Do Husbands Want to be Shorter than their Wives?
The Hazards of Inferring Preferences from Marriage Market Outcomes

Martha Bailey
University of Michigan
baileymj@umich.edu

Ariel Binder
University of Michigan
ajbinder@umich.edu

David Lam
University of Michigan
davidl@umich.edu

Extended abstract prepared for submission to the
2017 meeting of the Population Association of America

Date of draft: September 29, 2016

Do Husbands Want to be Shorter than their Wives?
The Hazards of Inferring Preferences from Marriage Market Outcomes

Abstract

Differences in spousal characteristics on dimensions such as age, height, and income are heavily studied in social science research. Researchers often draw inferences about preferences and norms based on observed distributions of spousal traits. For example, a "male taller" norm has been inferred from the fact that fewer wives are taller than their husbands than would occur with random mating. We show theoretically and empirically that it is difficult and potentially misleading to infer preferences from marriage market outcomes. If a gender gap in trait distributions exists, many different preferences will produce strong positive sorting on that trait in marriage market equilibrium. Applying these results to income differences between spouses, we show that what appears to be a norm against wives earning more than their husbands is consistent with a wide set of preferences, including a preference for equality of spousal incomes.

I. Introduction

Do men prefer to be taller than their wives? Do women prefer to be taller than their husbands? Would men and women prefer to have identical incomes between husbands and wives? What age difference between spouses is ideal from the perspective of husbands and wives? Differences in the characteristics of spouses has often been a focus of social science research. Much of this research attempts to understand the preferences of men and women about the ideal characteristics of spouses or about the ideal difference in spousal characteristics. For example, a number of researchers have investigated the extent to which there is a "male-taller" norm in various populations (e.g. Gillis and Avis, 1980; Stulp et al., 2013) and the extent to which preferences about height may affect other choices such as inter-ethnic marriage patterns (Belota and Fidrmuc, 2010). A large literature looks at income differences between spouses, investigating the extent to which there are preferences regarding spousal income differences and the impact of spousal income differences on outcomes such as time allocation decisions, consumption decisions, and marital stability.

One piece of evidence that is often cited in this research is the actual differences in spousal traits that are observed in a given population. For example, Stulp et al. (2013) show that the proportion of couples in which the husband is taller than the wife in a UK sample is greater than the proportion that would result from random matching, interpreting this as evidence of a male-taller norm. Bertrand, Kamenica, and Pan (2015) argue that there is a social norm against wives earning more than their husbands, with one important piece of evidence being a drop-off in the density of wife's share of marital income at 50%.

In this paper we argue that is very difficult and potentially quite misleading to infer preferences about spousal characteristics from observed pairings in the marriage market. The challenge comes from the fact that the underlying distributions of spousal characteristics will impose constraints on the feasible set of possible outcomes. When these constraints are combined with a tendency for positive or negative assortative mating on some trait, the observed matches in the marriage market may look very different than might have been expected based on preferences. We show, for example, that preferences in which men prefer to be shorter than their wives could produce exactly the same equilibrium set of marriages as preferences in which men prefer to be taller than their wives.

Beginning with simple models of marital sorting on height, we demonstrate theoretically that many different sets of preferences can produce the same set of observed pairings in the marriage market. We then develop much more general models that apply to a wide range of distributions and preferences, and consider the implications of these results for analysis of income differences between spouses. Using data on incomes of husbands and wives in U.S. Census data, we show that we can come very close to reproducing the actual income differences between wives using a model of marital sorting in which there is no preference or social norm related to husbands earning more than wives. Observed spousal income differences appear to be largely explained by the fact that men have higher average earnings than women, combined with a strong tendency for positive assortative mating on income. While this does not mean that there is no social norm related to husbands earning more than wives, our results suggest that great care must be taken in drawing inferences about preferences from the observed marriage market outcomes.

II. Becker's theory of marriage and a simple model of sorting on height

Our theoretical discussion requires that we be able to make predictions about how men and women are sorted in a marriage market. We build on Becker's (1973) economic theory of marriage, which provides well-known predictions about assortative mating on traits. Consider a man M and a woman F who are considering marriage. We assume they marry if and if only if it makes both better off compared to alternatives. Denote the "output" of the marriage by Z_{mf}. For now assume output can be divided $Z_{mf} = m_{mf} + f_{mf}$, where m_{ij} indicates what man i consumes when married to woman j. Although this may not be a minor assumption, since "household public goods" like children – or the income difference between spouses – cannot literally be divided in this way, Lam (1988) shows that the model can be applied to the case of household public goods under the assumption of transferable utility. Because output (or utility) can be divided up between husbands and wives, it is possible for men to make offers to potential wives (and women to make offers to potential husbands) of some division of output. This means that a man can in principle use "side payments" to attract a particular wife, and a woman can use side payments to attract a particular husband, making that person better off than they would have been with some other partner. In Section V below we will also consider the case of fully non-transferable utility. In the full version of the paper we will discuss the realism of transferability assumptions consider the implications of relaxing them.

Suppose we have a set of n women and n men, with marital output between woman i and man j denoted by Z_{ij}, and we consider all possible sortings of men and women. Drawing on results from other matching models in mathematics and economics, Becker showed that a competitive equilibrium in the marriage market will be the set of assignments that maximizes the sum of output across all marriages. The argument is a standard argument about the Pareto optimality of competitive markets. If an existing set of pairings does not maximize total output, then there must be at least two couples for which we could switch partners and increase total output. Given this, there must be an incentive for the individuals in those couples to capture that increase by a set of new matches and new division of output. This will be illustrated below for a simple example of two couples sorting on height.

Becker applied this very general result to the case of sorting on some trait *A*, where we will consider woman *f* to have a trait value A_f and man *m* to have trait value A_m, where *A* might be height, age, education, income, etc. We will characterize marital output (which might be some measure of joint marital happiness) as a function of the values of *A* for each partner $Z_{mf} = Z(A_m, A_f)$. Becker showed that the marriage market equilibrium will be characterized by positive assortative mating on *A* if

$$\frac{\partial Z(A_m, A_f)}{\partial A_m \partial A_f} > 0. \tag{1}$$

There will be positive assortative mating if the cross-partial in (1) is positive, and negative assortative mating if the cross-partial is negative. A positive cross-partial derivative can be interpreted as implying that the value of *A* for the husband and wife are complements, while a negative cross-partial implies they are negative. If, for example, having a higher educated husband raises the impact of the wife's education on marital output, then we will tend to see positive assortative mating on education. We will draw on the result in (1) extensively below.

Some of the key theoretical points can be demonstrated with a very simple model of sorting by height in the marriage market. Denote female height by H_f and male height by H_m. Suppose there are two women: F_1 is 60" tall and F_2 is 66" tall. There are two men: M_1 is 66" tall and M_2 is 72" tall. There are two possible ways to create two couples, 1) F_1M_1, F_2M_2, which is positive assortative matching on height, and 2) F_1M_2, F_2M_1, which is negative assortative matching on height. We are interested in what we can say about which sorting will be observed in a marriage market equilibrium.

In order to find the marriage market equilibrium, we need to describe how the heights of couples affect marital utility. Assume that people get utility from their individual consumption and some bonus that comes from being married. The gains from marriage take the very simple form of some bonus *K* (representing, say, economies of scale in consumption or benefits of household public goods) that is offset by some penalty that depends on the height difference between spouses. *K* can be thought of in monetary or consumption units, representing in the simplest example the amount of money the couple saves by being married. The penalty associated with the height difference between couples can also be given a monetary interpretation, representing the amount of additional consumption that would be required to compensate for the disutility from a given height difference between spouses.

We will consider various alternative cases for the loss function associated with the height difference between spouses. For the first case, suppose that all men and women agree that the ideal marriage is one in which the husband is 6" taller than his wife. Couples in which this is not the case experience some loss of utility that increases at an increasing rate as the height difference between spouses increases. A simple example is a quadratic loss function:

$$Z(H_m, H_f) = K - (H_m - H_f - 6)^2. \tag{2}$$

If the husband is 6” taller than the wife then there is no loss of utility from marriage. If the husband is the same height as the wife then the loss is 36. As a concrete and very literal example, this could mean that the couple would need an additional $36 worth of consumption to make them as happy as a couple with the ideal height difference. If the husband is 6” shorter than the wife then the penalty is 144. With these payoff functions, we can consider the two possible sortings of couples. If the taller man marries the taller woman and the shorter man marries the shorter woman, then each husband is 6” taller than his wife, generating a total marital utility of *2K* (zero penalty in either marriage). If we switch partners, then one couple (same height) has a penalty of 36 and the other couple (taller man and shorter woman) has a penalty of 144. Total marital utility is obviously highest with perfect rank-order sorting, and this is the competitive equilibrium we would expect to observe. If we started with the alternative sorting, everyone could be made better off by switching partners. If we observe the perfect rank-order sorting equilibrium and conclude that everyone prefers that husbands are taller than their wives, our inference would be correct.

Now consider a different payoff function in which the ideal couple is one in which the husband and wife have equal heights, with a penalty for height differences that is increasing in the difference.

$$Z(H_m, H_f) = K - (H_m - H_f)^2 \qquad (3)$$

With perfect rank-order sorting the total penalty is now 36 + 36 = 72, since each couple is 6” from the ideal height difference. In the alternate sorting we can create one ideal couple of equal heights, generating a penalty of zero. But the other couple (the tall man and the short woman) has a height difference of 12”, creating a penalty of 144 (which we can think of as 72 per spouse). Perfect rank-order sorting produces higher total marital utility (lower total penalties). This is because of the convex penalty function, which penalizes very large differences in height. The logic in terms of a competitive marriage market is as follows: Suppose we began with the sorting in which one couple has equal heights while the other couple has a 12” height difference. The individuals in the mismatched couple, F_1 and M_2 see that they would each be much happier if they could switch partners and have a 6” height difference instead of a 12” height difference. The question is whether F_1 would be able to induce M_1 to switch from F_2 to her. Her penalty would decline from 72 (half of 144) to 18 (half of 36) if she changed partners. The penalty for M_1 would increase from 0 to 18 (half of 36) if he switched partners. Clearly F_1 can more than compensate M_1 for changing, making him a side payment of at least 18, leaving herself better off after the switch. The exact same story can be told for M_2 inducing F_2 to switch to him. Every person will be better off after the re-sorting, so the positive assortative mating equilibrium is the one we should observe.

The resulting sorting of spouses with the preferences in (3) is exactly the same as the sorting with the preferences in (2) – the sorting with positive assortative matching on height. In this second case we would be drawing an incorrect inference if we interpreted the equilibrium as resulting from a preference for men to be taller than their wives. In fact the preference is for equal heights, and the distribution of heights allows for such a case. The reason we do not see it is because creating that match leads to another match of extremely unequal heights.

Taking this case even further, consider a payoff function in which the ideal couple is one in which the wife is 6" taller than her husband, with, once again, a penalty for deviations from the ideal that is increasing in the difference.

$$Z(H_m, H_f) = K - (H_f - H_m - 6)^2 \tag{4}$$

With perfect rank-order sorting the total penalty is 144 + 144 = 288, since each couple is 12" from the ideal height difference. In the alternate sorting the total penalty is 36 + 324 = 360. Once again it is positive assortative matching that produces the maximum total payoff across all marriages. If we started with negative assortative matching, a process of renegotiation analogous to the one just described should lead to a re-sorting. We will therefore expect that positive sorting will be observed as the equilibrium outcome. This then, is the interesting case in which the underlying preferences are that men prefer to be shorter than their wives. We never observe this in the actual marital outcomes, however. The reason is that the convex payoff function pushes the equilibrium toward a sorting that has small average differences between spouses. It is better to have everyone slightly off from the ideal rather than have some couples that are close to the ideal and other couples that are very far from the ideal.

III. More general models of differences in spousal characteristics

The case discussed above of two men and two women sorting on heights is very simple. The basic conclusions of the model apply to much more general cases, however. In the full version of the paper we will show that the model generalizes to cases with a large number of women and men covering a large range of heights. As long as men are on average taller than women and as long as the penalty function to height differences between spouses is convex in the height difference (that is, the penalty to a 2" height difference is more than twice as large as the penalty to a 1" height difference), we will tend to observe strong positive assortative mating on height, with the same set of matches implied by a wide range of preferences. Extending the result further, we will show that as long as the male distribution of heights stochastically dominates the female distribution, it will still be the case that a preference for having women taller than their husbands will produce the same equilibrium sorting in the marriage market as a preference for having men taller than their wives, and both will be indistinguishable from a preference for having husbands and wives of the same height.

These results obviously extend to differences in other characteristics such as income. Income has the additional complication that, unlike height, it is not an exogenous trait. The incomes of husbands and wives will be affected by decisions about labor supply and investments in human capital. These issues will be discussed in the paper. But assortative mating also plays a fundamental role in determining the income difference between spouses. Our results imply that if the male income distribution stochastically dominates the female income distribution, and if there is a tendency for strong positive assortative mating on income, it will tend to be relatively rare for women to earn more than their husbands. This tendency will exist even if the underlying norm is to have equal incomes between husbands and wives, or even for wives to earn more than their husbands. The paper will also discuss the fact that many other factors may lead to strong positive assortative mating on income or income-related characteristics. Lam (1988), for example, demonstrates that there will tend to be positive assortative mating on income whenever the gains from marriage results from household public goods, such as children. Individuals may have no preferences regarding the difference between spousal incomes at all, but the equilibrium set of pairings in the marriage market may look as if there is a norm that a husband should earn more than his wife.

IV. Literature on spousal income differences[1]

Since the large strides of the 1970s and 80s toward gender equality of education and labor market outcomes, social scientists have taken an interest in understanding the rise of dual-earner households and the distribution of income differences between spouses (see, e.g., Winkler, 1998; Brennan, Barnett and Gareis, 2001; Raley, Mattingly, and Bianchi, 2006). In a recent article, Bertrand, Kamenica, and Pan (2015) (hereafter referred to as BKP) put forth a provocative argument that patterns of household decision-making are consistent with a societal norm that the wife should not out-earn her husband. Their primary evidence consists of analysis of the distribution of the share of household income earned by the wife across a variety of Census Bureau samples and administrative data. In all samples they find that the distribution turns drops sharply and discontinuously at 0.5, the point at which the wife starts to out-earn the husband (Figure 1). Examining a variety of surveys, they supplement these discontinuities with other findings, such as: in marriage markets in which women are likelier to out-earn men, marriage rates are lower; when the wife's full earning potential exceeds her husband's she is less likely to work full-time; and when the wife does out-earn the husband the marriage is less stable and likelier to end in divorce. Despite another recent paper failing to reject the hypothesis in younger cohorts that divorce risk is unchanged when the wife outearns the husband (Scwartz and Gonalons-Pons, 2016), BKP assemble a convincing case that a husband-as-primary-earner social norm exists and inhibits further progress toward gender equitability in the household.

Figure 1. Share of household income earned by the wife in a sample of administrative income data, 1990-2004 (BKP Figure I)

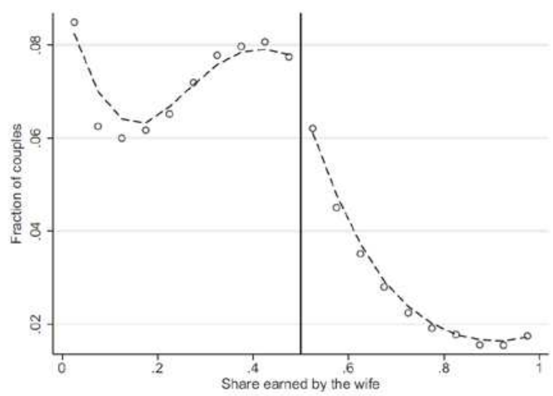

FIGURE I
Distribution of Relative Income (SIPP Administrative Data)

We do not dispute BKP's supplementary results nor argue that such a social norm does not exist. Nevertheless, in light of the above theoretical discussion, it may be misguided to make

[1] To be expanded in ongoing work.

a conclusion about agents' preferences toward spousal income differences based on the observed marriage market distribution of income differences. In the next section we present preliminary results indicating that so long as the marriage market is characterized by positive sorting on potential income (which as we showed above is consistent with a wide array of underlying preferences), it is possible to simulate distributions of the share of income earned by the wife that are nearly identical to the observed distribution, including the sharp drop at 0.5.

V. Preliminary Empirical Results

The theoretical results discussed in Sections II and III indicate that to the extent that individuals in the marriage market care at all about the income gap with their eventual spouses, there should be a strong tendency toward positive sorting on (potential) income in marriage market equilibrium. This sorting is driven by the fact that one side's income distribution (males') stochastically dominates the other side's (females'). In this section we present preliminary results demonstrating that if we calibrate the male and female income distributions according to Census data, and assume positive sorting on potential income, we can very closely replicate the empirical distribution of the wife's share of earned household income. As in BKP's analysis of Census data, our simulated distributions all exhibit modes between 0.4 and 0.5 and sharp drop-offs in mass thereafter. Using a Monte Carlo procedure combined with the test proposed by McCrary (2008), we estimate the sizes of hypothetical discontinuities in the densities at 0.5, and find them to be close in magnitude to the estimated discontinuity in observed data, albeit slightly smaller. Our evidence does not invalidate the notion that a "social penalty", which discontinuously applies if the wife out-earns the husband, exists in the match utility function. However, it does indicate that gender differences in the wage structure combined with any match utility function that delivers positive sorting on potential income (for example because of public consumption goods) can sufficiently explain the surprisingly low incidence of wives out-earning their husbands.

We start with a sample of couples drawn from the 2000 Census 5% sample. Following BKP, we restrict the sample to couples aged 18-65, process earned income variables following the procedure outlined in the paper's appendix, and keep only spouses reporting positive income. Note that wives often reduce their working hours or exit the labor force to raise young children, and re-enter the full-time workforce with lower earning potential (e.g. Mincer and Ofek, 1982). Because our simple treatment of the income process and marital sorting will not address the dynamics of household fertility and how they interact with labor supply decisions, we further restrict the sample to relatively young couples (aged 18-40) without children. We obtain a sample of 109,570 couples and calculate the share of household income earned by the wife (*wifeshare*) for each couple. Figure 2 plots a rough, 20-bin histogram of the distribution of *wifeshare* along with a line connecting the heights of the bins. As in BKP, the distribution exhibits a sharp drop at 0.5. According to McCrary's test the estimate of the sharp drop is 9.6% (with a standard error of 1.3%). This serves as the benchmark for our subsequent empirical exercises.

In our first exercise, we take observed income to be exogenous. Denote the income (log income) of individual i of gender g as Y_i^g (y_i^g), where g is m for males and f for females.

Assume log-normality of observed income; that is, $y_i^m \sim N(10.35, 0.75)$ and $y_i^f \sim N(10.00, 0.87)$, where the numbers correspond to the observed means and standard deviations of log income by gender in our sample. We simulate a sample of 50,000 males and 50,000 females from these distributions and match them not according to observed income rank but rather the rank of observed income perturbed with noise. That is, for each individual we assign $W_i^g = Y_i^g + u_i$, where *u* is white noise, and pair up males and females according to their rank of *W*. In this representation, *W* can be thought of as permanent income and the white noise captures transitory income shocks.[2] Calibrating the standard deviation of *u* to 16,000, Figure 3 displays a simulated distribution of *wifeshare*. It is remarkably similar to the actual distribution displayed in Figure 2, although different samples from the distribution alter slightly the location of the mode and the severity of drop-off thereafter. To take account of this sampling variability, we estimate the discontinuity at 0.5 by employing a Monte Carlo procedure. Simulating a sample of 10,000 males and 10,000 females, we estimate the discontinuity and its associated *t*-statistic via the McCrary test and repeat for 1,000 replications. With this conservative bootstrapping procedure we do not quite achieve statistical significance at conventional levels (the average *t*-statistic is 1.3), but the average discontinuity estimate is 7.1%--close to what is observed in the Census data. With this very simple income process and sorting mechanism, which is consistent with a wide array of underlying preferences, we are able to replicate the curious sharp drop in the distribution of *wifeshare* at 0.5.

Figure 2. Census 2000, Couples Aged 18-40 w/o Kids

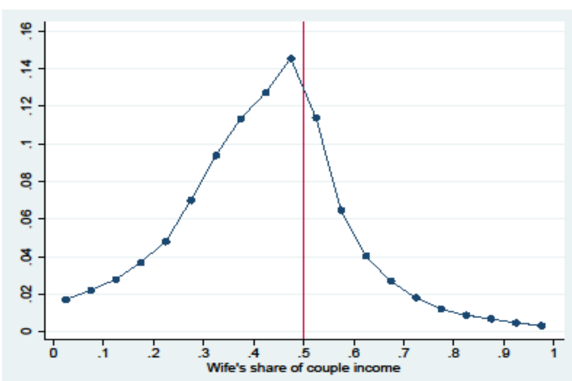

Figure 3. Simulated Sorting on Income + Noise

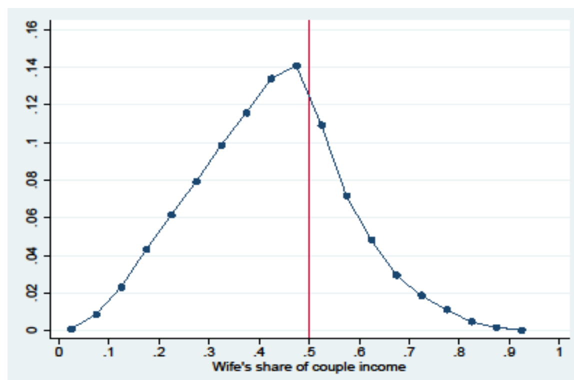

The results of the first exercise are striking given how little structure was imposed on the environment. However, it is a stretch to assume that the wife's income is exogenous, even despite focusing on a sample of childless wives of prime working age. For a variety of reasons, including the existence of social norms, the wife may reduce her labor supply to meet household objectives, even if children are not in the picture. In this exercise we endogenize the wife's

[2] Common specifications of earnings processes (e.g. Moffit and Gottschalk, 2002) assume transitory white-noise shocks enter log-linearly, rather than linearly as we have assumed. An alternative way to intepret the sorting of the marriage maket on income + noise might be the presence of a search friction whose magnitude is independent of the level of income (i.e. fixed costs). In ongoing work we are more carefully thinking about the types of earnings processes and search cost structures that are consistent with positive sorting on income + noise in marriage market equilibrium.

income via a simple labor supply model and explore the model's predictions about the distribution of *wifeshare*.

We assume that, for a given male m and female f, the match output function is given by

$$Z_{mf} = Z\left(Y_m, Y_f, P\right) = \frac{\left(0.61(Y_m + Y_f P)\right)^{1-\gamma}}{1-\gamma} - \psi P,$$

where Y_m and Y_f denote each spouse's (exogenous) full income, P is the wife's labor supply decision (constrained to be in the unit interval), γ is the CRRA parameter, and ψ is the disutility incurred by the household if the wife works. This simple specification of household utility has been used in recent work investigating determinants of wives' labor supply (e.g. Attanasio, Low, and Sanchez-Marcos, 2008). It assumes household consumption of earned income is a public good with congestion; the 0.61 is a McClements scale calibration capturing consumption economies of scale in marriage.[3] We depart from the framework of Section II and now assume fully *non*-transferable utility. With this assumption positive sorting on full income occurs in marriage market equilibrium so long as each member's full income positively affects match output.[4] It is trivial to show that this holds here (regardless of the wife's eventual labor supply decision). Once again assuming that each individual's full income is the sum of his or her permanent income and a transitory shock, positive sorting on full income + noise will arise in equilibrium.

After marriage, the wife takes her husband's and her own full income as given and chooses $P \in [0,1]$ to maximize the above utility function. Assuming an interior solution, she optimally chooses

$$P^* = \frac{\frac{1}{0.61}\left(\frac{\psi}{0.61 Y^f}\right)^{-\frac{1}{\gamma}} - Y^m}{Y^f};$$

if P^* lies outside of the unit interval, the appropriate corner solution applies.

To use the above model to draw valid conclusions about the distribution of *wifeshare* in marriage market equilibrium, we must reasonably calibrate it. Outside of the calibration we impose $\gamma = 1.5$, a standard value estimated in the macro literature. We assume log-normally distributed full incomes and allow the work disutility parameter, ψ, to be heterogeneous in the population and negatively correlated with Y_f.[5] The model thus contains 8 parameters to be calibrated, and we calibrate them by targeting 8 moments in our Census 2000 sub-sample: the means and standard deviations of male and female log observed income, the observed mean gender earnings ratio conditional on earning positive income (P^*>0), the observed mean gender earnings ratio conditional on full-time work (defined in the data as at least 1600 hours worked in the last calendar year; defined in the model as P^*>0.95), the female labor-force participation rate (defined in the data as the share of wives working positive hours in the last calendar year), and

---

[3] To illustrate, suppose P=1 and Y_m=Y_f. Then the couple enjoys a higher level of consumption in marriage than either member would as single.

[4] Starting from perfectly positive sorting, it is easy to show that no two individuals can become better off by dissolving their current matches and matching with each other. The inability of individuals to make transfer payments means we no longer need the cross-partial assumption on the match output function to generate positive sorting on the given trait in marriage market equilibrium.

[5] Imposing a negative correlation, as has been estimated in the literature (e.g. Eckstein & Lifshitz, 2011), ensures that positive sorting on potential income is not disturbed.

the female full-time labor-force participation rate. Table 1 summarizes the calibration—overall the model does a very good job of replicating the targets in the data.

With the calibrated model we can now simulate the distribution of *wifeshare*, and do so in the right graph of Figure 4 using a sample size of 50,000 couples. Once again, the simulated distribution is extremely close to what is observed in Census data (displayed in the left graph), and the distribution turns down sharply just before 0.5. Performing the same Monte Carlo version of McCrary's test above, we estimate a smaller and statistically insignificant discontinuity of 5.0%. Despite not quite replicating the size of the discontinuity observed in Census data, both exercises otherwise deliver extremely accurate representations of the distribution of *wifeshare* in marriage market equilibrium.

Table 1. Model Calibration

Description	Parameter	Calibrated Value
Mean male log income	μ^m	10.35
Standard deviation of male log income	σ^m	0.75
Mean female log full income	μ^f	10.16
Standard deviation female log full income	σ^f	0.70
Mean disutility of work	ψ	.0019
Standard deviation of disutility of work	σ^ψ	$\psi/2$
Correlation, disutility of work and female log full inc	ρ	-0.4
Standard deviation of transitory income shock	σ^u	13,000
Targets in the data	Data	Model
Mean male log observed income	10.35	10.35
Standard devation male log observed income	0.75	0.75
Mean female log observed income	10.00	9.98
Standard deviation female log observed income	0.87	0.87
Mean gender earnings ratio, all	0.74	0.71
Mean gender earnings ratio, full-timers only	0.80	0.79
Female labor-force participation rate	0.88	0.91
Female full-time labor-force participation rate	0.67	0.67

Figure 4. Comparison of Census 2000 Data (left) to Model Simulation (right)

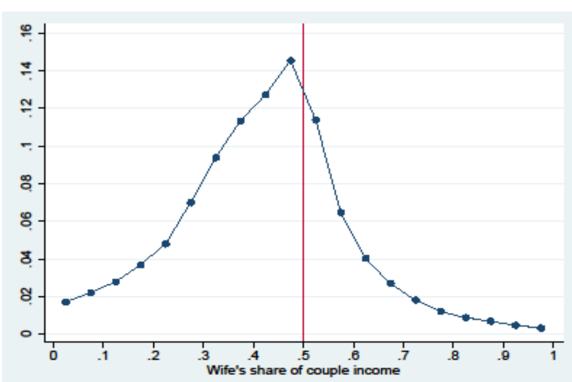

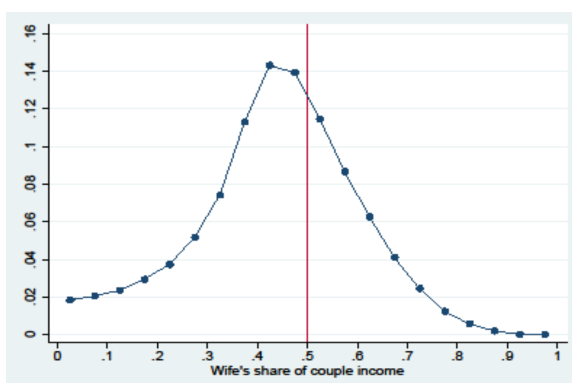

References

Attanasio, O., Low, H., & Sanchez-Marcos, V. (2008). “Explaining Changes in Female Labor Supply in a Life-Cycle Model,” *American Economic Review, 98*, 1517-1552.

Becker, G. (1973). “A Theory of Marriage: Part I,” *Journal of Political Economy, 81*(4), 813-846.

Belota, M., & Fidrmuc, J. (2010). “Anthropometry of Love: Height and Gender Asymmetries in Interethnic Marriages,” *Economics & Human Biology, 8*(3), 361-372.

Bertrand, M., Kamenica, E., & Pan, J. (2015). “Gender Identity and Relative Income within Households,” *Quarterly Journal of Economics*, 571-614.

Brennan, R. T., Barnett, R. C. and Gareis, K. C. (2001). “When She Earns more than He Does: A Longitudinal Study of Dual-Earner Couples,” *Journal of Marriage and Family, 63*(1): 168-182.

Eckstein, Z., & Lifshitz, O. (2011). “Dynamic Female Labor Supply,” *Econometrica, 79*, 1675-1726.

Gillis, J. S., & Avis, W. E. (1980). “The Male-Taller Norm in Mate Selection,” *Personality and Social Psychology Bulletin, 6*(3), 396-401.

Lam, D. (1988). “Marriage Markets and Assortative Mating with Household Public Goods: Theoretical Results and Empirical Implications,” *Journal of Human Resources, 23*(4), 462-487.

McCrary, J. (2008). “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test,” *Journal of Econometrics, 142*, 698-714.

Mincer, J., & Ofek, H. (1982). “Interrupted Work Careers: Depreciation and Restoration of Human Capital,” *Journal of Human Resources, 17*, 3-24.

Moffit, R., & Gottschalk, P. (2002). “Trends in the Transitory Variance of Earnings in the United States,” *Economic Journal, 112*, C68-C73.

Raley, S. B., Mattingly, M. J., & Bianchi, S. M. (2006). “How Dual are Dual-Income Couples? Documenting Change from 1970 to 2001,” *Journal of Marriage and Family, 68*(1), 11-28.

Scwartz, C. R., & Gonalons-Pons, P. (2016). “Trends in Relative Earnings and Marital Dissolution: Are Wives Who Outearn Their Husbands Still More Likely to Divorce,” *RSF: The Russell Sage Foundation Journal of Social Sciences, 2*(4), 218-236.

Stulp, G., Buunk, A. P., Pollet, T. V., Nettle, D., & Verhulst, S. (2013). “Are Human Mating Preferences with Respect to Height Reflected in Actual Pairings?” *PLoS ONE, 8*(1). doi:10.1371/journal.pone.0054186

Winkler, A. E. (1998). “Earnings of Husbands and Wives in Dual-Earner Families,” *Monthly Labor Review, 121*(4), 42-48.