDP Growth			0.123			-1.689***
			(0.16)			(-2.68)
Controls	No	No	Yes	No	No	Yes
R2	0.279	0.331	0.335	0.233	0.376	0.432
Observations	1980	1980	1980	1782	1782	1782

Panel B: Returns on perceived and actual mispricing

	(1)	(2)
Perceived Mispricing	-0.140***	-0.141***
	(-3.19)	(-3.19)
Mispricing		0.966
		(0.51)
R2	0.259	0.259
Observations	1782	1782

This table shows how expected returns vary with the information displayed. In columns 1 through 6 the answer given for percentage mispricing is regressed on each of the variables indicated in the left most column. All regressions include subject fixed effects. Standard errors are clustered by subject, with t-statistics in parentheses.