Costly Portfolio Adjustment*

Yosef Bonaparte† and Russell Cooper‡

January 11, 2011

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 an annual basis, less than 71% 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 and the consequent inactivity in portfolio adjustment 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 elasticity of intertemporal substitution obtained from log linear regressions of consumption growth.

*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, Cambridge University, the EIEF in Rome, Oxford University, the Toulouse School of Economics, the Finance Division at the UBC Sauder School of Business and at the Pennsylvania State University for comments. 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, University of Texas at Austin and the European University Institute, russellcoop@gmail.com

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

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

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, usually the return on a stock portfolio. 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 the portfolio earning $R_{t=1}^i$ in (1), 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]

[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 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 of the portfolio whose return is measured in the Euler equation, or its approximation. So, if there is infrequent costly adjustment of stock portfolios, then (1) will not hold when R^i_{t+1} is the measured real return on that portfolio. Yet, (1) will hold for assets without adjustment costs.

Second, the inaction at the household level will not necessarily disappear when looking at aggregated data. The inaction in household adjustment may well effect the estimation based upon aggregated data.

The nature and extent of the bias to 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 the 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), 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

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 an average of the 1995, 1998, 2001 and 2005 cross-sections.

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 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.

For our baseline results and in a number of extensions, we estimate the curvature of the household utility at a little over **3.5**.[7] Thus in our model, the curvature of the utility function is not the inverse of the coefficient on the interest rate in the standard consumption growth regression.

In addition, we find the presence of non-convex adjustment costs, due to both fixed trading costs and time costs at the household level. Finally, for our baseline we estimate a household discount factor of about 0.88, reflecting a high average portfolio return.

Throughout these extensions, the response of consumption growth to variations in the

[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.

[6]The approach of Cooper, Haltiwanger, and Willis (2010) 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]This the estimate of γ where household utility is represented by $u(c) = \frac{1}{1-\gamma}c^{1-\gamma}$.

interest rate, one of our key moments, varies considerably. But, the estimated curvature of the utility function is remarkably robust to the use of different interest rate series and to the inclusion in the estimation of households not participating in asset markets.

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 its portfolio. These costs take two forms: direct financial trading costs and the opportunity cost associated with the time involved in trading. An explicit model of these time costs, through search, is provided in Bonaparte and Fabozzi (2010). For our analysis, these costs are interpreted as costs of trading rather than costs of thinking about trading. [8] If the household chooses not to adjust, then consumption is equal to labor income. All capital income, such as unrealized capital gains and dividends, are assumed to be reinvested without cost.[9]

Our focus is on the consumption savings choice rather than portfolio rebalancing *per se*. Accordingly, we restrict attention to an optimization problem in which the household must adjust all the elements of its portfolio simultaneously. If there are fixed costs of adjustment which apply across all assets, such as information gathering costs, then it will be optimal for households to adjust an entire portfolio. In what follows we refer to the household as having a single asset.

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, an existing asset of s_{-1} and the return from the previous period was R_{-1}. Total financial wealth equals $R_{-1}s_{-1}$ at the start of a period.

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

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

[8]Models of rational inattention are discussed, *inter alia*, in Brunnermeier and Nagel (2008), Abel, Eberly, and Panageas (2009) and Alvarez, Guiso, and Lippi (2010) .

[9]We study a version of the model in which reinvestment is costly as well.

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

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

The consumption level is given by

$$c = R_{-1}s_{-1} + y \times \psi - s - C(s_{-1}, s). \tag{4}$$

Here s is the current purchase of the asset. We allow the household to be in debt but there is a limit, as indicated in (3), on indebtedness imposed in the estimation. 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 reflecting heterogeneity in 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.[10] Hence

$$s = R_{-1}s_{-1}. \tag{6}$$

This is in the argument for the transition of shares in (5).

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

[10] 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, explored below, would have the actual shares remain constant and allow consumption to absorb the return on the existing portfolio.

variations, though, are met by variations in consumption, not in the holdings of assets if the household chooses "no adjustment".

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 time opportunity costs of trade separately from actual trading costs. Our data are described in detail in the Appendix.

3.1 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 substitution. We estimate the preference parameters γ, along with the trading cost ψ and the discount factor β using simulated method of moments, described below.

We could also add a direct utility cost of trading.[11] We have instead chosen to model the cost of adjustment as a time cost, through the inclusion of ψ in (4).

Cost of Adjustment A novel part of the analysis 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 some asset i is bought and $j = s$ if it is sold. We assume:

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

[11] 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.

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 directly estimate these transactions costs for the buying and selling of common stocks. These estimates come from the data set used by 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.[12] 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^i_0	56.10	61.44
	(0.05)	(0.061)
Linear ν^i_1	0.0012	0.0014
	(1.63e-06)	(1.93e-06)
Quadratic ν^i_2	$-1.01e^{-10}$	$-1.28e^{-10}$
	(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

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.[13] Because of this fixed cost, even if $\psi \geq 1$, the model will still produce inaction.

[12]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.

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

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.

These estimates were obtained for the trade of a given asset i. Since agents who choose to trade may trade more than a single asset in a given year, the fixed cost per trade may understate the actual fixed cost incurred. We consider below the effects of larger fixed costs.

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, ψ in (4).

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 shock and a household specific shocks. These income processes along with the return process are represented as AR(1) processes.

The return process measures the return on stock holdings and includes both dividends and capital gains and has an average annual real return of 5.5% and a serial correlation of 0. We term this the "stock" return process in the discussion that follows. The data used to construct these processes are described in detail in the Appendix.[14]

3.2 Simulated Method of Moments Estimation: Results

Given these estimates of transactions costs, we use a simulated method of moments technique to estimate the remaining parameters, (γ, ψ, β).[15] There are three moments used to identify the parameters. The appendix provides details on the data used to calculate these moments. 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.

[14] As discussed, for example, in Guvenen (2007) and Browning, Ejrnaes, and Alvarez (2009), the estimation of the income process is by itself an open issue. Given our focus on portfolio adjustment costs, we have chosen a rather parsimonious model of the income process. The serial correlation of the annual real return is not statistically different from zero for the post-WWII and more recent samples.

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

There are two important ways in which this moment may not quite fit with the model's assumption that the only means of smoothing consumption is through costly portfolio adjustment. First, households obviously hold some liquid assets which can be used for smoothing consumption. For example, a household paid monthly can smooth consumption over a month by spending out of a paycheck deposited into a checking account. There are no portfolio adjustment costs incurred. In this sense, the 71% overstates the amount of inaction and thus the ability of a household to smooth consumption. By looking at annual rather than, say monthly data, we are smoothing over some of these higher frequency movements in consumption and income.

On the other hand, some portfolio adjustment reflects portfolio rebalancing without direct effects on consumption. In this case, the 71% understates the amount of inaction related to the use of portfolio adjustment for consumption smoothing.

Reality lies somewhere between the inaction rate used as a moment here and the assumption of costless portfolio adjustment used in the standard literature. Our results indicate that including these costs in the household optimization matter for inferences about preference parameters. We return to this issue in our robustness discussion.

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).[16] 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 stock returns for the 1967-92 period. To be precise, this is the growth rate of aggregated consumption not the average growth rate across households. The consumption series pertains

[16]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. In contrast to Carroll (2001) and Attanasio and Low (2004) we do not infer the degree of risk aversion directly from this parameter but rather infer it indirectly through the SMM estimation of the structural parameters.

to households who are direct participants in asset markets.

The estimate is $\alpha_1 = 0.0878$ with a standard error of 0.0021. This estimate is larger than the estimates reported elsewhere in the literature since our sample consists of asset market participants only.

The third moment is the median financial wealth to income ratio. The value of this ratio is 1.03 in our data. This moment is quite sensitive to the average return on the portfolio and is informative about β.

The parameter vector (γ, ψ, β) is estimated using Simulated Method of Moments. To do so, we solve the dynamic programming problem of an individual household. We then create a panel data set with 500 households and 500 time periods.[17] Households differ because of idiosyncratic income shocks. 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, β) which brings the simulated and actual moments as close as possible.

Table 2 summarizes our results for four different specifications of the return process. Since our model assumes there is a single asset, it leaves open exactly which return process is appropriate. We look at four return processes. These represent the last four columns of the table. The first set of rows are the moments and the second set report our parameter estimates. For the moments, one row reports the estimate of α_1 from the data and the following row reports the estimates from the simulated model. As we use different return processes, the estimate of α_1 is different for each column.

The first is the stock market return process that includes both dividends and capital gains. This is the appropriate return process if we interpret our single asset as a composite of stocks. We term this our baseline model.

The column labeled "Baseline" reports results for the estimation of γ, β and ψ. The curvature of the utility function is 3.51, the discount factor is 0.88 and the trading cost is about 2.6% of annual income. As we shall see, these parameter estimates are remarkably robust across alternative models.

At these parameter values, the sensitivity of consumption growth to the interest rate is

[17] For estimation, the state space included 5 values of return, 3 idiosyncratic income shocks, 3 aggregate income shocks and 500 elements in the asset space. We imposed a lower bound on assets of equal to 20% of the natural borrowing constraint, about $5,000 of unsecured debt. This constraint is binding in about 1.5% of our observations.

0.110, slightly above the estimate of α_1 in the data. The adjustment rate is about 72% and the wealth income ratio matches its data counterpart.

At $\psi = 0.974$ the cost of adjustment at the average income level is about $1,870. But, this overstates the adjustment costs actually paid since the choice of whether to adjust or not is endogenous and this calculation ignores this selection.

The estimate of β of 0.880 may seem relatively low. This is driven largely by the high mean return on the interest process of nearly 5.5%.

Moments	**Data**	**Baseline**	**Stocks and Bonds**	**S&P**	**CD**
Portfolio adjustment rate	0.71	0.72	0.68	0.702	0.75
α_1 data		0.088	0.149	0.083	0.571
α_1 sim.		0.110	0.068	0.0756	0.03
wealth income ratio	1.03	1.03	0.99	1.03	1.07
Estimated Parameters					
γ		3.51	3.81	3.79	3.42
ψ		0.974	0.964	0.967	0.976
β		0.880	0.895	0.916	0.924
fit		0.0006	0.0085	0.0001	0.298

Table 2: Moments and Parameters

The fit is close to zero for the just identified model.[18] The model though slightly overstates the sensitivity of consumption growth to interest rate movements. This motivates our study of other measures of asset returns. As we shall see, the estimates in the other specifications are generally very close to the baseline parameter values.

For each of these alternative processes, we estimate an AR(1) representation of the return. This is used to form the conditional expectation of the agent's optimization problem. This return is also used to estimate α_1 in (9). Consequently, Table 2 presents return specific estimates of α_1 both for the actual data and for the simulated data.

The second set of results, labeled "Stocks and Bonds" computes an annual return based on a portfolio composed of stocks (60%) and bonds (40%). This return had a mean of 3.44%

[18] We use an identify matrix and thus the fit is the sum, over the three moments, of the squared difference between the simulated and data moments. Table 3 provides information on the elasticities of the moments with respect to the parameters.

and the serial correlation of 0.086. This same process was used to estimate $\alpha_1 = 0.149$. Evidently consumption growth is more responsive to the return on this composite portfolio than the baseline.

The parameter estimates for this process are quite similar to the baseline. The estimate of γ is just above 3.8 and the adjustment cost is slightly larger. The low value for β again reflects the high portfolio return.

A third alternative is to calculate real returns based upon the S&P index. This measure of return was used by Vissing-Jorgensen (2002). The mean return was 1.016%, the serial correlation was zero and the standard deviation of the return is 0.16 over the sample period. One main difference between this and the process including dividends is in the average return, which is considerably lower for the S&P.

For this return process, the estimate of $\alpha_1 = 0.083$. With this moment along with inaction and the wealth to income ratio, the estimated values are $\gamma = 3.79$, $\psi = 0.967$ and $\beta = 0.916$. The estimates of (γ, ψ) are both close to the baseline. Given the lower mean return, the estimate of β is higher.

The fit for this return process is better than the baseline. This model is able to capture the muted response of consumption growth to the interest rate. But the estimated model does not match the adjustment rate as well as the baseline.

The final process studied was the return on a CD. The mean return for this process is 2.15%. In contrast to the other processes, the CD return had a very high serial correlation of 0.879. This model fits the moments very poorly. For this asset, it might be that the direct trading costs are negligible. The estimates without the financial trading costs were essentially identical. This indicates that the CD return is a not a good proxy for the joint return on the portfolio.

In all cases, the estimated degree of risk aversion is quite modest with γ ranging between 3.4 and 3.8. 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.

3.3 Sensitivity of Moments to Parameters

Table 3 shows how the moments respond to variations in parameters. These numerical derivatives are computed in the neighborhood of the Baseline parameter estimates. The three parameters are listed as rows and the three moments are the columns. An entry is an elasticity: the percentage change of a moment with respect to the percentage change in a parameter. So, for example, variations in β have large effects on both the wealth income ratio and the adjustment rate. The α_1 parameter is influenced most by variations in the adjustment cost, ψ.

Table 3: Elasticities of Moments to Parameters

parameter	adjustment rate	α_1	wealth income ratio
γ	2.7108	0.5193	4.5495
ψ	6.8047	6.9963	1.5632
β	37.2208	3.1711	45.5894

4 Robustness

Here we study the robustness of our results. In particular, we consider an alternative model of inaction, an estimation with shareholders and non-shareholders and a specification with much larger trading costs. A key result is that the estimates of the three parameters are very robust across these alternative models.

4.1 An alternative model of Inaction

Thus far inaction implied no assets 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 and hence consumption.

Specifically, this alternative model defines inaction as no changes in shares $s = s_{-1}$. In this case, $c = (R_{-1} - 1)s_{-1} + y$.

The column labeled “Alt. ” in Table 4 shows our results for this specification. The fit of this model is not nearly as good as in the baseline. The adjustment rate and aggregate EIS are both larger than the data and wealth income ratio is slightly low.

However, the estimates of (γ, ψ, β) are close to the baseline estimates. Evidently, this model of inaction implies that non-adjustment is less costly so that the cost of adjustment, ψ is slightly lower to get closer to the portfolio adjustment rate.

Moments	**Data**	**Baseline**	**Alt.**	**Part.**	**Big Costs**	**SC**	**Big Shocks**
Port. adj. rate	0.71	0.72	0.74	0.70	0.69	0.71	0.70
α_1 data		0.088	0.088	0.0138	0.088	0.088	0.088
α_1 sim		0.110	0.153	0.041	0.120	0.084	0.071
wealth income ratio	1.03	1.03	1.01	1.04	1.03	1.03	1.03
Est. Parms.							
γ		3.51	3.45	3.40	3.26	3.15	3.52
ψ		0.974	0.970	0.970	0.986	0.974	0.985
β		0.880	0.886	0.884	0.886	0.889	0.885
fit		0.0006	0.0052	0.0009	0.0012	0.00005	0.00034

Table 4: Moments and Parameters: Robustness

4.2 Market Participation

Thus far we have focused on the behavior of stock market participants. But many studies of the response of consumption growth to the interest rate measure the consumption of all households not just stock market participants.[19] We extend our analysis to that situation by adding in households who do not participate in stock markets.

To do so, we first reestimate (9) using aggregate consumption growth. The estimated value of α_1 was 0.0138. This is very small and not significantly different from zero, as is the case in many papers in this literature.

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. Accordingly, this estimate of α_1 is lower than the estimate obtained for stock market participants only.

[19]This issue of stock market participation is the focus, for example, of Alan (2006).

As this 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.[20]

Following the calculations in Vissing-Jorgensen (2002), we assume that 50% of the households participate in financial markets.[21] The estimate of α_1 from the simulated data comes from aggregating consumption across the households who participate in asset markets and those who do not and then calculating consumption growth rates.

Our finding are reports in the column labeled "Part." in Table 4. The parameter estimates are again very close to those reported for the baseline.[22] The fit is not as good as the baseline. Still, the model seems capable of matching the evidence of both participants alone and more aggregated data that includes non-participants.

4.3 Larger Fixed Trading Costs

As noted earlier, our measure of transactions costs is for a given asset while our model assumes adjustment of a portfolio of assets. From the SCF, the average number of trades by an agent is about 6 per year. Thus the fixed cost of trade reported in Table 1 may be too small. The row labeled "Big Costs" in Table 4 reports results when the fixed costs of trading is $1,000 rather than the value of around $60 reported in Table 1. This is enough of a payment to cover 16 trades, well above the average. In this way, the results bracket two extremes.

The next to last column of Table 4 shows that the estimate of γ for this model is a bit lower than the baseline. Since the estimation uses the baseline return process, the estimate of β is also close to the baseline. But, with the fixed financial costs of trade about 20 times higher than the baseline, the estimated value of ψ is larger to bring the adjustment frequency close to the data moment. The fit of this model is not as good as the baseline.

[20] 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.

[21] 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.

[22] Compared to Alan (2006) study of endogenous asset market participation, our discount factor is about the same and we find more curvature in the utility function.

4.4 Serial Correlation in Interest Rates

As noted earlier, we impose an iid process for interest rates. Both Carroll (2001) and Attanasio and Low (2004) discuss in the context of their simulation results the importance of serially correlated interest rate shocks for the identification of household risk aversion from the coefficients of (9). Though our identification is indirect, rather than direct, it is still of interest to study the effects of serial correlation on our estimates.

To do so, we use the parameterization of Attanasio and Low (2004) which assumes a serial correlation of 0.21 for an annual process. The column labeled "SC" in Table 4 summarizes our results. First, note that the fit is an order of magnitude better for this model, compared to the baseline case of iid interest rates. Second, the parameters are not far from the baseline though the curvature of the utility function is lower than the baseline estimate.

While this specification of the return process leads to a better fit of the model, it does so at a cost. For this specification, the beliefs of agents are not consistent with the underlying return process. We use this insight in Bonaparte and Cooper (2010) to look at the effects of overconfidence on household portfolio turnover.

4.5 Larger Income Shocks

The model only allows for consumption smoothing through costly portfolio adjustment. This may be appropriate for large income shocks, such as job losses and health shocks. But households might be able to smooth consumption at relatively little cost in the face of relatively small income variations.

To study this, we simulate the model allowing only large income shocks. To do so, we approximate the idiosyncratic income process using two states, where the high state is more than twice the low state. The interest rate process is still approximated using five states with the return ranging from 0.896 to 1.21. Thus interest rate movements are relatively large.

The results are presented in the column labeled "Big Shocks" in Table 4. The parameter estimates are close to the baseline and the fit remains very good. This suggests that our baseline results were largely driven by responses to larger shocks so that the limited ability to smooth consumption in face of smaller income shocks was not important for our estimation.

5 Inspecting the Results

In this section, we highlight the forces underlying our results. We first look at the role of aggregation across households. We then study the contribution of the adjustment costs.

5.1 Role of Aggregation

The model was estimated using the growth rate of average consumption. There might be interesting differences between the behavior at an individual level and in the aggregate. At the individual level, there are non-convexities leading to a discrete choice in adjustment that are smoothed over through the aggregation exercise. So the Euler equation that underlies (9) might be a better approximation to smoothed aggregate behavior than a version of this equation estimated using data at the individual household level. Further, heterogeneity in income at the individual level is smoothed when one studies aggregate consumption growth.

Moment	**Single Agent**	**Panel**	**Average Growth**	**Baseline**
α_1	0.037	0.043	0.044	0.110

Table 5: Effects of Aggregation and Adjustment Costs

We study the role of aggregation by looking at four simulation results, summarized in Table 5. These simulations are all using the baseline parameters: $\gamma = 3.51$, $\psi = 0.974$ and $\beta = 0.880$. Table 5 reports the estimates of α_1 from various simulations using these parameter values.

The column labeled "Single Agent" looks at the time series (400 periods) of a single household. The estimate of α_1 comes from the consumption and adjustment choices of a single household, there is no aggregation whatsoever. The estimate of α_1 is slightly lower than the baseline and not significantly different from zero.

The column labeled "Panel" shows the portfolio adjustment rate and α_1 estimated from a simulated panel data. Instead of aggregating across households to produce measures of aggregate consumption growth, here the simulated data and hence the moments come from the time series of individual households. For this analysis, 100 households were randomly selected from our simulated data. For this experiment, like the single agent case, the estimate of α_1 is lower than the baseline model.

The column labeled "Average Growth" calculates aggregate consumption growth by averaging across consumption growth at the household level. This is not the same aggregate used to create the baseline. The estimate of α_1 in this case is essentially the same as that obtained from the estimates at the level of the individual household and in the panel data.

The final column reproduces the baseline results. Here the response of consumption growth to interest rates is larger than the other treatments.

Clearly aggregation matters for matching the aggregate EIS in the data. As shown here, the differences are due to the manner in which aggregate consumption growth is calculated: as growth of average consumption (as in the baseline) or as average consumption growth.[23] The difference in these measures is from weighting. Yet the differences are not large: aggregation is apparently not a key issue in understanding the mapping from the moments to structural parameters.

5.2 Role of Adjustment Costs

Table 6 highlights the effects of adjustment costs on our results. To do so, we report the simulated moments with and without the various adjustment costs of the model. These are simulations using the baseline parameters, there is no further estimation involved.

Moments	Baseline	No Costs	Only $C(\cdot)$	Only ψ	$\psi = 0.95$
Portfolio Adjustment rate	0.72	0.97	0.93	0.72	0.62
α_1	0.110	0.112	0.115	0.110	0.096
wealth income ratio	1.03	1.09	1.09	1.03	0.993
fit	0.0006	0.072	0.052	0.0006	0.0098

Table 6: Moments and Parameters

The first column is the baseline as a basis for comparison purposes. The column, labeled "No Costs" presents moments for the model without any trading costs. That is, the trading costs summarized in Table 2 are set to zero and $\psi = 1$. The portfolio adjustment is of course

[23] Attanasio and Weber (1995) discuss the effects of these two forms of aggregation and report that their point estimate of the EIS is 0.214 using average consumption growth and increases to 0.452 using the growth of average consumption.

much higher than the baseline.[24]

Note that the inverse of α_1 remains far from γ, the curvature of the utility function. This is consistent with the findings of Carroll (2001), though we are using an infinite horizon model and looking at time series variation.

The column, labeled "Only $C(\cdot)$", reports moments using the transactions costs reported in Table 2. The adjustment rate is a bit lower due to these costs. But the costs coming through the estimated trading costs are not large enough to explain the observed inaction and this motivates an additional cost of adjustment through $\psi < 1$.

The column, labeled "Only ψ", reports moments for a model where the direct trading cost is ignored so that $C(\cdot) \equiv 0$. The moments are essentially identical to the baseline.

One feature of these simulations worth noting is that the aggregate EIS (α_1) is not that sensitive to variations in trading costs. However, from Table 3, the elasticity of α_1 with respect of ψ is about 7.0. This same elasticity is about 4.2 when the derivatives are evaluated at $\psi = 1$. So it is not the case that variations in ψ have no effect on this moment.

To further illustrate this, a simulation with $\psi = 0.95$ is shown in the last column of the table. In this case α_1 falls to 0.096. If we set $\psi = 0.90$, then α_1 falls further, to 0.07. Thus the response of consumption to the interest rate does depend upon portfolio adjustment costs.

The parameter estimate of $\alpha_1 = 0.110$ in the baseline model reflects the effects of interest rate movements on two margins: the extensive margin of adjust/no adjust and the intensive margin of how much to adjust. These two effects move in opposite directions. Because of the adjustment costs, households do not always adjust their portfolios which dampens the response of consumption growth to the interest rate. But, conditional on adjustment, households may react more to interest rate movements than they would otherwise.

Looking at the simulated data, the fraction of household's adjusting their portfolio depends upon the realized interest rate. The adjustment rate is 69% in the lowest interest rate state and rises monotonically to 75% in the highest interest rate state.

On the intensive margin, at the level of an individual household, we estimated α_1 conditional on portfolio adjustment. That is, (9) was estimated using observations of consumption growth and interest rates conditional on adjustment in period $t+1$. The resulting estimate was $\alpha_1 = 0.131$, considerably higher than the estimate of $\alpha_1 = 0.037$ reported in Table 5. The

[24]The fact it is not literally 1 reflects the discreteness of the state space.

latter estimate comes from pooled data over periods of adjustment and non-adjustment while the estimate of $\alpha_1 = 0.131$ captures the intensive margin by conditioning on adjustment.

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.

Our focus here has been on the consumption/saving margin. To study portfolio rebalancing requires a more complex model. 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. Adding these additional assets, as well as durable goods which entail their own adjustment costs, such as housing, remains for future work.

7 Appendix

7.1 Data Sources for the Driving Processes

We need both interest rate and income processes to solve the household optimization problem. These are computed directly from data on interest rates and income.

Interest Rate Processes The stock return process came from an AR(1) representation of the S&P data on real returns deflated by the CPI, 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 5.49% over the 1967-1992 period. The serial correlation of the return is near zero (though not very precisely estimated) and the standard deviation is 0.159.[25]

Household Income We use data from the PSID to estimate a household income process. Our sample period is from 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).

We decompose 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.

7.2 Definition of the moments

Adjustment The annual measure comes from the response to question X3928 in the Survey of Consumer Finances where responders who own stocks through a broker asked "Over the past year, about how many times did you or anyone in your family living here buy or sell stocks or other securities through a broker?" We assign a value of 1 if the household report at least one action of buy or sell and zero otherwise.

Consumption Growth Response to Interest Rates For this moment, we follow the literature on intertemporal consumption and study the response of consumption growth to interest rate movements. Households are classified as stock market participants if they report

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

ownership of stocks. As in Vissing-Jorgensen (2002), the consumption data for shareholders is from the CEX. We convert the data there to create annual rather than semi-annual measures of consumption growth. In our robustness, we study the response of average consumption growth to interest rates. For that analysis, we use data on aggregate real consumption expenditures from Robert Shiller available at `http://www.econ.yale.edu/~shiller/data.htm`. The interest rates process was described above.

Wealth Income ratio The SCF also report the total financial wealth (in a variable named FIN), which is the sum of all liquid assets such as stocks, bonds, mutual funds, and other quasi-liquid retirement accounts. Housing is not included. The SCF also reports the labor earnings for each household. For example, for the 1998 SCF wave, Variable X5702 reports information about: "In total, how much income from wages and salaries did you receive in 1997, before deductions for taxes and anything else?" Given this information, we generate a variable that measures the median ratio of wealth to income.

7.3 Estimating the transaction costs parameters

We utilize household account data set of more than 3 million observations used by Barber and Odean (2000) to calibrate the transaction costs parameters. The trading costs are estimated in a quadratic linear regression in which 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. Table 2 reports the estimated parameters.

References

Abel, A., J. Eberly, and S. Panageas (2009): "Optimal Inattention to the Stock Market with Information Costs and Transactions Costs," NBER Working Paper #15010.

Alan, S. (2006): "Entry Costs and Stock Market Participation over the life cycle," *Review of Economic Dynamics*, 9, 588–611.

Alvarez, F., L. Guiso, and F. Lippi (2010): "Durable Consumption and Asset Management with Transaction and Observation Costs," NBER Working Paper #15835.

Attanasio, O., and H. Low (2004): "Estimating Euler Equations," *Review of Economic Dynamics*, 7(2), 406–435.

Attanasio, O., and G. Weber (1995): "Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103(6), 1121–1157.

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 (2010): "Rationalizing Trading Frequency and Returns," NBER Working Paper #16022.

Bonaparte, Y., and F. Fabozzi (2010): "Household Search Choice: Theory and Evidence," forthcoming *Applied Economics*.

Browning, M., M. Ejrnaes, and J. Alvarez (2009): "Modeling Income Processes with lots of Heterogeneity," *Review of Economic Studies*, 33(1), 1–33.

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 (2010): "Euler-Equation Estimation for Discrete Choice Models: A Capital Accumulation Application," NBER Working Paper #15675.

GUVENEN, F. (2007): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?," *American Economic Review*, 97(3), 687–712.

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.