

NBER WORKING PAPER SERIES

RATIONALIZING TRADING FREQUENCY AND RETURNS:
MAYBE TRADING IS GOOD FOR YOU

Yosef Bonaparte
Russell Cooper
Mengli Sha

Working Paper 25838
<http://www.nber.org/papers/w25838>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2019

We are grateful to the National Science Foundation for supporting this research. This paper is a major revision of one circulated earlier as NBER Working Paper #w16022 entitled “Rationalizing Trading Frequency and Returns: Trading is Good for your Wealth” by Yosef Bonaparte and Russell Cooper. We thank seminar participants at the Federal Reserve Banks of Dallas and Kansas City, Cambridge University, the CCER at Peking University, the EIEF in Rome, Melbourne University, Oxford University, the Pennsylvania State University and the Toulouse School of Economics for comments and suggestions on earlier work that underlies this paper. Guozhong Zhu provided outstanding research assistance on the earlier paper. Comments and suggestions from seminar participants at the European University Institute and Universitat Bonn are greatly appreciated. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Yosef Bonaparte, Russell Cooper, and Mengli Sha. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Rationalizing Trading Frequency and Returns: Maybe Trading is Good for You
Yosef Bonaparte, Russell Cooper, and Mengli Sha
NBER Working Paper No. 25838
May 2019
JEL No. E03,E21,G11

ABSTRACT

Barber and Odean study the relationship between trading activity and returns. They find that households who trade more have a lower net return than other households. They argue that these results cannot emerge from a model with rational traders and instead attribute these findings to overconfidence. In contrast, we find that household financial choices generated from a dynamic optimization problem with rational agents and portfolio adjustment costs can produce trading and return patterns that closely mimic these facts. Adding various forms of irrationality, modelled as beliefs about income and return processes that are not data based, do not improve the ability of the model to explain the patterns of turnover and net returns. Irrationality can improve the ability of the model to match a larger set of moments, including these turnover and net return moments coupled with those that capture the wealth to income ratio and portfolio composition.

Yosef Bonaparte
University of Colorado Denver
1475 Lawrence St.
Denver, CO 80202
bonaparte20@gmail.com

Mengli Sha
Department of Economics
The Pennsylvania State University
mengli.sha0124@gmail.com

Russell Cooper
Department of Economics
European University Institute
Villa La Fonte
Via delle Fontanelle 18
I-50014 San Domenico di Fiesole (FI)
ITALY
and NBER
russellcoop@gmail.com

1 Motivation

Barber and Odean (2000) find that households with higher stock turnover have a lower net return. They interpret this as evidence households are overconfident and thus not rational. According to Barber and Odean (2000):

Our most dramatic empirical evidence supports the view that overconfidence leads to excessive trading ... On one hand, there is very little difference in the gross performance of households that trade frequently with monthly turnover in excess of 8.8 percent and those that trade infrequently. In contrast, households that trade frequently earn a net annualized geometric mean return of 11.4 percent, and those that trade infrequently earn 18.5 percent. These results are consistent with models where trading emanates from investor overconfidence, but are inconsistent with models where trading results from rational expectations.

This paper studies the implications of an optimizing model with costly portfolio adjustment for the relationship between frequency of trade and asset returns. We investigate two explanations for the findings of Barber and Odean (2000).¹ The first looks at the choice of rational agents faced with costs of portfolio adjustment. The second allows for different forms of irrationality, including overconfidence.

For rational, optimizing households, it seems natural to consider the differences in net returns as reflecting two forces: trading costs and a selection effect through household choice of whether to adjust their portfolio. Trading costs drive a wedge between gross and net returns. Household choice, both on the extensive (to adjust or not) and intensive (turnover conditional on adjustment) margins, creates an endogenous relationship between asset returns and portfolio adjustment.

We ask whether the presence of fixed and variable portfolio adjustment costs can generate the observed differences in returns based upon the frequency of trade. Our approach is to specify a dynamic optimization problem of a household and estimate its parameters. The uncertainty in the model comes from income shocks, which are partly household specific, as well as a stochastic return on the household stock portfolio. We generate simulated data from the estimated model to study the relationship between portfolio adjustment and returns.

Following the suggestion of Barber and Odean (2000), we also study a series of models which relax the assumption of perfect rationality to model overconfidence. This is in line with the literature that generally models overconfidence as mis-calibration, i.e., overestimating the

¹See Barber and Odean (2001) for additional evidence, including a gender breakdown, and discussion of overconfidence.

precision of information about the price of a financial security (e.g., Kyle and Wang (1997), Odean (1998) and Gervais and Odean (2001)). We specifically consider models in which traders over-estimate the volatility of labor income, over-estimate the serial correlation in returns or respond to signals that are not informative about returns.

To be clear, by “irrational” we mean to capture the behavior of utility maximizing households who hold beliefs about exogenous processes, such as returns, which are not consistent with data. Our evaluation of these alternatives is based upon the optimizing behavior built upon erroneous beliefs. The implied decision rules are taken to the same data as those based upon rational choice.

Barber and Odean (2000) conclude with a powerful statement:

Our central message is that trading is hazardous to your wealth.

This conclusion reflects their finding that net returns are lower for agents who trade more actively without earning higher gross returns. This trading behavior is subsequently viewed as irrational.

We do not concur. In contrast to these claims, we find that a model with rational households **is** capable of matching the observed differences between gross and net returns as a function of trading frequency. The introduction of irrational traders does **not** lead to any improvement in model fit. When the turnover moments are supplemented with those that capture the wealth income ratio and portfolio composition of investors, then the model fit is improved by the irrational households. Interestingly, those gains arise mainly from beliefs that assign more variability to income than estimated in the data rather than any irrational views about the determinants of stock prices.

To be clear, our model and estimation does not capture all of the elements associated with to the households in the Barber and Odean (2000). First, we do not have complete information on these households, such as their income process and total financial holdings. Second, for computation reasons, our model does not contain the full range of assets available to those traders.

Nonetheless our ability to match the turnover and net return moments and in the various robustness exercises makes clear that the patterns detected by Barber and Odean (2000) for their special sample does not necessarily require irrational behavior. Nor, from our analysis, is it apparent that irrationality improves the model’s ability to fit these turnover and net return moments.

2 Household Optimization

Here we briefly review the model, drawing upon Bonaparte, Cooper, and Zhu (2012) and Cooper and Zhu (2015), that is the basis of the household optimization problem.² The household is infinitely lived and has two assets: bonds and stocks. Bonds are costless to adjust, and have a certain return. Stocks yield a random return, higher on average than bonds. Stock holdings, by assumption, are costly to adjust.

The key to the model is the household choice of whether to adjust its portfolio or not. Adjustment is costly due to the presence of fixed and variable trading costs. The household may choose not to incur these costs, in which case consumption smoothing is achieved solely through adjustment in bond holdings. If the household adjusts, then it incurs a cost of portfolio adjustment. In this way, the model generates two types of turnover: the discrete choice of whether to adjust and the continuous choice of how much to adjust conditional on having incurred fixed adjustment costs.

To be clear, portfolio turnover refers to trade between stocks and bonds.³ The model does not include multiple stocks and thus rebalancing of the components of a stock portfolio is excluded. This implies that the model understates actual turnover and the associated costs of rebalancing. Missing these trades makes it more difficult to match the high turnover, low net return moments highlighted by Barber and Odean (2000).

Let $\Omega = (y, S, R^s)$ represent the state of the household where y is current labor income, $S = (b, s)$ is the current value of the holdings of bonds and stocks respectively and R^s is the stochastic return on stocks. The return on bonds, R^b is deterministic. A household chooses between (i) portfolio adjustment and (ii) no portfolio adjustment. This choice is given:

$$v(\Omega) = \max\{v^a(\Omega), v^n(\Omega)\} \quad (1)$$

for all Ω .

A household choosing to **adjust** selects the amount of stocks and bonds to solve:

$$\begin{aligned} v^a(\Omega) &= \max_{b' \geq 0, s' \geq 0} u(c) + \beta E_{\Omega'|\Omega} v(\Omega') \\ \text{s.t.} \\ c &= \psi y + R^b b + R^s s - b' - s' - C(s, s') - F. \end{aligned} \quad (2)$$

²An earlier version of the paper had only a single asset, as in Bonaparte and Cooper (2009). With that specification, portfolio adjustment and adjusting the margin between consumption and savings was not distinct.

³In section 4.4, we study an alternative measure of turnover.

In this problem, there is no borrowing and short sales of stocks are not allowed.

There are three forms of adjustment costs in the model. There is a time cost of stock adjustment represented by $\psi \leq 1$ in (2).⁴ Some of the estimation allows $\psi > 1$ so that agents generate a “utility gain” from trading. Second, the model allows a fixed cost of adjustment, F . This is distinct from the time cost of ψ as there is no interaction with income.⁵ Finally, the model includes direct trading costs, explained further below, captured by $C(s, s')$. In addition to the frequency of adjustment these costs also generate a demand for bonds and thus impact the stock share.

If the household chooses **not to adjust** its portfolio, then the trading and opportunity costs are avoided. There is re-optimization over bond holdings alone. The household chooses bonds to solve:

$$v^n(\Omega) = \max_{b' \geq 0} u(c) + \beta E_{\Omega'|\Omega} v(\Omega')$$

s.t.

$$c = y + R^b b - b' \tag{3}$$

$$s' = R^s s. \tag{4}$$

Here we assume that if there is no portfolio rebalancing, any return on stocks is automatically put into the stock account, i.e. $s' = R^s s$. In the robustness analysis, we relax this assumption so that stock returns are instead deposited into the bond account.

The policy functions generated by household optimization include an extensive margin (adjust, no adjust) and an intensive margin indicating the magnitude of the adjustment. Due to the adjustment costs, the model can produce both inaction in portfolio adjustment as well as large turnover rates. The incentive for portfolio adjustment comes from large shocks to income and returns. A large positive return shock may create a large enough wealth gain that households choose to rebalance their portfolios. This would generate a positive correlation between gross returns and trades. Likewise, a large adverse return shock might also cause financial wealth to fall so that rebalancing is worthwhile. In this case, a negative correlation between trading and gross return is created. Further, a large negative shock to income will also create an incentive to rebalance, independent of current returns. Added to this is the differential between measured gross and net returns created by the presence of trading costs, $C(\cdot)$. These types of responses to shocks form the link from the policy functions generated by the household optimization problem to the evidence of Barber and Odean (2000).

⁴Bonaparte, Cooper, and Zhu (2012) discusses other specifications of adjustment costs.

⁵We will study the two cases of $\psi > 0$ and $F > 0$ separately so that identification of the adjustment cost is not an issue.

3 Trading Costs and Exogenous Processes

The goal of the analysis is to study the relationship between portfolio turnover and return. To do so, we must estimate the parameters of the household choice problem. The estimation uses a simulated method of moments approach based upon this model of dynamic household choice. Solving the household optimization problem requires the specification of trading costs and exogenous processes.

3.1 Trading Costs

Bonaparte and Cooper (2009) estimate trading costs, $C(\cdot)$, directly from the data set used by Barber and Odean (2000). The data set provides information on common stock trades of about 78,000 households through a discount brokerage firm from January 1991 to December 1996.

Assume:

$$C^b(s_{-1}, s) = \nu_0^b + \nu_1^b(s - s_{-1}) + \nu_2^b(s - s_{-1})^2 \quad (5)$$

if the household buys an asset, $s > s_{-1}$. If instead the household sells, $s < s_{-1}$, then

$$C^s(s_{-1}, s) = \nu_0^s + \nu_1^s(s_{-1} - s) + \nu_2^s(s - s_{-1})^2. \quad (6)$$

Bonaparte and Cooper (2009) use the monthly household account data from Barber and Odean (2000) to estimate these parameters.⁶ The trading costs, measured in dollars, are estimated in a regression where the dependent variable is the commission and the independent variables are trade value (the price of the share times the quantity of share) and trade value squared per stock. Bonaparte and Cooper (2009) report the estimates in Table 1.

Though the linear and quadratic terms are statistically significant, the main cost of adjustment is the fixed cost per trade. While this cost may seem high relative to currently advertised fixed trading costs, it is still small compared to the average trade of a household in the data set of about \$12,500.

These estimates of trading costs do not include the bid-ask spread which, according to Barber and Odean (2000) are about 0.31% for purchases and 0.69% for sales. These additional costs are added to the linear terms reported in Table 1 when the trading costs are integrated into the household optimization problem.

⁶Details on the estimation can be found in Bonaparte and Cooper (2009). Through this procedure, we are able to decompose the commission costs reported in Table 1 of Barber and Odean (2000) into fixed and variable components.

Parameter	Buying	Selling
Constant ν_0^i	56.10 (0.05)	61.44 (0.061)
Linear ν_1^i	0.0012 (1.63e-06)	0.0014 (1.93e-06)
Quadratic ν_2^i	$-1.01e^{-10}$ (2.88e-13)	$-1.28e^{-10}$ (9.26e-13)
Adj. R^2	0.251	0.359
Number of Observations	1,746,403	1,329,394

Table 1: Estimated Trading Costs

3.2 Income and Returns

The income process for stockholders is annual. It is estimated from the Panel Study of Income Dynamics (PSID). The serial correlation is 0.842 and the standard deviation of the innovation is 0.29.⁷ Importantly, there is not sufficient information in the data set used by Barber and Odean (2000) to estimate the income process for individuals in that sample. Nor is it possible to extract a sample of households who directly own stock from the PSID to perhaps mimic those in the Barber and Odean (2000) sample.

As the frequency of the household choice problem is monthly, it is necessary to convert the annual income process to this higher frequency. This is done in two ways, distinguished by the presence of unemployment risk.

The first simply converts the annual process into a monthly one without adding any higher frequency unemployment risk. In this case, the monthly serial correlation is 0.9858 and the standard deviation of the innovation to income is 0.0904.

The second adds unemployment risk to the income process. As emphasized in Carroll (1992), it is important to recognize that, particularly, at the monthly frequency, households face significant risk of job loss. Thus, this second specification adds a zero labor income state to the process. These flows between employment and unemployment are taken from the Bureau of Labor Statistics.⁸

Specifically, each month an employed agent becomes unemployed with probability $\delta = 0.014$. Each month an unemployed agent finds a job with probability 0.27. The replacement rate for an unemployed agent is set at 40% of average income. If an unemployed agent finds a job, the wage is assumed to be the mean of the income process. Given these flows and the

⁷This is the same process as used in Bonaparte, Cooper, and Zhu (2012) and is described in the Appendix of that paper.

⁸Specifically, these probabilities characterizing this additional state are calculated from the seasonally adjusted flows taken from <https://www.bls.gov/webapps/legacy/cpsflowstab.htm> for 1990 to 2009.

estimated annual serial correlation and standard deviation of the innovation at the annual level, the monthly income process is estimated through a simulated method of moments approach, discussed in the Appendix. From this analysis, the monthly serial correlation is estimated at 0.9959 and the standard deviation of the income innovation is estimated at 0.0839, conditional on employment.

The real stock return, which includes capital gains and dividends, is measured as the S&P index monthly return from 1967-94. It is taken from CRSP (<http://wrds-web.wharton.upenn.edu/wrds/index.cfm>) The average monthly return is set at 1.0061 with a standard deviation of 0.0439. The estimated serial correlation of annual returns is not significantly different from zero.

4 Estimation Approach and Results

The estimation approach, given this choice of moments, minimizes the distance between actual and simulated moments:

$$J = \min_{(\Theta)} \left(\frac{M^s(\Theta) - M^d}{M^d} \right)' W \left(\frac{M^s(\Theta) - M^d}{M^d} \right). \quad (7)$$

Here W is the identity matrix as the moments come from a variety of data sets so that computing a variance-covariance matrix is not feasible.

Given the parameters, a simulated panel data set with 4000 households and 4000 time periods is created from the solution of the household's dynamic optimization problem. The simulated moments are calculated from this panel, just as in the actual data.⁹ Households differ because of idiosyncratic income shocks which generates differences in trading patterns and returns.

4.1 Moments

The analysis assumes the household choices are made on a monthly basis. There are a few advantages from studying choices at such a high frequency. First, by comparing to results from a related annual model, we examine the effects of time aggregation. Second, the monthly model allows a direct link to the high frequency data from Barber and Odean (2000). Third, the higher frequency choice model allows us to follow the estimation of the

⁹Our results do not change if the size of the simulated panel is increased. For these results, the (coarse) fine state space had $(20 \times 25 \times 10 \times 3) 100 \times 200 \times 10 \times 3$ elements. The solution entailed piecewise cubic hermite interpolation, with convergence for the value function defined on the fine grid.

elasticity of intertemporal substitution for stock market participants from Vissing-Jorgensen (2002).

For this analysis, there are two types of moments. The first set captures the relationship between turnover and net stock return, as in Barber and Odean (2000). The second set are frequently studied to summarize household financial choices. Together, the moments are chosen to be informative about the parameters to be estimated, $\Theta \equiv (\beta, \gamma, \psi)$.

Turnover and Return Moments The first set of moments come from the Barber and Odean (2000) study of household trading activity and return. These are monthly moments and include the portfolio turnover rate as well as the net stock return. From their Table V, Barber and Odean (2000) calculate monthly turnover as the average of sales and purchases. Though our model has only a single stock, it is able to match this average. The inclusion of the lowest turnover rate captures inaction in portfolio adjustment. Here portfolio turnover is defined as the absolute difference between end of period and beginning of period stock wealth, divided by the beginning of period stock wealth: specifically, turnover for household i is

$$T_i \equiv \left| \left(\frac{s'_i - s_i R^s}{s_i R^s} \right) \right|. \quad (8)$$

Note that the turnover rate depends on the households stock wealth at the start of the period, $s_i R^s$, inclusive of current stock returns. Later we discuss a portfolio based measure of turnover.

We calculate the net return for household i , denoted R_i^n , **on the stock portfolio** as

$$R_i^n = \frac{s_i R^s - C(s'_i - s_i R^s)}{s_i}. \quad (9)$$

This is not the net return on an individual trade but rather the net return on the entire stock portfolio. In our setting, it is impossible to compute the return on a particular trade without imposing some arbitrary accounting rule to assign trading costs to net returns of a particular purchase or sale. We return below to alternative measures of this return.

Clearly, there is a mechanical relationship between high turnover and low net return. If the household, for example, buys stocks, then using the specification of the trading cost function, $C(s'_i - R^s s_i)$, the net return becomes:

$$R_i^n = R^s - \frac{\nu_0^b}{s_i} - \frac{\nu_1^b(s'_i - R^s s_i)}{s_i} - \frac{\nu_2^b(s'_i - R^s s_i)^2}{s_i}. \quad (10)$$

From this calculation, the linear part of the cost function implies that the net return is lower when trades are large and when the initial stock holdings is relatively low. But the fixed

cost and the quadratic cost impact this relationship. In particular, low initial stock holdings imply a large differential between gross and net returns. Thus it is not solely turnover that determines this differential, but the state of the household undertaking the trade as well. Below we discuss how much of the low net returns is associated with high turnover through the linear part of the cost function.

For our analysis, we study the quintiles of the turnover rate as shown in Table 3. As in Barber and Odean (2000), we compute the time series average of the (cross sectional) average monthly net (of direct trading costs) return differential on the portfolio of the lowest and highest turnover quintile of -0.0046 . This is a monthly differential in return. If, for example, one portfolio earns 1.0146 per month and another earns 1.0046, then the difference in returns is 13.34% over a year.

To be precise, these are not exactly the measures of net return used by Barber and Odean (2000). Our return measure is based upon a portfolio not a single trade. The differences and consequences are explored below.

Financial Choice Moments The second set of moments, taken directly from Bonaparte, Cooper, and Zhu (2012), capture the financial choices of households in terms of the response of consumption to interest rate variations, portfolio composition, and savings.¹⁰ These moments are not taken from the Barber and Odean (2000) study but are from other sources. They are included to impose additional discipline on the parameter estimation and to provide insights into the representativeness of that sample.

The first moment is the elasticity of intertemporal substitution (EIS). The moment comes from Vissing-Jorgensen (2002) and is obtained from a regression of average consumption growth of stock market participants on the interest rate. The point estimate is 0.299. This moment is particularly informative about γ . Note though that in our model with infrequent adjustment, the inverse of the EIS is not necessarily equal to γ .

The second moment is the mean financial wealth to income ratio, which is 2.43 in the data. As discussed in Bonaparte, Cooper, and Zhu (2012), this is the sample average of the mean financial wealth to income from the Survey of Consumer Finance for the 1989-2007 period. This moment is quite informative about β .

The third moment is the stock share in financial wealth. This share reflects, in part, the gains to liquidity. Thus both the cost of stock adjustment, ψ , and the variability of income are important for matching this moment. It is also taken from the Survey of Consumer Finance for the same period.

¹⁰See Bonaparte, Cooper, and Zhu (2012) for further details on the data moments. The appendix provides additional details on the calculation of all moments.

The moments differ in terms of frequency. The first moment, calculated by Vissing-Jorgensen (2002) uses, monthly consumption data to estimate the EIS over 6 month intervals. This procedure is replicated in our calculation of the simulated moment. The second and third moments are calculated using annual data. Thus the monthly simulated data is time aggregated to produce an annual income measure by household. The stock holdings and financial wealth are computed on a 12 month basis to compute the second and third moments.

4.2 Results

The parameter estimates and moments are presented for a number of cases.¹¹ There are two monthly income processes, one without unemployment risk and a second including that risk. There are three sets of moments matched: (i) the financial moments, (ii) the turnover moments and (iii) all of the moments. Finally, the baseline estimation restricts $\psi \leq 1$ so that trading is costly. We also allow ψ to be unrestricted.

Table 2: Parameter Estimates

case	β	γ	ψ	J
	$\psi \leq 1$			
<i>No Unemp. Risk</i>				
turnover	0.9817	1.4070	1.0000	0.1594
fin. choice	0.9931	1.1100	0.9773	0.1249
all	0.9907	0.8838	0.9999	1.1123
<i>Unemp. Risk</i>				
turnover	0.9895	1.1072	0.9982	0.4413
fin. choice	0.9928	1.0429	0.9953	0.1086
all	0.9917	1.0873	0.9947	0.8569
	ψ unrestricted			
<i>No Unemp. Risk</i>				
turnover	0.9814	1.4063	1.0006	0.1340
all	0.9923	0.8911	1.3233	0.6849
<i>Unemp. Risk</i>				
turnover	0.9864	0.9517	1.2497	0.4324
all	0.9919	1.0908	1.3380	0.4952

This table reports estimated parameters for the various cases: β is the discount factor, γ is relative risk aversion and ψ is the fraction of income remaining after portfolio adjustment. J is the difference between model moments and data moments as described in equation (7).

¹¹Households with less than \$60 in stocks were excluded from the simulated data for the calculation of the financial moments as these households would not appear in the Barber and Odean (2000) data set.

Table 3: Data and Model Moments

case	EIS	WI	Sh	t1	t2	t3	t4	t5	DR
Data	0.2990	2.4300	0.6840	0.0019	0.0124	0.0289	0.0598	0.2149	-0.0046
$\psi \leq 1$									
<i>No Unemp. Risk</i>									
turnover	na	na	na	0.0016	0.0115	0.0231	0.0456	0.1842	-0.0040
fin. choice	0.2857	2.3864	0.9238	na	na	na	na	na	na
all	0.2392	0.9883	0.9332	0.0018	0.0101	0.0206	0.0341	0.2063	-0.0021
<i>Unemp. Risk</i>									
turnover	na	na	na	0.0024	0.0099	0.0188	0.0434	0.2042	-0.0029
fin. choice	0.3042	2.3628	0.9082	na	na	na	na	na	na
all	0.2446	1.7563	0.9182	0.0021	0.0084	0.0174	0.0357	0.2283	-0.0026
ψ unrestricted									
<i>No Unemp. Risk</i>									
turnover	na	na	na	0.0016	0.0115	0.0224	0.0442	0.1919	-0.0040
all	0.2599	1.8789	0.9427	0.0017	0.0090	0.0190	0.03297	0.2308	-0.0034
<i>Unemp. Risk</i>									
turnover	na	na	na	0.0020	0.0098	0.0215	0.0376	0.1244	-0.0044
all	0.2356	2.2249	0.9070	0.0019	0.0091	0.0195	0.0397	0.2405	-0.0037

This table reports data and simulated moments. “EIS” is the elasticity of intertemporal substitution, “WI” is the wealth to income ratio and “Sh” is the stock share, “ti” is the turnover rate for quintile “i” and “DR” is the difference in the net return between the highest and lowest turnover rate quintiles.

For this case of rational households, i.e. households whose beliefs are consistent with the exogenous estimated income and return processes, the parameter estimates are given in Table 2 and the moments in Table 3. The row labeled “turnover” focuses on the turnover and returns that motivate this study. The row “fin. choice” matches the moments summarizing household financial choices alone, leaving aside turnover and returns. The “all” row matches both the financial choice and turnover moments. These last two cases are motivated below.

These cases are further distinguished by whether the restriction $\psi \leq 1$ is imposed. The panel “ $\psi \leq 1$ ” imposes this constraint and the panels indicated by “ ψ unrestricted” relaxes it. The magnitude of ψ is of interest to determine whether the low return associated with high turnover is explained by a “utility gain” from trading.

Baseline The baseline estimates, with no unemployment risk and $\psi \leq 1$, are reported in the first block of the two tables. Focusing on the model’s ability to match the turnover and return moments, there are a couple of key points regarding parameter estimates from Table 2.¹² First, the estimated monthly discount factor is 0.9817, reflecting the monthly frequency

¹²Section 4.3 provides a detailed discussion of the mapping from parameters to moments.

of the optimization problem. This translates into an annual discount factor of 0.8012.¹³ For our estimation, a higher annual discount factor, say, of 0.96 would significantly reduce turnover as well as the differential in return between low and high turnover households.¹⁴

The estimated degree of risk aversion is 1.407. The estimate of γ is far from the inverse of the responsiveness of consumption growth to the interest rate, estimated at 0.299. To emphasize an important point, in this environment of costly adjustment, the estimated EIS bears no direct relationship to the degree of relative risk aversion.

The estimate of ψ is essentially one. That is, the constraint $\psi \leq 1$ is binding and the only costs of trading are those associated with the direct trading costs outlined above.

Table 4: Simulated Moments

case	EIS	WI	Sh	t1	t2	t3	t4	t5	DR
Data	0.2990	2.4300	0.6840	0.0019	0.0124	0.0289	0.0598	0.2149	-0.0046
Est. Financial	na	na	na	0.0006	0.0042	0.0102	0.0249	0.6830	-0.0062
Est. Turnover	0.0418	0.2680	0.9525	na	na	na	na	na	na

This table reports simulation results. “EIS” is the elasticity of intertemporal substitution, “WI” is the wealth to income ratio and “Sh” is the stock share, “ti” is the turnover rate for quintile “i” and “DR” is the difference in the net return between the highest and lowest turnover rate quintiles.

The moments in Table 3 reveal our first main result: **the estimated model nearly replicates the turnover and net return moments.** This includes, *DR*, the differential in monthly returns between the high and low turnover traders which is -0.0040 in the estimated model. Further, the turnover rate for the highest group is nearly the 21.49% from the data.

It is important to understand the mechanics operating here given the intuition of Barber and Odean (2000) that these patterns are inconsistent with the choices of rational agents. There are two key features of the model that generate these patterns: adjustment costs and idiosyncratic shocks. The adjustment costs matter directly as the trading cost generates a wedge between gross and net returns. Further, these adjustment costs create an incentive for both inaction in portfolio adjustment as well as large turnover once trade occurs. The idiosyncratic shocks generate dispersion in both stock holdings and turnover and thus create variation in turnover rates across households.

Looking at the highest turnover quintile, a couple of features are apparent. First the stock holdings of the high turnover group are considerably smaller than the low turnover

¹³This relatively low discount rate appears in other studies, such as Bonaparte, Cooper, and Zhu (2012), in order to match the average wealth income ratio across households.

¹⁴Intuitively, for the lower discount factor there is more inaction in adjustment and thus conditional on change there is more turnover.

group. This is consistent with Barber and Odean (2000) but the difference in holdings in the simulated data is larger. From their Table 5, the highest turnover group has holdings of about 2/3 that of the low turnover group. In our simulation, this ratio is 0.59. In fact, many of the high turnover group have stock holdings of only between 0 and 10 times monthly income, considerably lower than the mean financial wealth to income ratio.

Second, though not as dramatic, the high turnover group has slightly lower average income. In fact, many of the low turnover households have high income and high stock holdings. Finally, the trades of the high turnover group are extremely dispersed in response to idiosyncratic shocks and financial wealth differences. Some households are buying while others are selling. As these are high turnover households, these trades are substantial relative to their stock holdings.

Note that in this case the only adjustment costs are those estimated in Table 1. At these baseline parameters, the key is the fixed cost. Eliminating it reduces the return differential to -0.0009 . From (10), with only a fixed cost, the return differential remains large as long as the level of financial wealth is not excessive.

Matching Financial Choice Moments The row labeled “fin. choice” studies how well the model can fit the more traditional financial moments discussed above. The study of these moments is motivated by two questions. First, how representative are the households who match the turnover moments? Second, how well can the estimated model match financial moments as well as the turnover moments?

An initial perspective on this is to estimate the parameters to match the financial moments alone, denoted “fin. choice”. With these estimated parameters we can inspect the implications for the turnover moments. From Table 2, the estimated discount factor is much higher (the annual rate is 0.92) in order to match the financial moments while the risk aversion parameter is about the same. From Table 3, the estimated model matches the mean wealth to income ratio but the model has too high of a stock share. In this case, the estimated cost of adjustment, $\psi = 0.9773$ does not generate enough of a demand for liquidity to match the stock share.¹⁵

From Table 4, we can see the implications of this model, estimated to match the financial moments, on the turnover moments. The row labeled “Est. Financial”, displays these simulated turnover moments. Clearly at these parameters, the turnover moments are not well matched. In particular, the turnover rates for the first 4 quintiles are much lower than either the baseline or the data. And the turnover in the last quintile is extremely high, as is

¹⁵As in Bonaparte, Cooper, and Zhu (2012), this moment can be matched quite well using an annual model. Time aggregation seems to be the reason for not matching the moment at this higher frequency.

the reduction in the net return for this group. These moments are driven by the relatively low value of ψ which produces more inaction in portfolio adjustment punctuated by infrequent bursts of large turnover. These patterns are more extreme than those found in the Barber and Odean (2000) data.

We next study the issue of representativeness by a simulation in which the financial choice moments are generated by the baseline parameter estimates, i.e. those selected to match the turnover moments alone. Looking at the row of Table 4, labeled “Est. Turnover”, the baseline model is not able to match the financial choice moments. The EIS is very low, the mean wealth to income ratio is only 10% of the data moment and the stock share is much higher than the data. An interpretation of this is that the set of households in the Barber and Odean (2000) study are not representative of overall population.

The last exercise is to estimate parameters to match both the financial and turnover moments, labeled “all” in Tables 2 and 3. Clearly, the model is not capable of fitting the turnover and financial choice moments together. At the estimated parameters, the wealth income ratio is substantially below the data moment and the stock share is too high. If the only moments matched are the “financial choice” moments, the stock share is still too high.¹⁶ Further, this specification does not match the turnover moments as well as the baseline, particularly the differential in return.

The tension in matching all of these moments again points to the special character, in terms of income processes and wealth, of the Barber and Odean (2000) sample.¹⁷ We return to matching these additional moments in our discussion of irrational agents.

Unemployment Risk The block “With Unemp. Risk” introduces unemployment risk into the monthly income process. As discussed above, each household faces unemployment risk as well as the possibility of reentry into the labor force. This additional income risk should impact the precautionary saving of the household as well as creating a demand for liquidity. The model matching the turnover moments as well as “all” moments was re-estimated to include this additional source of uncertainty.

Looking first at the turnover moments alone, clearly the fit of the model is substantially worse than the baseline. One interpretation is that households in the Barber and Odean (2000) data do not face the unemployment risk of the population as a whole. Consistent with this interpretation is that the model estimated to match “all” moments does fit better

¹⁶Cooper and Zhu (2015) also misses the share considerably. Though not reported, the fit of these moments is closer when the decision period is annual rather than monthly. Clearly time aggregate plays a role here as well.

¹⁷It would be interesting to extract the Barber and Odean (2000) sample from our simulated data. But there is not enough information on income and total financial wealth to do so.

with the unemployment risk present.

Returning to the theme of this exercise, these models also generate large turnover rates and a differential in returns that is consistent with the data facts. But the magnitude of the return differential is not as large as in the baseline model.

Allowing Gains from Financial Trades: $\psi > 1$ The last case relaxes the constraint of $\psi \leq 1$. Note that this not only increases utility from trading but also provides additional resources to the household. This case is therefore difficult to interpret but it is instructive.

Looking at the turnover moments alone, if there is no unemployment risk, then the model allowing $\psi > 1$ fits the moments slightly better. The point estimate is $\psi = 1.0006$ and this small increase in the “cost” of trading helps to match the highest turnover quintile. Clearly though allowing this gain to trading is not key to matching the turnover moments alone.

For the other experiments, the model fit is better once the $\psi \leq 1$ constraint is relaxed, as indicated in the bottom half of these tables. From Table 2, allowing $\psi > 1$ has a large impact on the ability of the model to match both the financial and turnover moments. The best fit for all the moments requires $\psi = 1.338$ with the inclusion of unemployment risk.¹⁸ The large value of this parameter allows us to come closer to matching the wealth to income ratio. With unemployment risk, matching the turnover moments also shows an improved fit relative to the constrained case.

Summary The ability of the model to match the turnover moments stands in contrast to the arguments of Barber and Odean (2000). The model with rational agents creates large turnover as well as the net return differential documented in their study. These agents make these choices to promote their well-defined self-interest. We later introduce various forms of irrationality to determine if it is possible to improve upon these results.

4.3 Identification

Table 5 reports the response of the turnover moments to variations in the three parameters. The reported elasticities are calculated at the baseline estimates and simulated moments, based upon a 1% decrease in the parameters.¹⁹

There are a couple of points illustrated by these calculations. First, the moments are very sensitive to changes in the discount factor and the adjustment cost, (β, ψ) . The response to

¹⁸Thus some households trade a small amount each period to obtain this utility gain. As these trades are small, they appear in the lowest turnover category.

¹⁹The model is non-linear. Thus the magnitude and in some cases the sign of the moment change can depend on both the magnitude and direction of the parameter change.

Table 5: Elasticity of Moments to Parameter Values

parm.	t1	t2	t3	t4	t5	DR
β	-286.08	-53.42	-65.99	-36.54	44.40	43.00
γ	3.39	-0.19	4.01	4.76	0.56	8.58
ψ	100.00	97.36	56.08	44.85	-52.65	20.02

This table reports the elasticity of moments with respect to parameters for the baseline model, based upon a 1 % decrease.

γ is smaller but not insignificant, as some of the higher turnover rate moments are sensitive to variations in the risk aversion.

Second, some of the internal mechanisms of the model are revealed by these elasticities. Again, these are complicated due to the rich nature of the choice model as well as the moments being matched.

Looking at the adjustment cost, a reduction in ψ increases inactivity and thus causes the turnover rates in the lowest turnover quintiles to fall. But, when adjustment occurs, the turnover rate is higher. So the larger adjustment cost increases the mean turnover rate in the $t5$ category. This leads DR to fall since the net return is lower for the large turnover group.

The response to a reduction in β is large and complicated. As households become less patient, the adjustment rate falls in response to income shocks.²⁰ Consequently, the turnover rates fall as well. But the response to a return shock is somewhat different, in part because the gross return is paid to the stock account. Thus consuming from this flow requires turnover. Also, as β falls, all else the same, the level of financial wealth is lower which increases the turnover rate since stock holdings are in the denominator of that rate.

From Table 5, a reduction in β decreases the turnover in the largest quintile and also the difference in the net return. But, a reduction in β actually **increases** the turnover in the lowest 4 quintiles. Although the turnover in these quintiles is lower with the lower discount factor, the level of stock holdings is considerably lower and thus the turnover rate increases.

Finally, as households become more risk averse, consumption smoothing is more important. This increases the frequency of high turnover rates as seen in Table 5.

In addition to these calculations, the estimation entailed multiple starting values searching for the best fit. This would uncover local and global identification problems.

²⁰At the baseline parameters and the average stock return, a 1% increase in β increases the adjustment rate from about 20% to over 80%. Interestingly, these adjustment rates are lower when the stock return takes either its highest or lowest value.

4.4 Robustness

Here we study the robustness of our findings to alternative adjustment costs and measures of turnover rates and net return. In addition, we change our assumption on the treatment of stock earnings so that they are deposited into the bond rather than the stock account. The focus is on the model’s ability to match the turnover and net return differential moments for the case of $\psi \leq 1$.

To better understand these results, Table 6 presents simulations based upon perturbations from the baseline parameters as well as re-estimation results. Table 7 presents the new parameter estimates.

Table 6: Robustness: Data and Model Moments

case	EIS	WI	Sh	t1	t2	t3	t4	t5	DR
Data	0.2990	2.4300	0.6840	0.0019	0.0124	0.0289	0.0598	0.2149	-0.0046
Baseline	0.0418	0.2680	0.9525	0.0016	0.0115	0.0231	0.0456	0.1842	-0.0040
	Simulation								
BO	0.0418	0.2680	0.9525	0.0017	0.0116	0.0232	0.0456	0.1824	-0.0026
Port.	0.0418	0.2680	0.9525	0.0082	0.0224	0.0421	0.0831	0.2137	-0.0029
High Income	0.0769	0.5007	0.9564	0.0018	0.0092	0.0186	0.0380	0.1461	-0.0026
Low Cost	0.0318	0.1851	0.9594	0.0028	0.0109	0.0205	0.0422	0.1174	-0.0010
No Reinvestment	0.0231	0.1743	0.9374	0.0000	0.0000	0.0006	0.0093	0.0977	-0.0019
	Estimation								
F	na	na	na	0.0018	0.0117	0.0219	0.0425	0.2150	-0.0046
BO	na	na	na	0.0020	0.0104	0.0215	0.0424	0.2302	-0.0026
Port	na	na	na	0.0019	0.0106	0.0216	0.0505	0.1839	-0.0025
High Income	na	na	na	0.0019	0.0102	0.0192	0.0385	0.2297	-0.0037
Low Cost	na	na	na	0.0020	0.0109	0.0214	0.0442	0.2199	-0.0015
No Reinvestment	na	na	na	0.0005	0.0070	0.0188	0.0302	0.0834	-0.0011

This table reports data and simulated moments. “EIS” is the elasticity of intertemporal substitution, “WI” is the wealth to income ratio and “Sh” is the stock share, “ti” is the turnover rate for quintile “i” and “DR” is the difference in the net return between the highest and lowest turnover rate quintiles.

4.4.1 Fixed Adjustment Costs

For this experiment, the opportunity cost of adjustment is replaced by a fixed cost, denoted F . Here F is relative to mean income and is restricted, $F \geq 0$.²¹ For the re-estimation, there is no evidence of a fixed cost of adjustment, i.e. $F = 0$. The estimated discount factor and risk aversion are close to the baseline estimates. The fit is, in fact, slightly better than the

²¹As this is a change in the model itself, there are no simulation results possible at baseline parameters.

baseline in terms of matching the turnover moments alone.²² This result clearly indicates the robustness of our findings to other models of adjustment costs.

Table 7: Robustness: Parameter Estimates

case	β	γ	ψ	F	J
baseline	0.9817	1.4070	1.0000	na	0.1594
F	0.9814	1.4033	na	0.0000	0.1459
BO	0.9765	1.4344	0.9998	na	0.3672
Port	0.9857	1.7243	0.9931	na	0.3309
High Income	0.9818	1.3579	1.0000	na	0.3108
Low Cost	0.9809	1.4326	0.9963	na	0.5945
No Reinvestment	0.9927	2.0090	0.9999	na	2.0145

This table reports estimated parameters for the various cases: β is the discount factor, γ is relative risk aversion and ψ is the fraction of income remaining after portfolio adjustment, F is a fixed adjustment cost. J is the difference between model moments and data moments.

4.4.2 Alternative Measure of Net Return

As noted earlier, our measure of returns differs from that of Barber and Odean (2000). Their measure of net return is influenced by the trading pattern. To mimic their measurement, the revised net return is given by:

$$R_i^n = \frac{R^s s_i - C_b(s_{i,-1}, s_i) - C_s(s_i, s_{i'})}{s_i}. \quad (11)$$

For example, if a purchase in the last period is followed by a sale in the current period, then $C_b(s_{i,-1}, s_i)$ is the cost of buying last period and $C_s(s_i, s_{i'})$ is the cost of selling this period. Clearly, this measure is more closely linked to the net return on a trade rather than a portfolio.

From the simulation of the baseline with this alternative measure of turnover, Table 6 indicates that the turnover rates are the same in all quintiles as the baseline and the return differential falls. The re-estimation of the model yields a fit that is not as close as the baseline model. In this case the household's discount factor is estimated to be lower and the risk aversion higher than the baseline. The moments fit deteriorates largely due to excessive turnover in the $t5$ group and a lower DR than the baseline.

²²In keeping with the results reported in Bonaparte, Cooper, and Zhu (2012) about matching financial moments, we maintain the opportunity cost specification as our baseline.

4.4.3 Portfolio Return and Turnover

As is clear, our model has a single stock. Yet households generally hold more than a single stock combined with mutual funds, etc.²³ To mimic this other dimension of portfolio turnover in our model, we treat bonds as a second asset in calculating returns and turnover. As we shall discuss later, this has an interesting and significant impact for models of irrationality.

The treatment labeled “Port.” looks at the household’s total financial wealth, not just stocks. The turnover rate is calculated based on the sum of the absolute changes in the bond and stock holdings. The gross return on the portfolio is a weighted average of bond and stock returns. The net return subtracts stock trading costs.

A simulation using baseline parameters but this alternative definition of the portfolio yields much higher turnover rates than in the baseline. But the differential return is lower. This is a consequence of counting all trades, not just rebalancing, as turnover.

From the re-estimation, the fit is not as good as the baseline. Both the discount factor and the degree of risk aversion are estimated to be higher than in the baseline. The fit is worse because of a lower return differential.

4.4.4 High Income

The traders in the Barber and Odean (2000) study have an average income of about \$72,000 over the sample, substantially larger than the PSID average. Further, these households are direct stockholders, again making them atypical, and presumably reflecting their higher income and wealth status. As noted earlier, we cannot estimate the income process of these households. Nor can we estimate the income process for direct stock holders. In this section, we take a step in that direction by estimating the income process for relatively high income households in the PSID and then re-estimating the model.

The households were selected, as in Cooper and Zhu (2015) by education attainment in excess of 16 years. These higher education households had an average annual income of nearly \$68,000 over the 1978-97 period, closer to the Barber and Odean (2000) households. The monthly income process for these households had about the same persistence as the baseline sample but a higher standard deviation, 0.1151 compared to 0.0904.

Table 6 shows that at the baseline parameters, these higher income households have lower turnover rates and also a lower return differential compared to the baseline. Interestingly, the wealth to income ratio from these high income households is almost double the baseline. From Table 7, the parameter estimates for this case are close to the baseline, with substantial turnover and return differential, though the fit is not as good.

²³Barber and Odean (2000) report that an average household in their sample holds 4.3 stocks.

4.4.5 Low Cost

The case of “low cost” reduces the fixed costs of buying and selling to 10% of their estimated values. The motivation of this is simply that currently trading costs are considerably lower and it is interesting to see if the patterns of turnover and return differential remain.

Not surprisingly, the lower trading costs, at the baseline parameters, leads to higher turnover rates in the lowest quintile and lower turnover rates in the highest quintile. Consequently the return differential is lower.

If the model is re-estimated, at the new parameters turnover in the highest quintile is restored. But the return differential remains very low. The fit with the new estimation is not nearly as good as the baseline. The parameters are close to the baseline though the trading cost, $(1 - \psi)$, is positive and thus a bit higher than the baseline to compensate for the reduction in the fixed cost of trading.

4.4.6 Deposit of Stock Returns in Bond Account

In the baseline model, it is assumed that even if there is no portfolio adjustment, stock returns are deposited in the stock account. As a consequence, the stock share increases in high return states unless the adjustment cost is paid. When there is adjustment, there are stock sales in high return states.

An alternative is to assume that stock returns are paid to the bond account. With that assumption, portfolio rebalancing will go in the other direction in response to a high return shock. The effect of this modification on the moments will depend, in part, on the non-linear nature of the decision rules.

From the simulation at the baseline parameters, this modification creates a lot of portfolio adjustment inaction and even reduces the turnover in the $t5$ category. Further the return differential is lower.

When the parameters are re-estimated, the fit is still not nearly as close as the baseline. Households are estimated to be more patient and more risk averse. Still the turnover rates and return differential are lower than in the baseline and the data.

5 Irrational Beliefs

The presence of transactions costs are almost enough to both replicate the patterns of turnover and net returns found by Barber and Odean (2000). Prompted by the discussion of overconfidence in Barber and Odean (2000), we turn to alternative explanations that relax the assumption of perfect rationality. To be clear, as seen above, these irrational beliefs

are not needed to match the finding of a substantially lower return for high turnover traders. This analysis includes, in various forms, the favored explanation of overconfidence by Barber and Odean (2000).

For this exercise, we build upon the baseline model, i.e. the one without unemployment risk and ψ restricted. We present simulation results at the baseline parameters, perturbed by individual sources of irrationality.

Table 8 presents simulation results along with a measure of how well the alternative models fit the turnover moments. For discussion, the table also presents the financial moments, though these are not included in the fit measure. The simulations are intended to clarify how the behavior of households depends on their beliefs.²⁴

In a second exercise, the parameters were re-estimated, allowing for these alternative forms of irrationality. Those results are reported in the bottom part of Table 8 and the estimates are in Table 9. For this re-estimation, we study the baseline case of “no unemployment risk”. But in this case, we focus on matching all moments. **Importantly, there were no improvements in fit by the addition of irrational agents if only the turnover moments were matched.**

Table 8: Data and Model Moments: Irrational Beliefs

	EIS	WI	Sh	t1	t2	t3	t4	t5	DR	J
Data	0.2990	2.4300	0.6840	0.0019	0.0124	0.0289	0.0598	0.2149	-0.0046	na
Baseline	0.0418	0.2680	0.9525	0.0016	0.0115	0.0231	0.0456	0.1842	-0.0040	0.1594
	Simulation									
$\tilde{\mu}_R$	0.2934	1.4102	0.9911	0.0030	0.0104	0.0161	0.0218	0.0361	-0.0003	2.5272
$\tilde{\sigma}_R$	0.0338	0.2469	0.9601	0.0005	0.0062	0.0160	0.0306	0.1152	-0.0033	1.5257
$\tilde{\rho}_R$	0.0410	0.2641	0.9504	0.0005	0.0073	0.0202	0.0369	0.1419	-0.0033	1.1446
σ_y	0.0459	0.2844	0.9514	0.0033	0.0111	0.0201	0.0366	0.1246	-0.0036	1.0210
δ	0.0751	0.5610	0.9214	0.0000	0.0019	0.0126	0.0339	0.2873	-0.0056	2.3835
	Estimation									
Rational	0.2392	0.9883	0.9332	0.0018	0.0101	0.0206	0.0341	0.2063	-0.0021	1.1123
$\tilde{\rho}_R$	0.2375	0.9801	0.9272	0.0020	0.0106	0.0218	0.0385	0.2167	-0.0021	1.0213
σ_y	0.2294	1.1399	0.9192	0.0021	0.0097	0.0197	0.0364	0.2288	-0.0044	0.7747
z opt.	0.4654	2.1482	0.9013	0.0021	0.0080	0.0214	0.0451	0.2577	-0.0020	1.0426

This table reports data and simulated moments for irrational beliefs. For the treatments, σ_y allows excess volatility income, δ is the perceived unemployment risk, $\tilde{\mu}_R$ is the mean return, $\tilde{\sigma}_R$ is the standard deviation of the return and $\tilde{\rho}_R$ allows serial correlation in the stock return. For the simulations, $(\sigma_y, \tilde{\mu}_R, \tilde{\sigma}_R)$ are all increased by 10%. The value of $\tilde{\rho}_R$ is set at 0.10. δ is increased by 10% relative to the estimate of 0.014. The bottom block displays moments from estimation exercises that lead to improvements in the fit due to the inclusion of irrational agents.

²⁴Consistent with this, Table 11 presents the response of moments to parameters in the case of irrational beliefs about the volatility of income.

5.1 Returns

Consider the following process for the beliefs of agents about returns:

$$R_t = \tilde{\mu}_R + \tilde{\rho}_R R_{t-1} + \varepsilon_t \quad (12)$$

where ε is normally distributed with a mean of 0 and a standard deviation of $\tilde{\sigma}_R$.²⁵ The mean return is denoted $\tilde{\mu}_R$ and the serial correlation is $\tilde{\rho}_R$. These are perceived parameters that may not coincide with the true process for returns. Indeed, our interest is in studying the relationship between beliefs and the true process for trading strategies and portfolio returns.

From (12), the specification permits three types of deviations through the: (i) mean, (ii) standard deviation and (iii) persistence of the return process. Beliefs about a positive serial correlation in the return captures the frequently noted belief in stock market “momentum”.

The first three rows of Table 8 report simulation results in which the mean and the standard deviation of the return process were increased by 10%. The case of $\tilde{\rho}_R$ is the perceived serial correlation at 0.10 rather than its estimated value of 0.00.

Increasing the mean perceived return has a large impact on the moments. The turnover rate increases in the lowest quintile and falls considerably in the highest quintile relative to the baseline. The return differential almost disappears. For the financial moments, the higher mean return increases both the wealth to income ratio as well as the stock share.

Increasing the perceived standard deviation of returns, i.e. the riskiness of stocks, reduces turnover in all categories. The return differential is also lower.

When returns are believed to be serially correlated, the realized return has an effect on current wealth and on the distribution of future returns. The latter effect is like a signal. If $\tilde{\rho}_R$ is positive, then households are more confident that current returns provide information about future returns. At $\tilde{\rho}_R = 0.10$, as in the case of increased riskiness of returns, the turnover rates are reduced in all categories and the return differential is lower as well.

These simulations indicate that household behavior depends on their beliefs. Thus inferring beliefs from observed choices is feasible.²⁶ With this in mind, the model was re-estimated allowing perceptions of the return process to differ from the data.

There is one case in which the re-estimation led to an improvement in the fit of all moments: misperception in the serial correlation of returns. From Table 9, the estimated

²⁵Here we start from a standard AR(1) model and draw on the discussion in DeLong, Shleifer, Summers, and Waldmann (1991). While there are numerous papers in the literature using the concept of overconfidence, there are relatively few which point to a particular model of overconfidence. Gervais and Odean (2001) study overconfidence in a learning model that is beyond the scope of our study. Guiso and Jappelli (2006) study the effects of overconfidence on information acquisition.

²⁶For this estimation, we tried multiple starting values to ensure identification.

belief about the serial correlation of stock returns is 0.0265. From Table 8 this leads to a fall in the fit by 0.09 points. This increase in the serial correlation increased the turnover rates, particularly in the higher quintiles. Still the wealth to income ratio remained quite low relative to the data.

5.2 Income

The case of income volatility allows the agent to have a view of the income process that differs from that estimated from the PSID. If agents perceive the income process to be more volatile than it is, then this will create a demand for liquidity. Further, it motivates more trade and thus increases turnover rates, as in the data.

We study, through simulation, two deviations from rational beliefs. First, we allow an agent's belief of the standard deviation of the idiosyncratic income innovation, denoted σ_y , to increase by 10%. Second, we increase beliefs about unemployment risk by 10% relative to the estimate of 0.014.²⁷

The moments for these cases are given in Table 8. Relative to the baseline, the increase in σ_y increases the turnover in the lowest quintile and reduces it in the highest quintile. The effect on the return differential is minimal. Increasing the risk of unemployment increases the highest turnover rate and also increases the return differential. Neither of these alternatives improves the fit. Note also that neither improves the fit of the financial moments, particularly with regards to the stock share.

Table 9: Irrational Beliefs: Parameter Estimates

case	β	γ	ψ	σ_y	$\tilde{\rho}_R$	p_{zg}
Rational	0.9907	0.8838	0.9999	0.0904	0	0
$\tilde{\rho}_R$	0.9907	0.8807	0.9999	na	0.0265	na
σ_y	0.9910	0.9313	0.9886	0.1669	na	na
z opt.	0.9890	1.0995	0.9850	na	na	0.0595

Only cases of irrationality that improve the model fit are reported in this table. The cases reported here are with no unemployment risk and matching all moments. The last three elements in the first row are the actual estimates of the parameters from the data.

There was one case in which the re-estimation led to an improvement in model fit. From Table 8, a model in which agents misperceive income risk leads to a substantial improvement in fit by about 30%. From Table 9, the estimated value of σ_y is almost twice that of the

²⁷For this case, we increase the unemployment risk relative to the estimated case with unemployment risk.

baseline. The fit improves largely because of the increase in the wealth to income ratio and the increase in the return differential.

5.3 Noisy Advice

Here we consider another form of irrational beliefs associated with a signal provided, say, by a financial advisor about future returns. From our specification of the stock return process, future returns are not predictable. But an agent may be induced to believe the advice of an advisor, leading to excessive turnover and relatively low net returns.

To study this formally, assume there is an iid signal, denoted z , that the household believes is correlated with future returns. The discrete choice is again given by (1) and the options of adjustment and non-adjustment given by (2) and (3) respectively with the modified state vector of (Ω, z) .

This choice is given:

$$v(\Omega, z) = \max\{v^a(\Omega, z), v^n(\Omega, z)\} \quad (13)$$

for all (Ω, z) . The revised options are:

$$\begin{aligned} v^a(\Omega, z) &= \max_{b' \geq 0, s' \geq 0} u(c) + \beta E_{\Omega'|\Omega, z} \int_{z'} v(\Omega', z') dG(z') \\ \text{s.t.} \\ c &= \psi y + R^b b + R^s s - b' - s' - C(s, s'). \end{aligned} \quad (14)$$

if adjustment. If the household does not adjust, it solves

$$\begin{aligned} v^n(\Omega, z) &= \max_{b' \geq 0} u(c) + \beta E_{\Omega'|\Omega, z} \int_{z'} v(\Omega', z') dG(z') \\ \text{s.t.} \\ c &= y + R^b b - b' \quad (15) \\ s' &= R^s s. \quad (16) \end{aligned}$$

In these expressions, $G(z')$ is the cdf of z' . The conditional expectation in these expressions highlights that the sole role of z is to provide information about Ω' .

As households believe z is informative about future returns, their decisions will depend on this random variable. This source of irrationality is similar to overconfidence about the serial correlation about returns but realizations of z only influence household beliefs, not their budget sets. This allows z to be, at least in the mind of the household, a predictor of future stock returns, $R^{s'}$.

Table 10: Data and Simulated Model Moments: Noisy Advice

	EIS	WI	Sh	t1	t2	t3	t4	t5	DR	J
Data	0.2990	2.4300	0.6840	0.0019	0.0124	0.0289	0.0598	0.2149	-0.0046	
Baseline	0.0418	0.2680	0.9525	0.0016	0.0115	0.0231	0.0456	0.1842	-0.0040	0.1594
	z opt.									
Stocks	0.0542	0.3426	0.9527	0.0012	0.0092	0.0203	0.0414	0.1873	-0.0042	0.4096
Port.	0.0542	0.3426	0.9527	0.0067	0.0205	0.0404	0.0786	0.2020	-0.0029	7.2063
	z pess.									
Stocks	0.0389	0.2537	0.9383	0.0018	0.0116	0.0222	0.0448	7.7505	-0.0432	1300.1
Port.	0.0389	0.2537	0.9383	0.0078	0.0214	0.0393	0.0816	0.2797	-0.0030	10.644

This table reports data and simulated moments for irrational beliefs. “EIS” is the elasticity of intertemporal substitution, “Ad” is the adjustment rate, “WI” is the wealth to income ratio and “Sh” is the stock share, “ti” is the turnover rate for quintile “i” and “DR” is the difference in the net return between the largest and smallest turnover rate cells. For the simulations, $p = 0.01$. The case “Stocks” calculates turnover from stock trades, as in the baseline model, while “Port.” looks at turnover of the portfolio.

To implement this, assume $z \in \{0, 1\}$. When $z = 0$, households believe that R^s is an iid process, as specified in the baseline model. Alternatively, $z = 1$ leads the households to put all weight on the lowest (highest) realization of $R^{s'}$. The simulation sets p , the probability that $z = 1$ for the two cases in which this realization of z is viewed as either extremely good or bad news, at 1%.

Table 10 presents the simulation results. For each type of signal, there are two calculations of the turnover moments. The case “Stocks” calculates turnover from stock trades, as in the baseline model, while “Port.” looks at turnover of the portfolio as discussed in sub-section 4.4.3. The latter case is of interest since the news may lead to both excessive bond as well as stock trades.

As is clear from the last column of the table, introducing advice of this form does not improve the fit of the model. There is a huge asymmetry between the optimistic and pessimistic cases. When the signal is about future bad returns, the turnover rate is extremely large. Essentially agents who receive bad news in the current period sell their stock holdings. In the subsequent period, they rebuild their portfolio, which leads to a large turnover rate. The resulting differential in return is almost 10 times that of the baseline. As $p = 0.01$, the turnover is only affected in the top quintile. When the measure of turnover is portfolio, this large adjustment disappears since the stock holdings do not fall close to zero upon the good news.

The results are different when the news is good, so that weight is put on the highest return state rather than the lowest. In this case, the turnover rates in the lowest quintiles fall slightly but the large trades appearing in the case of the pessimistic shock does not occur.

These simulations again show that choices depend on these signals. Going a step further, the model was re-estimated allowing for these signals. The estimation included the likelihood of the news.

In the case of “good news”, labeled “z opt.” in Tables 8 and Table 9, introducing advice did improve the model fit by 0.07 points. The estimated probability of “good news”, denoted p_{zg} , was nearly 6%. This led to a reduction in the estimated value of ψ . Relative to the (rational) baseline, this led to higher turnover in all but one of the quintiles and a substantially higher wealth to income ratio.

6 Conclusion

The goal of this paper was to assess the claim made by Barber and Odean (2000) that the patterns of returns as a function of portfolio turnover was consistent with overconfident agents and inconsistent with rational traders. The approach was to study the implications of a household dynamic optimization problem, emphasizing both the extensive and intensive margins of portfolio choice. In our model, portfolio adjustment costs created both inaction and large turnover and also drove a wedge between gross and net returns.

Parameters are estimated to match moments turnover rates and net return patterns. We estimated models with both rational and irrational agents. For the latter, we study specifications in which agents either hold beliefs about income or returns that are not consistent with the data. We also study cases in which agents receive signals of future returns that are false.

We reach two conclusions. First, models with rational agents can match the turnover moments emphasized in Barber and Odean (2000). Second, introducing various forms of irrationality do influence household choices, but does not improve the fit of the model with respect to the turnover and net return moments.

References

- BARBER, B., AND T. ODEAN (2000): “Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors,” *The Journal of Finance*, 55(2), 773–806.
- BARBER, B. M., AND T. ODEAN (2001): “Boys will be boys: Gender, overconfidence, and common stock investment,” *The Quarterly Journal of Economics*, 116(1), 261–292.

- BONAPARTE, Y., AND R. COOPER (2009): “Costly Portfolio Adjustment,” NBER Working Paper # 15227, revised 2010.
- BONAPARTE, Y., R. COOPER, AND G. ZHU (2012): “Consumption Smoothing and Portfolio Rebalancing: The Effects of Adjustment Costs,” *Journal of Monetary Economics*, 59, 751–68.
- CARROLL, C. D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 23(1992-2), 61–156.
- COOPER, R., AND G. ZHU (2015): “Household Finance: Education, Permanent Income and Portfolio Choice,” *Review of Economic Dynamics*, 20, 63–89.
- DELONG, J., A. SHLEIFER, L. SUMMERS, AND R. WALDMANN (1991): “The survival of noise traders in financial markets,” *The Journal of Business*, 64(1), 1–19.
- GERVAIS, S., AND T. ODEAN (2001): “Learning to be overconfident,” *Review of Financial Studies*, 14(1), 1.
- GUISSO, L., AND T. JAPPELLI (2006): “Information Acquisition and Portfolio Performance,” mimeo.
- KYLE, A. S., AND F. A. WANG (1997): “Speculation duopoly with agreement to disagree: Can overconfidence survive the market test?,” *The Journal of Finance*, 52(5), 2073–2090.
- ODEAN, T. (1998): “Volume, volatility, price, and profit when all traders are above average,” *The Journal of Finance*, 53(6), 1887–1934.
- VISSING-JORGENSEN, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110(4), 825–853.

7 Appendix

This appendix provides additional information about the calculation of the income process and other moments.

7.1 Income

Notation

The following uniform notation is used throughout this note. For income, denote the absolute level by Y , the log by y , and the residuals by \tilde{y} . The superscript a indicates annual and m

indicates monthly data.

Annual Income

From Bonaparte, Cooper, and Zhu (2012), henceforth BCZ, the annual income process for stock holders obtained by the following steps (as described in the BCZ appendix): pool all the observations together, regress income on age, age^2 , education attainment, gender and marital status. Take the residuals from the regression and use the residuals to run the AR(1) process

$$\tilde{y}_t^a = \rho^a \tilde{y}_{t-1}^a + \epsilon_t^a. \quad (17)$$

The persistence of the income shock is estimated to be $\rho^a = 0.84224$, and the standard deviation of the innovation is $\sigma_\epsilon = 0.29027$.

Monthly Income

From Bureau of Labor Statistics CPS survey, <https://www.bls.gov/webapps/legacy/cpsflowstab.htm>, data on flows and levels are obtained to calculate the monthly probability of being separated from employment, the probability of finding a job conditional on being unemployed. Our definition of unemployment includes unemployment and not in the labor force. The probabilities are calculated from averaging over monthly flows over stocks. For example, the probability of being unemployed is $\frac{\text{flow}(e \rightarrow u + e \rightarrow n)}{\text{employmentlevel}}$.

We make the following assumptions on the monthly income process. Suppose a person was unemployed last month, then with prob p_{ue} , he is employed this month, and receives the average income, 6.²⁸ In logs, $y_{ue}^m = \log(6)$. With prob $1 - p_{ue}$, this person is still unemployed, he receives the unemployed benefit, with is 0.4 times the average monthly income, so this translates to $y_{uu}^m = \log(0.4 \times 6)$.

Suppose a person was employed last month, then with probability p_{ee} , he is still employed this month and receives

$$\tilde{y}_t^m = \rho^m \tilde{y}_{t-1}^m + \epsilon_t^m. \quad (18)$$

In levels, this corresponds to

$$y_t^m = \mu_m(1 - \rho^m) + \rho^m y_{t-1}^m + \epsilon_t^m \quad (19)$$

Since we assume that the innovation ϵ_t follows normal distribution of mean 0 and standard deviation σ_ϵ , Y^m follows log normal distribution of parameters (μ_m, σ_y^m) , where $\sigma_y^2 = \frac{\sigma_\epsilon^2}{1 - \rho^2}$.

²⁸This corresponds to the monthly average income of \$6000, annual \$72000 for stock holders.

Hence,

$$6 \equiv \mathbb{E}Y^m = e^{\mu_m + \frac{\sigma_y^2}{2}} \rightarrow \mu^m = \log(6) - \frac{\sigma_y^2}{2} \quad (20)$$

With prob $1 - p_{ee}$, this person gets unemployed and get $y_{eu}^m = \log(0.4 \times 6)$.

7.1.1 Estimation

Given these monthly flows into and out of unemployment, it is necessary to estimate the parameters, $(\rho^m, \sigma_\epsilon^m)$, of the monthly income process. This is essentially a SMM exercise with the annual parameters, $(\rho^a, \sigma_\epsilon^a)$ as moments to match. The following algorithm was used for this purpose:

1. guess a vector of parameters $\rho^m, \sigma_\epsilon^m$
2. simulate a panel of level of monthly incomes, aggregate to a panel of annual incomes (level)

$$y^a = \log\left(\sum_{i=1}^{12} \exp(y_i^m)\right) \quad (21)$$

3. use the simulated annual income panel to do the following AR(1) regression

$$y_t^a = \mu_a(1 - \rho^a) + \rho^a y_{t-1}^a + \epsilon_t^a \quad (22)$$

and calculate the variance of the residual.

4. compare the ρ^a and σ_ϵ^a with the estimates from BCZ and go back to step 1 if not close enough.

As indicated in the text, the serial correlation of the monthly income process conditional on employment is estimated to be 0.9959, and the standard deviation of the innovation is 0.0839 under the mean recovery case.

7.2 Financial Moments

7.2.1 Measurement

Elasticity of Intertemporal Substitution (EIS) is obtained by the following regression

$$\frac{1}{H} \sum_{i=1}^H (\ln C_{t+1}^i - \ln C_t^i) = \beta_0 + \sigma R_{t+1} + u_{t+1}$$

which corresponds to Equation (9) in Vissing-Jorgensen (2002). This is a semiannual moment as t corresponds to 6 months. In the simulation, 6-month windows are constructed, within which half-year consumption profiles are aggregated from monthly consumption profiles. For each household i at time t , the growth rate is calculated, and then at each t we average over these households to get the average log difference. The return R_{t+1} is defined as the product of the monthly returns in 6 months.

Wealth Income Ratio (WI) is an annual moment. In the simulation, 12-month windows are created. 12 monthly incomes are aggregated to get the annual income while the wealth at time t is the end-of-the-period sum of stock and bond holdings. The WI is obtained by dividing the sum of annual wealth positions over the sum of all annual incomes. Then an average over time is taken.

Stock Share (SH) is a triennial moment. In the simulation, 3-year windows are created. At the end of each t , SH is calculated as the sum of stockholdings of all households over the sum of the financial wealth of all households. Then an average over time is taken.

Turnover and Net Return To mimic the data set in BO, a six-year window is constructed in the simulation. In each month, households are put into different quintiles according to their turnover rates, the average turnover rate in each quintile q_t^i is calculated as well as the average net return $Rnet_t^i$. Finally a time average is taken to get the mean turnover and mean net return in each quintile. The difference in net return is calculated by subtracting the mean net return of the highest quintile from that of the lowest quintile.

7.3 Elasticity of Moments to Parameters: σ_y Case

Table 11: Elasticity of Moments to Parameter Values

parm.	EIS	WI	Sh	t1	t2	t3	t4	t5	DR
β	-95.2134	-93.1816	-11.4822	-100.0000	-100.0000	-99.9197	-86.7350	71.1842	372.6229
γ	0.6658	-0.3164	-0.1408	7.8129	2.3014	0.7020	1.4694	16.0455	-1.5627
ψ	-8.4888	-3.7727	0.6946	-43.8412	-19.4917	-17.8923	-8.8820	17.7616	15.3085
σ_y	1.1440	-1.1773	-0.4873	10.3505	4.3297	2.3675	3.0928	19.5675	5.2176

This table reports the elasticity of moments with respect to parameters for the baseline model, based upon a 1 % decrease.