Costly Portfolio Adjustment*

Yosef Bonaparte[†] and Russell Cooper[‡]

August 1, 2009

Abstract

This paper studies the dynamic optimization problem of a household when portfolio adjustment is costly. The analysis is motivated by the observation that on a monthly basis, less than 10% of stockholders typically adjust their portfolio of common stocks. We use this, and related observations, to estimate the parameters of household preferences and portfolio adjustment costs. We find significant adjustment costs, beyond the direct costs of buying and selling assets. These adjustment costs imply that inferences drawn about household risk aversion and the elasticity of intertemporal substitution are biased: household risk aversion is lower compared to other estimates and it is not equal to the inverse of the elasticity of intertemporal substitution.

1 Motivation

Since the contribution of Hansen and Singleton (1982), most studies of household intertemporal consumption choice base inferences upon an Euler equation of the form

$$u'(c_t^i) = \beta E_t[R_{t+1}^i u'(c_{t+1}^i)] \tag{1}$$

*We are grateful to the National Science Foundation for supporting this research. We thank seminar participants at the Federal Reserve Bank of Dallas, the Federal Reserve Bank of Kansas City, the CCER at Peking University, Melbourne University and at the Pennsylvania State University for comments. Guozhong Zhu provided outstanding research assistance.

[†]The University of Southern Mississippi, College of Business, bonaparte20@gmail.com

[‡]Department of Economics, University of Texas at Austin and the European University Institute, cooper@eco.utexas.edu

where $i = 1, 2, ...I$ is an index for the household, c_t^i is period t consumption, R_{t+1}^i is the return in period $t+1$ for household i, where $t = 1, 2, ...$ is an index of time.

Estimation of preference parameters proceeds along two lines. One approach, following Hansen and Singleton (1982), imposes (1) on the representative household, using observations of consumption and asset returns on a monthly basis. Other studies have allowed some heterogeneity in the consumption paths across households and more recently in asset returns as well.[1] These studies focus on estimation of the curvature of the household's utility function.

A second approach, following Hansen and Singleton (1983) and used in Hall (1988), rests upon an approximation of (1). With this approach, the researcher estimates the relationship between consumption growth and the (log of the) real interest rate. The coefficient on the interest rate is often interpreted as the elasticity of intertemporal substitution (EIS). With some additional structure, the inverse of the EIS is the degree of risk aversion.[2]

Hall (1988) summarizes his findings as " ... supporting the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero." This result of a near-zero EIS is consistent with many other studies of this important link between consumption growth and interest rates. Under some specifications of the utility function this low estimate of the EIS implies a high level of household risk aversion.

This paper challenges the interpretation of these findings. Equation (1) is a necessary condition for optimality **if** the household reoptimizes by adjusting its portfolio each period. Yet, on a monthly basis, only 8.6% of households owning stocks adjust their portfolio of common stocks. The fraction rises to 71.0% on an annual basis.[3]

The inaction in portfolio adjustment has two implications. First, it is not appropriate to use (1) as a basis for estimation at the individual level since this condition assumes continuous adjustment. Second, the inaction at the household level will not necessarily disappear when looking at aggregated data. The inaction in household adjustment may well

[1] See Bonaparte (2008) and the discussion therein.

[2] As noted in Hall (1988), papers in the literature differs in terms of describing the estimation in terms of the EIS or household risk aversion.

[3] This annual measure comes from the response to question X3928 in the Survey of Consumer Finances. The monthly inaction comes from question X7193. Here action means the purchase or sale of stocks or other securities through a broker. This adjustment rate is for an average of the 1995, 1998, 2001 and 2005 cross-sections.

effect the estimation based upon aggregated data.

The nature and extent of the bias introduced on estimates of the intertemporal elasticity of substitution is not clear *a priori*. Though the inaction seems to naturally pull the estimated EIS lower, when adjustment occurs it can be more sensitive to interest rate movements. We re-estimate household risk aversion using a model which explicitly recognizes costs of portfolio adjustment. We find that degree of risk aversion is considerably lower than one would infer from previous studies.

These findings come from the analysis of dynamic optimization of a household with costly portfolio adjustment. These costs come from the buying and selling of assets as well as the time cost of portfolio adjustment.

The model is related to but, in important ways, is distinct from the literature on asset market participation. The literature, stemming from the insights of Mankiw and Zeldes (1991), recognizes that only a subset of households are directly involved in asset markets. Thus conditions such as (1) hold for only the households actively involved in assets markets. As argued in Vissing-Jorgensen (2002), the neglecting the asset market participation decision may bias estimates of the elasticity of intertemporal substitution downward. For some specifications of utility, this leads to an upward bias in the estimated degree of risk aversion.

Our observations go beyond the participation decision to recognize that even the subset of households directly owning assets and hence involved in asset markets do not adjust their portfolio's each month let alone each year. This is a distinct reason why (1) may not hold. Moreover, ignoring the underlying adjustment costs which rationalize the inaction in portfolio composition may also lead to a miscalculation in the value of asset market participation itself.[4]

Luttmer (1999) also studies the implications of trading costs in the intertemporal optimization problem of a household. His result is a lower bound on these costs for given values of the curvature of the period utility function and a parameter for an external habit. Our approach and results are different. We obtain some trading costs directly from the data and provide estimates of both a preference parameter and a direct opportunity cost of trading.[5]

Given inaction in portfolio adjustment, the model is estimated using a simulated method

[4] The literature on asset market participation, such as Heaton and Lucas (1996), looks at the role of transactions costs for the participation decision, without focusing on portfolio adjustment.

[5] Brunnermeier and Nagel (2008) find that the household asset market participation decision is sensitive to wealth levels.

of moments procedure, rather than the standard approach of estimating parameters from *ex post* errors associated with (1).[6] To implement this approach, we take some key moments from the data and find parameter values for the household's problem which reproduce these moments in simulated data. From this exercise, we obtain estimates of household utility functions as well as the time cost of trading. Additional trading costs are estimated directly from the data.

Following Hansen and Singleton (1983), Hall (1988) and numerous other researchers, we too regress consumption growth on the interest rate. But in our model, the coefficient on the interest rate is not an estimate of the household EIS. Instead, we use that coefficient and others as moments to infer the household degree of risk aversion.

2 Household Behavior

We model the household through the specification of a dynamic discrete choice problem. The household can either choose to adjust its portfolio or not. If it chooses to adjust, then it incurs transactions costs but is able to change the composition of its portfolio.[7] If the household chooses not to adjustment, then its consumption is equal to labor income. All capital income, such as unrealized capital gains and dividends, are assumed to be reinvested without cost.

The specification of this discrete choice model is consistent with the measure of inaction in the data. That is, the statistics reported above about the fraction of inaction in household portfolio adjustment were measures of household trades. Inaction meant there were no purchases and no sales of assets by a household. Nonetheless portfolio values do change along with asset prices even when there is no direct adjustment by the household. Income variations, though, are met by variations in consumption, not in the holdings of assets.

[6]The approach of Cooper, Haltiwanger, and Willis (2008) which estimates parameters from Euler equations at the level of an individual in the presence of non-convex adjustment costs could be used here as well.

[7]These costs take two forms. Brunnermeier and Nagel (2008) conclude that portfolio costs must be supplemented by " ... a cost of giving any attention at all to the portfolio..." and this additional cost is captured in our model.

2.1 Dynamic Optimization Problem

Formally, $v(y, s_{-1}, R_{-1})$ is the value of the household's problem when it has current income y, has an existing portfolio of s_{-1} and the return vector from the previous period was R_{-1}. Here $s_{-1} = (s^1_{-1}, s^2_{-1}, ...s^I_{-1})$ is the vector of asset holdings which incorporate capital gains and purchases from the past. The return vector R_{-1} provides information about returns over the next period. Total financial wealth equals $\sum_i R^i_{-1} s^i_{-1}$.

The value of the household's problem is given by the maximum over the options of adjusting or not:

$$v(y, s_{-1}, R_{-1}) = \max\{v^a(y, s_{-1}, R_{-1}), v^n(y, s_{-1}, R_{-1})\} \tag{2}$$

for all (y, s_{-1}, R_{-1}). If the household chooses to adjust its portfolio, then the value of the problem is:

$$v^a(y, s_{-1}R_{-1}) = \max_s u(c) + \beta E_{R,y'|R_{-1},y} v(y', s, R). \tag{3}$$

The consumption level is given by

$$c = \sum_i R^i_{-1} s^i_{-1} + y \times \psi - \sum_i s^i - C(s_{-1}, s). \tag{4}$$

Here s^i is the purchase of shares of asset i. In this expression, the function $C(\cdot)$ captures the cost of portfolio adjustment. It is a key part of the model in order to match observations of inaction in the adjustment of portfolios by households.

The other cost of adjustment is parameterized by ψ. It effects consumption and thus current utility through a reduction in labor income, assuming $\psi < 1$. We call this a time cost of trade to distinguish it from the trading costs paid to a broker. One interpretation of this cost is lost income due to time spent on portfolio adjustment. In this case, this cost is distributed across the population with household labor income.

If there is no adjustment, the value of the problem is:

$$v^n(y, s_{-1}, R_{-1}) = u(y) + \beta E_{R,y'|R_{-1},y} v(y', s, R). \tag{5}$$

When there is no adjustment in the portfolio, the household consumes only its labor income and the cost of adjustment is absent. The proceeds from its existing portfolio are assumed to be costlessly reinvested.[8] Hence

$$s^i = R^i_{-1} s^i_{-1} \tag{6}$$

for $i = 1, 2, ...I$. This is in the argument for the shares next period in (5).

3 Quantitative Analysis

As noted earlier, (1) will generally not hold for this optimization problem since adjustment is intermittent. Thus we do not try to estimate parameters directly from an Euler equation or its log-linear approximation.

Our approach is to calculate pertinent moments from data and to estimate parameters of the utility function from a simulated method of moments approach. The trading costs of portfolio adjustment are estimated directly from observations outside of the dynamic programming model. This allows us to identify the utility costs of trade separately from actual trading costs.

For the quantitative analysis, we set $I = 1$ so that there is a single asset. In this case, the decisions of portfolio adjustment and intertemporal consumption are the same. Thus we interpret the asset as a composition of multiple assets with a single cost of adjustment. We do this for computational reasons since the state space of the multi-asset problem can be too large.

3.1 Data

We use data from a couple of sources for our study. First, as explained below, we use the monthly household account data set, as in Barber and Odean (2000), to estimate trading costs. So, in contrast to other studies, our trading costs are estimated directly from the data. Second, we use data from the Survey of Consumer Finances to measure inaction in portfolio adjustment. This is a broader measure of inaction than the one coming from the study of Barber and Odean (2000). Third, we use data on aggregate real consumption expenditures from Robert Shiller available at `http://www.econ.yale.edu/~shiller/data.htm`. Fourth,

[8] This means that any capital gains are converted into new shares since the price of a share is kept fixed at unity. An alternative formulation would have the actual shares remain constant and allow consumption to absorb the return on the existing portfolio.

we use data from the PSID to estimate a household income process.[9] The precise use of these data in identifying paramters is explained in detail below.

3.2 Parameterization

Here we present our specification of functional forms. Some of the parameters are estimated outside of the dynamic programming problem and those estimates are provided here as well.

Utility Function We assume a utility function of the CRRA form: $u(c) = \frac{1}{1-\gamma}c^{1-\gamma}$. For this specification, our estimate of risk aversion is the inverse of the elasticity of intertemporal subsitution.

We could also add a direct utility cost of trading.[10] We have instead chosen to model the cost of adjustment as a time cost, through the inclusion of ψ in (4).

We estimate the preference parameters γ, along with the trading cost ψ using simulated method of moments, described below. As our data is annual and we set $\beta = 0.96$ in our baseline estimation.

Cost of Adjustment A key part of the model is the parameterization of the adjustment cost function. We assume that the cost function, $C(s_{-1}, s)$ is separable across assets so that $C(s_{-1}, s) = \sum_i C^j(s^i_{-1}, s^i)$ where $j = b$ if the asset is bought and $j = s$ if it is sold. We assume:

$$C^b(s^i_{-1}, s^i)) = \nu^b_0 + \nu^b_1(s^i - s^i_{-1}) + \nu^b_2(s^i - s^i_{-1})^2 \tag{7}$$

if the household buys asset i so that $s^i > s^i_{-1}$. If instead the household sells asset i, then

$$C^s(s^i_{-1}, s^i)) = \nu^s_0 + \nu^s_1(s^i_{-1} - s^i) + \nu^s_2(s^i - s^i_{-1})^2. \tag{8}$$

We use the monthly household account data to estimate these transactions costs for the buying and selling of common stocks. These estimates come from the data set used by

[9]Our sample period from the PSID is 1967-1993. Households labor incomes are deflated using the CPI obtained from BLS. The sample selection includes the following: (i) male,(ii) between 20 and 64 years old, (iii) not from the SEO sample, (iv) real hourly labor earnings between \$2 and \$400, (v) work between 520 hours (10 hours per week) and less than 5110 hours (14 hours a day, everyday) and (vi) The households have more than 19 years of observations (out of 25 years of the sample).

[10]As in much of the discrete choice literature, such as Rust (1987), we could make this additional additive and stochastic so that each household faces a choice specific shock.

Barber and Odean (2000). It contains information on common stock trades of about 78,000 households through a discount brokerage firm from January 1991 to December 1996.

In our sample, we have over three million observations where in each observation a stockholder (trader) reports: trade date, buy or sell, quantity of shares transacted, commission (in dollars value), cusip identifier and the price.[11] We estimate these trading costs from 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. If a stockholder bought different stocks in a given month, the stockholder reports the commission, quantity and price on each one of these stocks separately.

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

Table 1: Estimated Trading Costs

To get a sense of magnitudes, the average purchase in our sample has a value of about \$11,000. The cost of this trade is about \$70.00. For trades of this size, the quadratic term is negligible.

The estimates, along with standard errors in parentheses, are given in Table 1. Note there is a fixed cost of trading of around \$60.00.[12] Because of this fixed cost, even if $\psi \geq 1$, the model will still produce inaction.

To be clear, the estimates of these transactions costs is one piece of the adjustment cost function. The other component is the cost of trade, the ψ in (4).

[11] This includes trades which entail a negative commission using our estimates. For these calculations, the $(s^i - s^i_{-1})$ terms in $C^s(s^i_{-1}, s^i))$ are values not quantities.

[12] For comparison, Heaton and Lucas (1996) only allow a quadratic cost of adjustment.

Income and Returns The dynamic optimization problem of the household includes conditional expectations over its income and the return on assets. The income variable itself is the product of a common shocks and a household specific shocks. These income processes along with the return process are represented as AR(1) processes.

The income process came from the PSID. We decomposed income into a common and a household specific shock. 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 idiosyncratic shock is both more persistent and more volatile. The average level of real income is set at $72,000 annually to mimic the income levels in the household account data.

The return process came from an AR(1) representation of the data on real returns, including dividends, from Robert Shiller, available at `http://www.econ.yale.edu/~shiller/data.htm`. The inclusion of dividends in this return measure is consistent with our model of inaction where dividends are costlessly reinvested. For that process, the mean return is 7.51% over the 1967-1992 period. The serial correlation of the return is 0.264 (though it is not very precisely estimated) and the standard deviation is 0.1531.[13]

Market Participation As explained below, we will be using estimates of the EIS from aggregated data as a moment in our estimation. Since the aggregated data include both stock market participants as well as non-participants, we add to our model a group of households who are not active in financial markets. The consumption of these households is set equal to their income.[14] Following the calculations in Vissing-Jorgensen (2002), we assume that 50% of the households participate in financial markets.[15]

[13] The moments of the return process are quite sensitive to the choice of sample period. We chose this period to be close to the sample period for the other processes and moments.

[14] The model could in principle be extended to allow these households some smoothing outside of direct stock ownership. The addition of an additional asset though requires an additional state variable and is thus computationally quite demanding.

[15] This is consistent with the more recent findings of Bucks, Kennickell, Mach, and Moore (2009) that about 50% of households own stock directly or indirectly.

3.3 Simulated Method of Moments Estimation: Results

Given these estimates of transactions costs and our calibration of $\beta = 0.96$, we use a simulated method of moments technique to estimate the remaining parameters, (γ, ψ).[16] There are two moments used to identify the parameters. We talk about other implications of our just identified model below.

The model is estimated at the annual frequency. The first moment comes from the SCF data set where in an average year 71% of households adjustment their portfolio. This moment is very informative about the cost of trade, ψ.

The second moment comes from the sensitivity of consumption growth to returns estimated from the log-linear approximation of a consumption Euler equation, drawing upon Hansen and Singleton (1983).[17] To create this moment, we regress the log of consumption growth on the log of the return:

$$log(\frac{c_{t+1}}{c_t}) = \alpha_0 + \alpha_1 \times log(R_{t+1}) + \zeta_{t+1}. \tag{9}$$

At this stage, we offer no structural interpretation of the moment, α_1. We call this coefficient the **aggregate EIS** to distinguish it from the EIS at the household level.

We estimate α_1 from data on real consumption growth of non-durables and services and the Shiller measure of return for the 1967-92 period. The estimated value of α_1 was 0.0138. This is very small and not significantly different from zero. Similar results are reported, for example, by Vissing-Jorgensen (2002).

Since (9) is estimated from aggregated data, it averages over stock market participants as well as non-participants. Further, it averages over portfolio adjusters and non-adjusters. Accordingly, the sensitivity of consumption growth to the return reflects both the participation decision as well as the adjustment decision. A low value of α_1 is indicative of both non-participation and inaction in portfolio adjustment. Thus, in simulation, α_1 is responsive to variations in both γ and ψ.

To be clear, we do not interpret α_1 as an estimate of the household EIS nor its inverse as

[16] In contrast, Luttmer (1999) estimates parameters by looking at deviations from observed outcomes using trading rules.

[17] Despite the warnings contained in Carroll (2001), this is nonetheless a useful moment for inferring the underlying curvature of the utility function. In simulation, we see that this regression coefficient is sensitive to variations in γ.

an estimate of risk aversion. Instead we use our model of household choice to infer household risk aversion using α_1 as a moment.

To estimate the parameters, we solve the dynamic programming problem of an individual household. We then create a panel data set with 500 households and 500 time periods.[18] Households differ because of idiosyncratic income shocks. An additional 500 households who are non-participants in asset markets were added to the data with the restriction that their consumption was equal to their income. We compute moments from the simulated data set in exactly the same manner that the moments were calculated in the actual data. The estimation finds the parameters (γ, ψ) which brings the simulated and actual moments as close as possible.[19]

Moments	**Data**	No Adjustment Costs	Only γ	Both (γ, ψ)	Boundary
Portfolio Adjustment rate	0.71	0.962	0.953	0.894	0.790
α_1	0.0138	0.691	0.622	0.534	0.53
Estimated Parameters					
γ		1.0008	1.00072	1.00067	1.0007
ψ		1	1	0.992	0.980
fit		0.523	0.423	0.305	0.279

Table 2: Moments and Parameters

The second column, labeled "No Adjustment Costs", estimates the curvature of the utility function without any trading costs. That is, the trading costs summarized in Table 1 are set to zero and $\psi = 1$. The estimate of γ is around 1.0, which is the log specification. The value of the aggregate EIS estimated from the simulated data is considerably higher than the actual data moment of 0.014. Though the model has no costs of trading, there is still some portfolio inaction. This is in part due to the fact that wealth is constrained to remain on a grid.

The third column, labeled "Only γ", reports an estimate of γ for the case with $\psi = 1$, using the transactions costs reported in Table 1. The curvature of utility function is again very close the log specification of $\gamma = 1$. The fit, which is the sum of the squared deviations

[18] The state space included 5 values of return, 3 idiosyncratic income shocks, 3 aggregate income shocks and 200 elements in the asset space. We imposed a lower bound on assets of zero.

[19] This is a just-identified exercise and an identity matrix was used as a weighting matrix.

between sample and simulated moments, is 0.423, better than the fit without the trading costs. The estimate of α_1 in the simulated data is 0.622, though γ is around 1. The adjustment rate for the estimated model was 95.3%, much higher than the sample rate. This motivates us to add an additional cost of adjustment through $\psi < 1$.

The fourth column, labeled "Both (γ, ψ)", reports the case where both γ and ψ were estimated. The fit in this case is much better as ψ is estimated to be 0.992, less than 1. As a consequence the adjustment rate is much lower. Also, due to the additional inaction, consumption growth is less responsive to interest rate movements. As in the previous estimates, the curvature of the utility function is close to the log specification.

At $\psi = 0.992$ the cost of adjustment at the average income level is \$576. But, this overstates the adjustment costs actually paid since, as discussed below, higher income households are less likely to adjust their portfolio.

For this estimation, the zero lower bound on wealth was hit in about 5% of the observations.[20] The last column reports estimates when the lower bound is hit about 11% of the time. The fit is better owing to the increased inaction which itself is due to hitting the zero wealth bound. This case is here too for comparison with other estimates reported in our discussion of robustness.

In all cases, the estimated degree of risk aversion is quite modest with γ close to unity. Importantly, one cannot infer the EIS nor the degree of risk aversion directly from the inverse of α_1. Put differently, given our estimate of γ, the estimate of α_1 close to zero does not reflect a high degree of risk aversion. Rather, it reflects the inaction of households in their portfolio adjustment as well as the presence of non-participants in aggregated data.

4 Robustness

Here we consider the robustness of our results. In particular, we consider an alternative model of inaction, other measures of return and estimation with shareholders alone.

[20]That is, the estimation was done with a side constraint of hitting this lower bound not more than 5% of the observations to minimize the effects of the bound on our results.

4.1 Another Version of Inaction

As noted earlier, the meaning of inaction in this portfolio adjustment problem is not entirely clear. We have thus far interpreted inaction as meaning no trades so that any adjustment in the value of a portfolio arose through changes in asset valuations not direct trades.

There is, however, another model of inaction one might consider. In the alternative model, inaction is interpreted as no change in asset holdings. Thus unrealized capital gains and paid dividends are treated as current income.

Specifically, this alternative model defines inaction as no changes in shares $s^i = s^i_{-1}$ for $i = 1, 2, ...I$. In this case, $c = \sum_i (R^i_{-1} - 1)s^i_{-1} + y$.

The column labeled "Alt. inaction" in Table 3 shows our results for this specification. The estimated parameters are a slightly higher estimate of γ compared to the baseline and a substantial cost of adjustment with $\psi = 0.8348$. The fit of this model is not nearly as good as the baseline.

Moments	**Data**	**Baseline**	**Alt. inaction**	**S&P**	Shareholders
Portfolio adjustment rate	0.71	0.894	0.849	0.692	0.697
α_1	0.0138	0.534	1.03	0.0849(0.057)	0.263(0.104)
Estimated Parameters					
γ		1.00067	1.343	3.75	4.34
ψ		0.992	0.901	0.939	0.919
fit		0.305	1.055	0.0011	0.025

Table 3: Moments and Parameters: Robustness

4.2 S&P

Thus far we have used the Shiller return process. This measures the real Standard and Poor's (S&P) annual return including dividend. An alternative is to calculate real returns based upon the S&P index.[21] The model was re-estimated using the S&P return. To do so, we estimated the return process over the period 1967 to 1992. The mean return was 1.0161, the serial correlation is essentially 0 and the standard deviation of the return is 0.16. The

[21] The data also come from Robert Shiller's web page, `http://www.econ.yale.edu/~shiller/data.htm`.

main difference between this and the process including dividends is in the average return, which is must lower for the S&P measure of return.

This process was used to solve the household's dynamic optimization problem. As before, the (γ, ψ) parameters were found to brings the moments calculated in simulated data as close as possible to the real data counterparts.

As indicated in fifth column of Table 3, the estimate of $\alpha_1 = 0.057$ is different due to using the S&P return rather than the Shiller return process. With this moment along with inaction, the estimated values are $\gamma = 3.75$ and $\psi = 0.939$. The estimate of γ is much higher than in the baseline. For this estimation, the lower bound on wealth was hit in about 11% of the observations.

Interestingly, the fit is considerably better than the model estimated with the Shiller return measure. On that basis, one might conclude that this interest rate process was closer to that obtained by shareholders. Yet, this process is not consistent with our model of inaction.

4.3 Shareholders Only

Our results thus far mix both shareholders and non-shareholders in creating the simulated data. This is appropriate since both types of households are present in aggregated data underlying the estimate of α_1.

Here we consider the sample of stockholders alone drawing upon Vissing-Jorgensen (2002). For stockholders, we estimate an aggregate EIS of 0.104 for shareholders.[22]The S&P return was used in this regression and thus we use that process in our model as well.

We use that moment along with the adjustment rate of 0.71 to estimate (γ, ψ) for stockholders alone. Accordingly, we did not include non-participants in the calculation of the consumption growth rate used to estimate α_1.

The results are shown in the last column of Table 2. The adjustment rate in the simulated data is around 0.697. Here the value of ψ is lower than in the baseline, around 0.92. And the value of γ is again considerably larger than in the baseline, $\gamma = 4.34$. In terms of fit, this model fits the sample much better than the baseline model.

For this estimate of γ, the aggregate EIS is about 0.23, which is considerably lower than

[22]This is using annual CEX data and no instruments in the regression.

the baseline model. This is still larger than the estimate of α_1 despite the fact that there are no non-participants on the data. Thus this difference between the aggregate EIS and α_1 reflects the inaction induced by the adjustment costs.

5 Inspecting the Results

In this section, we highlight the forces underlying our results. We first look at aggregation across households. We then look at the policy functions of a household using our baseline estimates.

5.1 Role of Aggregation

We study the role of aggregate by looking at three simulation results, summarized in Table 4. These simulations are all using the baseline parameters: $\gamma = 1.00067$ and $\psi = 0.992$.

The column labeled "Single Agent" shows the portfolio adjustment rate and α_1 estimate from simulated data. The adjustment rate is the same as in the baseline since all of the households in the model are *ex ante* identical. Hence a long simulation of a single agent will yield the same average adjustment rate as one would calculate from a panel.

That statement is not true for α_1. For a single agent, the α_1 moment is very small compared to the shareholders and baseline cases. In the case of a single agent, the consumption growth used in (9) to estimate α_1 reflects the impact of both aggregate and idiosyncratic income shocks. Due to the presence of adjustment costs, individual consumption growth is not very responsive to interest rate movements.

The moments produced by the single agent are closer to those from the actual data than the baseline. This is seen by the lower value of fit than the baseline. But recall that the moments themselves were produced from aggregated not individual data.

But aggregated consumption is considerably more responsive to interest rates, as seen in the column labeled "Shareholders". To obtain these results, we simulated a panel using baseline parameters. We then regressed the growth of average consumption of shareholders alone on the log of the interest rate. This aggregation across heterogeneous households implies a much larger response in the growth of average consumption to variations in the interest rate. We interpret this as a consequence of smoothing over inaction by aggregation.

The final column reproduces the baseline results. These results differ from "Shareholders" due to the inclusion of both non-participants. The presence of these agents lowers the response of consumption growth to interest rates.

Moments	**Single Agent**	**Shareholders**	**Baseline**
Portfolio adjustment rate	0.894	0.894	0.894
α_1	0.022	0.736	0.534
fit	0.034	0.556	0.305

Table 4: Effects of Aggregation

5.2 Who Trades

It is useful to have a sense of the discrete choice made at the individual household level. One way of doing so is to summarize the portfolio adjustment choice through a simple probit regression. The results are reported in Table 5.

Variable	marginal effect
Asset Holdings (millions)	0.000497
	0.0000383)
Income (millions)	-0.00486
	0.000388)

Table 5: Probit Regression Results: Simulated Data

From these results, an increase in wealth will increase the likelihood of adjustment: an increase of asset holdings by \$100,000 increases the adjustment probability by about one-half of a percentage point. But an increase in income actually reduces the probability of adjustment: an increase of income by \$10,000 dollars decreases the probabilty to adjust by about one-half of a percentage point. This is presumably due to the role of ψ as an adjustment cost.

6 Conclusion

Our goal in this paper was to understand the implications of infrequent portfolio adjustment for estimates of household risk aversion and the elasticity of intertemporal substitution. To do so, we formulated and estimated a model incorporating fixed and quadratic costs of portfolio adjustment. The estimation was through a simulated method of moments approach to match a parameters linking interest rates to consumption growth from aggregated household data and inaction of portfolio adjustment.

We found support for costs of portfolio adjustment directly from trading data. Our estimates find a utility loss associated with trading and a risk aversion estimate that is much lower than the one inferred from estimation based upon the representative agent model with continuous adjustment.

One limitation of our model is the absence of multiple assets. An alternative formulation would include asset specific trading cost to distinguish between assets which are easily adjusted, such as a savings account, and those which require more decision time, such as complex options trades. If, for some of the assets, there are no costs of adjustment at all, including both direct and time costs, then the agent could use that asset for the purpose of consumption smoothing in each period. For this single asset, (1) will hold. Adding these additional assets, as well as durable goods such as housing, remains for future work.

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.

Brunnermeier, M. K., and S. Nagel (2008): "Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-evidence on Individuals," *American Economic Review*, 98(3), 713–36.

Bucks, B., A. Kennickell, T. Mach, and K. Moore (2009): "Changes in U.S. Family Finances 2004 to 3007: Evidence from the Survey of Consumer Finance," mimeo, Board of Governments, Division of Research and Statistics.

Carroll, C. (2001): "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *Advances in Macroeconomics*, 1(1), 1–38.

Cooper, R., J. Haltiwanger, and J. Willis (2008): "Euler-Equation Estimation for Discrete Choice Models: A Capital Accumulation Application," mimeo.

Hall, R. (1988): "Intertemporal substitution in consumption," *The Journal of Political Economy*, pp. 339–357.

Hansen, L., and K. Singleton (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 91, 1269–1286.

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.

Heaton, J., and D. Lucas (1996): "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104, 443–487.

Luttmer, E. (1999): "What Level of Fixed Costs can Reconcile Asset Returns and Consumption Choices?," *Journal of Political Economy*, 107, 969–997.

Mankiw, N., and S. Zeldes (1991): "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, 29, 97–112.

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Analysis of Harold Zurcher," *Econometrica*, 55(5), 999–1033.

Vissing-Jorgensen, A. (2002): "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110(4), 825–853.