Risk, return, and sentiment in a virtual asset market*

Maurizio Montone[1] and Remco C.J. Zwinkels[†2]

[1]Utrecht University

[2]Vrije Universiteit Amsterdam and Tinbergen Institute

December 2021

Abstract

The joint-hypothesis problem casts doubt on the results of market efficiency research. Specifically, it is hard to assess to what extent financial markets reflect economic fundamentals or mispricing. To address this issue, we study price formation in a large virtual asset market where fundamentals are predetermined and publicly known. We find that a number of well-established determinants of returns from the real world also affect asset prices in this market, despite the absence of systematic risk. The results suggest that prices in real financial markets include a substantial behavioral component, which is likely underestimated in canonical asset pricing tests.

Keywords: Asset pricing; Market efficiency; Natural experiment.

JEL Classification: G14; G41; C93

*We thank William Forbes, Clifton Green, David Hirshleifer, and seminar participants at Kingston University London, Utrecht University, the Behavioral Finance Working Group Conference at Queen Mary University of London, the ECASE seminar series at Erasmus University Rotterdam, and Vrije Universiteit (VU) Amsterdam for helpful feedback and suggestions. We are also grateful to Jonathan Jansen and Finn de Jong for excellent research assistance, and Jonathan Moed who was first to explore the idea of using price data from the FIFA internal player market to conduct asset pricing research.

†Corresponding author. De Boelelaan 1105, 1081HV Amsterdam, Netherlands. Email: r.zwinkels@vu.nl. Phone: +31 20 59 85220.

1. Introduction

The joint-hypothesis problem casts doubt on the results of market efficiency research (Fama (1970, 1991)). Specifically, it is hard to assess to what extent financial markets reflect economic fundamentals or mispricing. To identify the former, a common approach is to construct portfolios that capture sources of systematic risk that are otherwise latent (see, e.g., Fama and French (2004)). Notable examples are portfolios of stocks ranked on size or book-to-market (Fama and French, 1992), among many others. It is not entirely clear, however, whether such factors also represent mispricing to some extent (see, e.g., Daniel and Titman (1997); Stambaugh et al. (2012)).

A different approach is to try and identify mispricing directly, by modeling investor biases through a proxy called "sentiment". The idea is to capture instances in which economic agents hold unduly optimistic or pessimistic beliefs, i.e., not based on the facts at hand (Baker and Wurgler, 2006). This process typically involves the creation of an index. The most prominent examples in this sense are the indices of consumer sentiment (Carroll et al., 1994), and investor sentiment (Baker and Wurgler (2006, 2007); Baker et al. (2012)). Nonetheless, the issue with these measures is that they may reflect economic fundamentals (see, e.g., DeVault et al. (2019) for an excellent discussion).

In this paper, we propose a novel solution to the joint-hypothesis problem. We identify a credible and ex-ante distinction between economic fundamentals and mispricing, and quantify the relative impact of the two on asset returns. Specifically, we consider a large virtual asset market in which fundamentals are predetermined and publicly known, which allows us to directly and independently assess the impact of investor behavior on price formation. As such, this is a unique setting to carry out asset pricing tests.

We find three main empirical results. First, a number of factors that are known to price returns, such as size and book-to-market, represent mispricing rather than risk in this market. Second, investor sentiment has an independent and substantial impact on returns. Third, investor attention plays an important role in the price formation process. The results suggest that prices in real-life financial markets include a substantial behavioral component, which we estimate to account for about one-third of the overall asset volatility. Interestingly, the magnitude of this effect is close to that of the decrease in returns on asset pricing anomalies that follow academic publications (McLean and Pontiff, 2016).

Experimental studies also propose a similar approach, by recreating a simplified asset market in which fundamentals are predetermined and publicly known. The main finding of this line of research is that prices deviate from fundamentals, but eventually reach the rational equilibrium.[2] However, the structure of such markets features a small number of assets (typically one or two), participants (up to a few dozen), and trading sessions (30-60 max). As a result, this setup is not suitable for canonical asset pricing tests. Also, sentiment can only be defined endogenously.

On the other hand, betting markets also offer the opportunity to study market efficiency (Sauer, 1998). With respect to experimental markets, this setting features a large number of assets and participants. However, fundamentals are privately known (if at all), assets have a short time span (for example, a sports week), and transactions are zero-sum games (unlike stocks). These features make it hard to study standard asset pricing issues, and also make the definition of sentiment rather narrow if compared to that of other financial markets.[3]

In our study, we overcome these limitations. We analyze the price formation process from an in-play market for soccer players in the online video game FIFA 19, which shows many similarities in structure and demographics to real-life financial markets. In addition, it is a closed system in which the economic fundamentals are captured by a set of "ratings", i.e., predetermined scores for a number of player characteristics. These ratings are known to all participants, defined prior to the start of the market, and orthogonal to the physical world. To the best of our knowledge, this paper is the first to study the price dynamics of a video game's in-play market.

FIFA is a soccer simulator video game, and one of the most popular games worldwide. FIFA 19, the version of the video game we study, was sold approximately 20 million times and played by 36 million gamers. It is especially popular in Central and South America,

[2]See Forsythe et al. (1982, 1984), and Friedman et al. (1984), for relatively short-lived assets, and Smith et al. (1988) for relatively long-lived assets.

[3]For example, bettors tend to exhibit a preference for underdogs (Ali, 1977; Snyder, 1978; Asch et al., 1982; Asch and Quandt, 1987; Ziemba and Hausch, 1987; Golec and Tamarkin, 1991; Direr, 2011), or for specific teams (Avery and Chevalier, 1999; Kuypers, 2000; Levitt, 2004; Forrest and Simmons, 2008).

Eastern and Southern Europe, and the Middle East.[4] Although there are more sports-simulator video-games available, also with player markets, FIFA is by far the most popular. We specifically focus on the game mode FIFA Ultimate Team (FUT), which allows gamers to play against each other online. To participate in the game, users can exchange virtual FIFA money, known as "coins", and set up a squad by signing virtual soccer players.

At the beginning of the season, gamers participate in a primary market for players. FIFA itself provides the initial supply by selling random sets of players with mixed abilities, known as "packs", for a fixed fee. Afterwards, gamers can exchange players on a secondary market. The latter is a continuous market, with an open limit-order book. As such, it is similar to a stock exchange, with a fixed supply of assets in the short term.[5] Therefore, it also constitutes the focus of our study. The main difference with respect to an equity market is that short selling is not allowed, which constitutes a limit to arbitrage (Shleifer and Vishny, 1997).[6]

Attesting to the high utility gamers extract from playing video games, total leisure demand is very sensitive to innovations in leisure luxuries such as gaming computer use (Aguiar et al. (2021)). In addition to the utility of playing the game itself, however, gamers have two important financial incentives. First, they earn coins for each match they win. Second, they can sell their team to other users through an active market for FIFA accounts.[7] Specifically, every gamer is given a ranking in a general leaderboard depending on the performance against their opponents. The higher the position, the larger the monetary value of their team. These monetary and reputational incentives are key features of the game that make it comparable to actual financial markets.[8]

Finally, it is also important to notice that the demographics of gamers are comparable to those of real investors. The Entertainment Software Assocation (ESA) states that the average gamer is 34 years old, male, and college educated. Sports games such as FIFA

[4]See, e.g., Gamstat.com.

[5]The primary market remains open, but becomes less important after the start of the season.

[6]However, shorting is also fairly limited in the stock market itself (Jones and Lamont, 2002). A large fraction of institutional investors are prohibited from taking short positions (see, e.g., Chen et al. (2002)), and there are also additional mechanisms that induce pessimistic investors not to trade (see, e.g., Diether et al. (2002); Antoniou et al. (2016)).

[7]See, e.g., Playerauctions.com.

[8]Attesting to the relevance of the financial component of the game, there is even an active online community including several prolific influencers where gamers exchange tips on how to maximize returns from trading players (see, e.g., Futchief.com).

are especially popular amongst the age group from 34 to 54. Although it is impossible to observe the exact demographics of stock market participants, some studies suggest that they are quite similar to those reported above.[9]

To determine players' expected returns in the transfer market, we develop a simple theoretical model. For simplicity and without loss of generality, we consider a virtual transfer market with two types of soccer players, with high-skill and low-skill, respectively. As in the game, high-skill players are not only more talented than low-skill players, but also present in lower supply. As a result, they sell at a higher equilibrium price. Since ratings are fixed and predetermined, the two player types are uncorrelated assets.

Gamers are risk-averse, fully rational, and exhibit mean-variance utility over final wealth. The optimal portfolio of soccer players solves the following trade-off. On the one hand, high-skill players increase the quality of the team. On the other hand, their expensive nature reduces the size of the squad gamers are able to afford. In turn, a smaller squad provides less insurance against player injuries, suspensions, and fatigue during the game. Gamers then make two choices. First, they determine the optimal portfolio composition of high- and low-skill players. Second, they decide on the optimal amount of coin reserves to hold as a liquidity buffer, in case new players are needed.

In equilibrium, we show that all assets in this economy are priced by the security market line. Since fundamentals are predetermined and uncorrelated, however, there is no covariance risk across assets, and then the beta of any player simply reflects the ratio between the player's volatility and the volatility of the market portfolio. The mechanism is as follows. Players with lower ratings are more volatile, as there is more uncertainty associated with their performance.[10] As a result, they command a higher risk premium. The source of risk, however, is specific to the player, and therefore completely idiosyncratic.

This is an important result for two reasons. First, beta pricing makes portfolio analysis

[9] For example, Grinblatt et al. (2011) find an average age of approximately 40 years for the entire population of male Finnish retail investors. Furthermore, they show that stock market participation increases with IQ, suggesting that college-educated people are more likely to participate in the stock market. Bauer et al. (2009) use data from a large discount broker in the Netherlands, and find an average age of 45 years and a proportion of men of 75%. These estimates are also similar to those from the sample of Barber and Odean (2001), retrieved from a large U.S. discount broker.

[10] For example, players with a high "passing" score are more likely to pass the ball accurately to a teammate than those with a low score. Therefore, the outcome of the pass is more uncertain in the latter case.

meaningful in this market. Players with a higher beta, defined as above, should yield higher returns. Second, the absence of systematic risk makes factor-model extensions inapplicable to this market. Factors such as size and book-to-market may correlate with latent state variables in the real world (Fama and French, 1995, 2004). In this market there are no state variables, and then neither factor should have any impact on returns.

In the empirical analysis, we consider daily and weekly prices for the entire trading season of FIFA 19. The data set is retrieved from FutBin, and spans the period from September 2018 to September 2019, for a sum total of 361 daily observations per player.[11] To avoid potential liquidity biases, we only consider players above a certain quality level and that are still active in their real-life soccer careers. Of the 20,941 players available, then, the filtering results in 1,994 unique players.

In the first set of empirical tests, we construct a number of factor-mimicking portfolios that represent the cornerstone of modern asset pricing, such as those based on size and book-to-market (Fama and French, 1992), market beta (Frazzini and Pedersen, 2014), short-term mean reversion (Lo and MacKinlay, 1990), and volatility (Blitz and van Vliet, 2007). Interestingly, we find evidence for each of these asset pricing anomalies in our sample, even though there is by construction no change in the players' fundamentals during the sample period. In light of this, such anomalies seem to constitute genuine mispricing.

To get a sense of the magnitude, consider the book-to-market anomaly. The monthly Sharpe ratio of the long-short portfolio on U.S. value and growth stocks is 0.16.[12] In our weekly portfolios, which represents the closest analog to real-world monthly returns, we obtain an estimate of 0.49. The larger number reflects the absence of fundamental news, which reduces trading and thus the standard deviation of returns. To make a more instructive comparison, we transform these measures into coefficients of variation, and obtain estimates of 5.81 and 2.05, respectively. Their ratio suggests that mispricing may account for around one-third of the total volatility of the real-life value premium. We obtain similar estimates

[11] The latest edition of FIFA is released in September of each year. Gamers can still play the game and trade players after a new edition becomes available, but market liquidity and data availability decrease substantially.

[12] The sample period is from July 1926 through December 2019, and the data is retrieved from Kenneth French's website. For the sake of comparability with the FIFA market, we consider equal-weighted returns. This choice also addresses the concern that value-weighting partly conceals mispricing patterns (Baker and Wurgler, 2006).

when considering the size premium.

These results are based on univariate portfolio sorts. For robustness, we also estimate the effect of the characteristics introduced above on returns through both Fama-MacBeth regressions and fixed-effects regressions. We find that all characteristics are significantly related to expected returns, and continue to carry the same signs. The results are also robust to restricting the sample to the most liquid period, i.e., by excluding the first and last two months of trading. We also find similar results when excluding "superstar" players, defined as those with a rating above 90, and when using weekly instead of daily returns.

In the second set of tests, we introduce a measure of investor sentiment that is perfectly orthogonal to the players' value. We consider the daily news sentiment index from Buckman et al. (2020), and study its relation with asset prices. Consistent with previous studies, we find that investor sentiment is associated with a contemporaneous increase in asset returns followed by mispricing correction (Baker and Wurgler, 2006, 2007; Baker et al., 2012), and amplifies the magnitude of asset pricing anomalies (Stambaugh et al., 2012). The novelty of our findings lies in the fact that they stem from a measure of sentiment that is genuinely uncorrelated with asset fundamentals. As a result, we lend support to the credibility of the findings from the existing investor sentiment literature.

Following Da et al. (2011), we also test whether players that receive more attention are associated with a contemporaneous increase in returns. To this end, we alternatively define attention as Google search intensity for a given player, the cumulative number of real-life goals scored by a player until the period under consideration, and a measure of whether a player is performing exceptionally well in real life, defined as the instance in which they get selected in the FIFA "Team of the Week". It is important to note that none of these measures are related to the actual performance of players in the game, because the latter is entirely based on predetermined ratings. Nonetheless, we find evidence that all three measures of attention are associated with higher returns.

In a related analysis, we also explore the effect of attention-grabbing player characteristics on valuations. Specifically, we consider players from the four main European soccer leagues (i.e., England, Spain, Germany, and Italy), from a top team (i.e., one that ended up among the top three in the prior season), and whose on-field position is the most advantageous to score goals (i.e., striker). We find a significant contribution of all three attributes to player pricing. Combined, players that play at a top team, and in a main league, are up to 4.5% overpriced on average. All in all, these results confirm that market participants are prone to the attention bias, as attention-grabbing events from real life irrationally affect pricing in the game.

Finally, we study the dynamics of player returns by estimating a $GARCH(1,1)$ time-series model, introducing an extended set of control variables in the level equation for each player individually. We find a Monday effect in player returns (Keim and Stambaugh, 1984), as well as a reconfirmation of the positive effect of sentiment changes on returns (Baker and Wurgler, 2006), and a small negative daily autocorrelation in returns (Lo and MacKinlay, 1990). In addition, we find volatility clustering due to the GARCH effect, with coefficients that are qualitatively similar to those typically found for financial assets (Bollerslev, 1987). In addition, we find evidence for the leverage effect when estimating a more sophisticated $GJR-GARCH$ model, due to Glosten et al. (1993a). Overall, the dynamics of player returns show great statistical similarity to returns from real financial markets.

The remainder of the paper is organized as follows. Section 2 describes the video game with its transfer market. Section 3 introduces the theoretical model. Section 4 describes the data. Section 5 presents the empirical results. Section 6 concludes.

2. FIFA

2.1. General

The FIFA video games are a series of soccer simulators. A new version has been put on the market each consecutive year since the release of FIFA International Soccer, or "FIFA 94", in the weeks leading up to Christmas of 1993. The games are developed by Electronic Arts (EA) under their "EA Sports" label, which is also the creator of other notable sports simulators such as Madden NFL (American Football) and NHL (Hockey). These games have been released on multiple platforms, such as Xbox, PlayStation, PC, and Nintendo Wii. The FIFA series is the most popular video game produced by EA, as well as one of the most popular video games worldwide.

Global annual sales of the 2019 version of FIFA, known as "FIFA 19", amount to approximately 20 million units, the majority of which are for the PlayStation console. These units are sold at an average price of around $60, which implies sales revenues of roughly $1.2 billion from copies of the game alone. It should be noted that this estimate excludes any in-game additional downloadable content, which is made available by EA for a fee, and accounts for two-thirds of EA's total revenues.

EA is very protective of its non-financial information, and therefore does not usually share any statistics on usage or demographics. As a notable exception, on the occasion of the FUT five-year anniversary, EA stated in 2014 that the total number of unique users participating in FUT was 21,849,017. Users played around 264,000 matches per day over that year, and exchanged virtual soccer players a whopping 1.46 billion times.

Although the specific demographics of FIFA gamers are unknown, the Entertainment Software Association (ESA) publishes detailed demographics of gamers for the Unites States.[13] The average user is 33 years old, and roughly two-thirds of gamers are adults. A breakdown into age categories shows that 21% of gamers are under 18, 40% between 18 and 35, 18% between 36 and 49, and 21% above 50. Sports video-games such as FIFA are especially popular among gamers between 34 and 54 years old, of whom 54% are male, and 52% college educated.

Overall, there is a large number of participants in the FUT market. Furthermore, the age distribution suggests that market participants are mostly adults, male, and highly educated. In light of this, the population of gamers seems to be not too different from that of financial market participants.

2.2. FIFA Ultimate Team and the market

In this paper, we focus on FIFA 19. Specifically, we consider the game mode FUT and its internal asset market. In this feature, users compose their own teams and compete against each other in online matches. To participate in the game, gamers need to earn virtual FIFA money, known as "coins", and set up a squad of virtual soccer players.

At the beginning of the season, gamers participate in a primary market for players. FIFA

[13] See the *2019 Essential Facts About the Computer and Video Game Industry* published by the ESA.

itself provides the initial supply by selling random sets of players with mixed abilities, known as "packs", for a fixed fee. When a user buys a pack, all the players contained therein are added to their squad. Much like in real-life sports, users then need to pick a formation, select the best player for each position on the field, and come up with a starting line-up of 11 elements including a goalkeeper.

After setting up the initial squads, gamers can exchange players on a secondary market. This is a continuous user-driven exchange, with an open limit-order book. Users trade for two purposes. First, they can modify the composition of their squad by replacing players that do not suit their tactics with others that do. Second, they can try to sell their players for a profit, in order to earn more coins. There are elaborate guides available on the internet that provide information on how to trade effectively within the FUT market, with the goal of optimizing profits from trades. Many of these guides refer to the basic principles of microeconomics and the dynamics of financial markets, and encourage users to expand their knowledge on these topics.

In addition to the utility of playing the game, gamers have two important financial incentives. First, they earn coins for each match they win. Second, they can sell their team to other users for real-world currency through an active market for FIFA accounts. Specifically, every gamer is given a ranking in a general leaderboard depending on the performance against their opponents. The higher the position, the larger the monetary value of their team.

The players in FIFA are characterized by a number of attributes that reflect their real-life soccer skills. For example, the "pace" score determines how fast a player is, and the "shot" score indicates the accuracy of the player's shooting efforts. Other relevant characteristics are "passing", "dribble", "defense", and "physique". These attributes are fixed, predetermined, and known to all gamers. When considered all together, they determine a player's overall rating, expressed as an integer number between zero and 100. In turn, the rating represent the player's "fundamentals". Hence, all fundamentals are publicly known in this market.[14]

It is important to note that the market is fragmented across different gaming platforms.

[14]Note that players can achieve performances slightly above or below their ratings depending on whether one or more teammates share their nationality, or play in the same club or league in real life. This effect, known as "chemistry", assigns an overall chemistry score to the team (between 1 and 100), and colors (green, yellow, or red) for pairwise chemistry between any two players on the team. The magnitude of the effect is limited, however, and should not affect our analysis in any obvious way.

For example, trading activity on the Xbox market is isolated from trading activity on the PlayStation or PC markets. Furthermore, the market runs for one year only. When a new edition of FIFA is launched, the market essentially stops running. Accounts cannot be transferred to the new edition, and each gamer starts with a small default endowment of coins and players. This feature makes the total supply of players exceptionally low at the start of the year, which induces gamers to turn to the primary market to build their squad.

3. Model

Players in FUT can be thought of as finitely-lived dividend-paying assets. Adding a player with a higher rating to a team increases the quality of the team, and thereby increases the probability of winning matches and the accompanying pay-off in the form of FIFA coins. Hence, the earned FIFA coins can be interpreted as dividends. Dividends are therefore stochastic (the outcome of a match is uncertain), and the expected value of dividends is conditional on the player's rating (a higher rating increases the probability of winning).

Based on the mechanics of the market described above, we derive a simple theoretical model that determines the expected returns for players. We consider a two-period game. At time 0, agents start with a zero endowment of players and need to build a squad. To participate in the transfer market, they need to purchase game-specific coins through cash. The thus attained amount of coins constitutes agent j's initial wealth, denoted by w_{0j}. The transfer market is characterized by perfect information, so there is nothing to learn from market prices.

For simplicity and without loss of generality, players can be of two types: high-skill or low-skill. High-skill players have a higher overall rating, defined as a numerical evaluation of their abilities, than low-skill players. Since ratings are fixed and predetermined, these player types are uncorrelated assets. The supply of high-skill players in the game is only a fraction of the supply of low-skill players. Due to the lower supply and higher ratings, high-skill players sell at a higher price.

At time 1, agents compete with each other in a series of games, and earn coins for each victory. At the end of each period, all teams are ranked in a leaderboard according to their overall performance. The final value of agent j's team, denoted by $\tilde{w}_{1j}$, is determined by the position of the team in the leaderboard, plus the amount of coins earned and/or saved (if any). This overall value can be exchanged for real cash.

Agents are risk-averse, and exhibit mean-variance utility over $\tilde{w}_{1j}$. The optimal portfolio of players solves the following trade-off. On the one hand, high-skill players increase the quality of the team. On the other hand, their high price reduces the size of the squad agent j is able to afford. In turn, a smaller squad provides less insurance against player injuries and suspensions during the game.

Agents then make two choices. First, they determine the optimal portfolio composition of high- and low-skill players. Second, they decide on the optimal amount of coin reserves to hold as a liquidity buffer, in case new players are needed. Holding coins earns no returns, i.e., the risk-free rate is zero.

Investor j solves:

$$\max_{\{\alpha_{Hj},\alpha_{Lj}\}} E[u_j(\tilde{w}_{1j})] = E(\tilde{w}_{1j}) - \frac{\gamma_j}{2} var(\tilde{w}_{1j}), \tag{1}$$

subject to:

$$\tilde{w}_{1j} = w_{0j}\left(1 + \sum_{i=L,H} \tilde{r}_i \alpha_{ij}\right), \tag{2}$$

where γ_j is the coefficient of absolute risk-aversion, and α_{ij} is the fraction of wealth invested in asset i. Note that high-skill players exhibit a higher first and second moment of returns (i.e., $\bar{r}_H > \bar{r}_L$ and $\sigma^2_H > \sigma^2_L$).

In Appendix A, we show that all assets in this economy are priced by the security market line. The betas, however, have a different interpretation than usual. Since there is no covariance risk across assets, the beta of asset i simply reflect the ratio between the asset i's volatility and the volatility of the market portfolio.

4. Data

We retrieve price data for all players traded in the FUT market of FIFA 19 from FutBin, a third-party website that aggregates FUT statistics. The sample period is from September 22, 2018, to September 19, 2019. In total, we obtain 361 daily observations for 20,941 players, coming from 156 countries, and playing at 632 clubs in 43 national leagues. Since the market is continuous, there are no opening or closing prices. For this reason, the prices made available by FutBin are recorded at midnight GMT.[15]

For liquidity reasons, we only consider the prices from the PlayStation market. We limit the analysis to currently active players with the highest ratings, labeled as the "gold" category, excluding famous players from the past that are no longer active in real life, known as "icons". Although we cannot observe turnover, we do observe that the prices of many players are stale on some days. In light of this, we eliminate all players with zero daily returns. In so doing, we obtain a final sample of 1,994 unique players with price observations for each of the 361 trading days.

To assess the relevance of goal-scoring players in the game, we introduce a dummy variable that takes on value one when a player is a striker.[16] We also observe "Team of the Week" (TOTW) selections, a weekly recurring event that takes place each Wednesday, where EA selects 23 players that are currently performing exceptionally well in real life. This selection is based on a relatively large set of variables, and fully carried out by EA.[17]

Finally, we create dummy variables indicating the prestige of the players' real-life club and league. A "Top Club" is defined as a club that ended among the top three of the domestic league at the end of the previous season (i.e., in 2018), whereas "Top League" includes the English, Spanish, German, and Italian leagues, which represent the most important national competitions in Europe according to the soccer governing bodies.

[15] Given that FIFA is predominantly played in Europe, this also represents a time of day with a fairly low level of volatility and liquidity.

[16] Specifically, this category includes players denoted as "Striker", "Left Forward", "Right Forward", "Center Forward", and "Center Attacking Midfielder".

[17] For the occasion, EA releases special versions of these exceptionally-performing players. These players temporarily exhibit additional skills, and trade separately from their standard version (so both are simultaneously available in the market). However, special players are rare, traded less, and only exist for a limited period of time, so we discard them from the analysis.

Figure 1 plots the equal-weighted average price across the 1,994 players in the sample, normalized to 100 on day 1, along with daily log returns.[18] As hypothesized in Section 3, there is an overall negative price trend. On the other hand, prices exhibit a steep increase at the beginning of the season, followed by a gradual negative trend afterwards. Specifically, the index loses approximately 40% of its value over the course of the year. The initial increase in price is caused by supply shortages, which induces gamers to turn to the primary market. The subsequent negative trend is explained by the decreasing horizon over which players can be utilized or traded. A final striking observation is the crash in prices of over 20% in August, which most likely reflects the thin trading that takes place towards the end of the market.

As in real financial markets, daily returns are volatile and centered around zero. In addition, there appear to be two clusters of high volatility and extreme returns: one at the start of the season (September), and another one at the end (August).

The high volatility of returns reflects the fact that there are at least three reasons for trading in this market. First, gamers may want to replace players that become unavailable due to injuries or suspensions (for example, when they receive a red card during the game). Second, gamers earn coins for each match they win, and can afford to sign better players as a result. Third, sophisticated traders may try to profit from players' mispricing.

Table 1 presents some descriptive statistics for prices and returns. In Panel A, we find a wide price range, from 350 coins to 2.5 million coins, with a high skewness to the right. Volatility is also substantial, as the standard deviation reaches the price level at the 75^{th} percentile. In Panel B, we find that average daily returns are very close to zero, but exhibit a wide range from +281% to -317%.[19] The distribution shows considerable excess kurtosis, as is the case for financial markets (Fama, 1965), and positive skewness, which is a pattern that is also observed for stock returns (Albuquerque, 2012).

In Table 1, Panel C, we consider player characteristics. Ratings, as defined by EA, range from 75 to 94.[20] The highest ranked players in the game are Cristiano Ronaldo and Lionel Messi (both 94), and the number of players per rating category is inversely related to the

[18] A rating-weighted index follows a similar pattern.

[19] A lower bound of less than -100% is due to the fact that we use log returns.

[20] The relatively high lower-bound is due to the fact that we focus on "gold" players.

Figure 1: Player prices and returns over time

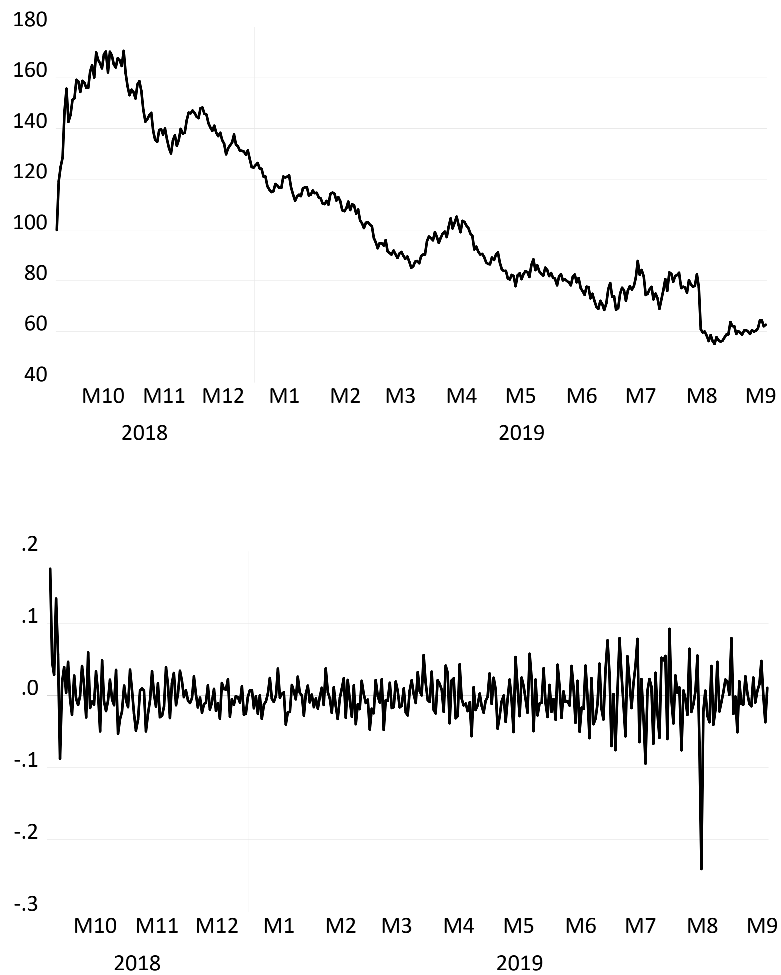

This figure plots the time series of daily average prices (top), and average log returns (bottom), of the 1,994 players included in our sample. The level of the price index is normalized to 100 on day 1, and the sample period ranges from September 22, 2018 to September 19, 2019.

Table 1: Descriptive statistics

	Min.	25%	50%	75%	Max.
	Panel A: Prices				
Mean	401.42	631.42	813.11	1,388.22	155,0434
Median	395.00	558.00	726.50	1000.00	1,587,479
Maximum	654.00	1,986.50	3,169.50	8,173.00	2,521,475
Minimum	350.00	400.00	406.00	700.00	483,542
Std. Dev.	23.19	199.53	322.97	1,054.55	644,541
Skewness	-0.52	1.73	2.68	3.76	13.33
Kurtosis	1.45	6.93	12.79	22.71	216.02
	Panel B: Returns				
Mean	-0.0096	-0.0008	0.0000	0.0013	0.0105
Median	-0.0169	0.0000	0.0012	0.0086	0.0460
Maximum	0.1223	0.7153	0.9491	1.2080	2.8104
Minimum	-3.1760	-1.2239	-0.9658	-0.7316	-0.1241
Std.Dev.	0.0283	0.1638	0.2176	0.2638	0.5096
Skewness	-6.8529	-0.2622	0.0107	0.2762	4.8639
Kurtosis	2.9233	5.0719	6.8660	10.7345	107.7861
	Panel C: Attributes				
	Rating	Top Club	Top League	Striker	TOTW
Mean	78.126	0.294	0.578	0.139	0.008
Median	77.000	0.000	1.000	0.000	0.000
Maximum	94.000	1.000	1.000	1.000	1.000
Minimum	75.000	0.000	0.000	0.000	0.000
Std.Dev.	3.280	0.456	0.494	0.346	0.087
Skewness	1.393	0.902	-0.315	2.088	11.338
Kurtosis	4.927	1.814	1.099	5.360	126.557

This table reports the descriptive statistics of the cross section of individual player prices (Panel A), returns (Panel B), and attributes (Panel C). In Panels A and B, we calculate the descriptive statistics individually for each player (over the time series), and for each statistic we report the minimum, first quartile, median, third quartile, and maximum of the cross-sectional distribution across players. In Panel C, we report aggregate descriptive statistics with a breakdown into relevant player attributes. "Rating" is the overall rating score of the player in the game; "Top Club" is a dummy that takes on value one when a player's real-life team ended among the top three of the domestic league at the end of the previous season; "Top League" is a dummy variable that takes on value one if the player's real-life team is in the English, Spanish, German, and Italian leagues, which represent the most important national competitions in Europe according to the soccer governing bodies; "Striker" is a dummy variable indicating attacking players; and "TOTW" indicates the selection players by EA into the team of the week, based on exceptional real-life performance.

rating. The percentage of players from top leagues and top clubs is relatively high, and equal to 57.8% and 29.4%, respectively. This is explained by the fact that we exclude lower-rated players, which in turns implies a lower probability of playing for a top team or in a top league. Strikers represent 13.9% of the players in the game, and TOTW selections account for 0.8% of the player/day observations. Note that the TOTW dummy is the only player characteristic that is time-varying.

5. Empirical analysis

Our empirical analysis proceeds as follows. We start by estimating the relation between prices and fundamentals in Subsection 5.1. Subsequently, we proceed with the analysis of returns in Subsection 5.2, sentiment in Subsection 5.3, attention in Subsection 5.4, and additional properties of returns in Subsection 5.5.

5.1. Prices

As a preliminary analysis, we estimate the relation between player prices and ratings. To this end, we acknowledge the fact that the supply of players is inversely related to their ratings. By construction, the relation between supply and rating is concave. Furthermore, since top players are in high demand and limited supply, prices should be a convex function of ratings. Hence, the test equation is as follows:

$$log(P_{i,t}) = \alpha + \beta_1 Rating_i + \beta_2 Rating_i^2 + \eta_t + \varepsilon_{i,t}, \tag{3}$$

in which $P_{i,t}$ is the price of player i in period t, $Rating_i$ is the rating of player i, and η_t represents time fixed effects (either daily or weekly).

The estimates are in Table 2, Panel A. In column (1), we start out with a simple linear model without the squared term. We find a positive and highly significant relation between a player's price and its rating. In column (2), we estimate the full equation. Consistent with our priors, we observe a positive estimated coefficient on $Rating^2$, which indicates a convex relation. Also, introducing the squared term improves the model fit from 56.6% to 69.5%, which suggests that the relation between prices and ratings is indeed non-linear.

Table 2: Prices and ratings

	Panel A: Full sample				Panel B: November-July			
	Linear	Convex	Time FE	FMB	Linear	Convex	Time FE	FMB
Rating	0.253***	-3.954***	-3.954***	-3.954***	0.253***	-4.096***	-4.096***	-4.096***
	(35.200)	(-19.009)	(-19.005)	(-85.475)	(34.706)	(-19.183)	(-19.180)	(-90.061)
$Rating^2$		0.026***	0.026***	0.026***		0.027***	0.027***	0.027***
		(19.973)	(19.970)	(90.195)		(20.112)	(20.108)	(95.313)
$Adj.R^2$	0.566	0.695	0.717	0.723	0.589	0.732	0.752	0.755
Obs	719,834	719,834	719,834	719,834	544,362	544,362	544,362	544,362

	Panel C: Rating <90				Panel D: Weekly data			
	Linear	Convex	Time FE	FMB	Linear	Convex	Time FE	FMB
Rating	0.238***	-4.989***	-4.989***	-4.989***	0.247***	-4.022***	-4.022***	-4.022***
	(33.082)	(-29.309)	(-29.304)	(-90.354)	(30.774)	(-17.108)	(-17.110)	(-32.100)
$Rating^2$		0.033***	0.033***	0.033***		0.027***	0.027***	0.027***
		(30.438)	(30.433)	(94.077)		(17.963)	(17.965)	(33.860)
$Adj.R^2$	0.517	0.664	0.690	0.699	0.550	0.685	0.704	0.712
Obs	714,780	714,780	714,780	714,780	103,688	103,688	103,688	103,688

This table presents the estimation results of Equation (3). Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively. Standard errors are two-way clustered; t-statistics in parentheses.

Then we perform two additional tests. In column (3), we re-estimate the test equation with time fixed-effects, to allay the concern that a time trend such as the one observed in Figure 1 may potentially drive the results. Reassuringly, we find similar estimates. In column (4), finally, we estimate Fama-MacBeth regressions. The estimates are analogous to those from the time fixed-effects regressions, but the model fit increases to 72.3% due to the time-variation in the β coefficients.

Next, we perform a number of robustness checks. In Panel B, we re-estimate the test equation in the subsample of observations from November 1 through July 31. The later start-date roughly captures the moment in which the primary market has exhausted its function as the main supplier of players, and the earlier end-date excludes the period of thin trading that takes place towards the end of the season. The results show that the estimated coefficients remain qualitatively similar, with an even higher goodness of fit.

In Panel C, we address the concern that the convex relation between prices and ratings may be driven by a small number of star players. To this end, we exclude players with a rating above 90. Reassuringly, we find similar results. In Panel D, finally, we estimate the model using weekly rather than daily data, in order to reduce the effect of potential noise in daily prices. The estimates are virtually unchanged.

Overall, the results show that up to 75% of the variation in prices is explained by ratings, and the relation is quite stable across time and players. To a large extent, then, the market seems to be pricing players correctly. To look further into this issue, we turn to the analysis of returns.

5.2. Returns

Market beta

We start the analysis of returns by estimating a fifty-day rolling market-beta for each player:

$$r_{i,t} = \alpha + \beta r_{m,t} + \varepsilon_{i,t}, \tag{4}$$

where $r_{i,t}$ is the log-return of player i on day t, and $r_{m,t}$ is the log-return on the market index, defined as an equal-weighted index of all players in the sample, and normalized to 100 on the first day of the sample period.[21] Subsequently, we rank players on their market beta and form decile portfolios, rebalanced either daily or weekly.[22]

The results are in Table 3, Panel A. Based on the portfolio sorts, we find that the market beta is not priced in the cross-section of player returns. Also, there is no discernible pattern across deciles. The 10-1 spread is an insignificant 0.6%. Below, we shed further light on this pattern in our analysis of investor sentiment.

Size and book-to-market

The empirical asset pricing literature has proposed a number of factors to explain the cross-section of asset prices. The most well-known and influential attempt in this respect is arguably the Fama and French (1993) three-factor model, which essentially augments the CAPM with a size and a book-to-market factor. Next, we construct these additional factors in our sample, and include them in the analysis.

We start with size. Although we do not have data on the supply of players in the game, we can measure size using the price level $P_{i,t}$ itself. The reason is as follows. The supply of high-skill players is a fraction of the supply of low-skill players ($n_H < n_L$). However, the highly convex relation between prices and ratings more than compensates this effect ($P_H > P_L$), thus making high-skill players "large caps" ($n_H P_H > n_L P_L$). Also, the supply of player i as a proportion of the total supply of players, n_i/N, is fixed.[23] Therefore, the time series variation of size is only determined by the player's market price.

In light of these considerations, we rank stocks based on the price level, sort them in decile portfolios, and rebalance them either daily or weekly. The returns on these portfolios are in Table 3, Panel B. Except for portfolios D1 and D10, returns monotonically decrease with size. The 10-1 spread is a significant 5.5% with daily rebalancing, and a marginally insignificant 2.6% with weekly rebalancing. For a better comparison with equity markets, and specifically with the SMB factor from Fama and French (1993), we also calculate the top-50-minus-bottom-50 spread. With daily (weekly) rebalancing, this spread equals a large

[21] Note that the risk-free rate is zero in this market.

[22] The results that follow are similar when replacing raw returns with CAPM alphas.

[23] This feature is hard-wired in the game, as the probability that a pack contains a player with a given rating is fixed and decreases with the rating.

Table 3: Returns on portfolios sorted on beta, size, and book-to-market

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D10-D1	T50-B50
Panel A: Beta												
Daily	0.001	0.002	0.002	0.003	0.002	0.001	0.001	-0.001	-0.002	-0.005	-0.006	-0.003
	(0.188)	(0.658)	(0.474)	(0.808)	(0.354)	(0.152)	(0.131)	(-0.099)	(-0.236)	(-0.560)	(-0.697)	(-0.697)
Weekly	0.001	0.001	0.002	0.002	0.001	0.001	0.001	0.001	0.000	-0.001	-0.002	-0.001
	(0.093)	(0.157)	(0.156)	(0.203)	(0.041)	(0.052)	(0.045)	(0.030)	(-0.019)	(-0.057)	(-0.131)	(-0.105)
Panel B: Size												
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D10-D1	T50-B50
Daily	0.033	0.038	0.036	0.028	0.011	-0.010	-0.022	-0.029	-0.052	-0.022	-0.055***	-0.057***
	(6.649)	(6.019)	(4.853)	(3.833)	(1.585)	(-1.471)	(-3.069)	(-3.923)	(-8.271)	(-6.995)	(-9.434)	(13.681)
Weekly	0.016	0.016	0.014	0.011	0.007	0.001	-0.006	-0.013	-0.026	-0.011	-0.026	-0.024***
	(1.140)	(0.900)	(0.723)	(0.596)	(0.396)	(0.050)	(-0.326)	(-0.707)	(-1.584)	(-1.308)	(-1.656)	(6.634)
Panel C: Book-to-Market												
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D10-D1	T30-B30
Daily	-0.070	-0.039	-0.026	-0.009	0.007	0.018	0.026	0.031	0.033	0.034	0.104***	0.078***
	(-12.046)	(-5.911)	(-4.130)	(-1.411)	(1.009)	(2.757)	(4.153)	(5.595)	(6.830)	(8.623)	(19.437)	(18.355)
Weekly	-0.031	-0.017	-0.008	-0.002	0.004	0.008	0.011	0.015	0.015	0.015	0.046***	0.034***
	(-2.115)	(-1.086)	(-0.501)	(-0.099)	(0.222)	(0.450)	(0.651)	(0.933)	(1.074)	(1.302)	(3.335)	(9.207)

This table presents the average daily returns on portfolios sorted on market beta (Panel A), size (Panel B), and book-to-market (Panel C), with daily and weekly rebalancing. D10-D1 denotes the decile 10 minus decile 1 portfolio; T50-B50 denotes the top 50 minus bottom 50 percentile portfolio. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively, with t-statistics in parentheses.

and statistically significant 5.7% (2.4%). This inverse relation between market capitalization and returns is consistent with the evidence from equity markets.

Next, we look into the book-to-market ratio. A well-known empirical regularity of financial markets is that stocks with a high book-to-market ratio outperform stocks with a low book-to-market ratio. This spread, known as the value premium, has generated a heated debate on whether it constitutes a genuine measure of risk, or also represents mispricing. In the analysis that follows, we shed light on this issue. Although players in FUT do not formally have a "book" value, we can calculate the fundamental value based on the model from Equation (3). Specifically, we calculate the book-to-market value per player as:

$$BM_{i,t} = (\alpha + \beta_1 Rating_i + \beta_2 Rating_i^2) - ln(P_{i,t}), \tag{5}$$

using the estimated coefficients α, β_1, and β_2 from the Fama-MacBeth regressions in Table 2, such that the overpricing measure does not suffer from a look-ahead bias.[24] Then we sort players on $BM_{i,t}$, and create decile portfolios. Again, we rebalance the portfolios either daily or weekly.

The returns on these portfolios are in Table 3, Panel C. We find a monotonically increasing pattern, as returns are positively associated with the book-to-market ratio. The 10-1 spread portfolio earns a significant average return of 10.4% with daily rebalancing, and 4.6% with weekly rebalancing. For a better comparison with equity markets, the HML factor from Fama and French (1993) to be precise, we also calculate the top-30-minus-bottom-30 spread. With daily (weekly) rebalancing, this spread equals a significant 7.8% (3.4%). Again, these results are consistent with those from the real world.

To get a sense of the magnitude, consider the book-to-market anomaly. Using U.S. stock data from July 1926 through December 2019, the monthly Sharpe ratio of the equal-weighted long-short portfolio on value and growth stocks with 30% breakpoints is 0.16. In our equal-weighted weekly portfolios, which represents the closest analog to real-world monthly returns, we obtain an estimate of 0.49. The larger number reflects the absence of fundamental news, which reduces trading and thus the standard deviation of returns. To make a more instructive comparison, we transform these measures into coefficients of variation, and obtain estimates

[24]Note also that the goodness of fit of these regressions is quite high (above 70%).

of 5.81 and 2.05, respectively. The ratio between the latter and the former is 0.35, which suggests that mispricing may account for around one-third of total volatility of the real-life value premium.

The results are similar when considering the top and bottom 10% breakpoints, and when we replace the book-to-market ratio with cashflow-to-price or earnings-to-price. We also analyze the HML factor from Fama and French (1993), which is defined in terms of value-weighted returns and includes a second sort on size. The monthly Sharpe ratio of HML is 0.10, which implies a coefficient of variation of 9.81. Compared with the coefficient of variation from the FIFA sample (2.05), the estimates suggest that mispricing may account for around one-fifth of total volatility of the value-weighted book-to-market spread.

Next, we consider the size anomaly. Using 10% breakpoints, the monthly Sharpe ratio of the real-world long-short portfolio on small and large stocks is 0.11, which implies a coefficient of variation of 9.40. Its counterpart from the FIFA sample exhibits a Sharpe ratio of 0.23, and thus a coefficient of variation of 4.27. The ratio between the latter and the former coefficient of variation equals a whopping 0.45, which implies that mispricing may characterize about half of the volatility of the real-world size premium for extreme decile portfolios.

We also analyze the SMB factor from Fama and French (1993), which is defined in terms of value-weighted returns and uses the 50% breakpoint. SMB exhibits a monthly Sharpe ratio of 0.06, which implies a coefficient of variation of 16.09. Its counterpart from our sample exhibits a coefficient of variation of 2.82, which constitutes 27% of the magnitude of the real-world coefficient. Specifically, the number suggests that a little more than one-fourth of the real-world value-weighted size premium may reflect mispricing.

Mean reversion and volatility

Apart from the three factors from Fama and French (1993), we can construct two more factors in this market: short-term mean reversion (Lo and MacKinlay, 1990), and low volatility (Blitz and van Vliet, 2007). The typical momentum strategy considers a lookback period of three months to a year (Jegadeesh and Titman, 1993). That is not possible in the FUT market because it only runs for one year. Therefore, here we focus on short-term mean reversion

(Lo and MacKinlay, 1990). Specifically, we sort players on their return from the previous day (week), form decile portfolios, and rebalance them daily (weekly).[25] To test for the low volatility effect in our data, we rank players on weekly return volatility, sort them in decile portfolios, and rebalance them either daily or weekly.

The results for mean reversion are in Table 4, Panel A. We find that the short-term mean-reversion effect is quite strong. The long-short portfolio yields returns of 7.0% with daily rebalancing, and 6.8% with weekly reblancing. The results are mainly driven by deciles 8 to 10, which exhibit particularly low returns.[26]

The results for volatility are in Table 4, Panel B. We find a strong low volatility premium, with an almost monotonic decrease in returns across deciles. Although the magnitude of the premium is not as large as for the other factors, the long-short portfolios still yields 5.0% with daily rebalancing, and 1.6% with weekly rebalancing.

Characteristics

Next, we study the direct effect of the characteristics introduced above on returns through Fama-MacBeth regressions. First, we include one characteristic at a time. The estimates are in Table 5, columns (1) to (5). The results are similar to those from the portfolio sorts. We find that market beta is not significantly related to the cross-section of player returns, book-to-market has a positive effect on returns, whereas size, past returns, and volatility exert a negative effect.

In column (6), we include all five factors simultaneously to investigate whether their cross-correlations help explain the results. We find that all coefficients keep their sign and significance, although the magnitude decreases slightly for all factors. In column (7), we find similar results when we restrict the sample to the trading period from November through July. In column (8), we also find similar estimates when excluding players with a rating above 90 from the analysis. Finally, in column (9), we obtain similar estimates when using weekly data, actually with slightly higher statistical significance, presumably due to the fact

[25]Experiments with a lookback and holding period of one month gives qualitatively similar results, but with slightly lower excess returns.

[26]In fact, the spreads are mostly driven by deciles 9 and 10 for most factors. This might be a result of the short sale constraints in this market, causing negative information to be impounded into prices more slowly (Beber and Pagano, 2013).

Table 4: Returns on portfolios sorted on past returns and volatility

	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D10-D1
Panel A: Mean reversion											
Daily	0.004	0.013	0.018	0.019	0.018	0.014	0.006	-0.003	-0.018	-0.066	-0.070***
	(0.569)	(2.215)	(3.296)	(3.828)	(3.851)	(2.915)	(1.191)	(-0.568)	(-2.635)	(-6.884)	(-8.145)
Weekly	0.018	0.015	0.012	0.011	0.009	0.006	0.003	-0.004	-0.015	-0.051	-0.068***
	(2.544)	(2.297)	(2.067)	(1.981)	(1.621)	(1.230)	(0.506)	(-0.664)	(-2.371)	(-7.278)	(-10.871)
Panel B: Low volatility											
	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	D10-D1
Weekly	0.003	0.005	0.005	0.004	0.002	0.001	0.000	0.000	-0.004	-0.012	-0.016**
	(1.069)	(1.565)	(1.037)	(0.711)	(0.358)	(0.083)	(-0.004)	(0.014)	(-0.569)	(-1.469)	(-2.117)
Daily	0.010	0.014	0.015	0.011	0.008	0.004	0.001	-0.006	-0.016	-0.037	-0.050***
	(4.492)	(4.642)	(3.737)	(2.332)	(1.415)	(0.586)	(0.093)	(-0.829)	(-1.869)	(-4.177)	(-5.671)

This table presents the average daily returns on portfolios sorted on previous returns (Panel A), and volatility (Panel B), with daily and weekly rebalancing. D10-D1 denotes the decile 10 minus decile 1 portfolio. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively, with t-statistics in parentheses.

Table 5: Fama-MacBeth regressions of returns

Dependent: Ret	1	2	3	4	5	6	7 (M11-M07)	8 (<90)	9 (weekly)
Beta(-1)	-0.001					0.000	0.001	0.000	-0.001
	(-0.726)					(0.022)	(0.726)	(0.037)	(-0.239)
Log(BM(-1))		0.460***				0.340***	0.351***	0.376***	0.538***
		(19.056)				(15.874)	(14.642)	(14.730)	(5.495)
Log(P(-1))			-0.017***			-0.008***	-0.008***	-0.005**	-0.029***
			(-8.251)			(-5.005)	(-4.407)	(-2.460)	(-4.502)
Ret(-1)				-0.120***		-0.092***	-0.074***	-0.093***	-0.372***
				(-8.834)		(-7.540)	(-5.373)	(-7.668)	(-19.277)
Log(Vol(-1))					-0.016***	-0.009***	-0.011***	-0.009***	-0.019***
					(-4.502)	(-3.984)	(-4.223)	(-3.910)	(-4.018)
$Adj.R^2$	0.048	0.046	0.041	0.060	0.053	0.172	0.181	0.176	0.285
Obs	618,140	618,140	618,140	618,140	618,140	618,140	522,428	604,810	87,736

This table presents the results of Fama-MacBeth regressions of player returns on a number of characteristics. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively, with t-statistics in parentheses.

that weekly data contain less noise.

For robustness, we repeat the analysis by estimating fixed-effects regressions. Specifically, we control for time-invariant player-specific characteristics and potential time trends.[27] The estimates are in Table 6. Reassuringly, we find that the results are similar to those from the Fama-MacBeth analysis. The main difference between the two sets of findings is that the coefficient of market beta is negative and significant in fixed-effects regressions when including all regressors simultaneously, which is consistent with the low-beta anomaly (Frazzini and Pedersen, 2014). Furthermore, the significance of size decreases somewhat when using weekly data.

Overall, the returns on players in FUT exhibit dynamics similar to equities both in the time series and the cross section. Given that the fundamental values are publicly known, and there is no state variable in this economy, return factors seem to reflect mispricing rather than risk.

5.3. Sentiment

Another way to detect mispricing is to try and identify investor biases directly, by capturing instances in which economic agents hold unduly optimistic or pessimistic beliefs. Studies of this sort, however, are also sensitive to the joint-hypothesis problem, as it is hard to rule out the possibility that sentiment contains information on economic fundamentals.

In this paper, we consider a measure of sentiment that is perfectly orthogonal to the asset market under consideration. Specifically, we take the news sentiment measure from Buckman et al. (2020). The two main advantages of this measure are that it captures the general level of sentiment in the media at large, not just financial news, and it is available at the daily frequency for our entire sample period. We first study the direct effect of sentiment on stock returns (Baker and Wurgler, 2006, 2007), and then the interaction effect between sentiment and the characteristics introduced above (Stambaugh et al., 2012).

The results are in Table 7. We find that sentiment has an important effect on the price formation process in FUT. In column (1), the direct contemporaneous effect of changes in

[27]Note that player fixed-effects capture the time-invariant n_i/N, which allows us to interpret P as a genuine measure of size.

Table 6: Fixed-effects regressions of returns

Dependent: Ret	1	2	3	4	5	6	7 (M11-M7)	8 (<90)	9 (weekly)
Beta(-1)	-0.002					-0.004**	-0.005***	-0.004**	-0.009
	(-1.316)					(-2.490)	(-2.886)	(-2.551)	(-1.471)
log(BM(-1))		1.068***				0.818***	0.785***	0.861***	2.049***
		(22.000)				(7.832)	(5.988)	(7.474)	(6.156)
log(P(-1))			-0.126***			-0.035**	-0.057***	-0.030*	-0.038
			(-19.146)			(-2.537)	(-3.270)	(-1.964)	(-0.705)
Ret(-1)				-0.124***		-0.056***	-0.027	-0.056***	-0.271***
				(-9.026)		(-3.894)	(-1.598)	(-3.871)	(-10.313)
log(Vol(-1))					-0.014***	-0.003**	-0.006***	-0.003*	-0.019***
					(-10.580)	(-2.054)	(-3.831)	(-1.941)	(-3.303)
$Adj.R^2$	0.195	0.259	0.255	0.218	0.211	0.258	0.281	0.260	0.336
Obs	591,607	674,062	674,062	674,062	674,062	591,607	505,319	587,267	84,074
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Player FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

This table presents the results of panel regressions of player returns on a number of characteristics. All models include both player and time fixed effects. Standard errors are clustered at both the player and the time level. Statistical significance at the 10%, 5% and 1% level is denoted as ** and ***, respectively, with t-statistics in parentheses.

Table 7: Sentiment

Dependent: Ret	1	2	3	4	5	6	7
dSENT	0.070***						
	(3.705)						
dSENT(-1)		-0.303***					
		(-16.503)					
dSENT(-1)*Beta(-1)			-0.074***				
			(-6.250)				
dSENT(-1)*log(BM(-1))				-1.172***			
				(-4.139)			
dSENT(-1)*log(P(-1))					0.170***		
					(13.297)		
dSENT(-1)*Ret(-1)						-1.453***	
						(-9.941)	
dSENT(-1)*log(Vol(-1))							-0.035***
							(-3.160)
$Adj.R^2$	0.066	0.066	0.258	0.258	0.258	0.258	0.258
Obs	591,607	591,607	591,607	591,607	591,607	591,607	591,607
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	No	No	Yes	Yes	Yes	Yes	Yes
Player FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

This table presents the results of panel regressions of player returns on a number of characteristics and changes in sentiment ("dSENT"), defined as the daily news sentiment index from Buckman et al. (2020). The direct effects of the factors are included in the model, but excluded from the table for the sake of brevity. All models include both player and period fixed effects. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively, with robust t-statistics in parentheses.

sentiment on returns is positive and highly significant. In column (2), lagged changes in sentiment are negatively related to returns. Combined, these two dynamics suggest that the effect of sentiment shocks on returns is reversed in the subsequent period, as in Baker and Wurgler (2006).

In columns (3) to (7), we find that sentiment strongly mediates the relation between characteristics and returns. In column (3), we find that the CAPM relation becomes more negative in times of high sentiment. This pattern is consistent with the idea that sentiment traders undermine an otherwise positive mean-variance trade-off during high sentiment periods (Yu and Yuan, 2011), overbidding for high-beta stocks especially in the presence of leverage constraints (Frazzini and Pedersen, 2014). Overall, then, the insignificant results from the previous unconditional analysis of market beta seem to hide important conditional patterns.

In columns (4) and (5), we find that high sentiment is causing a slower mispricing correction of the size and book-to-market spreads, a well-known effect in equity markets (see, e.g., Baker and Wurgler (2006)). In columns (6) and (7), the results suggest that past returns and volatility have a stronger effect on returns in times of high sentiment, which likely reflects the well-known asymmetric trading pattern across up and down market states.

5.4. Attention

Apart from sentiment, attention is another important phenomenon in the behavioral asset pricing literature (Da et al., 2011). An often-used exogenous measure of attention for a particular item is its search intensity on Google, as measured by Google Trends. In one of the first applications of Google Trends in finance, Da et al. (2011) find that higher search volume is associated with higher prices, among others.

In this paper, we analyze the effect of attention on player prices by relating a player's search volume to price formation in FUT. The hypothesis is that gamers push up the price for players that come under the spotlight.[28] To test for this, we augment the model from the previous subsection with player attributes that are likely to attract attention.

[28]Note that attention could theoretically also push prices down. However, given the short sale constraints in this market, we expect the net effect to be positive.

For each player, we obtain weekly search-intensity data from Google Trends. Google does not provide the absolute number of searches per item, but rather a measure that is normalized over time and space. In light of this, the results depend on the exact time period and geographical area. We select the period from September 1 2018 through September 1 2019, and set the geographical area to global. The search intensity per player is only relative to the player itself, so we cannot compare search intensity across players.[29] The inclusion of player fixed-effects, however, provides an econometric solution to this issue.

In addition to the Google Trends index, the set of attention-grabbing attributes we include is the (cumulative) number of goals a player scores in real life in the national league, and an indicator of whether the player is performing exceptionally well in real life as assessed through the TOTW selection.[30] Since Google Trends data and TOTW selections are weekly, and league matches are typically also weekly events, we carry out this analysis at the weekly frequency.

The results are in Table 8. We find evidence consistent with our conjecture. In column (1), an increase in attention as measured by search volume on Google is associated with an increase in returns. In column (2), the inclusion in the TOTW selection exerts a positive and significant effect on returns. in column (3), we observe a positive association between returns and goals scored in real life by a player.

In columns (4) to (6), we perform a few additional tests. We show that the attention effects of search volume, form, and goals do not capture the same type of attention, as all three effects survive each other's inclusion. Furthermore, the results are similar when restricting the sample to the period from November through July, and when excluding star players from the analysis.

Next, we study the relation between attention and valuations, where the latter is defined as the market-to-book value of a player ($MB_{i,t}$). Specifically, we expect players that attract high levels of attention to exhibit high valuations. To test for this, we run a panel regression of the market-to-book ratio on a number of attention-grabbing attributes, including a set of dummy variables for top leagues, top clubs, and strikers. We also consider interaction

[29]Google Trends does provide an option to compare search terms, but this is only possible with up to five terms. Thus we cannot compare the 1,994 players in our sample at the same time.

[30]Goal-scoring statistics are from Gracenote.

Table 8: Attention

Dependent: Ret	1	2	3	4	5 (M11-M07)	6 (<90)
dGT	0.014**			0.010	0.015**	0.011
	(2.120)			(1.533)	(2.165)	(1.551)
TOTW		0.415***		0.416***	0.414***	0.424***
		(18.612)		(18.414)	(17.039)	(18.360)
Goals			0.005**	0.006**	0.006**	0.006**
			(2.205)	(2.231)	(2.290)	(2.216)
Beta(-1)	-0.009	-0.009	-0.009	-0.010	-0.012*	-0.010
	(-1.500)	(-1.529)	(-1.485)	(-1.572)	(-1.788)	(-1.596)
log(BM(-1))	1.963***	2.027***	2.040***	1.925***	1.788***	2.005***
	(5.640)	(6.031)	(6.112)	(5.484)	(4.626)	(5.278)
log(P(-1))	-0.061	-0.040	-0.040	-0.065	-0.095*	-0.056
	(-1.119)	(-0.739)	(-0.739)	(-1.199)	(-1.740)	(-0.945)
Ret(-1)	-0.264***	-0.271***	-0.270***	-0.264***	-0.278***	-0.264***
	(-10.276)	(-10.339)	(-10.284)	(-10.269)	(-10.745)	(-10.287)
log(Vol(-1))	-0.021***	-0.019***	-0.019***	-0.021***	-0.028***	-0.021***
	(-3.880)	(-3.293)	(-3.317)	(-3.890)	(-7.128)	(-3.833)
$Adj.R^2$	0.342	0.344	0.337	0.350	0.351	0.351
Obs	82,425	84,074	84,074	82,425	71,259	81,823
Player FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

This table presents the results of panel regressions of player returns on a number of characteristics and attention-grabbing attributes. "dGT" represents changes in the Google Trends index for a given player; "TOTW" is a dummy that takes on value one when a player is selected in the EA team of the week; "Goals" represents the cumulative number of goals a player has scored in real life up to period under consideration. All models contain both player and period fixed effects. Standard errors are clustered at both the player and the time level. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively, with t-statistics in parentheses.

terms between these variables, for example because a striker that plays for a top club is likely to attract more attention than a striker that plays for a mid-table club, other things being equal. Since these are time-invariant player characteristics, we do not include player fixed-effects in this analysis.

The results are in Table 9. In column (1), we find that players active in a top league have on average a significant 1.5% higher valuation level than players active in other leagues. In column (2), players from top teams exhibit on average a 1.4% higher valuation than other players. In column (3), the average valuation of strikers is an insignificant 0.4% lower than that of non-strikers. These findings become slightly stronger when including all variables in the model simultaneously, in column (4).

In column (5), we study the interaction effects between these attributes. First, the direct effect of being a striker becomes positive, although still insignificant. Furthermore, we find that all interactions are important. A player active in both a top league and a top club exhibits an additional increase in valuation of 0.9%. The positive striker effect decreases for strikers in top leagues (-1%). The same holds for a striker from a top club (-2%). Hence, only strikers in lesser-known teams enjoy a valuation premium. Finally, the triple interaction of being a striker from a top club in a top league leads to a valuation increase of 2.0%. The total effect of being a striker for a top club in a top league is an overvaluation of 4.5%. Non-strikers in top clubs in top leagues exhibit an average overvaluation of 4.9%.[31]

Overall, attention-grabbing attributes significantly affect player prices. Whereas the direction of causality in financial assets can be unclear, in this case the attributes under consideration are purely exogenous to rational price formation. Thus, we can be reasonably certain that it is the attention for players that increases prices.

5.5. Return properties

Returns on financial assets exhibit a number of stylized properties, such as skewness and kurtosis, as discussed in Section 4, but also volatility clustering (Engle, 1982). In this section, we take a closer look at the statistical properties of returns in the FUT market.

[31] Note that Striker or interactions including Striker are not always significant stand-alone. An F-test shows that combined, however, including the Striker dummy including its interactions adds to the explanatory power of the model.

Table 9: Attention and valuations

Dependent: log(M/B)	1	2	3	4	5	6 (M11-M07)	7 (<90)	8 (Weekly)
Top League	0.015***			0.021***	0.019***	0.020***	0.020***	0.020***
	(5.541)			(7.216)	(5.362)	(5.937)	(5.562)	(5.101)
Top Club		0.014***		0.021***	0.020***	0.021***	0.020***	0.018***
		(4.365)		(6.221)	(4.005)	(4.328)	(4.059)	(3.660)
Striker			-0.004	-0.003	0.006	0.004	0.006	0.005
			(-1.384)	(-0.863)	(0.961)	(0.719)	(0.995)	(0.683)
League*Club					0.009	0.007	0.008	0.010
					(1.282)	(1.048)	(1.165)	(1.380)
League*Striker					-0.010	-0.010	-0.010	-0.009
					(-1.244)	(-1.295)	(-1.330)	(-1.156)
Club*Striker					-0.020**	-0.019**	-0.021**	-0.020**
					(-2.016)	(-1.970)	(-2.082)	(-2.016)
League*Club*Striker					0.020	0.019	0.017	0.020
					(1.346)	(1.335)	(1.163)	(1.424)
Rating	0.064***	0.076***	0.069***	0.069***	0.074***	0.074***	0.002***	0.073***
	(4.022)	(4.907)	(4.406)	(4.430)	(4.536)	(4.539)	(0.120)	(4.561)
Rating^2	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	(-4.130)	(-5.007)	(-4.471)	(-4.601)	(-4.701)	(-4.698)	(-0.229)	(-4.729)
$Adj.R^2$	0.018	0.015	0.010	0.029	0.032	0.039	0.022	0.032
Obs	704,311	704,311	704,311	704,311	704,311	532,623	699,257	101,452
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

This table reports the results of panel regressions of player valuations, defined as the market-to-book ratio, on a number of characteristics and attention-grabbing attributes. Top League is a dummy variable indicating whether the player plays in one of the five top leagues; Top Club is a dummy variable indicating whether the player plays for a top club, which is a club that ended in the top three of the league in the previous year; Striker is a dummy indicating whether the player is a striker as indicated by one of the forward positions in FIFA. Statistical significance at the 10%, 5%, and 1% level is denoted as *, **, and ***, respectively. Standard errors are two-way clustered; t-statistics in parentheses.

Figure 2: Autocorrelation

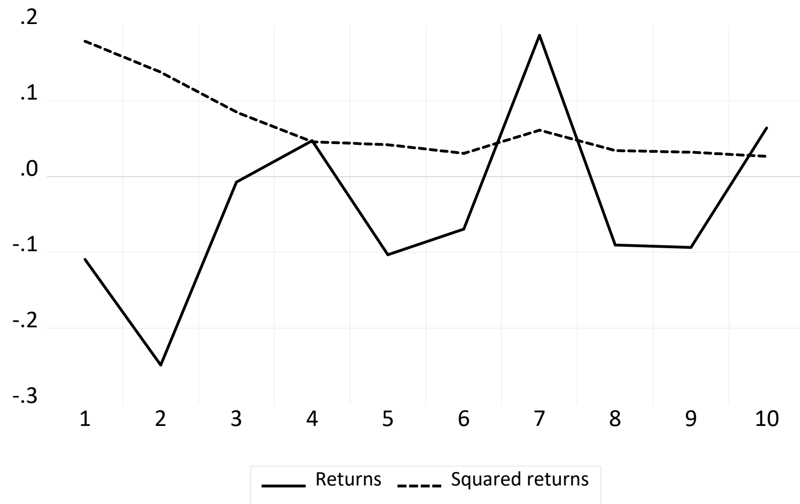

This figure shows the average autocorrelation in player returns and squared returns for the first 10 daily lags.

In Figure 2, we display the average autocorrelation in returns and squared returns over the 1,994 players in our sample for the first ten lags. The pattern is similar to that from equity markets: slightly negative at short lags, and quite random afterwards. The absolute magnitude of these correlations, however, is somewhat larger than their equity counterparts. The peak on day seven is particularly interesting, and could be an indication of day-of-the-week effects.

The autocorrelation pattern in squared returns is also comparable to that from equities, as it exhibits a long memory process. There is a very gradual decay in average autocorrelations. The difference with respect to equity returns is that the correlations appear to be somewhat lower on average for the shorter lags.

For a more formal analysis, we estimate the following time-series model for each player in our sample:

$$\begin{gathered} Ret_t = c + \Sigma d_i D_i + e_i X_{i,t-1} + \varepsilon_t \\ \sigma^2_{t+1} = \omega + \alpha\varepsilon^2_t + \beta\sigma^2_t, \end{gathered} \tag{6}$$

where D_i is a day-of-the-week dummy for day i, where we pick Saturday as our reference category; $X_{i,t}$ is a vector of control variables that include market beta, log book-to-market, log size (price), past returns, log volume, and changes in sentiment. We estimate Equation (6) for each individual player in our sample, then calculate the average of the estimated coefficients and assess their cross-sectional significance.

The results are in Table 10. We find a number of interesting patterns. In this pure time-series setting, the effect of market beta on returns is positive and significant also unconditionally. On the other hand, all the other signs are as expected. In addition, there are strong day-of-the-week effects. Returns on Monday are especially low, and equal to -17.3% relative to Saturday.[32] Tuesday and Wednesday returns are higher, with 0.6% and 0.5%, respectively, relative to Saturday. The high returns for Saturday, Tuesday, and Wednesday are explained by the fact that real-life European soccer matches typically take place on these days, which might capture an attention effect as the players in our sample are active in real life on these days.[33]

In column (1), we estimate a $GARCH(1,1)$ model. Interestingly, the returns to players in FUT display GARCH effects, or volatility clustering, similar to equity returns. Specifically, α and β are positive and highly significant, with $\beta > \alpha$, even though the absolute magnitude of both coefficients is slightly lower than their equity counterparts. In column (2), we estimate a GJR-GARCH model (Glosten et al., 1993b), i.e., an asymmetric volatility model. As in equity returns, here we also observe an asymmetry between positive and negative returns, as indicated by the significant γ coefficient. However, we find that positive returns add more to volatility than negative returns, which is the opposite of the leverage effect in equity markets. This might be caused by the absence of short selling in FUT, causing downturns to be more damped than upturns.

6. Conclusion

In this paper, we propose a novel approach to the joint-hypothesis problem. We analyze the price formation process from a large virtual market for soccer players in the online video game FIFA 19, which shows many similarities in structure and demographics to real-life financial

[32]The Monday effect is a widely documented feature in asset returns (see, e.g., French (1980) and Keim and Stambaugh (1984)).

[33]League matches take place on the weekend, and Champions League games on Tuesdays and Wednesdays.

Table 10: Return dynamics

	GARCH	GJR-GARCH
Beta(-1)	0.002***	0.002***
	(5.880)	(6.240)
log(BM(-1))	0.462***	0.435***
	(17.453)	(22.918)
log(P(-1))	-0.202***	-0.203***
	(-40.320)	(-50.473)
Ret(-1)	0.058***	0.066***
	(16.497)	(22.095)
log(Vol(-1))	-0.003***	0.026***
	(-0.611)	(7.382)
dSENT(-1)	-0.100***	-0.105***
	(-5.781)	(-7.405)
Sunday	-0.018***	-0.013***
	(-13.397)	(-13.214)
Monday	-0.193***	-0.153***
	(-48.221)	(-51.868)
Tuesday	0.042***	0.036***
	(25.403)	(29.479)
Wednesday	-0.012***	-0.005***
	(-7.653)	(-4.590)
Thursday	-0.158***	-0.130***
	(-60.867)	(-63.536)
Friday	-0.008***	-0.006***
	(-3.884)	(-4.667)
ω	0.007***	0.007***
	(34.704)	(38.998)
α	0.265***	0.459***
	(49.502)	(24.535)
γ		-0.275***
		(-11.436)
β	0.559***	0.582***
	(72.082)	(94.062)

This table reports the estimation results of Equation (6). The coefficients are first estimated for each player individually, and then aggregated across players. Statistical significance at the 10%, 5%, and 1% level is denoted as ** and ***, respectively, with cross-sectional t-statistics in parentheses.

markets. One remarkable feature of this market is that fundamentals are predetermined and publicly known, which allows us to directly assess the impact of investor behavior on price formation independently of economic fundamentals. To the best of our knowledge, this paper is the first to study the price dynamics of a video game's in-play market.

We find that a number of factors that are known to price returns, such as size and book-to-market, represent mispricing rather than risk in this context. Also, investor sentiment and attention have an independent and substantial impact on asset prices. The results suggest that prices in real-life financial markets include a substantial behavioral component, which is likely underestimated in canonical asset pricing tests. Specifically, we estimate that mispricing may constitute anywhere from one-fifth to one-half of the spreads observed in equity markets.

References

Aguiar, M., M. Bils, K. K. Charles, and E. Hurst (2021). Leisure luxuries and the labor supply of young men. *Journal of Political Economy 129*(2), 337–382.

Albuquerque, R. (2012). Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns. *The Review of Financial Studies 25*(5), 1630–1673.

Ali, M. M. (1977). Probability and utility estimates for racetrack gamblers. *Journal of Political Economy 85*(4), 803–815.

Antoniou, C., J. A. Doukas, and A. Subrahmanyam (2016). Investor Sentiment, beta, and the cost of equity capital. *Management Science 62*(2), 347–367.

Asch, P., B. G. Malkiel, and R. E. Quandt (1982). Racetrack betting and informed behavior. *Journal of Financial Economics 10*(2), 187–194.

Asch, P. and R. E. Quandt (1987). Efficiency and profitability in exotic bets. *Economica 54*(215), 289–298.

Avery, C. and J. Chevalier (1999). Identifying investor sentiment from price paths: The case of football betting. *Journal of Business 72*(4), 493–521.

Baker, M. and J. Wurgler (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance 61*(4), 1645–1680.

Baker, M. and J. Wurgler (2007). Investor sentiment in the stock market. *The Journal of Economic Perspectives 21*(2), 129–151.

Baker, M., J. Wurgler, and Y. Yuan (2012). Global, local, and contagious investor sentiment. *Journal of Financial Economics 104*(2), 272–287.

Barber, B. and T. Odean (2001). Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics 116*(1), 261–292.

Bauer, R., M. Cosemans, and P. Eichholtz (2009). Option trading and individual investor performance. *Journal of Banking & Finance 33*(4), 731–746.

Beber, A. and M. Pagano (2013). Short-selling bans around the world: Evidence from the 2007–09 crisis. *The Journal of Finance 68*(1), 343–381.

Blitz, D. C. and P. van Vliet (2007). The volatility effect. *The Journal of Portfolio Management 34*(1), 102–113.

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *The Review of Economics and Statistics 69*(3), 542–547.

Buckman, S., A. H. Shapiro, M. Sudhof, and D. Wilson (2020). News sentiment in the time of covid-19. *FRBSF Economic Letter 2020-08*, 1–5.

Carroll, C., J. Fuhrer, and D. Wilcox (1994). Does consumer sentiment forecast household spending? *American Economic Review 84*(5), 1397–1408.

Chen, J., H. Hong, and J. C. Stein (2002). Breadth of ownership and stock returns. *Journal of Financial Economics 66*(2-3), 171–205.

Da, Z., J. Engelberg, and P. Gao (2011). In search of attention. *The Journal of Finance 66*(5), 1461–1499.

Daniel, K. and S. Titman (1997). Evidence on the characteristics of cross sectional variation in stock returns. *The Journal of Finance 52*(1), 1–33.

DeVault, L., R. W. Sias, , and L. T. Starks (2019). Sentiment metrics and investor demand. *The Journal of Finance 74*(2), 985–1024.

Diether, K. B., C. Malloy, and A. Scherbina (2002). Differences of opinion and the cross section of stock returns. *The Journal of Finance 57*(5), 2113–2141.

Direr, A. (2011). Are betting markets efficient? Evidence from European football championships. *Applied Economics 45*(3), 343–356.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*(4), 987–1007.

Fama, E. F. (1965). The behavior of stock-market prices. *The Journal of Business 38*(1), 34–105.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance 25*(2), 383–417.

Fama, E. F. (1991). Efficient capital markets: II. *The Journal of Finance 46*(5), 1575–1617.

Fama, E. F. and K. R. French (1992). The cross-section of expected stock returns. *The Journal of Finance 47*(2), 427–465.

Fama, E. F. and K. R. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics 33*(1), 3–56.

Fama, E. F. and K. R. French (1995). Size and book-to-market factors in earnings and returns. *The Journal of Finance 50*(1), 131–155.

Fama, E. F. and K. R. French (2004). The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives 18*(1), 25–46.

Forrest, D. and R. Simmons (2008). Sentiment in the betting market on Spanish football. *Applied Economics 40*(1), 119–126.

Forsythe, R., T. Palfrey, and C. Plott (1982). Asset valuation in an experimental market. *Econometrica 50*(3), 537–567.

Forsythe, R., T. Palfrey, and C. Plott (1984). Futures markets and informational efficiency: A laboratory examination. *The Journal of Finance 39*(4), 955–981.

Frazzini, A. and L. H. Pedersen (2014). Betting against beta. *Journal of Financial Economics 111*(1), 1–25.

French, K. R. (1980). Stock returns and the weekend effect. *Journal of Financial Economics 8*(1), 55–69.

Friedman, D., G. Harrison, and J. Salmon (1984). The informational efficiency of experimental asset markets. *Journal of Political Economy 92*(3), 349–408.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993a). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance 48*(5), 1779–1801.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993b). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance 48*(5), 1779–1801.

Golec, J. and M. Tamarkin (1991). The degree of inefficiency in the football betting market. *Journal of Financial Economics 30*(2), 311–323.

Grinblatt, M., M. Keloharju, and J. Linnainmaa (2011). IQ and stock market participation. *The Journal of Finance 66*(6), 2121–2164.

Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance 48*(1), 65–91.

Jones, C. M. and O. A. Lamont (2002). Short-sale constraints and stock returns. *Journal of Financial Economics 66*(2-3), 207–239.

Keim, D. B. and R. F. Stambaugh (1984). A further investigation of the weekend effect in stock returns. *The Journal of Finance 39*(3), 819–835.

Kuypers, T. (2000). Information efficiency: An empirical study of a fixed odds betting market. *Applied Economics 32*(11), 1353–1363.

Levitt, S. (2004). Why are gambling markets organized so differently from financial markets? *Economic Journal 114*(495), 223–246.

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

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

Sauer, R. D. (1998). The economics of wagering markets. *Journal of Economic Literature 36*(4), 2021–2064.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Smith, V. L., G. L. Suchanek, and A. W. Williams (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica 56*(5), 1119–1151.

Snyder, W. W. (1978). Horse racing: Testing the efficient markets model. *The Journal of Finance 33*(4), 1109–1118.

Stambaugh, R. F., J. Yu, and Y. Yuan (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics 104*(2), 288 – 302.

Yu, J. and Y. Yuan (2011). Investor sentiment and the mean-variance relation. *Journal of Financial Economics 100*(2), 367–381.

Ziemba, W. T. and D. B. Hausch (1987). *Dr. Z's beat the racetrack.* William Morrow, New York.

Appendix A

The first two moments of the distribution of final wealth are:

$$E(\tilde{w}_{1j}) = (1 + \bar{r}_H \alpha_{Hj} + \bar{r}_L \alpha_{Lj})\, w_{0j}, \tag{A.1}$$

$$var(\tilde{w}_{1j}) = \left(\sigma_H^2 \alpha_{Hj}^2 + \sigma_L^2 \alpha_{Lj}^2\right) w_{0j}^2, \tag{A.2}$$

where α_{ij} is the ratio between demand for asset i, a_{ij}, and current wealth w_{0j}. Note that the covariance between the two assets is zero (i.e., $\sigma_{HL} = 0$). The two first-order conditions yield:

$$\bar{r}_H = \gamma_j \sigma_H^2 \alpha_{Hj} w_{0j}, \tag{A.3}$$

$$\bar{r}_L = \gamma_j \sigma_L^2 \alpha_{Lj} w_{0j}. \tag{A.4}$$

Solving out, the optimal investments are:

$$\alpha_{Hj} w_{0j} = \frac{\bar{r}_H}{\gamma_j \sigma_H^2} \equiv a^*_{Hj}, \tag{A.5}$$

$$\alpha_{Lj} w_{0j} = \frac{\bar{r}_L}{\gamma_j \sigma_L^2} \equiv a^*_{Lj}, \tag{A.6}$$

which implies the same portfolio formation $(\frac{\alpha^*_{Hj}}{\alpha^*_{Lj}})$ for all investors.

To derive equilibrium returns, take the first-order condition and sum up across all investors:

$$\sum_{j=1}^{M} \tau_j \bar{r}_H = \sigma_H^2 \sum_{j=1}^{M} a^*_{Hj}, \tag{A.7}$$

$$\sum_{j=1}^{M} \tau_j \bar{r}_L = \sigma_L^2 \sum_{j=1}^{M} a^*_{Lj}, \tag{A.8}$$

where $\tau_j \equiv \frac{1}{\gamma_j}$ represents investor j's risk tolerance.

We can divide by all investors (M), and apply the equilibrium conditions $a^*_{Hj} \equiv \bar{a}_H$, $a^*_{Lj} \equiv \bar{a}_L$:

$$\tau_M \bar{r}_H = \sigma_H^2 \bar{a}_H, \tag{A.9}$$

$$\tau_M \bar{r}_L = \sigma_L^2 \bar{a}_L, \tag{A.10}$$

where τ_M is the average level of risk tolerance. Rearranging, we obtain:

$$\bar{r}^*_H = \frac{\bar{a}_H}{\tau_M} \sigma_H^2, \tag{A.11}$$

$$\bar{r}^*_L = \frac{\bar{a}_L}{\tau_M} \sigma_L^2. \tag{A.12}$$

Now, define market returns as:

$$\tilde{r}_M = \frac{\bar{a}_H}{\bar{a}_M}\tilde{r}_H + \frac{\bar{a}_L}{\bar{a}_M}\tilde{r}_L, \tag{A.13}$$

where $\bar{a}_M \equiv \bar{a}_H + \bar{a}_L$. Note that the covariance between the returns on asset i and market returns is:

$$cov(\tilde{r}_H, \tilde{r}_M) \equiv \sigma_{HM} = \frac{\bar{a}_H}{\bar{a}_M}\sigma_H^2, \tag{A.14}$$

$$cov(\tilde{r}_L, \tilde{r}_M) \equiv \sigma_{LM} = \frac{\bar{a}_L}{\bar{a}_M}\sigma_L^2, \tag{A.15}$$

Then we can rewrite returns as:

$$\bar{r}_H^* = \frac{\bar{a}_M}{\tau_M}\sigma_{HM}, \tag{A.16}$$

$$\bar{r}_L^* = \frac{\bar{a}_M}{\tau_M}\sigma_{LM}. \tag{A.17}$$

But then the relation must hold also for the market portfolio:

$$\bar{r}_M^* = \frac{\bar{a}_M}{\tau_M}\sigma_M^2, \tag{A.18}$$

which implies:

$$\frac{\bar{a}_M}{\tau_M} = \frac{\bar{r}_M^*}{\sigma_M^2}. \tag{A.19}$$

Using this result, we can finally rewrite returns as:

$$\bar{r}_H^* = \sigma_{HM}\frac{\bar{r}_M^*}{\sigma_M^2} \equiv \beta_H \bar{r}_M^*, \tag{A.20}$$

$$\bar{r}_L^* = \sigma_{LM}\frac{\bar{r}_M^*}{\sigma_M^2} \equiv \beta_L \bar{r}_M^*, \tag{A.21}$$

which represents the security market line. Note that betas are just a function of volatility:

$$\beta_H = \frac{\frac{\bar{a}_H}{\bar{a}_M}\sigma_H^2}{\left(\frac{\bar{a}_H}{\bar{a}_M}\right)^2 \sigma_H^2 + \left(\frac{\bar{a}_L}{\bar{a}_M}\right)^2 \sigma_L^2}, \tag{A.22}$$

$$\beta_L = \frac{\frac{\bar{a}_L}{\bar{a}_M}\sigma_L^2}{\left(\frac{\bar{a}_H}{\bar{a}_M}\right)^2 \sigma_H^2 + \left(\frac{\bar{a}_L}{\bar{a}_M}\right)^2 \sigma_L^2}. \tag{A.23}$$