

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.

© 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
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, Italy 50133
and University of Texas
and also NBER
russellcoop@gmail.com

Rationalizing Trading Frequency and Returns*

Yosef Bonaparte[†] and Russell Cooper[‡]

May 15, 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, bonaparte20@gmail.com

[‡]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). First, drawing 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. In particular, 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. The 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.

Our approach is to specify a dynamic optimization problem of a household. 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 follow the approach in Bonaparte and Cooper (2009) to estimate the financial costs and time cost of portfolio adjustment, along with the household utility function.¹ We generate simulated data from the parameterized 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 and the 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

¹As discussed below, we use a slightly lower discount factor, relax the borrowing constraint and estimate using shareholders only to better match the set of traders captured in the Barber and Odean (2000) study.

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 in 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 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 of Bonaparte and Cooper (2009). 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.

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.² Total financial wealth this period is $R_{-1}s_{-1}$.

The household chooses between the options of adjusting or not:

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

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

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

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

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

²For the purpose of this exercise, the household has a single financial asset.

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, $s = s_{-1}R_{-1}$.³

2.2 Quantitative Analysis

Bonaparte and Cooper (2009) estimate trading costs, $C(\cdot)$, directly from the data set used by Barber and Odean (2000) and then use a simulated method of moments approach to estimate other parameters. As in Bonaparte and Cooper (2009) assume $u(c) = \frac{1}{1-\gamma} c^{1-\gamma}$. Thus the key parameters to estimate are the curvature of utility and the time cost of adjustment, (γ, ψ) .

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.⁴ The trading costs 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 current

³Bonaparte and Cooper (2009) discuss and estimate alternatives to this model of no adjustment. This specification fits the data best.

⁴Details 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

trading costs, it is still small compared to the average trade of a household in the data set of about \$12,000.

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

Income and returns are modeled as AR(1) processes. Household income variable is the product of a common shock and a household specific shock. As discussed in Bonaparte and Cooper (2009), the serial correlation of the common component was estimated to be 0.58 and the standard deviation of the innovation to aggregate income was estimated to be 0.028. For the idiosyncratic component, the serial correlation was 0.90 and the standard deviation of the innovation was 0.238.

The return process measures real returns, including dividends, and comes from Robert Shiller, available at <http://www.econ.yale.edu/~shiller/data.htm>. The mean return is 5.49% over the 1967-1992 period with a serial correlation of nearly zero and a standard deviation of 0.1633.⁵

2.2.3 Moments

Bonaparte and Cooper (2009) use two moments to identify the parameters (γ, ψ) . The first

⁵This period was chosen to best overlap with the data on the income process and the data used to estimate trading costs. Since the standard error on the serial correlation of the return process was so high, we set this to zero in the estimation.

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. A similar result is reported by Vissing-Jorgensen (2002).⁶ The response of consumption growth to interest rate movements is larger here than in other studies due to the focus on a sample of shareholders.

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 (γ, ψ) using two moments: the frequency of inaction and the estimate of α_1 from (7). For that exercise, $\beta = 0.96$ was imposed as well as a restriction of $s \geq 0$.

The estimation approach and results are modified slightly for this study of gross and net returns. Since a key issue in calculating average adjustment costs and thus net returns is the level of trade, it was important that the model do a good job of matching the average trade in the data, of approximately \$12,000. To match this additional feature of the data, we set $\beta = 0.94$ and place a negative lower bound on s . Further, the estimation of α_1 in (7) comes from a sample of stock market participants only. Despite these differences, the estimates of (γ, ψ) are similar to those reported in Bonaparte and Cooper (2009).

The moments from the data and the parameter estimates are summarized in Table 2.⁷ The moments here are closer to the actual data moments than the estimates from Bonaparte and Cooper (2009). The average trade in the simulated data was about \$18,000, still large relative to the data.

⁶Households are classified as stock market participants if they report ownership of stocks. As in Vissing-Jorgensen (2002), the data is from the CEX. We convert the data there to create annual rather than semi-annual measures of consumption growth.

⁷The 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.

Moments	Data	Model
Portfolio Adjustment rate	0.71	0.745
α_1	0.0878	0.338
Estimated Parameters		
γ		1.052
ψ		0.991
fit		0.064

Table 2: Moments and Parameters

The data moments are summarized in the second column. The estimates are presented in the third column. Interestingly, there is evidence of adjustment costs beyond the trading costs. Further, utility is quite close to the log case.⁸ The fit of 0.064 is calculated as the sum of the squared differences between the simulated and data moments.

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.

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.

⁸Bonaparte and Cooper (2009) study the robustness of these estimates to different measures of return and other specifications of the model.

The third measure nets out both the financial cost and the income loss due to ψ .⁹

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 and fifth columns compute return for the lowest and highest quintiles of the turnover rate distribution.¹⁰

Average Return	Adjusters	Non-adjusters	Lowest Turnover	Highest Turnover
Gross	1.067	1.031	1.028	1.108
Net FC	1.064	1.031	1.028	1.106
Net All	1.060	1.031	1.028	1.104

Table 3: Returns and Trading

The table reveals that neither the transactions cost explanation of the net return differential between adjusters and non-adjusters nor the selection story fits the facts. For both cases, the key observation is that the gross return for adjusters, 1.067, exceeds that of the non-adjusters of 1.031. This is inconsistent with the findings of Barber and Odean (2000).

First, consider the transactions costs explanation. 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 \$12,000 and is about \$18,000 in the simulated data. Thus the fixed cost is much too small in both the actual and simulated data to explain the differential in returns.

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

⁹To 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 $t - 1$ total portfolio.

¹⁰The turnover rate is the absolute value of the net change in the portfolio divided by its initial value.

can become large differences in annual returns. If an agent incurs 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.¹¹

Second, 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 is creating a higher, not lower, return for the adjusters. This is because agents are adjusting 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.¹² The gross return of the high turnover group exceeds that of the low turnover group. Further, the adjustment costs are not large enough to overturn that ordering for net returns.

As a final point, note that 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

¹¹Exploring this would require the estimation and simulation of our model at a much higher frequency which is difficult due to β near 1.

¹²The 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.

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.

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 1.052 to 5.0. We find households trade more often but the differential in the gross rate of return remains: the gross return for the adjusters is nearly 6 percentage points higher than the gross return for non-adjusters. But, as before, the difference between gross and net returns 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 $\psi = 1.01$, the gross return for adjusters is almost 6 percentage points above that of the non-adjusters.

Finally, we looked at adjustment costs in excess of those estimated in Table 1. For one experiment, the fixed cost of trading was increased to \$1000. In another, the linear term was increased to 5%. While both experiments influenced the moments, they 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).

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:

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

where ε is normally distributed with a mean of 0 and a standard deviation of $\tilde{\sigma}$.¹³ 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.¹⁴ When returns are 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

Our findings for the various experiments are summarized in Table 4. To be clear, these calculations are all at the baseline parameters allowing for only one form of overconfidence at a time. 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

¹³Here 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.

¹⁴See Daniel, Hirshleifer, and Subrahmanyam (2001) for a discussion of models of overconfidence in which agents overestimate the informativeness of signals.

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	1.031	1.028	NaN	1.029	1.03	1.032	1.143
Adjusters							
Gross	1.067	1.063	1.057	1.068	1.069	1.067	1.023
Net FC	1.064	1.061	1.056	1.065	1.066	1.064	1.020
Net All	1.060	1.060	1.056	1.061	1.062	1.060	1.017
Highest Turnover							
Gross	1.108	1.126	1.195	1.148	1.156	1.109	1.106
Net FC	1.106	1.124	1.194	1.146	1.154	1.107	1.103
Net All	1.104	1.123	1.194	1.144	1.152	1.105	1.102
Lowest Turnover							
Gross	1.028	1.03	0.93	1.039	1.03	1.15	1.177

Table 4: Models of Overconfidence

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.

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, all households adjust.¹⁵

¹⁵Recall that we did not re-estimate the model parameters for each specification of overconfidence.

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 high different. 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.

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” case where the returns for the lowest turnover group are considerably 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 0.9853, so that there is no monotone relationship between returns and turnover rates.

Daniel, Hirshleifer, and Subrahmanyam (2001) also discuss the implications of this last form of overconfidence for asset trades. They focus on the intensive margin. 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, over-confident 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 74% in the high return state but only 61% 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.

Though these models of overconfidence do not match the findings of Barber and Odean (2000), they do have other interesting implications. For example, when $\tilde{\rho} = 0.3$, the overall fit of the model is much improved: it falls to 0.010 compared to the baseline of 0.064. The improvement of the fit comes from the reduction in the α_1 to -0.0039 from 0.338. With more perceived persistence in the interest rate, the amount of saving is more responsive and thus

specification	High Return	Low Return
baseline	0.74	0.61
OC	0.44	0.92

Table 5: Adjustment Rate Dependence on Gross Return

consumption less responsive to interest rate movements.

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.

There are a number of avenues for future work. Throughout this exercise, we assume that agents have a common return. Based upon Bonaparte (2008) this is not the case and perhaps this added dimension of heterogeneity will bring about additional selection effects.

In addition, 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. (2008): “Estimating the Elasticity of Intertemporal Substitution with Household-Specific Portfolios,” mimeo.

-
- BONAPARTE, Y., AND R. COOPER (2009): “Costly Portfolio Adjustment,” NBER Working Paper # 15227.
- 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.
- VISSING-JORGENSEN, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110(4), 825–853.