

NBER WORKING PAPER SERIES

RATIONALIZING TRADING FREQUENCY AND RETURNS

Yosef Bonaparte
Russell Cooper

Working Paper 16022
<http://www.nber.org/papers/w16022>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2010

We are grateful to the National Science Foundation for supporting this research. 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. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Yosef Bonaparte and Russell Cooper. 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
Yosef Bonaparte and Russell Cooper
NBER Working Paper No. 16022
May 2010, Revised November 2010
JEL No. E21,G11

ABSTRACT

Barber and Odean (2000) study the relationship between trading frequency and returns. They find that households who trade more frequently have a lower net return than other households. But all households have about the same gross return. They argue that these results cannot emerge from a model with rational traders and instead attribute these findings to overconfidence. Using a dynamic optimization approach, we find that neither a model with rational agents facing adjustment costs nor various models of overconfidence fit these facts.

Yosef Bonaparte
Robert Day School of Economics and Finance
Claremont McKenna College
500 East Ninth Street
Claremont, CA 91711-6400
yosef.bonaparte@claremontmckenna.edu

Russell Cooper
Department of Economics
European University Institute
via della Piazzola, 43
Firenze, 50133
ITALY
and University of Texas
and also NBER
russellcoop@gmail.com

Rationalizing Trading Frequency and Returns*

Yosef Bonaparte[†] and Russell Cooper[‡]

November 17, 2010

Abstract

Barber and Odean (2000) study the relationship between trading frequency and returns. They find that households who trade more frequently have a lower net return than other households. But all households have about the same gross return. They argue that these results cannot emerge from a model with rational traders and instead attribute these findings to overconfidence. Using a dynamic optimization approach, we find that neither a model with rational agents facing adjustment costs nor various models of overconfidence fit these facts.

1 Motivation

Barber and Odean (2000) find that households who adjust their portfolio more frequently 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

*We are grateful to the National Science Foundation for supporting this research. 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.

[†]Robert Day School of Economics and Finance, Claremont McKenna College, 909-607-0756, ybonaparte@cmc.edu

[‡]Department of Economics, the European University Institute and Department of Economics, University of Texas at Austin, russellcoop@gmail.com

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). The first looks the choice of rational agents faced with costs of portfolio adjustment. The second allows for overconfidence.

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 will 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. Thus both ingredients are necessary to match the observations assuming agents are rational.

Building upon Bonaparte and Cooper (2009), 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 portfolio. We generate simulated data from the estimated model to study the relationship between portfolio adjustment and returns.

Second, following the suggestion of Barber and Odean (2000), we study a series of models which relax the assumption of perfect rationality to model overconfidence. We consider models in which traders over-estimate the mean of the return process, under-estimate the variance or over-estimate the serial correlation.

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 then viewed as irrational.

We do not concur. It is certainly possible for gross and net returns of traders to be below those of non-adjusters in an optimizing framework. The question is quantitative: do the costs of trade, the processes for income and returns and the rational choices of households generate the pattern of gross and net returns found in the data?

We find that none of these specifications are capable of matching the observed differences between gross and net returns as a function of trading frequency. While there are differences

between returns earned by adjusters and non-adjusters, these are reflected in both the gross and net returns earned by traders. In contrast, the facts presented by Barber and Odean (2000) highlight differences in net but not gross returns.

Our results come from two sources. First, the estimated adjustment costs are not large enough to explain the observed differences between gross and net returns. Second, at the estimated parameters, the selection effects coming from the household choices on both the extensive and intensive margins are very powerful so that, in our baseline model, adjusters earn a higher net (and thus gross) return than non-adjusters. In the model with overconfidence, there are cases in which adjusters earn less than non-adjusters. But even here, the selection effect is the dominant factor as the difference appears in both the gross and net returns.

While trading is costly, how hazardous it actually is remains an open issue.

2 Household Behavior: Model and Estimates

Here we briefly review the model and estimates of Bonaparte and Cooper (2009) that is the basis of the household optimization problem we study. 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 is equal to its labor income. 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 adjustment conditional on having incurred fixed adjustment costs.

2.1 Household Optimization

Denote by $v(\Omega)$ the value of the household's problem in state $\Omega \equiv (y, s_{-1}, R_{-1})$ where y is current income, s_{-1} is holding of the single asset from the previous period and R_{-1} is the return from the previous period.¹ Total financial wealth this period is $R_{-1}s_{-1}$.

The household chooses between the options of adjusting or not:

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

for all Ω . If the household chooses to adjust, then

$$v^a(\Omega) = \max_s u(c) + \beta E_{R, y' | R_{-1}, y} v(\Omega') \quad (2)$$

¹For the purpose of this exercise, the household has a single financial asset.

where s is the new holding of the asset and Ω' is the future value of Ω . Household consumption is

$$c = R_{-1}s_{-1} + y \times \psi - s - C(s_{-1}, s). \quad (3)$$

There are two costs of adjustment here. The first, given by the function $C(\cdot)$, represents direct trading costs. The second, parameterized by ψ , represents the lost income due to time spent on portfolio adjustment. As households will experience different realizations of income they will have different adjustment costs.

If there is no portfolio adjustment, then

$$v^n(\Omega) = u(y) + \beta E_{R,y'|R_{-1},y} v(\Omega'). \quad (4)$$

In this case, $c = y$ and $\Omega = (y', s_{-1}R_{-1}, R)$ so that gross proceeds from the existing portfolio create the portfolio for the current period without any cost, $s = s_{-1}R_{-1}$.²

The policy functions generated by household optimization include an extensive margin (adjust, no adjust) and an intensive margin indicating the magnitude of the adjustment. Note that optimizing households base their choices on realized income and returns. Their choice of portfolio turnover reflects net rather than gross returns.

2.2 Quantitative Analysis

Bonaparte and Cooper (2009) estimate trading costs, $C(\cdot)$, directly from the data set used by Barber and Odean (2000). A simulated method of moments approach is used to estimate other parameters.

2.2.1 Trading Costs

Bonaparte and Cooper (2009) 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

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

They use the monthly household account data Barber and Odean (2000) to estimate these parameters.³ The trading costs are estimated in a regression where the dependent

²Bonaparte and Cooper (2009) discuss and estimate alternatives to this model of no adjustment. This specification fits the data best.

³Details 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.

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.

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

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 current 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 0.3% 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.

2.2.2 Income and Returns

The appendix of Bonaparte and Cooper (2009) contains more detailed information on the data. Income and returns are modeled as AR(1) processes. Household income variable is the product of a common shock and a household specific shock.

We consider a couple of return processes. The first, as in Bonaparte and Cooper (2009), measures real returns, including dividends, and comes from Robert Shiller, available at <http://www.econ.yale.edu/~shiller/data.htm>. The mean return is 5.51% over the 1967-1992 period with a serial correlation of essentially zero and a standard deviation of 0.1589. We refer to this as the “baseline”.

The second measure of returns is computed directly from the household data used in Barber and Odean (2000).⁴ This return measure is not used in Bonaparte and Cooper

⁴For the return process, we employ the Barber and Odean (2000) data. Our sample selection includes

(2009) but is used here to try to mimic the returns faced by the households whose behavior generated the differences in gross and net returns described by Barber and Odean (2000). For that process, the mean return is about 17.6% over the sample, the serial correlation is 0.107 and the standard deviation of the innovation is 0.2725. We refer to this case as “HH” since it is based upon the average of household returns in the Barber and Odean (2000) data.

2.2.3 Moments

Bonaparte and Cooper (2009) use three moments to identify the parameters (β, γ, ψ) . All moments are from annual data.

The first moment comes from the SCF data set where in an average year 71% of households adjustment their portfolio. The second moment comes from the estimation of the log-linear approximation of a consumption Euler equation, drawing upon Hansen and Singleton (1983):

$$\log\left(\frac{c_{t+1}}{c_t}\right) = \alpha_0 + \alpha_1 \times \log(R_{t+1}) + \zeta_{t+1}. \quad (7)$$

The estimate of $\alpha_1 = 0.0878$ is from data on real consumption growth of non-durables and services of stock market participants and the Shiller measure of return for the 1967-92 period.

The final moment is the median financial wealth to income ratio, which is 1.03 in the data. This moment is quite informative about β .

2.2.4 Estimation Results

Estimation involves the solution of an agents dynamic optimization problem, (1)-(4), and the creation of a simulated panel data set with 500 households and 500 time periods. Households differ because of idiosyncratic income shocks which generates differences in trading patterns and returns.

Parameter estimates are obtained by minimizing the distance between the simulated and actual moments. Bonaparte and Cooper (2009) estimated (β, γ, ψ) using three moments: the frequency of inaction, the estimate of α_1 from (7) and the median financial wealth to income ratio.

The moments from the data and the parameter estimates are summarized in Table 2.⁵ The column labeled “Data” shows the actual moments and the column labeled “Baseline” is

households who hold stocks consecutively for the whole period from February 1991 through December 1996. The selected sample contains 1,079,877 observations for 14,478 households; each household has 71 monthly consecutive observations.

⁵The weighting matrix was the identify matrix for this exercise so that the fit is simply the sum of the squared differences between the actual and simulated moments.

from the estimated model. The column labeled “HH” are results based upon the household return process.

Moments	Data	Baseline	HH(data)	SC
Portfolio Adjustment rate	0.71	0.72	0.61 (0.52)	0.74
α_1	0.088	0.110	-	0.0969
wealth income ratio	1.03	1.03	3.80 (3.80)	1.04
Estimated Parameters				
γ	-	3.51	-	3.51
ψ	-	0.974	0.834	0.974
β	-	0.88	0.797	0.88
fit	-	0.0006	0.0091	0.001

Table 2: Moments and Parameters

Baseline The data moments are summarized in the second column. The estimates are presented in the third column.⁶ The moments are pretty well matched though the value of α_1 is not quite as low as in the data. The degree of risk aversion is considerably less than the inverse of α_1 . Interestingly, there is evidence of adjustment costs beyond the trading costs: $\psi < 1$. The fit of 0.0006 is calculated as the sum of the squared differences between the simulated and data moments.

HH Returns The results using the return process estimated from household data are presented in the “HH” column. The portfolio adjustment rate and the wealth income ratio moments were recomputed for this sample and are indicated in the table as 0.52 and 3.80 respectively. Due to data limitations for these households, we cannot re-estimate α_1 .

Given the portfolio adjustment rate and wealth income ratio as moments, we re-estimated (ψ, β) to match these two parameters, fixing $\gamma = 3.51$. The estimated β is much lower for this sample, reflecting the high average return on this portfolio. Further the cost of adjustment, ψ , is much lower as well to generate the lower adjustment rate of 52% in the household sample. The fit is not as good with this sample as it is for the baseline model. Still, this sample will provide additional information on the differential of returns for adjusting and non-adjusting households.

⁶Bonaparte and Cooper (2009) study the robustness of these estimates to different measures of return and other specifications of the model.

3 Returns and Trading

Given this model and estimates, we now turn to the main point of the analysis: the relationship between returns (both gross and net) and trading patterns. As noted above, Barber and Odean (2000) find an inverse relationship between net returns and trading frequency but no significant differences in gross returns.

We use our model to evaluate this evidence. In our model all agents have rational expectations so that any difference in returns associated with trading frequency comes from the optimal choices of households. For these households, their choices on the extensive and intensive margins depend on perceived net returns.

In theory, there are two ways in which the model can link differences in net return to trading frequency. The first is direct: the presence of a trading cost will reduce the net return. The second, more subtle link, can be generated by selection: agents choose whether to adjust or not (the extensive margin) and how much to adjust (the intensive margin) in a state contingent manner. Thus optimal behavior on the part of agents will itself generate a relationship between returns and trading decisions on both the extensive and intensive margins.

Table 3 presents calculations of different return measures for the baseline model using the parameters reported in Table 2. The gross return is simply the annual return on the portfolio. The second return measure nets out the financial costs, from Table 1, of trading. The third measure nets out both the financial cost and the income loss due to ψ .⁷

The timing here is important. The return is realized from period $t - 1$ to t . The adjusters/no-adjusters distinction refers to the decisions taken in period t . Thus the table analyzes whether adjustment responds to higher or lower realized returns, as in the state vector of the optimization problem. In this sense, the dependence of the return on adjustment reflects the selection of whether to adjust or not.

The columns of the table relate to the extensive and intensive margins. The second and third columns of the table look at agents who adjust and those who choose not to adjust. The fourth, fifth and sixth columns compute return for the lowest, middle and highest quintiles of the turnover rate distribution.⁸

Baseline Looking first at the results from the baseline, the table reveals that neither the transactions cost explanation of the net return differential between adjusters and non-

⁷To be specific, the net return is the gross return minus the average cost of trading. The average cost of trading equals the total cost of trading divided by the quantity traded. For timing, the net return in period t is the gross return earned between periods $t - 1$ and t minus the trading costs incurred in period $t - 1$ as a fraction of the period $t - 1$ asset trade.

⁸The turnover rate is the absolute value of the net change in the portfolio divided by its initial value.

Return	Adjusters	Non-adjusters	Lowest Turn.	Middle Turn.	Highest Turn.
Baseline					
Gross	1.060	1.049	1.050	1.061	1.065
Net FC	1.058	1.049	1.050	1.060	1.061
Net All	1.025	1.049	1.050	1.049	0.977
HH					
Gross	1.189	1.151	1.149	1.066	1.263
Net FC	1.189	1.151	1.149	1.066	1.262
Net All	1.134	1.151	1.149	1.046	1.157

Table 3: Returns and Trading

adjusters nor the selection story fits the facts. For both cases, the key observation is that the gross return for adjusters, 1.060, exceeds that of the non-adjusters of 1.049. This is inconsistent with the findings of Barber and Odean (2000).

To understand these results, consider the transactions costs explanation for the differentials in returns. Looking at adjusters, the return net of transactions costs is only slightly lower than the gross return. From Table 1, the main adjustment cost is the fixed component of around \$60.00. For these costs to create a differential in return of 6 percentage points, as reported by Barber and Odean (2000), the average trade would have to be about \$1,000. In fact, the average trade in the data is about \$12,500 and is about \$13,800 in the simulated data. Thus the fixed cost is much too small relative to the trade size in both the actual and simulated data to explain the differential in returns.

The selection effect can create differences in net return based on trading frequency. But, that explanation runs into two problems. First, as can be seen from Table 3, the selection effect generates differences in gross as well as net returns. This runs counter to the evidence from Barber and Odean (2000). Second, the selection effect creates a higher, not lower, return for the adjusters. This is because agents adjust more frequently in high return states than in low return ones. We return to this point below.

Barber and Odean (2000) study the returns for low and high turnover households, thus focusing jointly on the extensive and intensive margins. Turning to the model's implications for the relationship between return and turnover, we see the same basic patterns. We split our data into quintiles based upon the absolute value of the turnover rate. Table 3 reports the returns for the lowest and highest turnover groups.⁹ The gross return of the high turnover

⁹The lowest group has zero turnover as does a fraction of the second lowest group. As the return for non-adjusters is not identical, the return for the lowest turnover group is not the same as the return for the non-adjusters.

group exceeds that of the low turnover group. Further, the adjustment costs are not large enough to overturn that ordering for net returns.

HH Returns The lower section of Table 3 shows the measured returns from the HH simulations. Recall that these simulations are based on a return process calculated for individual households from the Barber and Odean (2000) study.

As seen in Table 3, the results follow the same pattern as the baseline. The returns (gross and net) are higher for adjusters and for the highest turnover group, compared to the low turnover group. Despite using the return process in the data from the Barber and Odean (2000) study, the pattern of gross and net returns that they uncovered does not appear in the simulated data based upon optimizing behavior. Once again, the key is the selection effect since trade is slightly more frequent in high return states.

There is one important difference between our model and the underlying data studied by Barber and Odean (2000). Our model is specified and estimated at an annual frequency. In contrast, Barber and Odean (2000) have monthly data and their return differentials are geometric annualized monthly returns. Thus relatively small differences in monthly returns can become large differences in annual returns. If an agent trades frequently enough to incur a transactions cost each month which reduces the net return in that month by one-half of a percentage point, then the cost is over 5 percentage points on an annual basis. So while it might be that the transactions costs could be large enough at a higher frequency to explain the lower net return of adjusters, this point does not explain the differences in gross returns which reflects the selection effects.¹⁰

Alternative Parameters We consider a couple of extensions of our model to study the impact of variations in risk aversion and trading costs on returns. If agent's are more risk averse, then they are more likely to trade for the purposes of consumption smoothing. To study that, we increase γ from 3.51 to 10.0. We find households trade more often but the differential in the gross rate of return remains: the gross return for the adjusters is nearly 4 percentage points higher than the gross return for non-adjusters. But, as before, the difference between gross and net returns for traders is negligible.

Another interesting possibility is that agents actually prefer to trade so that $\psi > 1$. Though the estimate of ψ is less than one, it is interesting to see what return patterns are produced by this model. At $\psi = 1.01$, all households adjust their portfolios. This reflects the fact that the financial costs of trading are low enough that they do not by themselves

¹⁰In some simulations at the monthly frequency, the returns of adjusters exceed those of the non-adjusters, as in the results reported in Table 3. But due to the discount factor near unity and the lack of high frequency moments, estimation of the model is not possible.

create inaction: the gross and net returns are about the same for the adjusters.

However, there are interesting patterns of returns across the turnover rate. The net and gross return is a concave function of the turnover rate, reaching 9% for the middle quintile and falling to nearly 6.2% for the highest turnover group. When ψ is bigger than one, some of the trades are undertaken by high income households simply for “the joy of trading” even though returns are not high enough to warrant these trades. Still these are differences in both net and gross returns, not in net returns alone.

Finally, we looked at adjustment costs in excess of those estimated in Table 1. For this experiment, the fixed cost of trading was increased to \$1000. The gross return to adjusters exceeded that of the non-adjusting households while the net returns were about the same.

In another experiment, the linear term was increased to 5%. This additional cost also did not influence the basic findings of our model: gross returns and net returns are higher for adjusters (high turnover households) than non-adjusters (low turnover households).

Differences in Expected Utility As a final point, note that for all of these treatments this model has completely rational households. All trades are consistent with the maximization of discounted expected utilities. Investors who trade do so precisely because the expected utility from trading **exceeds** that of not trading.

In contrast, Barber and Odean (2000) say:

The two models yield different predictions about the gains of trading. The rational expectations model predicts that investors who trade more (i.e., those whose expected trading is greater) will have the same expected utility as those who trade less. The overconfidence model predicts that investors who trade more will have lower expected utility.

This theme that traders who trade more have the same expected utility as those who trade less is not a property of the model. The choice of whether to trade comes from the optimal choice of the agents in (1). There is no presumption of indifference. Further, Barber and Odean (2000) contend that in a model with rational investors (they refer to the model of Grossman and Stiglitz (1980)) more active traders will have higher gross returns but no difference in net returns. The gross return differential is a property of this model but net returns also differ across traders even though they are rational. A key aspect of our model is the heterogeneity across agents in income and wealth. In an optimizing model, these differences induce the utility differentials between adjusters and non-adjusters as well as the selection effects which underly the reported relationships between adjustment, turnover and returns.

4 Evaluating Models of Overconfidence

Given that the presence of transactions and opportunity costs are not enough to create the pattern of gross and net returns found by Barber and Odean (2000), we turn to their favored explanation: overconfidence. We model overconfidence in three ways: (i) a higher than actual return, (ii) a lower than actual standard deviation of the return and (iii) more persistence in the return process.

4.1 Models of Overconfidence

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

$$R_t = \tilde{\mu} + \tilde{\rho}R_{t-1} + \varepsilon_t \quad (8)$$

where ε is normally distributed with a mean of 0 and a standard deviation of $\tilde{\sigma}$.¹¹ The mean return is denoted $\tilde{\mu}$ and the serial correlation is $\tilde{\rho}$. This process 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.

4.2 Approach

From (8), our specification permits three types of deviations through the: (i) mean, (ii) standard deviation and (iii) persistence of the return process. For each of these specifications, we parameterize the deviation from the true process and solve the model assuming household's hold these beliefs. The realized net and gross returns are calculated from simulated data where the actual process is used for the exogenous returns process.

The first deviation from truth comes from excessive optimism about the mean of the process, so that $\tilde{\mu}$ is higher than the actual mean of the return process. The second deviation allows the household to believe that the standard deviation of the process is smaller than truth: i.e. $\tilde{\sigma}$ is less than the true standard deviation of the innovation. In this case, the household perceives less uncertainty than reality.

The final case allows a deviation between the perceived and actual serial correlation of the return process, $\tilde{\rho}$ may differ from the actual serial correlation parameter.¹² When returns are

¹¹Here 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.

¹²See Daniel, Hirshleifer, and Subrahmanyam (2001) for a discussion of models of overconfidence in which agents overestimate the informativeness of signals.

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. For our estimated return process, there is no evidence of serial correlation over the sample period. Thus our return process is iid and the current return provides no information about the future. If, however, $\tilde{\rho}$ is positive, then households are more confident that current returns provide information about future returns.

4.3 Results

Return	Baseline	Mean		Std.		Ser. Corr.	
		$\mu * 1.01$	$\mu * 1.05$	$Std * 0.9$	$Std. * 0.8$	$\rho = 0.1$	$\rho = 0.3$
Non-Adjusters							
Gross	1.0494	1.0495	1.0495	1.0488	1.0484	1.0697	1.0750
Adjusters							
Gross	1.06	1.06	1.0579	1.06	1.0602	1.0515	1.0502
Net FC	1.0583	1.0585	1.0574	1.0584	1.0584	1.0500	1.0465
Net All	1.025	1.0312	1.0477	1.0264	1.022	1.0196	0.9870
Highest Turnover							
Gross	1.0652	1.0608	1.0975	1.0666	1.0667	1.0670	1.0690
Net FC	1.0608	1.0569	1.0959	1.0625	1.0616	1.0629	1.0573
Net All	0.9768	0.9858	1.0628	0.9829	0.9669	0.9855	0.8749
Lowest Turnover							
Gross	1.0501	1.0505	1.0137	1.0536	1.0677	1.0685	1.0712

Table 4: Models of Overconfidence

Our findings for the various experiments are summarized in Table 4. To be clear, these calculations are all at the baseline parameters, including the stochastic processes for income and interest rates, allowing for only one form of overconfidence at a time.¹³ These deviations from truth are used to index the columns of the table. The first set of rows report the gross returns for adjusters and non-adjusters, which indicate whether portfolio adjustment occurred in a given period or not. The second set of rows reports returns for the highest and lowest turnover groups. As the lowest turnover group did not trade, only gross returns are reported.

¹³Hence in the simulations the distribution from which the returns are drawn differs from the beliefs of agents.

There are two types of net returns reported, corresponding to the objects calculated in Table 3. The first, “Net FC”, calculates the net return accounting for direct financial costs of the actual trade, using the estimates in Table 1. The second “Net All” uses both the direct financial costs as well as the opportunity cost, parameterized by ψ , to calculate trading costs.

We first discuss results based on the extensive margin of adjust or not. We then discuss the relationship between return and turnover rates.

The results from the baseline appear in the second column of the table. Recall that there is a gross rate of return differential in favor of the adjusters. This reflects the selection effect. The net return, though of course lower than the gross return, is still higher than the return of the non-adjusters.

The first form of misperception is in the mean of the return process, $\tilde{\mu} > \mu$, shown under the columns labeled “Mean”. If there is overconfidence about the mean of the return, then the adjustment rate increases. When the mean is viewed as higher than truth by a factor of 1.01, there is relatively little change in the returns though the adjustment rate is higher. At a mean return increased by a factor of 1.05, the same pattern of net and gross returns remained.

The second form of misperception is in the standard deviation of the innovation, $\tilde{\sigma} < \sigma$, shown under the columns labeled “Std.” If the perceived standard deviation is lower than truth, the gap in gross returns widens with the adjusters having the higher gross return. As in the other specifications, the costs of trade reduce the net return, but it still exceeds the return of the non-adjusters.

The final case of misperception in the serial correlation, $\tilde{\rho} > \rho = 0$, shown under the columns labeled “Ser. Corr” is quite different.¹⁴ Here the agents are more confident that the returns will persist and are responsive to this signal. In this case, the gross, and hence net return, to the adjusters is below the gross return of the non-adjusters. This is much closer to the pattern highlighted by Barber and Odean (2000). But, as before, this pattern appears in the differential on gross returns.

Though these models of overconfidence do not match the findings of Barber and Odean (2000), they do have other interesting implications. The column in Table 3 labeled “SC” presents the moments when households believe $\rho = 0.3$. In this case, the overall fit of the model, though not as good as the baseline of 0.0006, is 0.001 using the baseline parameters.¹⁵ The goodness of this fit comes from the reduction in α_1 from 0.11 to 0.0969. With more perceived persistence in the interest rate, the amount of saving is more responsive and thus consumption less responsive to interest rate movements.

¹⁴Daniel, Hirshleifer, and Subrahmanyam (2001) also discuss the implications of this last form of overconfidence for asset trades. They focus on the intensive margin.

¹⁵The fit can be improved further by additional estimation.

Instead of looking at the margin of adjust or not, it is informative to look at the patterns of returns for different turnover rates. The results for the highest and lowest turnover groups mimic the results based on adjust/no adjust. That is, except for the of misperceptions about the serial correlation, the gross returns are higher for the highest turnover group of agents. But this ordering is switched in the “Ser. Corr. =0.3” case where the returns for the lowest turnover group are slightly higher than the highest turnover group. In this case, however, it is important to note that the gross returns are lowest for the middle turnover group, at 1.0356, so that there is not a monotone relationship between returns and turnover rates in this case. In contrast, the relationship is monotone in the results of Barber and Odean (2000).

It is useful to see how overconfidence impacts on the likelihood of trade. To do so, compare the choices of whether to adjust or not for two types of households. The first, rational trader, has preferences and beliefs based on the baseline model. The second, overconfident trader, has the same preference but with beliefs that $\tilde{\rho} = 0.3$. We focus on the sensitivity of trade to variations in the gross return on the asset.

The results are summarized in Table 5. For the baseline model of a rational agent, the probability of making a trade is about 71% in the high return state but only 66% in the low return state. Thus, as noted in the discussion of Table 3, the gross return of adjusters is lower than the gross return of the non-adjusters.

But, in this model of overconfidence, the results flip. As is evident from Table 5, the adjustment rate is higher in the low return states for the overconfident agents. The perception of a persistent low returns leads them to adjust their portfolio.

specification	High Return	Low Return
baseline	0.71	0.66
OC	0.66	0.74

Table 5: Adjustment Rate Dependence on Gross Return

The results in Table 5 complement the findings on gross returns reported in Table 4. When the adjustment rate is higher in high return states, as in the baseline, then gross returns are higher for adjusters compared to non-adjusters. But, when the adjustment rates are lower in high return states, as in the case with a belief in a serial correlation of 0.3, then the realized returns of the adjusters are lower than those of the non-adjusters.

5 Conclusion

The goal of this paper was to assess the claim in Barber and Odean (2000) that the patterns of returns as a function of the frequency of trade was consistent with overconfident agents and inconsistent with rational traders. In our model, the frequency of trade translates into whether households incur a cost to adjust their portfolio or not.

Using parameter estimates which match moments of adjustment rate, the sensitivity of consumption growth to interest rate movements and the volume of trade, we found that the model of rational agents produced differential in both gross and net returns in which adjusters earned more than non-adjusters. Since agents are utility maximizing, trading choices are optimal *ex ante* though in any dynamic stochastic model, there can be *ex post* regret.

Models of overconfidence in either the mean, the standard deviation or the persistence of shocks, did not match the observations of Barber and Odean (2000) either. In particular, these models also created differences in net returns from differences in gross returns. This reflects the power of the selection effects which dominate trading costs in determining the pattern of net and gross returns.

We have focused on only a few of potentially many models of overconfidence. We are confident only that the models we have looked at are at odds with the data.

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.
- BONAPARTE, Y., AND R. COOPER (2009): “Costly Portfolio Adjustment,” NBER Working Paper # 15227, revised 2010.
- DANIEL, K., D. HIRSHLEIFER, AND A. SUBRAHMANYAM (2001): “Overconfidence, arbitrage, and equilibrium asset pricing,” *Journal of Finance*, pp. 921–965.
- 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.
- GROSSMAN, S., AND J. STIGLITZ (1980): “On the impossibility of informationally efficient markets,” *The American Economic Review*, 70(3), 393–408.

GUIISO, L., AND T. JAPPELLI (2006): “Information Acquisition and Portfolio Performance,” mimeo.

HANSEN, L., AND K. SINGLETON (1983): “Stochastic Consumption, Risk Aversion and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, 91(2), 249–265.