

# Quantifying Explanations for Black-White Wealth Inequality\*

David Low

NYU

March 2013

PRELIMINARY AND INCOMPLETE

## Abstract

The black-white wealth gap is large, and regressions of wealth on income and demographics can explain only a fraction of it. A standard Blinder-Oaxaca decomposition finds that the average black household, with \$45,467 in net wealth, would have \$67,975 (i.e. 49.5% more) if it were white but other characteristics were unchanged. Many explanations for this gap have been proposed, but few have been tested. This paper uses a structural model of consumption, savings, and portfolio choice over the lifecycle to quantify the importance of old and new explanations for the black-white wealth gap. Preliminary results suggest that income is important but not as central as previous regression-based studies have found, while a previously ignored factor - mortality - may play a modest role in the black-white wealth gap.

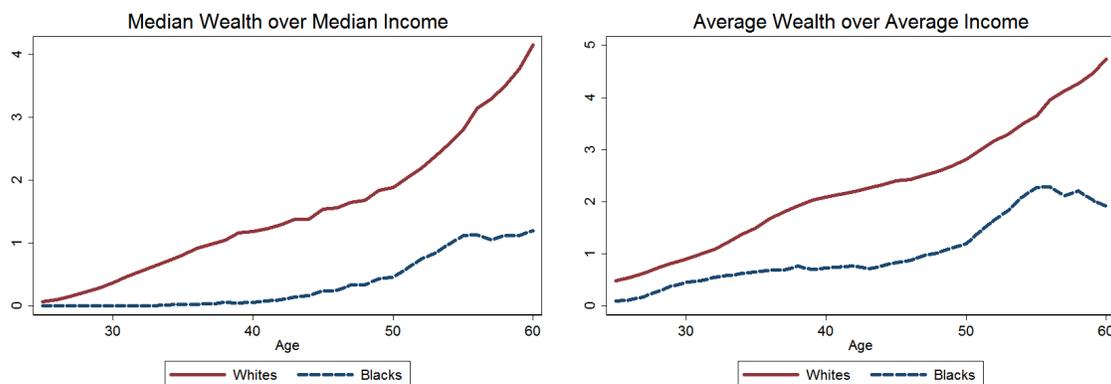
---

\*My advisor, Raquel Fernández, provided invaluable support with this paper. I would also like to thank Christopher Flinn, Kevin Thom, Ahu Gemici, Gianluca Violante, Joseph Mullins, David Evans, Andrew Demers, and Elliot Lipkowski for many very useful comments.

# 1 Introduction

Blacks in the United States have much less wealth than whites. Consider Figure 1, which plots the wealth to income ratio for whites and blacks over the lifecycle.

Figure 1: Wealth to Income Ratio



Many economists have used regressions to quantify the amount of the wealth gap that can be explained by income and other observables, e.g. Blau and Graham (1990), Altonji et al. (2000), Scholz and Levine (2004) and Altonji and Doraszelski (2005). The typical finding is that the amount of the wealth gap explained depends heavily on the counterfactual; giving whites the black distribution of observables reduces the wealth gap by much more than giving blacks the white distribution of observables. For example, Altonji and Doraszelski (2005) explain 79% of the wealth gap for married couples using the white regression, but only 25% of it with the black regression. The implications are discussed in more detail in Section 3, but the sheer size of the wealth gap means that even the white regression misses a great deal of wealth. For example, the average black household, with \$45,467 in net wealth, would have \$67,975 (i.e. 49.5% more) if it were white but other characteristics were unchanged.<sup>1</sup>

Sociologists (and the media) have offered a long list of potential explanations for this wealth gap, from financial illiteracy to discrimination in the housing market, but there has been very little quantitative work to establish how much of the wealth gap these factors actually explain. A rare exception is Menchik and Jianakoplos (1997), who estimate the additional effect of giving blacks the white distribution of bequests. They find bequests explain from as little as 1% - 2% of the gap (using NLSY data) to as much as 12%-19% of the gap (using SCF data). The upper bound suggests

<sup>1</sup>For comparison, the average white household has \$145,626 in wealth, and would have only \$75,816 (i.e. 47.9% less) if it were black but other characteristics were unchanged. The white regressions explain more in absolute terms, but in proportional terms - which are arguably more important - the results are more comparable.

that giving blacks the white distribution of income, demographics, and bequests can explain about half of the wealth gap.<sup>2</sup>

Other studies have identified differences between whites and blacks that could help explain the wealth gap, though they have not quantified by how much. Chiteji and Stafford (1999) find that, *ceteris parabus*, blacks are less likely to invest in high-return equities than whites, and argue that financial literacy passed down from parents to children may contribute to the black-white wealth gap. Charles and Hurst (2002) find that, conditional on other observables, blacks are less likely to be approved for a mortgage than whites, though they conclude the “effect [of the mortgage approval gap] on the race gap in housing transitions is small.” Bayer et al. (2012) conclude that blacks pay a statistically significant premium of 3% when buying houses.

Could these seemingly small differences in the property market translate into large observed differences in wealth, as is often assumed? Are there important interactions between inheritances, financial literacy, and the property market? There is no previous quantitative work to judge. Moreover, there are several other potential contributors to the wealth gap - including starkly different marriage, divorce, and mortality rates - that, to the best of my knowledge, have never been quantitatively tested. Even if each factor individually explains only a small portion of the wealth gap, do they combine to explain most or all of it? Or is the racial wealth gap truly a mystery?

A realistic structural model of asset accumulation could shed significant light on these questions, but to date no one has applied such a model to the study of black-white wealth inequality. That is the purpose of this paper. Structural models have often been used to study other social questions in economics. Keane and Wolpin (2010), in perhaps the paper closest in spirit to this one, explore differences in the education, marital status, fertility, employment, and welfare decisions of white, black, and Hispanic women. Gender inequality in the labor market has also proven a fertile ground for structural research. Van Der Klaauw (1996), Attanasio et al. (2008), Eckstein and Lifschitz (2011), Fernández and Wong (2011), Fogli and Veldkamp (2011) and Fernández (forthcoming) study historical trends in female labor force participation, while Albanesi and Olivetti (2009), Erosa et al. (2010), and Gayle and Golan (2012) examine the gender wage gap. This paper applies a familiar methodology to a new subject.

The rest of this paper is organized as follows. Section 2 discusses the data used in this paper, while Section 3 presents more detailed information on the wealth gap and explanations for it. Section

---

<sup>2</sup>The SCF numbers from Menchik and Jianakoplos (1997) do not control for other parental characteristics that may be correlated with bequests, such as lifetime income or financial literacy, and so they are probably picking up the effect of a number of variables beyond bequests. See Charles and Hurst (2003) for more information. They also have a much less extensive set of control variables than Altonji and Doraszelski (2005) or I use. In Appendix A, I find that the additional explanatory power of controls for parental income and education is negligible.

4 develops a realistic structural model of asset accumulation and portfolio choice over the lifecycle, which is estimated in Section 5. Section 6 uses this model to explore the contributions of various explanations for the black-white wealth gap, and Section 7 concludes.

This paper is preliminary. Sections 5, 6 and 7 are particularly incomplete, and should only be read as an indication of the objectives and future direction of this paper.

## 2 Data

The primary data source I use is the Panel Survey of Income Dynamics (PSID), a longitudinal survey that has followed families and their offshoots since 1968. The PSID is an attractive data source for several reasons, including its panel dimension and its detailed data on wealth and demographics. The PSID was annual until 1997, and biannual after; the analysis in this paper uses all available PSID waves. One major advantage of the PSID is its relatively good coverage of wealth data, which was collected in 1984, 1989, 1994, and every two years since 1999.

The main sample of the PSID was designed to be nationally representative, but the Survey of Economic Opportunity (SEO) subsample oversamples low-income households. Since sample size is an issue, the SEO data is included in my analysis in order to increase the number of black families observed. This means the data is not nationally representative, so family-level weights provided by the PSID are used whenever necessary. These weights are intended to make the sample nationally representative within but not necessarily across years, so I weight all years equally.

I keep only data from households with white or black heads, between the ages of 23 and 65. To remove outliers, observations with the head's labor income, wife's labor income, total family income, head's hourly wage, or any category of wealth in the bottom or top percentile are dropped.

All values are expressed in constant 2010 dollars. The closest empirical counterpart to the exogenous income process in the model is total family non-asset after-tax income. Therefore I subtract asset income from total family income reported in the PSID. An estimate of each's family's tax liabilities, from the NBER's TAXSIM program, is taken from Heathcote et al. (2010). Since these estimates are only provided through 2002, I estimate liabilities for future years with a regression of tax liabilities on a cubic in income, the number of children, marital status, and dummies for region of residence. Future versions of this paper will use the TAXSIM program to replace the estimates from this regression.

I divide the assets in the data into three categories, (1) liquid, (2) illiquid, and (3) home equity, which correspond to the asset categories in my model. Liquid wealth is defined as money in checking or savings accounts, minus nonmortgage debt. Home equity is calculated as the value of the home

minus the value of the mortgage outstanding on it. Illiquid wealth refers to any one of five asset categories: (1) farms or businesses, (2) stocks, (3) annuities or IRAs, (4) other real estate, or (5) other assets.<sup>3</sup> The only asset class in the PSID I exclude from my analysis is vehicles, as they depreciate quickly in value and are more similar to durable consumption than wealth.

The PSID does not have good data on interest rates or mortality, so I supplement it with data from the 2010 Survey of Consumer Finances (SCF) and the National Longitudinal Mortality Survey (NLMS), as described in Section 5.

### 3 Documenting the Wealth Gap

One obvious, and important, explanation for the black-white wealth gap is that blacks earn less than whites do. They are also less likely to get married when single, and more likely to get divorced or widowed when married. There are also many other significant observable differences between blacks and whites, such as family size and geographic location, that could potentially explain the wealth gap. A series of papers, including Blau and Graham (1990), Altonji et al. (2000), Scholz and Levine (2004) and Altonji and Doraszelski (2005), have explored how much of the wealth gap income and demographics can explain in a regression framework. Following the most recent and thorough of these, Altonji and Doraszelski (2005), I estimate a version of the following statistical model:

$$W = \beta_0 + \beta_1 Y + \beta_2 X + \epsilon \tag{1}$$

where  $W$  is net household wealth, including home equity, and  $Y$  and  $X$  are vectors of relevant income and demographic variables, respectively. The full list of control variables is extensive and very similar to those used in Altonji and Doraszelski (2005), so I provide it in the appendix.

I then estimate the parameters of the model separately by race and household type (i.e. married couple, single male, or single female). The standard way to determine the amount of the black-white wealth gap explained by these control variables is to take its Blinder-Oaxaca decomposition, which separates the gap into the portion predicted by the regression and the portion that is not. Define  $W_B^A$  as the predicted average wealth level of group  $B$ , using the coefficients from the regression on group  $A$ . The Blinder-Oaxaca decomposition of the racial wealth gap into its “explained” and “unexplained” components, using the regression on group  $A \in \{W, B\}$ , is then  $(W_W^A - W_B^A)/(W_W^W - W_B^B)$ . Table 1 presents the results. For comparison, results broken down by gender and marital status, alongside comparable results from Altonji and Doraszelski (2005), are provided in the appendix.

<sup>3</sup>Examples of other assets provided by the PSID are “bond funds, cash value in a life insurance policy, a valuable collection for investment purposes, or rights in a trust or estate”.

Table 1: The Wealth Gap

Whites		Blacks		Gap \$	% Explained	
$\hat{W}_W^W$	$\hat{W}_W^B$	$\hat{W}_B^W$	$\hat{W}_B^B$		White	Black
\$145,626	\$75,816	\$67,975	\$45,468	\$100,158	77.5%	30.3%

Notes:  $\hat{W}_S^W$  is the average predicted wealth level, according to the white regression, for sample group  $S$ , which can be  $W$  for whites or  $B$  for blacks.  $\hat{W}_S^B$  is the average predicted wealth level, according to the black regression, for sample group  $S$ . The gap is defined as  $\hat{W}_W^W - \hat{W}_B^B$ , while the part of the gap that is explained by regression  $S$  is defined as  $\hat{W}_W^S - \hat{W}_B^S$ .

The biggest lesson of Table 1 is that the white regression explains much more of the gap than the black regression does. Before exploring this point - which has been noted many times before, by Altonji and Doraszelski (2005) and others - two points are worth stressing.

First, at \$100,158, the sheer size of the wealth gap means that even the white regression misses a great deal of wealth. The average black household has \$45,468 in net wealth but “should” have \$67,975 (i.e. 49.5% more), according to the white regression model. Second, Table 1 almost certainly understates the magnitude of the wealth puzzle, since many of the control variables in Equation 1 may be endogenous.

Why does the white regression explain so much more of the gap than the black regression? The direct implication is that giving whites the black distribution of income and demographics would reduce the wealth gap by much more than would giving blacks the white distribution of income and demographics. Altonji and Doraszelski interpret this to mean that black wealth levels respond much less to changes in income and demographics than white wealth levels do.

Another and perhaps more interesting way of phrasing the same point is that the black-white wealth gap is mainly a middle- and upper-class phenomenon. Table 2 displays predicted wealth levels for both races, broken up by age-specific income terciles.

Table 2: The Wealth Gap by Income

Income Bracket	Whites		Blacks		Gap \$	% Explained	
	$\hat{W}_W^W$	$\hat{W}_W^B$	$\hat{W}_B^W$	$\hat{W}_B^B$		White	Black
Bottom 33%	\$83,129	\$38,017	\$24,046	\$19,382	\$63,747	92.7%	29.2%
Middle 33%	\$114,680	\$59,817	\$63,262	\$43,886	\$70,795	72.6%	22.5%
Upper 33%	\$191,825	\$102,077	\$118,912	\$74,457	\$117,368	62.1%	23.5%

Notes:  $\hat{W}_S^W$  is the average predicted wealth level, according to the white regression, for sample group  $S$ , which can be  $W$  for whites or  $B$  for blacks.  $\hat{W}_S^B$  is the average predicted wealth level, according to the black regression, for sample group  $S$ . The gap is defined as  $\hat{W}_W^W - \hat{W}_B^B$ , while the part of the gap that is explained by regression  $S$  is defined as  $\hat{W}_W^S - \hat{W}_B^S$ . Income terciles do not condition on race, but do condition on age, e.g. a household headed by a 23 year old whose income places it the top 33% of all households headed by 23 year olds is placed in the top income tercile. This is done to control for the correlation between income and wealth induced by age.

The key point is that income and demographics explain almost all of the considerable wealth gap for those on the bottom of the income distribution, but as income increases the unexplained portion of the gap gets larger, both in absolute and proportional terms. This means that potential

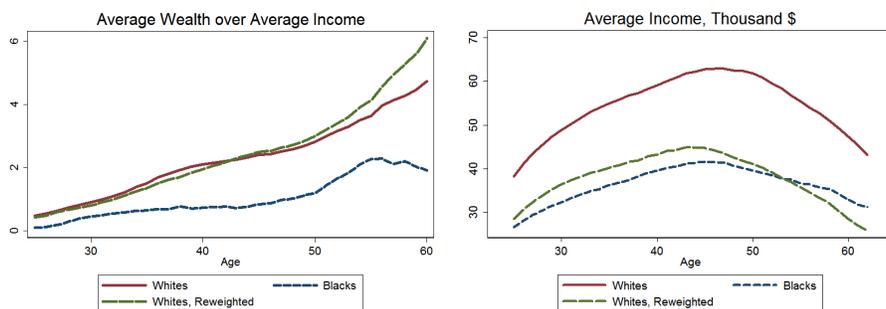
explanations for the wealth gap that apply primarily to low-income families, such as the conspicuous consumption argument advanced by Charles et al. (2009), are not promising candidates to explain the black-white wealth gap.

The higher estimated response of white wealth to income led Barsky et al. (2002) to suspect model misspecification could be at the root of the apparent black-white wealth gap. Specifically, if wealth is a convex function of income, and whites on average have higher income, then a linear regression will attribute a higher effect of income on wealth to whites even if both races have the same wealth function. Barsky et al. (2002) use nonparametric techniques to reweight the white income distribution so that it approximately matches the black income distribution, and find they eliminate 64% of the wealth gap. I replicate their methodology on my larger sample, and obtain a very similar result of 62.7%.

Does this mean that the income distribution, rather than just its mean, can explain the wealth gap? It is worth remembering that the counterfactual conducted by Barsky et al. (2002) - changing white characteristics to match black characteristics - is the same counterfactual performed by the white model in Table 1. It may therefore not be surprising that their conclusions are similar, especially to the extent that income and other demographic variables favoring wealth accumulation, such as education or marital status, may be correlated. Still, their point is worth pursuing.

Recall that Figure 1, presented at the beginning of this paper, controls for average income at each age by dividing average wealth by average income. It therefore controls for the first moment of the income distribution, but no more. The left panel of Figure 2 presents these results again, with an additional line that uses the same nonparametric technique as Barsky et al. (2002) to control for the entire income distribution.

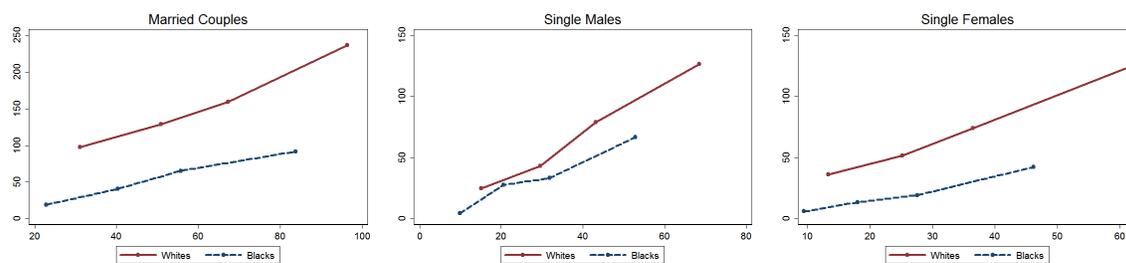
Figure 2: Changing the White Distribution of Income to Match the Black Distribution



Notes: New weights for white families are calculated so that the reweighted white distribution of income approximately matches the black distribution of income at each age. I use the same nonparametric kernel density technique as Barsky et al. (2002), with the same Epanechnikov kernel function and the same bandwidth.

Figure 2 strongly suggests, perhaps surprisingly, that the additional explanatory power provided by controlling for the entire income distribution is negligible. To understand why, note that the power of the convexity argument depends on three things: (1) a highly convex wealth function, (2) few blacks near the top of the income distribution, and (3) black and white wealth functions which almost agree over their common support. Figure 3 plots average wealth against average income for each of four race- and age-specific income quartiles. Note there is actually little evidence for any of these assumptions.

Figure 3: Average Wealth as a Function of Average Income, Thousand \$



Notes: The y-axis is wealth; the x-axis is income. Both axes are in thousands of dollars. Each income quartile includes households whose income falls within that quartile by race *and* age, e.g. a household headed by a black 23 year old whose income places it the top 25% of all households headed by black 23 year olds is placed in the top income quartile. This is done to control for the correlation between income and wealth induced by age.

Figure 3 also illustrates the size of the wealth gap. Note, for example, that the average high-earning (top 25%) married black household makes \$83,741 a year, and has \$91,782 in net wealth. The average low-earning (bottom 25%) married white household earns only \$31,088 a year, and yet has \$97,475 in assets - i.e. \$5,6931 *more*.

Besides Barsky et al. (2002), there is to the best of my knowledge only one other paper that claims to have made significant progress in explaining the black-white wealth gap. White (2007) constructs the first structural model of black-white wealth inequality and finds that he can explain the entire wealth gap. The driving force in his model is expected convergence between the income levels of infinitely-lived black and white dynaties; blacks expect (and receive) faster income growth than whites as the legacies of racism fade. This means that, at the equilibrium interest rate, blacks would like to borrow while whites would like to save. The model is elegant: it simultaneously matches the black-white income and wealth gaps, using only different initial conditions at Emancipation and segregated schooling. Unfortunately, it makes several strongly counterfactual predictions. A lifecycle version of the model would seem to predict that the wealth gap should increase most quickly at the beginning of the lifecycle, when expected income growth matters most. Figure 2 does not bear this prediction out. Furthermore, the key assumption in White (2007) - that black income levels are converging to white income levels - is questionable, as there is substantial evidence that black-white

income convergence has slowed or stopped.<sup>4</sup> Moreover, since the model depends on blacks borrowing against, and whites saving for, their children’s future, it would seem to predict a significantly smaller black-white wealth gap for childless families. The data, if anything, suggest the opposite.<sup>5</sup>

There is therefore little quantitative understanding of why the black-white wealth gap exists, despite an almost embarrassing number of potential explanations. The next section begins to address this problem, by building a realistic structural model intended to quantify the most promising explanations for the wealth gap.

## 4 The Model

This paper will eventually quantify the wealth effects of racial differences in:

1. Income and family size
2. Bequests
3. Property market terms
4. Mortality rates
5. Marriage and divorce rates
6. Financial literacy

This section constructs a realistic lifecycle model of asset accumulation that incorporates all of these factors. Note that even the explanations that have been studied in the literature before, (1) and (2), may have different implications in a structural model than in a regression. For example, a structural model can quantify the effect of racial differences in entire income *profiles*, rather than just current income and a measure of permanent income at each age.

Also note that almost all of the explanations listed above disproportionately affect incentives to save in illiquid assets or housing. The availability of liquid assets, which are an imperfect substitute, reduces the impact of these factors on wealth accumulation. Future versions of this paper will examine the consequences of endowing both blacks and whites with quasi-hyperbolic (a.k.a.  $\beta - \delta$ ) preferences. These preferences are time-inconsistent; future selves will consume more, and spend less, than the current self would like. Agents can save in housing or illiquid assets to constrain the

---

<sup>4</sup>One potential reason for the slowdown of convergence is the rise in the skill premium over the last few decades. For more information, see Chay and Lee (2000) and Neal (2006).

<sup>5</sup>Re-estimating Equation 1 for married couples, replacing all child and dependent variables with one dummy variable that is true if either the head or the spouse has ever had or adopted a child, yields an estimated increase in wealth of \$4,975 for whites and \$52,272 for blacks, a difference that is significant at the 5% level. In other words, the black-white wealth gap is on average \$47,297 smaller for couples *with* children.

consumption of future selves, but if housing or illiquid assets become less attractive - e.g. through the racial differences discussed above - then agents may simply decide to forgo saving altogether, rather than saving in liquid assets that will be squandered by future selves. Quasi-hyperbolic preferences should therefore magnify the effects of the factors listed above.<sup>6</sup> Future versions of this paper will present two sets of results, the first with exponential and the second with quasi-hyperbolic preferences. However, as this paper is preliminary and quasi-hyperbolic preferences complicate the exposition considerably, for now I restrict attention to exponential preferences.

The basic framework is a model of consumption, savings, and portfolio choice with many ingredients that are standard in the literature. Onto this I add several extensions that are designed to incorporate the explanations for the black-white wealth gap listed above.

## 4.1 The Basic Framework

Time is discrete. Consumers receive an exogenous, stochastic income flow  $\{y_t\}$  for their working lives, from periods  $t = T_{born}$  until retirement in period  $T_{retire}$ . After retirement, they receive an exogenous, nonstochastic income until certain death in period  $T_{die}$ .

Consumers value consumption  $c$ . They also care about the flow value of where they live, which can either be value derived from the house they own,  $\kappa h'$ , or the house they rent,  $r$ . To reflect the economies of scale in a household, consumption and housing are both deflated by family size  $e$ , which is a deterministic function of time. Period utility over consumption and housing is therefore given by  $u(\frac{c}{e}, \frac{\max\{r, \kappa h'\}}{e})$ .

Agents have access to three different kinds of assets. The first kind of assets are called “liquid assets”, denoted  $l$ , and are meant to correspond to checking and savings accounts. Savings in liquid assets earn a rate of return  $R^l$ . Liquid assets can also be borrowed (up to a limit defined below), at an interest rate that begins at  $R^l$  and increases continuously in the amount borrowed thereafter. The continuity assumption ensures that the Euler Equation is always necessary for an optimum, and is made for computational tractability. While it cannot necessarily be justified on theoretical grounds, standard measure-theoretic arguments show that it is quantitatively harmless, and it is certainly more realistic than an assumption often made in other lifecycle models, e.g. in Li et al. (2012) and Attanasio et al. (2012), that the interest rate on debt is *equal* to  $R^l$ .<sup>7</sup> Subject to the borrowing limit, liquid assets can be spent freely on consumption or savings.

The second kind of asset is “housing”, denoted  $h'$ , which is also modeled after much of the

<sup>6</sup>For a much more detailed discussion of quasi-hyperbolic preferences, see Laibson (1998), Harris and Laibson (2001), and Angeletos et al. (2001).

<sup>7</sup>Davis et al. (2006) find in a lifecycle savings model that matching the data requires borrowing costs that exceed the risk-free rate on savings; otherwise young agents will borrow too much to invest in high-return assets.

literature, e.g. Cocco (2004), Cocco et al. (2005), Yao and Zhang (2005), Li and Yao (2007), Yang (2009), Bajari et al. (2010), Attanasio et al. (2012), and Li et al. (2012). Since housing enters the utility function directly, a major benefit of owning a home is that the owner can live in it.

Another major benefit of owning a home is that it can be sold later. The price of housing at time  $t$ ,  $p_h(t)$ , increases at a constant rate  $R^h$ . Any time an agent moves, he must pay a proportion  $\phi_h$  of the value of both his new and old housing stock as a sunk cost, which is intended to represent both broker fees and moving costs. Therefore the total cost of adjusting from housing  $h$  to housing  $h'$  in period  $t$  is given by:

$$\psi_h(h, h', t) = p_h(t)(h' - h) + \phi_h[h' + h] \quad (2)$$

The final major benefit of owning a home is that it can be used as collateral. Agents can borrow up to a proportion  $\xi_h$  of the value of their home; they can also borrow up to a proportion  $\xi_y$  of their minimum income  $\underline{y}'$  next period. Therefore the borrowing constraint is given by:

$$l' \geq -\xi_y \underline{y}' - \xi_h h' \quad (3)$$

For simplicity and following much of the literature, I assume that the housing space is discrete, with  $nH + 1$  possible values.

$$h' \in \mathcal{H} \equiv \{0, H_1, \dots, H_{nH}\} \quad (4)$$

Finally, agents can buy “illiquid assets”, which are denoted  $i$  and behave very similarly to the illiquid assets in other models, e.g. Angeletos et al. (2001) and Kaplan and Violante (2012). With PSID data, illiquid assets will be identified as assets in any of 5 categories: (1) stocks, (2) IRAs, (3) businesses or farms, (4) real estate other than primary residence, and (5) other. The price of an illiquid asset at time  $t$  is given by  $p_i(t)$ , which increases at a constant rate  $R^i > R^l$  every period. Agents must pay a fixed cost,  $\Gamma$ , to buy or sell illiquid assets, as well as a proportion  $\phi_i$  of the value of illiquid assets that are traded. This means that the total cost of adjusting from  $i$  illiquid assets to  $i'$  illiquid assets in period  $t$  is given by:

$$\psi_i(i, i', t) = p_i(t)(i' - i) + \phi_i|i' - i| + \Gamma I_{i' \neq i} \quad (5)$$

I assume that the illiquid asset space is also discrete, with  $nI + 1$  possible values.

$$i' \in \mathcal{I} \equiv \{0, I_1, \dots, I_{nI}\} \quad (6)$$

The assumptions made above imply that the budget constraint is given by:

$$\frac{1}{R_l(l')}l' + c + r + \psi_i(i, i') + \psi_h(h, h') = l + y \quad (7)$$

Finally, I assume that illiquid assets become liquid for free at retirement, and can thereafter no longer be purchased. This assumption directly models the fact that many illiquid assets, such as 401(k)s, are in fact freely accessible after retirement. It is also intended to model in a simple way the complex financial planning that can occur with other illiquid assets, such as equities or businesses, to ensure they are available at or soon after retirement.<sup>8</sup> This assumption also eliminates a state variable for a substantial portion of the lifecycle, significantly speeding up the code.

## 4.2 Extensions

**Motivation:** Recall that the explanations for the black-white wealth gap that this paper will consider are racial differences in:

1. Income and family size
2. Bequests
3. Property market terms
4. Mortality rates
5. Marriage and divorce rates
6. Financial literacy

Explanations (1), (3), and (6) can be directly parameterized in the basic framework detailed above, but quantifying the impact of the others requires additional model ingredients. Introducing these ingredients is the goal of this subsection.

**Bequests:** First, agents face a small age, race, and gender specific probability  $\pi_b$  every period of receiving a large exogenous bequest every period. Bequests are folded in to the income process,  $\{y_t\}$ .

**Stochastic death:** Second, stochastic death can occur before the final period  $T_{die}$ , with an age, race, and gender-specific probability  $\pi_d(t)$ . A bequest motive provides a value function,  $B$ , on the event of death.<sup>9</sup>

---

<sup>8</sup>For example, consumers often shift equities from high-risk portfolios to low-risk ones as retirement approaches.

<sup>9</sup> $B$  can also be motivated with the assumption that the agent knows ahead of time that he will die, and has some amount of time to consume his remaining assets.

**Marriage & Divorce:** A much more significant addition to the baseline framework is exogenous marriage and divorce. Single males, single females, and married couples are each given their own income and family size process. Transition probabilities to and from marriage are gender- and age-specific. As in Fernández & Wong (2011), the continuation value of married couples is taken to be the weighted average of the individual continuation values.

An important question is how to model the transition of assets through marriage. Computing the value of marriage requires predicting the asset level of the spouse an agent will find. Fernández & Wong (2011) assume that singles know the characteristics of the person they will marry, if they marry; this allows agents to predict their asset level after marriage, at the expense of an additional state variable. Voena (2012) makes the even simpler assumption that agents marry other agents with identical asset positions. While computationally very tractable, this assumption does not respect gender differences in wealth, as Fernández & Wong (2011) point out.

I make an assumption that respects gender differences in wealth, but does not introduce another state variable. Let  $\Lambda(t)$  denote the ratio of median male wealth to median female wealth at age  $t$ . To maintain tractability while coming closer to matching gender differences in wealth, I assume that for every dollar in liquid assets, illiquid assets, and housing that a woman brings to a marriage, a man brings  $\Lambda(t)$  dollars. Since the housing and illiquid asset spaces are discrete (and a couple must choose which house to live in), I assume that the newlywed couple receives the house and illiquid assets of the current agent, but also receive the proceeds from selling a house worth  $\hat{\Lambda}(t)h$  and illiquid assets worth  $\hat{\Lambda}(t)i$ , where  $\hat{\Lambda} = \Lambda$  for women and  $\hat{\Lambda} = \Lambda^{-1}$  for men.

An even more difficult question is what happens to assets after divorce. There is surprisingly little quantitative research on the question, despite the large asset losses that could occur in divorce because of litigation costs and the difficulty of splitting some assets, like houses. The majority of economic models with divorce, including Fernández and Wong (2011) and Voena (2012), assume that it does not destroy any assets at all. Cubeddu and Ríos-Rull (2003) allow for a proportion of assets to be lost in divorce, but do not take a stand on the size of this proportion. The issue is further complicated by the fact that almost all previous papers with assets and divorce have only one asset class, instead of the three considered in this paper. A serious quantitative analysis of the costs of divorce, allowing for different costs incurred by different asset classes, is a fascinating but daunting topic that is beyond the scope of this paper.

Therefore I make assumptions intended to provide a reasonable lower bound on the amount of wealth lost in divorce. First I assume that the costs of splitting illiquid assets in litigation equals or exceeds the cost of liquidating those assets; this means the divorcing couple liquidates the assets themselves. The other assumption is that courts order a divorcing couple to sell their home, and to

split the proceeds and all remaining wealth equally. The agents are allowed to negotiate in order for one of the agents to keep the home, but only if the keeper pays the other half of the value of the home. Given this arrangement, if both agents would prefer to keep the home, each receives it with probability one half.<sup>10</sup>

Widowhood is also allowed in the model, but is much simpler and involves no loss of assets.

## 5 Estimation

This section details the procedure used to parameterize the model.

### 5.1 Estimation

Eventually, I will estimate the parameters of the model using the Simulated Method of Moments (SMM). Time constraints mean that, for the moment, this is not feasible, so to establish preliminary results I calibrate the model as described below. As is standard, the first stage of calibration sets the values of some parameters outside the model. The second stage of calibration chooses the values for other parameters so that the predictions of the model match empirical moments as closely as possible.

#### 5.1.1 First Stage

Many of the parameters used in the model have standard values in the literature, or can be estimated from the data without the use of a structural model. The values chosen for these parameters are described below.

**Demographics:** Households begin life at age 23 and live to a maximum age of 90. The mandatory retirement age is 65.

The effective family size (i.e. the consumption and housing deflator) is calculated for each observation as in Li et al. (2012). If  $N_c$  is the number of children in a household and  $I_m$  is an indicator for a household headed by a married couple, then the deflator for a household is  $(1 + I_m + .7 \cdot N_c)$ .<sup>7</sup> Age-, race-, and gender-specific empirical averages of this value are taken from the data and used as the deflator values facing agents in the model.

Similarly, age-, race-, and gender-specific marriage probabilities are taken as the empirical averages in the PSID, as are age- and race-specific divorce probabilities. Age- and race-specific

---

<sup>10</sup>Allowing for an efficient solution, such as Nash bargaining, would be computationally very expensive as the outcome would have to be computed at every point in the state space for married couples, even when divorce probabilities are low. It also seems very reasonable to assume that divorce proceedings do not lead to Pareto optimal outcomes.

gender wealth ratios are also taken from the data, as the ratio of median male wealth to median female wealth.

**Initial Distribution of State Variables:** At the beginning of life, agents receive an income shock drawn from the stationary distribution. Marital status, gender, and the level of liquid assets, illiquid assets, and home equity are drawn together by endowing the agent with the states of a randomly-chosen 23 year old individual in the PSID.

**Savings and Prices:** The real interest rate on liquid savings is set at 2.7 percent, the same value used in Li et al. (2012) and just slightly lower than the value used in Bajari et al. (2010). The real appreciation rate of home prices is set to 0, which again matches the value in Li et al. (2012) and is slightly lower than the .178 percent used in Bajari et al. (2010). The real interest rate on illiquid assets is set to match the average return on equities of 6.7 percent in Gomes and Michaelides (2005), which is very similar to the 7.0 percent that Kaplan and Violante (2012) calculate. Future versions of this paper will use the Survey of Consumer Finances (SCF) to set the rates of return on these assets.

The rental price of a home is set to 7.5% of its value, which is the rate used in Li and Yao (2007).

The maximum interest rate on debt is set to 5.8%, which is the empirical average rate paid on debt by respondents in the 2010 SCF. The interest rate on debt increases linearly from the rate on savings,  $R^l = 2.7\%$ , to this maximum rate, as the amount of debt increases from 0 to \$10,000, and is constant thereafter. This functional form is chosen both for its simplicity and to approximate a constant interest rate on debt.

**Debt:** The proportion of labor income that can be borrowed against,  $\xi_1$ , is set to .2 for now, following Heathcote, Storesletten, & Violante (2010). The proportion of home equity that can be borrowed against,  $\xi_2$ , is set to .8, following Bajari et al. (2010) and Li et al. (2012). Later versions of this paper may estimate the values of these parameters.

**Mortality:** Mortality data comes from the most recently available National Longitudinal Mortality Survey (NLMS), from 2008. This was intended to be a sample representative of the non-incarcerated U.S. population on March 1, 1983. Though it is now somewhat dated, I chose to use this dataset because of its good data on the income and demographics of the respondents.

The initial surveys of the NLMS occurred in 11 separate waves between 1979 and 1987. Respondents were followed for 11 years after their initial survey. If they died in this timeframe, their age of death

was recorded. The age-, race-, and gender-specific mortality rate is assumed to be the proportion of NLMS respondents of that age, race, and gender who were observed to die at a given age.

The bequest motive is turned off for now by setting the bequest function equal to 0 everywhere, but future versions of this paper may consider nontrivial bequest functions, like the one used in Li et al. (2012).

**Income:** Following much of the lifecycle literature, I assume that labor income follows a deterministic trend but is subject to transitory and persistent idiosyncratic shocks. Specifically,

$$\log(y_t) = g_t + z_t + \epsilon_t \tag{8}$$

where  $g_t$  is the deterministic component of income and  $\epsilon_t$  is the transitory shock. The persistent shock  $z_t$  follows an AR(1) process with persistence  $\rho$ ,

$$z_t = \rho z_{t-1} + \eta_t$$

$\epsilon_t$  and  $\eta_t$  are normal random variables with mean 0 and variances  $\sigma_\epsilon^2$  and  $\sigma_\eta^2$ , respectively. I estimate the parameters of the process ( $\rho$ ,  $\sigma_\epsilon^2$ , and  $\sigma_\eta^2$ ) by Minimum Distance Estimation, the method used by Guvenen (2009), Fernández and Wong (2011), Kaplan and Violante (2012) and many others. Because this process is so standard, I leave the formal details for the appendix. Table 3 presents the results.

Table 3: Estimated Parameters of the Income Process

$\hat{\rho}$	$\hat{\sigma}_\eta^2$	$\hat{\sigma}_\epsilon^2$
.965	.008	.056
(.018)	(.003)	(.010)

Notes: Standard errors are in parentheses.

These estimates are well in line with others in the literature, though they are all somewhat low. There are several potential explanations for my low estimates, as my sample differs in some details from others in the literature, most notably the inclusion of households headed by single males and especially single females. It is possible, for example, that singles face less income uncertainty than married couples, though this seems unlikely given the risk sharing that is often assumed to occur between married couples. An explanation I prefer is described in detail in Guvenen (2009); income profiles (absent any shocks) may in fact be heterogeneous, while the estimation procedure restricts

them not to be. This forces the estimation to treat heterogeneity as shocks.<sup>11</sup> Since my estimation procedure allows for profiles to differ by household type and race, it allows for more heterogeneity than usual, which in turn may reduce the portion of income variance that is attributed to shocks.<sup>12</sup>

Age-, race-, and marital status-specific income profiles are taken as the empirical median.

Income after retirement follows Fernández and Wong (2011). Specifically, let  $y_r$  denote an individual's income in the last period before retirement (ignoring temporary shocks). Let  $\bar{y}_r$  denote the empirical average of income at age 64. An individual's income after retirement is then 90% of  $y_r$  up until a threshold level of  $.38 \bar{y}_r$ , plus 32% of  $y_r$  until a higher threshold of  $1.59 \bar{y}_r$ , plus 15% of any remaining income after this point.

The chance of receiving any bequest at any age is set to 0 now. Future versions of this paper will calibrate the bequest process from the data.

**Utility:** The utility function is CES between consumption and housing, and CRRA over time. It is the same as in Li et al. (2012):

$$u(c_t, h_t, n_t) = \frac{(\omega (\frac{c_t}{n_t})^{\frac{\theta-1}{\theta}} + (1-\omega) (\frac{h_t}{n_t})^{\frac{\theta-1}{\theta}})^{\frac{(1-\gamma)(\theta)}{\theta-1}}}{1-\gamma} \quad (9)$$

This CES specification is close to the Cobb-Douglas specification used in most of the literature, e.g. Li and Yao (2007), but is more flexible in that it allows the elasticity of substitution between consumption and housing,  $\theta$ , to be different than one.<sup>13</sup> When  $\theta$  is less than one - as estimated by Bajari et al. (2010) and Li et al. (2012) - the elasticity of substitution between consumption and housing is low, so an increase in the price of housing leads to an increase in the total amount spent on housing. Accounting for this is especially important when assessing the importance of factors that may make it effectively more difficult for blacks to own homes. However, for the purposes of this draft, I set  $\theta = 1$ ; this means that the Euler equation is invertible in closed form, which speeds up the modified Endogenous Grid Method I use considerably. As it is a crucial parameter, future versions of this paper will estimate  $\theta$ .

---

<sup>11</sup>Heterogeneity in growth rates will be interpreted as very persistent idiosyncratic shocks. As detailed in Guvenen (2009), failing to account for this heterogeneity will bias  $\hat{\rho}$  upwards. The direction of the bias for  $\hat{\sigma}_\eta^2$  is not clear, and depends on the relative variance of profile heterogeneity and true persistent shocks. This may explain while accounting for heterogeneity moves Guvenen's  $\hat{\sigma}_\eta^2$  upwards and mine downwards.

<sup>12</sup>Allowing for only one profile, common to households of all types and both races, changes the estimates to  $\hat{\rho} = 0.969$ ,  $\hat{\sigma}_\eta^2 = 0.011$ , and  $\hat{\sigma}_\epsilon^2 = 0.0718$ , which are all much closer to typical values in the literature.

<sup>13</sup>See Bajari et al. (2010) and Li et al. (2012) for further discussion.

### 5.1.2 Second Stage

I calibrate six parameters inside my model. The first three are utility parameters: the discount factor,  $\beta$ , the coefficient of relative risk aversion  $\gamma$ , and the weight of consumption in the utility function,  $\omega$ . The other three parameters involve the costs of adjusting asset levels:  $\phi_h$  and  $\phi_i$ , the proportional costs associated with changing housing and illiquid asset levels, respectively, and  $\Gamma$ , which is the fixed price agents must pay to change their illiquid asset levels.

The moments I target are white lifecycle profiles in:

1. average wealth
2. average liquid wealth
3. average illiquid wealth
4. average home equity
5. median wealth
6. proportion of agents with a home
7. proportion of agents with illiquid assets

Results are reported in Table 4.

Table 4: Calibrated Parameter Values

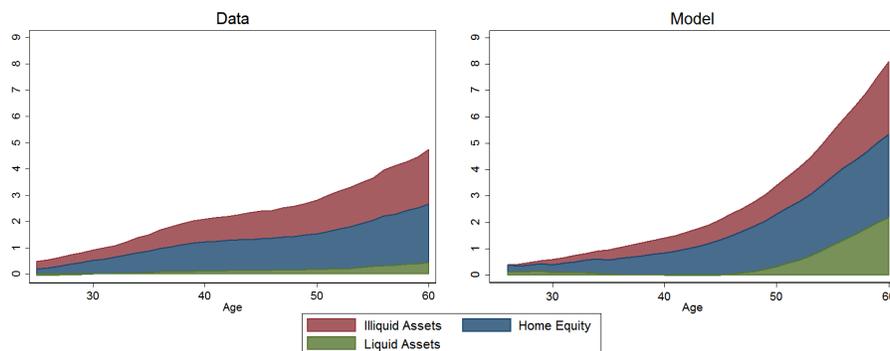
Parameter	Calibrated Value	Interpretation
$\beta$	.964	Discount factor
$\omega$	.810	Utility weight of consumption
$\gamma$	2.03	Coefficient of relative risk aversion
$\phi_h$	.146	Proportional cost of adjusting housing
$\phi_i$	.100	Proportional cost of adjusting illiquid assets
$\Gamma$	1.01	Fixed cost of adjusting illiquid assets

These values are all suspiciously close to the starting guesses. The minimization algorithm I use is the Nelder-Mead simplex algorithm, which is often not effective at finding global minima far from a starting guess. The algorithm instead probably settled in a local minimum. Future versions of this paper will employ a more robust minimization routine - such as iterated grid search starting from a Halton sequence of points, or the Markov Chain Monte Carlo (MCMC) approach detailed in Chernozhukov and Hong (2003) - to find the global minimum. For now, the results presented in Section 6 should be treated as very preliminary, and indicative only of the objectives and future direction of this paper.

## 6 Results

Figure 4 presents the average levels of liquid wealth, illiquid wealth, and home equity over the lifecycle for whites, from the data and as predicted by the model.

Figure 4: Average Wealth to Average Income Ratio



Under the current calibration, the model underpredicts wealth accumulation early in life, and significantly overpredicts wealth accumulation near retirement. There are two likely reasons why.

The first, as already mentioned, is the fact that the minimization routine did not find the global minimum; there are almost certainly different parameter values that can do a better job of matching the data.

The second is the current restriction on the elasticity of substitution between consumption and housing,  $\theta$ , to be 1. Bajari et al. (2010) and Li et al. (2012) estimate a much lower value of around .33. When  $\theta = 1$  young agents are willing to rent relatively small homes, so they can save in order to purchase a large house when possible and pay moving costs only once. When  $\theta$  is significantly lower, agents are unwilling to live in small homes but also prefer to avoid paying the money to rent large ones, so they move earlier and more often. This should increase wealth accumulation early in life, especially in home equity, and decrease it later.

Still, it may be interesting to examine the predictions of the calibrated model. Recall that when estimating the size of the wealth gap that can be explained, there are two main counterfactuals to consider. The first adjusts the environment faced by whites to that faced by blacks, while the second does the reverse. Results in this paper will be presented in the first counterfactual, for two main reasons. The first is practical; whites make up a much larger share of the population, so data on their wealth and processes is more trustworthy. The baseline calibration, against which all other results are presented, targets whites partly because it needs to be reliable. The second reason is that the interesting policy question is not why whites have so much wealth but why blacks have so little, which is answered by the first counterfactual.

Results in the literature are typically presented in terms of the percent of the wealth gap that certain factors can explain. Therefore, as an example of the kind of results that can be expected from future versions of this paper, Table 5 presents the percent of the wealth gap at age 60 that can be explained by different experiments. Since I intentionally do not calibrate a model to black data, my definition of the wealth gap is the difference between wealth predicted by the model at age 60, and 40.8% of this wealth - which is the empirical ratio of average black wealth to average white wealth at age 60. The results answer the question: how much do each of these experiments lower average wealth at age 60, expressed as a percent of the total wealth gap?

I conduct five experiments in total. The first three of these experiments each change one process to match the process estimated for blacks in the data, while leaving all other processes and parameters at the levels for whites. The fourth experiment increases the price at which agents can buy (but not sell) houses by 3%, which is the premium that Bayer et al. (2012) estimate blacks pay for houses. The last experiment makes all four of these changes simultaneously.

Table 5: Percent of Gap at Age 60 Explained

↓ Income	↓ Marriage & ↑ Divorce	↑ Mortality	↑ Home Prices	All
29.5%	17.9%	7.5%	2.6%	58.6%

Though clearly it is too early to take the results of Table 5 seriously, they are still interesting.

For example, note that while income seems to be the most important contributor to the wealth gap, my estimate of its effect is somewhat lower than in previous non-structural studies. Menchik and Jianakoplos (1997) estimate that differences in permanent income account for between 30 - 72% of the gap for singles, and 37 - 58 % of it for married couples, while Barsky et al. (2002) put the figure at 64%. However, Menchik and Jianakoplos (1997) have only a small number of other control variables, while Barsky et al. (2002) have none at all.<sup>14</sup> It seems likely that these estimates are picking up the effect of other variables, such as education and health, that are highly correlated with income and that also favor wealth accumulation. Other regression-based studies, such as Altonji and Doraszelski (2005), have a much more extensive set of control variables but modest sample sizes do not allow them to separate the effect of income from the effect of their other controls. This is one advantage of the structural approach of this paper. Income growth, which is lower for blacks than for whites, also plays a role in lowering the explanatory power of income by incentivizing young whites to borrow against expected future income. Previous non-structural studies could not account for this effect.

<sup>14</sup>They do control for age in a limited way by restricting their sample to 45-50 year olds.

It also seems that the effect of racial differences in mortality rates, which has never been quantified before, may be non-negligible. The value of 7.5% reported in Table 5 is probably an overestimate, since the model does not yet include a meaningful bequest motive. However, note that married agents in the model (who do most of the wealth accumulation) do care about they spouses they leave behind; it is only singles that do not value wealth after death. Introducing a bequest motive may therefore not lower this estimate as much as would otherwise be expected.

Note the very modest effect of increasing home prices reported in Table 5 is almost certainly an underestimate because of the current restriction on the elasticity of substitution between consumption and housing,  $\theta$ , to be 1. Currently agents in the model are more willing to adjust to an increase in house prices by lowering housing demand than they should be.

## 7 Conclusion

Though clearly preliminary, this paper presented the first structural estimates of the importance of several explanations for the black-white wealth gap - including some never before considered by the literature.

There are several explanations for the wealth gap that this paper does not yet consider but will. Racial differences in financial literacy can be represented by the model in a simple way by adjusting the fixed price,  $\Gamma$ , that agents must pay in order to adjust illiquid assets levels. This adjustment needs to be quantified carefully, however. The effect of racial differences in inheritance rates can also be estimated by the model by introducing inheritance processes calibrated to PSID data.

There are also potentially important barriers that blacks face in the property market that are not adequately reflected in the house price experiment conducted above. There is evidence, for example, that conditional on observables blacks have less access to credit and see a lower rate of return on their houses than whites do. Quantifying these differences carefully, and introducing them into the model to estimate their effects, is an important goal for future versions of this paper.

Quasi-hyperbolic preferences should also prove an interesting addition to the paper, if all of the other factors considered leave a significant portion of the wealth gap unexplained. Despite the considerable evidence in favor of quasi-hyperbolic preferences (e.g. see Angeletos et al. (2001)), they remain controversial, so results with standard exponential preferences will still always be presented.

One explanation for the black-white wealth gap this paper will probably never consider is racial differences in culture or preferences. One reason is that these differences are hard to quantify, but a more important one is that this does not seem to be a promising line of research. Historically, the question was whether or not blacks saved *more* than whites, e.g. Klein and Mooney (1953) and

Galenson (1972). Moreover, what little evidence there is on the subject suggests that blacks should, conditional on observables, probably accumulate more wealth than whites. Benjamin et al. (2010) present experimental evidence that non-immigrant blacks are more patient and more risk-averse than whites; similarly, Sahm (2007) finds in the PSID that blacks are more risk-averse than whites. Higher risk-aversion should lead blacks to accumulate more precautionary wealth, though it could also cause them to participate less in the stock market and hence earn a lower rate of return on their assets. The evidence is not conclusive, and whether blacks do have different savings preferences and whether these preferences matter quantitatively is an open question.

## References

- Albanesi, Stefania and Claudia Olivetti (2009) “Home Production, Market Production, and the Gender Wage Gap: Incentives and Expectations,” *Review of Economic Dynamics*, Vol. 12, pp. 80–107.
- Altonji, Joseph and Ulrich Doraszelski (2005) “The Role of Permanent Income and Demographics in Black/White Differences in Wealth,” *The Journal of Human Resources*, Vol. 40(1), pp. 1–30.
- Altonji, Joseph, Ulrich Doraszelski, and Lewis Segal (2000) “Black/White Differences in Wealth,” *Economic Perspectives*, Vol. 24(1), pp. 38–50.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001) “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, Vol. 15(3), pp. 47–68.
- Attanasio, Orazio, Hamish Low, and Virginia Sánchez-Marcos (2008) “Explaining Changes in Female Labor Supply in a Life-Cycle Model,” *American Economic Review*, Vol. 98:4, pp. 1517–1552.
- Attanasio, Orazio P., Renata Bottazzi, Hamish W. Low, Lars Nesheim, and Matthew Wakefield (2012) “Modelling the Demand for Housing over the Life Cycle,” *Review of Economic Dynamics*, Vol. 15, pp. 1–18.
- Bajari, Patrick, Phoebe Chan, Dirk Krueger, and Daniel Miller (2010) “A Dynamic Model of Housing Demand: Estimation and Policy Implications,” *Working Paper*.
- Barsky, Robert, John Bound, Kerwin Charles, and Joseph Lupton (2002) “Accounting for the Black-White Wealth Gap: A Nonparametric Approach,” *Journal of the American Statistical Association*, Vol. 97(459), pp. 663–673.

- Bayer, Patrick, Marcus D. Casey, Fernando Ferreira, and Robert McMillan (2012) “Price Discrimination in the Housing Market,” *NBER Working Paper 18069*.
- Benjamin, Daniel, James Choi, and A. Joshua Strickland (2010) “Social Identity and Preferences,” *American Economic Review*, Vol. 100, pp. 1913–1928.
- Blau, Francine and John Graham (1990) “Black-White Differences in Wealth and Asset Composition,” *The Quarterly Journal of Economics*, Vol. 105(2), pp. 321–339.
- Charles, Kerwin and Erik Hurst (2002) “The Transition to Home Ownership and the Black-White Wealth Gap,” *The Review of Economics and Statistics*, Vol. 84(2), pp. 281–297.
- Charles, Kerwin Kofi and Erik Hurst (2003) “The Correlation of Wealth across Generations,” *Journal of Political Economy*, Vol. 111(6), pp. 1155–1182.
- Charles, Kerwin Kofi, Erik Hurst, and Nikolai Roussanov (2009) “Conspicuous Consumption and Race,” *The Quarterly Journal of Economics*, Vol. 124(2), pp. 425–467.
- Chay, Kenneth Y. and David S. Lee (2000) “Changes in Relative Wages in the 1980s: Returns to Observed and Unobserved Skills and Black-White Wage Differentials,” *Journal of Econometrics*, Vol. 99, pp. 1–38.
- Chernozhukov, Victor and Han Hong (2003) “An MCMC Approach to Classical Estimation,” *Journal of Econometrics*, Vol. 115, pp. 293–346.
- Chiteji, Ngina S. and Frank P. Stafford (1999) “Portfolio Choices of Parents and Their Children as Young Adults: Asset Accumulation by African-American Families,” *The American Economic Review Papers and Proceedings*, Vol. 89(2), pp. 376–380.
- Cocco, João F. (2004) “Portfolio Choice in the Presence of Housing,” *The Review of Financial Studies*, Vol. 18(2), pp. 535–567.
- ao Cocco, Jo Francisco J. Gomes, and Pascal J. Maenhout (2005) “Consumption and Portfolio Choice over the Life Cycle,” *The Review of Financial Studies*, Vol. 18(2), pp. 491–533.
- Cubeddu, Luis and José-Víctor Ríos-Rull (2003) “Families as Shocks,” *Journal of the European Economic Association*, Vol. 1(2/3), pp. 671–682.
- Davis, Steven J., Felix Kubler, and Paul Willen (2006) “Borrowing Costs and the Demand for Equity over the Life Cycle,” *The Review of Economics and Statistics*, Vol. 88(2).

- Eckstein, Zvi and Osnat Lifschitz (2011) “Dynamic Female Labor Supply,” *Econometrica*, Vol. 79(6), pp. 1675–1726.
- Erosa, Andrés, Luisa Fuster, and Diego Restuccia (2010) “A Quantitative Theory of the Gender Gap in Wages,” *iMdea working papers series*.
- Fella, Giulio (2011) “A Generalized Endogenous Grid Method for Non-Concave Problems,” *Working Paper*.
- Fernández, Raquel (forthcoming) “Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century,” *American Economic Review*.
- Fernández, Raquel and Joyce Cheng Wong (2011) “The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women’s Work,” *NBER Working Paper 17508*.
- Fogli, Alessandra and Laura Veldkamp (2011) “Nature or Nurture? Learning and the Geography of Female Labor Force Participation,” *Econometrica*, Vol. 79(4), pp. 1103–1138.
- Galenson, Marjorie (1972) “Do Blacks Save More?” *The American Economic Review*, Vol. 62(1/2), p. 1972.
- Gayle, George-Levi and Limor Golan (2012) “Estimating a Dynamic Adverse-Selection Model: Labour-Force Experience and the Changing Gender Earnings Gap 1968-1997,” *Review of Economic Studies*, Vol. 79, pp. 227–267.
- Gomes, Francisco and Alexander Michaelides (2005) “Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence,” *The Journal of Finance*, Vol. 55(2), pp. 869–904.
- Guvenen, Fatih (2009) “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, Vol. 12, pp. 58–79.
- Harris, Christopher and David Laibson (2001) “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, Vol. 69(4), pp. 935–957.
- Heathcote, Johnathan, Fabrizio Perri, and Gianluca Violante (2010) “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967 - 2006,” *Review of Economic Dynamics*, Vol. 13(1), pp. 15–51.
- Kaplan, Greg and Gianluca Violante (2012) “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Mimeo*.

- Keane, Michael P. and Kenneth I. Wolpin (2010) “The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women,” *International Economic Review*, Vol. 51(3), pp. 851–892.
- Klein, L. R. and H.W. Mooney (1953) “Negro-White Savings Differentials and the Consumption Function Problem,” *Econometrica*, Vol. 21(3), pp. 425–456.
- Laibson, David (1998) “Life-cycle consumption and hyperbolic discount functions,” *European Economic Review*, Vol. 42, pp. 861–871.
- Li, Wenli, Haiyong Liu, Fang Yang, and Rui Yao (2012) “Housing over Time and over the Life Cycle: A Structural Estimation,” *Working Paper*.
- Li, Wenli and Rui Yao (2007) “The Life-Cycle Effects of House Price Changes,” *Journal of Money, Credit, and Banking*, Vol. 39(6), pp. 1375–1409.
- Menchik, Paul and Nancy Jianakoplos (1997) “Black-White Wealth Inequality: Is Inheritance the Reason?” *Economic Inquiry*, Vol. 35(2), pp. 428–442.
- Neal, Derek (2006) “Why Has Black-White Skill Convergence Stopped?” in Eric Hanushek and Finish Welsh eds. *Handbook of the Economics of Education*: Elsevier.
- Sahm, Claudia (2007) “How Much Does Risk Tolerance Change?” *Finance and Economics Discussion Series*.
- Scholz, John and Kara Levine (2004) “U.S. Black-White Wealth Inequality,” in K. Neckerman ed. *Social Inequality*: Russell Sage Foundation.
- Van Der Klaauw, Wilbert (1996) “Female Labour Supply and Marital Status Decisions: A Life-Cycle Model,” *The Review of Economic Studies*, Vol. 63(2), pp. 199–235.
- Voena, Alessandra (2012) “Yours, Mine, and Ours: Do Divorce Laws Affect the Intertemporal Behavior of Married Couples?” *Working Paper*.
- White, T. Kirk (2007) “Initial conditions at Emancipation: The long-run effect on black-white wealth and earnings inequality,” *Journal of Economic Dynamics & Control*, Vol. 31, pp. 3370–3395.
- Yang, Fang (2009) “Consumption Over the Life Cycle: How Different is Housing?” *Review of Economic Dynamics*, Vol. 12, pp. 423–443.
- Yao, Rui and Harold Zhang (2005) “Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints,” *The Review of Financial Studies*, Vol. 18(1), pp. 198–239.

## A Appendix: Blinder - Oaxaca Decomposition

This appendix details the control variables used in, and the results from, estimating Equation 1.

My control variables are chosen to mirror and slightly extend those used in Altonji and Doraszelski (2005). The individual-specific measure of expected income included in  $Y$  is the residual from race- and gender-specific regressions of labor income on a fourth-order polynomial in age, a dummy for whether the individual has children, the number of children, region of residence, education level, and dummies for residence in a metropolitan area, children, and year. This variable is forward-looking, and intended partly to capture expectations of future income growth. I also calculate a family-specific measure of permanent income as the average total nonasset income reported by that family in every year available; this variable is more backwards-looking, and is meant to estimate the total income a family has had available to accumulate assets, which may be particularly important if borrowing constraints are tight.

In the case of couples, my income controls are the expected income of the head, the expected income of the wife, the expected income of the family, temporary (year-specific) nonasset family income, these incomes squared, and interaction terms of the expected income of the head and wife with temporary family income. Income controls for singles are the same, except dropping terms for the wife.

For couples, demographic control variables  $X$  include fourth-order polynomials in the age of the head and of the wife, the number of children present in the household, a dummy for whether there is at least one child in the household, a dummy for whether support is given to dependents outside the household, the total amount of money given to dependents outside the household, the education levels of the household head and wife (six categories), the number of siblings of the head and of the wife, the number of previous marriages of the head and of the wife, the total number of children born to or adopted by the head and the wife, dummies for whether the head or wife report poor health, dummies for whether the head or wife are self-employed, the total number of hours worked by the head and the wife, dummies for the region of residence, a dummy for residence in a metropolitan area, and whether the head's previous marriage ended in widowhood or divorce. I also include a full set of year dummies. Control variables for singles are the same, without the controls for the wife.

Detailed results from this regression, broken down by household type, are provided in Table 6.

In terms of the proportion of the wealth gap explained, my results are broadly consistent with Altonji and Doraszelski. My regressions explain a smaller proportion of the gap for married couples, but slightly more of the gap for single men. Since I use several more waves of the PSID than they do, weight regression results by the population weights, and have a slightly larger set of control variables, I view my estimates as more reliable as theirs. Still, the key point - that across all three

Table 6: Actual and Predicted Wealth Levels by Income

	Baseline Results			Altonji and Doraszelski (2005)		
	Gap (2010 \$)	% Expl. W	% Expl. B	Gap (2010 \$)	% Expl. W	% Expl. B
Married Couples	\$105,393	60%	17%	\$205,480	79%	25%
Single Men	\$43,467	128%	38%	\$71,051	120%	31%
Single Women	\$55,950	67%	43%	\$81,941	103%	33%

Notes: In each box, the first column displays the size of the wealth gap. The second and third columns show the percent of this gap that is explained by the white and black models, respectively.

groups, the white model explains much more than the black model - remains.

However, Altonji and Doraszelski (2005) estimate the size of the gap to be much larger than I do. This is driven by the fact that I am more aggressive in dropping wealth outliers than they are. Recall that I drop all observations in the top 1% of the wealth distribution; Altonji and Doraszelski (2005) drop observations with residuals from a median regression of wealth on observables in the top .5%. If I instead drop only observations in the top .5% of the wealth distribution, my estimates of the size of the gap increase considerably, to \$140,897, \$60,032, and \$66,820 for married couples, single males, and single females, respectively. Note the fact that dropping more observations near the top of the wealth distribution does not increase the proportion of the gap explained by regressions is further evidence against the convexity argument of Barsky et al. (2002).

Finally, to investigate the additional explanatory power provided by the inheritability of wealth, as discussed in Charles and Hurst (2003), Table 7 presents results both from the baseline regression and from a regression that also includes controls for the education and economic status of the head's parents (and the wife's, if applicable). These controls are not used by Altonji and Doraszelski (2005).

Table 7: Actual and Predicted Wealth Levels by Income

	Baseline Results			Parental Controls		
	Gap	% Expl. W	% Expl. B	Gap	% Expl. W	% Expl. B
Married Couples	\$105,393	60%	17%	\$106,331	46%	21%
Single Men	\$43,467	128%	38%	\$65,360	109%	32%
Single Women	\$55,950	67%	43%	\$51,993	78%	37%

Notes: In each box, the first column displays the size of the wealth gap. The second and third columns show the percent of this gap that is explained by the white and black models, respectively.

These additional controls do not improve the fit of the regressions, which is why I exclude them from my baseline specification. This may seem surprising, but parental education and economic status are highly collinear with control variables that are already included in the regression.

## **B Estimating Parameters of the Income Process**

Details will be provided in later versions of this paper. See Guvenen (2009) for a very similar methodology.

## **C Model Solution**

The discrete housing and illiquid asset grids, combined with their nonconvex adjustment costs and the non-constant interest rate of debt, mean that the value function is not globally concave. My solution algorithm is based heavily on Fella (2011), who extends the Endogenous Grid Method to non-concave problems. This algorithm is accurate and fast enough to allow me to estimate my model with SMM, a procedure Bajari et al. (2010) go to great lengths to avoid using to estimate another model of housing.