COLLEGE CHOICE AS A COLLECTIVE DECISION*

Nick Huntington-Klein

Department of Economics

California State University, Fullerton

nhuntington-klein@fullerton.edu

April 18, 2017

Abstract

Although the choice between colleges can be thought of as being made collectively by a family, models of educational choice almost universally portray the decision as made by the student alone. Using a novel experimental method for identifying collective decision functions, I find that students have more influence than parents over the decision, but not exclusive control. Students care more than parents about classroom experience and future earnings. Ignoring the dual-agent nature of the decision can weaken predictions and lead to poorly-targeted policy designs.

Keywords: Education, collective choice, experiment, college

JEL Codes: I21, J24, D13

*The author was supported by Grant #R305B090012 from the Institute of Education Sciences during this project. I thank Patrick Bajari, Mark Long, Shane Murphy, Elaina Rose, William Zumeta, anonymous reviewers, and session participants at the University of Washington, UC Santa Barbara, CSU Fullerton, and the WEAI for their comments and suggestions. The experiment described in this paper was approved by the University of Washington IRB. The author declares that he has no relevant or material financial interests that relate to the research described in this paper.

1 Introduction

College attendance is an important determinant of human capital. The choices of whether or not to attend college, which college to attend, and what to major in all affect a student's performance in the labor market. Aggregated, these choices determine much of the supply of skilled labor. In this paper, I model the decision of which college to attend as a collective household choice made by students and parents, and estimate the parameters of that model using an experiment. This approach improves upon the standard approach to modeling higher education.

Recent years have seen major advances in the study of choice in higher education. Treatment of financial incentives now commonly models heterogeneity in field-specific ability, which bears heavily on choice (Stinebrickner and Stinebrickner, 2014). New work also seriously considers non-financial incentives. Preference for being in a particular educational environment may drive student choice even more strongly than financial incentives at the level of choosing whether or not to attend college (Bleemer and Zafar, 2015; Belfield et al., 2016), which college to attend (Jacob et al., 2017), and on the choice of major (Wiswall and Zafar, 2015).

Beyond expanding and modifying the list of incentives that drive educational choice, recent work emphasizes the uncertainty that students face in making their choice. There is exhaustive evidence that student beliefs about the financial costs and benefits of different educational options are not perfect (see Hastings et al., 2017, for a partial review of this literature). Further, students may not be initially aware of their own preferences and abilities, which leads to meaningful differences between ex-ante and ex-post educational returns (Altonji et al., 2012). Students may not even know what options are available to them, the "choice set problem." While the choice set problem has been long known, recent work has demonstrated the significant impact that it has on choice, especially for low-income students (Hoxby and Avery, 2013; Hoxby and Turner, 2013).

Despite significant progress, one prominent unrealistic feature remains in most studies of educational choice: college choice is nearly always presented as the choice of *a student*. In real household settings, both students and parents have some agency in the decision. Educational models acknowledge this, but assume only that parental characteristics have an influence on what is fundamentally a student's choice. For example, some include parental education, parental encouragement, or parental approval as a predictor, as is done in Zafar (2012). This approach still treats the decision as being that of a student.

In this paper, I develop a model of intra-household choice between colleges based on their attributes, which allows the decision to directly represent the preferences of both students and parents. Student and parent preferences over college attributes are elicited independently using a stated-choice conjoint analysis experiment in which college attributes are varied randomly. The household choice function consists of student and parent preferences combined using a bargaining parameter. I estimate the bargaining parameter by measuring how sensitive reported student choice is to parental opinion, exogenously varying "parental opinion of this college" as a college attribute presented to students, and vice verse. For example, if reported student choice is insensitive to parental opinion, but reported parental choice is sensitive to student opinion, then the weight of the household decision falls upon the student. This approach allows for student, parent, and household choice functions to be cleanly identified, a task that would be very difficult in observational data.

I gather a data set of 964 online participants. Separating the sample into students and parents by age, I find that student and parent priorities are significantly different. In particular, students respond more strongly, relative to parents, to how enjoyable the classes are, and to future earnings. The student emphasis on class enjoyability and future earnings may represent a difference in priorities compared to parents, or may represent the fact that students will directly experience those classes and earnings, which will affect parents only indirectly. Parents respond strongly to their child's opinion, and so the weight of bargaining power falls to the students. But, parents do still hold non-negligible bargaining power, contrary to the implicit assumption of student-only models.

In the wider literature on educational choice, collective choice between students and their parents is generally ignored, despite evidence that children and parents have a say in household decisions (Lundberg et al., 2009; Dauphin et al., 2011). Economists use college choice models to answer a number of interesting questions, such as the influence of consumption value or the uncertainty of outcomes on choice. These questions will be better understood if parental influence is properly incorporated into the model. Without a properly specified model, empirical results from choice data will mistake parental preferences for student preferences. Giustinelli and Manski (2017) review the current literature on family decision-making in education, and emphasize how differences in expectations held by parents and students mean that family decision-making should be considered an essential component of models of educational choice. Without a proper collective treatment of choice, we will misunderstand how educational incentives of interest actually translate into choices made by making incorrect inferences about individual choice functions, and will fail entirely to see mechanisms that operate only through a collective structure.

There is a small amount of previous work on collective student/parent decision making in educational choice. Using data from the United States, Kalenkoski (2008) presents a model of Nash bargaining and rejects that decisions are made as a unitary household and that parents have altruistic preferences. Giustinelli (2016) looks at parent-child interactions in the choice of high school curriculum in Italy. She finds that students tend to lead decision-making. A child's preference for a particular type of curriculum is the most important predictive factor in choice. Attanasio and Kaufmann (2014) use data from Mexico to look at the influence of both student and mother labor market expectations on the choice of whether to attend high school and whether to attend college. Both mother and student expectations matter.

This literature on student/parent educational decisions makes a distinction between the standard "unitary" model, where a household is treated as an individual, and a model in which multiple agents with different preferences come together to make a decision. The assumption of a unitary household is rejected in a long series of empirical papers on decision-making between husbands and wives, generally showing that income is spent differently depending on which household member receives it (Thomas, 1990; Lundberg et al., 1997; Duflo, 2003, e.g.). See Fortin and Lacroix

(1997) or Ermisch (2003) for a review of this literature.

The rejection of the household unitary model significantly improved understanding of intra-household choice, and emphasized the importance of targeting in labor and transfer policy. Education policy can benefit in the same way from a better understanding of targeting. Targeting becomes especially important in light of a new wave of college policies that provide outreach and information about college (Swail and Perna, 2002; Domina, 2009; Jensen, 2010; Hoxby and Turner, 2013; College Affordability and Transparency Center, 2014; Hastings et al., 2015; Bleemer and Zafar, 2015, e.g.). These policies necessarily must choose which members of a household to focus on. Without a proper collective model of college choice, these decisions will be made blind.

2 Model

In this section I present a collective model of college choice along the lines of Chiappori (1992), which assumes that the decision process is Pareto efficient but otherwise does not impose a particular model of bargaining on choice (Ermisch, 2003). I also show how certain parameters of that model can be identified experimentally.

Parents (p) and their student children (s) must choose a college option c from the set of available colleges $\mathbb{C}$.[1] They make this decision collectively, maximizing a combination of student and parent utilities over two periods. The first period covers the time spent in college, and the second covers time in the post-college labor market.

$$\max_{c\in\mathbb{C},z_1^s,z_1^p,z_2^s,z_2^p,L,B}[U^s(x_c,z_1^s)+\delta^s U^s(0,z_2^s)+\mu[U^p(x_c,z_1^p)+\delta^p U^p(0,z_2^p))]] \tag{1}$$

$$\text{s.t. } t_c+z_1^s+z_1^p=Y_1^s+Y_1^p+L, \tag{2}$$

$$z_2^s+(1+r)L\leq Y_{c2}^s+B, \text{ and } z_2^p+B\leq Y_2^p \tag{3}$$

[1] $\mathbb{C}$ also includes the option to not attend college at all.

where U^s and U^p are student and parent utility functions, μ is the weight of parental utility relative to student utility in the decision-making process, and δ^s and δ^p are student and parent discount factors.

Each college option c offers consumption value attributes x_c at a tuition cost of t_c, both of which only apply in the first period. College attributes x_c act here as public goods preferred by both students and parents. Parents do not directly experience many of these attributes, but they may directly care about their child's experience. In addition to tuition payments, parents and children must pay for and distribute other goods z to the student (z_1^s, z_2^s) and the parent (z_1^p, z_2^p).

Tuition and other goods are paid for using student and parent income Y_1^s, Y_{c2}^s, Y_1^p, and Y_2^p, and student loans L that must be repaid at an interest rate r. Second-period student income Y_{c2}^s depends on the college chosen in the first period due to college variation in networking opportunities and human capital accumulation. In the second period, students and parents may make transfers B between each other.

I do not consider credit constraints. It is a strong assumption that credit constraints never matter. However, it is likely that in the United States credit constraints are not a dominant factor in college choice (Dynarski, 2002; Carneiro and Heckman, 2002). In the experiment, subjects are instructed to assume that sufficient loans are available to cover tuition.

Assuming noncompetitive college markets, colleges are not priced hedonically. That is, college attributes x_c and Y_{c2}^s are not a direct function of price t_c. Rather than choosing a college with optimal levels of x_c and Y_{c2}^s given a price for each attribute, the family chooses c from available attribute baskets $\mathbb{C}$ given a set price t_c for each basket.

For each college $c \in \mathbb{C}$, the optimal allocation $(z_1^{s*}, z_2^{s*}, z_1^{p*}, z_2^{p*}, L^*, B^*|c)$ is found using first-order conditions from Equation 1 and the constraints in Equations 2-3. The goods allocation is a function of fixed parameters r, Y_1^s, Y_1^p, and Y_2^p and college-specific variables t_c, Y_{c2}^s, and x_c. The household's objective function given a specific college $c \in \mathbb{C}$ can then be based on indirect lifetime utility functions V^s and V^p, which aggregate indirect utility over both periods:

$$\max_{c\in\mathbb{C}}[V^s(x_c, t_c, Y^s_{c2}) + \mu V^p(x_c, t_c, Y^s_{c2})] \tag{4}$$

Equation 4 can be linked to household decision functions from the experimental results.

In the choice task, I characterize the solution to the maximization problem in Equation 4 as being noisily reported by either the student or the parent. In the experiment, subjects are presented with random variation in x_c, t_c, Y^s_{c2}, and either V^p or V^s depending on whether they are choosing as a student or a parent. Since V^p is exogenously given to students, and V^s is exogenously given to parents, this means that students and parents are, respectively, solving

$$\max_{c\in\mathbb{C}}[V^s(x_c, t_c, Y^s_{c2}) + \mu V^{p*}_c + \varepsilon^s_c] \tag{5}$$

$$\max_{c\in\mathbb{C}}\left[\frac{V^{s*}_c}{\mu} + V^p(x_c, t_c, Y^s_{c2}) + \varepsilon^p_c\right] \tag{6}$$

where the objective function for parents is divided through by μ to simplify later calculations.[2] V^{p*}_c and V^{s*}_c are exogenously given indirect utility for parents and students, respectively, and ε^s_c and ε^p_c are random decision shifters, assumed to follow a generalized extreme value distribution to allow for a multinomial logit choice specification.[3] These shifters can be taken either as specification error, or execution error in reporting the choice. The random decision shifters are

Given random variation in the college attributes and the other person's opinion, the multinomial logit estimation identifies up to a scale parameter the empirical decision functions $\hat{V}^s(x_c, t_c, Y^s_{c2}, V^{p*}_c)$ and $\hat{V}^p(x_c, t_c, Y^s_{c2}, V^{s*}_c)$ (Orme, 2006; Raghavarao et al., 2011). These empirical decision functions are theoretically equivalent to $V^s(x_c, t_c, Y^s_{c2}) + \mu V^{p*}$ and $\frac{V^{s*}}{\mu} + V^p(x_c, t_c, Y^s_{c2})$, as they appear in Equations 5 and 6.

[2] This assumes that $\mu > 0$, a necessary condition for the collective model to be appropriate, rather than a non-cooperative model. In the results section I discuss the possibility that $\mu \leq 0$ for some part of the sample.

[3] The model can also be identified nonparametrically. In work available from the author, I show that sample mean results are robust to the use of the nonparametric Hainmueller et al. (2014) estimator.

These empirical decision functions allow me to calculate the probability of choosing a given college c, and how changes in attributes affect the probability of choice. And, importantly, since V^{p*} and V^{s*} are given exogenously and held constant, the effect of changes in x_c, t_c, or Y^s_{c2} on the probability of choice work entirely through the indirect utility function of the person reporting the decision. For example, the marginal effect of x_c on the probability of student choice, calculated using $\hat{V}^s$, reveals information about the student's theoretical indirect utility V^s, but nothing about the parent's theoretical indirect utility V^p, since parental indirect utility is held constant in $\hat{V}^s$. And so, drawing from Equations 5 and 6,

$$\frac{\partial \hat{V}^i}{\partial a_c} \propto \frac{\partial V^i}{\partial a_c} \, \forall \, i \in \{s, p\}, a_c \in \{x_c, t_c, Y^s_{c2}\} \tag{7}$$

The derivative of the decision function, for the student or parent, is proportional to the derivative of that person's indirect utility function. The derivative is only proportional because the decision function is identified up to a scale parameter. Because the indirect utility function is used here rather than the utility function, these first derivatives do not separately identify the marginal utility associated with x_c, t_c, or Y^s_{c2} directly, but are functions of those marginal utilities as well as the fixed parameters of the model: $\delta^s, \delta^p, r, Y^s_1, Y^p_1$, and Y^p_2.

The effect of the other person's opinion on the reported household decision reflects the bargaining power that the other person holds in the decision. Drawing from Equations 5 and 6, the derivative of $\hat{V}^s$ with respect to exogenous parental opinion V^{p*}_c is μ and the derivative of $\hat{V}^p$ with respect to V^{s*}_c is $\frac{1}{\mu}$. To account for the differences in scale between the theoretical "other person's utility" and the measured variable, as well as the fact that the logit only identifies the model up to a scalar, I use the ratio of the two derivatives to derive information about the bargaining parameter μ.

The bargaining parameter estimate $\hat{\mu}$ is then

$$\hat{\mu} = \left(\frac{\partial \hat{V}^s}{\partial V_c^{p*}} \left(\frac{\partial \hat{V}^p}{\partial V_c^{s*}} \right)^{-1} \right)^{\frac{1}{2}} \tag{8}$$

The empirical model illustrates how a respondent's different empirical responses to college attributes can reveal, separately, information about their own preferences and about the bargaining parameter μ, and also shows how a theoretically-grounded estimate of μ can be derived from empirical decision functions.

3 Experimental Method and Data

I perform a choice-based conjoint analysis to identify the decision functions from the previous section ($\hat{V}^s$ and $\hat{V}^p$). The central approach of a choice-based conjoint analysis is to present subjects with a choice set of goods. Subjects choose their preferred option. The attributes of the goods are varied by the researcher, so it is straightforward to identify the decision function. Conjoint analysis is a popular tool in marketing, environmental, and health fields, but has only recently seen much use in labor or education economics, as in Wiswall and Zafar (2016).

Conjoint analysis allows the collective household decision function to be estimated without access to detailed observational data or actual student/parent pairs. Conjoint analysis has been previously used in the marketing literature to model joint choice (Krishnamurthi, 1988). A line of work by Aribarg and coauthors (Aribarg et al., 2002, 2010; Wang et al., 2013) uses conjoint analysis to identify a joint choice model in which learning another's preferences leads one to revise their own. Even within a more standard economic framework of bargaining problems where preferences are constant, conjoint analysis allows for a straightforward way of identifying a rather difficult model in order to understand important facts about parent and student preferences and how they might fit together to form a household decision.

I ask subjects to choose between four hypothetical colleges, each with five attributes. The subject may also choose a "No college" option.

The subject compares the colleges on the basis of (1) The earnings associated with a degree from the college, (2) How enjoyable they would find classes, (3) How they would rate the quality of the social life, (4) The annual tuition charged, and (5) The opinion of one's parents. There is an alternate "Parent" form of the question in which the subject is told to imagine that they are a parent with a child who is about to go off to college, and are instructed to choose for the child based on the child's ratings and opinions. Each of these attributes varies at four levels:[4]

- Earnings at age 25: $48,000, $54,000, $60,000, $66,000
- Enjoyability of classes:[5] Not a lot, A little, Fairly, Very
- Social life: Poor, Okay, Good, Great
- Annual tuition: $5,000, $10,000, $15,000, $20,000
- Parents'/child's opinion: Hate, Dislike, Like, Love

for a total of $4^5 = 1{,}024$ different possible college profiles. These are grouped randomly into sets of four and the subject chooses between them and the No college option. See Appendix A for the full instrument and instructions. The earnings and tuition levels chosen for inclusion in the experiment are necessarily somewhat arbitrary, but are chosen to represent believable ranges for the earnings of college graduates and tuition.

Outside of Parents'/child's opinion, which allows the collective model to be identified, the particular attributes included are chosen to address several of the major features of college choice. Earnings and tuition directly relate to the financial return of choosing a particular college. Enjoyability of classes and Social life are meant to capture the academic and social consumption value of attending a particular college. Consumption value is presented abstractly rather than concretely

[4] To avoid order effects, the order in which the five attributes are presented is varied across respondents.

[5] I use "enjoyability," a direct measure of consumption value, rather than the "quality" of classes. In trial runs of the experiment that used "quality," respondents were confused by colleges featuring high-quality classes that led to low earnings.

(using variables for the quality of the gym or the diversity of the student body, for example) because this both reduces dimensionality and allows the attributes to be presented as the student's own rating, whereas a gym might matter a lot for one student and not at all for another.

Since these attributes are assigned, several problems relevant to most work on college choice do not apply here. I do not need to distinguish between a college not chosen because it is not preferred, and a college not chosen because it was not in the choice set since the student was not accepted or did not consider it as a possibility. I also do not need to worry about the distinction between a college's actual attributes and what students believe those attributes to be. Indeed, respondents are told to assume that other aspects that influence student choice, such as distance to home, are constant across all options. This does assume that respondents are capable of holding non-stated attributes constant, and do not try to infer them from known attributes.

It is still possible that respondents may make their choice conditional on unstated attributes that they attempt to infer from the information available. If the impact of stated attributes on choice is conditional on some unstated attribute that respondents assume is correlated, coefficients will be biased. This problem could be addressed by having students state the probability of choosing each option rather than selecting one over all others (Manski et al., 1999). However, this would heavily increase the cognitive complexity of the choice task, and so I rely on the ability of respondents to hold other attributes constant.

Subjects are presented with a total of thirteen choice tasks. In six of these tasks, they are shown "Student" questions and told to choose as though they were a student. In another six, they are shown "Parent" questions and told to choose as though they are a parent. An additional warm-up question is given first; this response is discarded. Giving both forms of the task to all respondents allows me to check how these preferences change with age and so determine the plausibility of generalizing to younger respondents, and also encourages subjects to think about how their bargaining partner might make choices.

I collected data using the survey pool at SocialSci.com.[6] There are 964 survey participants, each of which provides twelve usable choice observations, for a total of 11,568 observations. In each observation, the respondent chooses between four random sets of college attributes and a No college option.

Demographic information about the respondents, asked after the choice tasks are completed, is shown in Table 1, breaking out the Student and Parent groups that make up much of analysis. The sample is not specific to actual college students or parents of college-age children. Students are any respondent aged 18-25 who is responding in the Student role. Parents are any respondent aged 39-47 responding in the Parent role. 39 and 47 are the 25th and 75th percentiles of mother's age for an 18-year-old in the National Longitudinal Survey of Youth 1997 cohort.

The sample is better representative of groups that are likely to attend college than of the general U.S. population. Respondents are more female, more white, and have higher levels of education than the U.S. population: more than 60% of the sample has a Bachelor's degree or above. Given the question at hand, this may be a more desirable group to represent than the total population. This representation is in line with the unpopularity of the No college option, chosen in only about 3% of cases, lower than the population rate of college non-attendance. However, the format also may bias respondents towards selecting a college by devoting most of the visual space to the different colleges; results should not be taken as informative about the choice between college and no college.

Consistent with high education levels and also with optimistic self-reporting, 83% of the sample reports "above average" academic ability. Self-reported academic ability is used to get a sense of academic confidence and affiliation, rather than as an objective measure of ability.

The student and parent groups are broadly similar demographically. The students are less

[6] SocialSci.com has been discontinued since the execution of this study. Contact the author for a copy of the recruitment and subject treatment protocols followed by the site. While participation in the survey pool and the survey was voluntary, SocialSci worked to actively recruit respondents in demographic groups underrepresented in the pool. All subjects were given a small amount of compensation (less than $5) for their participation.

Table 1: Respondent Background Information

Attribute	Full Sample Mean	Students	Parents
Male	.418	.406	.443
Female	.573	.578	.544
Neither Gender	.009	.016	.013
White	.733	.691	.763
Black	.087	.066	.105
Hispanic	.070	.066	.092
Asian	.090	.160	.035
Other Race/Ethnicity	.038	.035	.031
Age	36.791	22.012	43.048
Less than HS Degree	.004	.004	.004
HS Degree/GED	.081	.117	.066
Some College, No Degree	.217	.379	.171
1-2 Yr. Degree or Cert.	.088	.055	.123
Bachelor's Degree	.360	.383	.320
Master's Degree	.177	.055	.224
Professional Degree	.033	.004	.039
Doctorate	.039	.004	.053
Self-reported Academic Ability			
Far Below Average	.003	.000	.004
Much Less than Average	.005	.004	.004
A Little Less than Average	.024	.016	.026
Average	.138	.070	.180
A Little More than Average	.322	.359	.316
Much More than Average	.378	.418	.294
Far Above Average	.131	.133	.175
N	964	256	228

Other Race/Ethnicity includes unlisted races, respondents who wrote "mixed" but did not specify which races were mixed, and those who reported "don't know" or would not report. Student is respondents aged 18-25, and Parent is respondents aged 39-47. Full sample mean also includes respondents who were in neither age range.

white, black, or Hispanic relative to the parent group, but more Asian, or other races and ethnicities. They are slightly less male. The students have lower educational attainment, which is to be expected given that many of them are too young to have earned a bachelor's degree. They also see themselves as having higher academic ability than the parents do.

With minor exceptions, choices seem to be made with attention paid to the task. The median completion time for the entire survey is 7 minutes and 49 seconds, with a median time of 22 seconds spent on each choice task disregarding the warm-up task. Colleges that are strictly dominated by another in the choice set appear in 7.4% of choice tasks, but are only chosen in .3% of choice tasks. Observations in which a strictly dominated option is chosen are left in the data to avoid excluding true nonstandard preferences, and because such choices would be expected in small amounts given a choice process with some amount of randomness or execution error.

Given this data I estimate decision functions using a conditional logit specification. For most analyses, the model is fully saturated, with a binary variable representing each level of each college attribute and the other person's opinion, plus a binary variable indicating the No college option.

The model is estimated using the hierarchical Bayes (HB) algorithm (see Rossi et al., 2005; Sawtooth Software, 2009; Sermas, 2014). Rather than estimating population average parameters β, HB generates individual-level coefficients β_i with a Markov chain Monte Carlo algorithm.[7] The HB estimator assumes that each coefficient in the logit model follows a joint normal distribution $\beta_i \sim N(\alpha, \Sigma)$, where β_i is a vector of model coefficients for individual i, α is a vector of the mean for each coefficient, and Σ is a variance-covariance matrix of the coefficients. β_i parameters are updated in each iteration of the Markov chain with a random-walk Metropolis-Hastings algorithm. The algorithm is outlined in more detail in Rossi, Allenby, & McCulloch (Appendix A, 2005).

The HB estimator for the logit model allows for heterogeneity in preferences, which relaxes the independence of irrelevant alternatives (IIA) assumption in standard multinomial logit analysis. However, the normal distribution can give unintuitive estimates for a small portion of the sample; for example, a positive preference for higher tuition. So, I constrain coefficients such that lower tuition, or higher levels of any other attribute, have a positive effect on choice. Estimates without these constraints show similar results, and are available from the author.

HB takes advantage of full-sample choice behavior to inform individual-level estimates. Iden-

[7]For the saturated model I use 8,000 Markov chain iterations, and then another 2,000 draws to generate average coefficient values. The interacted linear model, used as an alternate specification at certain points in the paper, instead uses 100,000 iterations. Convergence logs and images are available from the author.

tifying individual coefficient estimates with only six choice tasks per person/role seems low, but HB is capable of effectively and accurately estimating individual-level heterogeneity in parameters even with few choice observations per respondent (Lenk et al., 1996; Sawtooth Software, 2009). HB performs similarly to finite mixture distributions in Monte Carlo analysis (Andrews et al., 2002) and is chosen here over the use of a finite mixture distribution because it directly models heterogeneity at the individual level, and is designed to produce point estimates of parameters at the individual level.

I take the individual-level coefficients β_i generated by the HB algorithm and compare estimates across respondent groups. These comparisons comprise the main results of the paper.

4 Results

4.1 Model Estimates

I estimate the model as fully saturated: each attribute enters as a series of dummies for each of the four levels of the attribute.[8,9] The model has considerable predictive power. I first estimate the model while leaving out the final choice task for each respondent and role, using estimates from only the first five observations to predict the sixth, out-of-sample observation. On average, the model assigns a choice probability of .777 to the out-of-sample option that is actually chosen, compared to a random baseline of about .25. In 86.4% of cases the option with the highest assigned probability is actually chosen.

In Figures 1a-1c I show the average coefficients over the whole sample, Students, and Parents,

[8]All 964 respondents are used in estimation, even though only individual coefficients for those in the 18-25 and 39-47 age groups are used in most analyses. The larger sample improves population distribution estimates, and thus individual coefficient draws.

[9]This model assumes no interactions between attributes. Average coefficients from a model with a full set of second-level interaction terms between attributes, which specifies the attributes linearly to avoid estimating 180 coefficients, is reported in Appendix B. Qualitative results are similar using the saturated or linear interacted model. For simplicity of presentation, superior fit, and to avoid the use of a linear specification of the attributes, I present results from the uninteracted model.

respectively, estimated using all choice tasks rather than leaving the last one out as in the validation exercise. Students are respondents aged 18-25 responding in the role of student, and Parents are those aged 39-47 responding in the role of parent.

Student and parent average coefficients have some notable similarities. In both cases, the enjoyability of classes ranks as one of the most important attributes, and earnings and social life are ranked very similarly to each other. Consistent with work on the relative influence of consumption value and earnings over other parts of the college decision (Alstadsæter, 2011; Wiswall and Zafar, 2015; Belfield et al., 2016), the consumption value inputs of the enjoyability of classes and social life receive a particularly strong weight as compared to earnings. A one-unit increase in the enjoyability of classes (e.g. going from "A little enjoyable" to "Fairly enjoyable") has more influence on choice than a rather large \$6,000 annual raise or \$5,000 annual tuition increase. Consistent with literature on financial aid (Dynarski, 2002) and debt aversion (Cunningham and Santiago, 2008), responses to tuition are strong relative to earnings. A decrease of \$5,000 in tuition has a similar influence on the decision as an increase of \$6,000 in annual earnings, implying aversion to debt or a rather large discount rate.

There are differences between students and parents as well. Parents place a large weight on the child's opinion relative to the other attributes, giving a one unit change of child opinion similar influence to a one unit of change of the enjoyability of classes. The weight on child's opinion is highly nonlinear. Going from a college that the student "Dislikes" to one that they "Like" in particular has a very large influence on parental choice. Students do not place as much weight on the parent's opinion, giving a one unit change of parental opinion less influence than \$6,000 of earnings or a one unit change in social life. Student weight on parental opinion also exhibits a Dislike/Like discontinuity but it is not as stark as for parents.

Table 2 uses the individual coefficients drawn for each respondent to show the average marginal effects for each level of each attribute, comparing students and parents.[10] Variance in parameters

[10]Marginal effects vary with the choice set. I average marginal effects over 10,000 randomly determined choice sets

Figure 1: Coefficients in Saturated Model

(a) All Responses

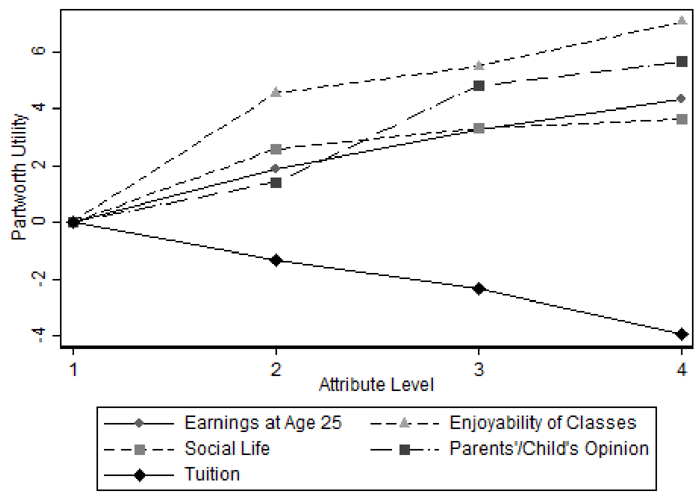

(b) Students

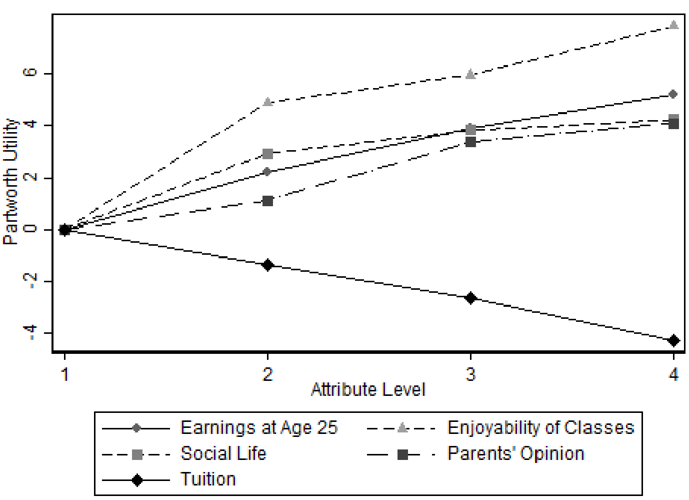

(c) Parents

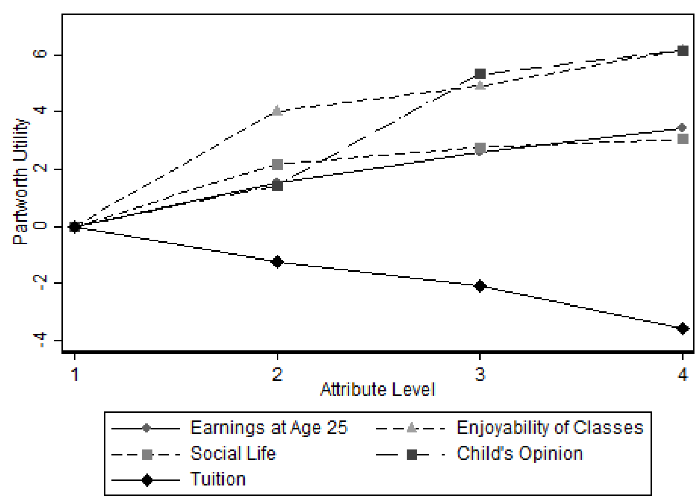

Table 2: Marginal Effects of Students vs. Parents

	Students [SD]		Parents [SD]		Difference (SE)	
Earnings at Age 25						
$48,000 to $54,000	.081	[.041]	.062	[.042]	-.019***	(.004)
$54,000 to $60,000	.085	[.063]	.053	[.052]	-.032***	(.005)
$60,000 to $66,000	.082	[.060]	.054	[.052]	-.028***	(.005)
Enjoyability of Classes						
Not a lot to *A little*	.153	[.062]	.132	[.076]	-.021***	(.006)
A little to *Fairly*	.067	[.059]	.051	[.050]	-.016***	(.005)
Fairly to *Very*	.128	[.089]	.080	[.068]	-.048***	(.007)
Social Life						
Poor to *Okay*	.120	[.055]	.093	[.057]	-.027***	(.005)
Okay to *Good*	.047	[.033]	.031	[.030]	-.016***	(.003)
Good to *Great*	.028	[.038]	.019	[.022]	-.009***	(.003)
Tuition						
$5,000 to $10,000	-.081	[.066]	-.077	[.081]	.004	(.007)
$10,000 to $15,000	-.066	[.054]	-.045	[.046]	.021***	(.005)
$15,000 to $20,000	-.064	[.04]	-.055	[.043]	.008**	(.004)
Parent's/Child's Opinion						
Really dislike to *Dislike*	.044	[.04]	.045	[.043]	.001	(.004)
Dislike to *Like*	.108	[.088]	.178	[.134]	.070***	(.010)
Like to *Love*	.043	[.028]	.062	[.048]	.019***	(.004)

Students are respondents aged 18-25 responding as students; Parents are respondents aged 39-47 responding as parents. Standard deviations are shown in brackets, and the standard error of the difference estimate is shown in parentheses. */**/*** indicates statistical significance at the 10%/5%/1% level.

may represent variance in direct preference for the attribute, or in any other model primitives like discount rate or parental income that determine demand for that attribute, as shown in Equations 1-3. The table also allows for student-parent comparisons on the same scale.

Students and parents respond differently to different college attributes. In each case, student choice responds more rapidly to changes in college attributes than parent choice does. Responses are most similar for tuition, for which parents and students respond very similarly to changes from $5,000 to $10,000 and from $15,000 to $20,000. The differences are largest for earnings and the enjoyability of classes. These differences likely partially reflect who is directly affected by these

per person per attribute.

college characteristics. Parents only indirectly experience earnings and the enjoyability of classes, but may share in paying tuition costs.

Both student and parent responses to linear increases of earnings and tuition are themselves relatively linear. Responses to the enjoyability of classes are weakest in the middle, moving from A little enjoyable to Fairly enjoyable. Social life has less marginal influence at higher levels; both parents and students want to avoid Poor social life, but further improvements are less important.

Parents put more emphasis on the student's opinion than the student puts on the parent's opinion, at least for the higher two marginal effects. The difference is especially strong for the change from Dislike to Like. The stronger response to student opinion implies that the bargaining parameter μ is less than 1, and that student opinion takes priority in household choice. However, students respond to parental opinion as well; μ is greater than 0, which is the implied value in a student-only model.

μ varies depending on how strongly each person feels. When considering colleges held in similar regard by both, estimates of μ based on these averages are $(.044/.045)^{.5} = .991$, $(.108/.178)^{.5} = .778$, and $(.043/.062)^{.5} = .832$, all three tilting bargaining power towards the student, although μ could be as high as $(.108/.045)^{.5} = 1.519$ in some cases if opinions are mismatched.

The exact value of μ depends on the choice set and the identities of the student-parent pair. As long as the student responds positively to the parent's opinion, $\mu > 0$ and the student-only model can be formally rejected. However, there are some respondents with very low responses to the other person's opinion, as shown in Figure 2.

There are quite a few observations where response to parental opinion is close to zero. For these, a non-cooperative model may better describe behavior, since these students report household choices made without much weight on their parents' opinions. Arbitrarily defining a "very low" response as .005 or less (as this would indicate that parental preference has very little meaningful effect on student choice), 2.0%, 4.7%, and 2.0% of the students have a very low response to parental opinion at the lowest, middle, and highest margins, respectively, with 7.0% of the sample

Figure 2: Distribution of Parent's Opinion Marginal Effects Among Students

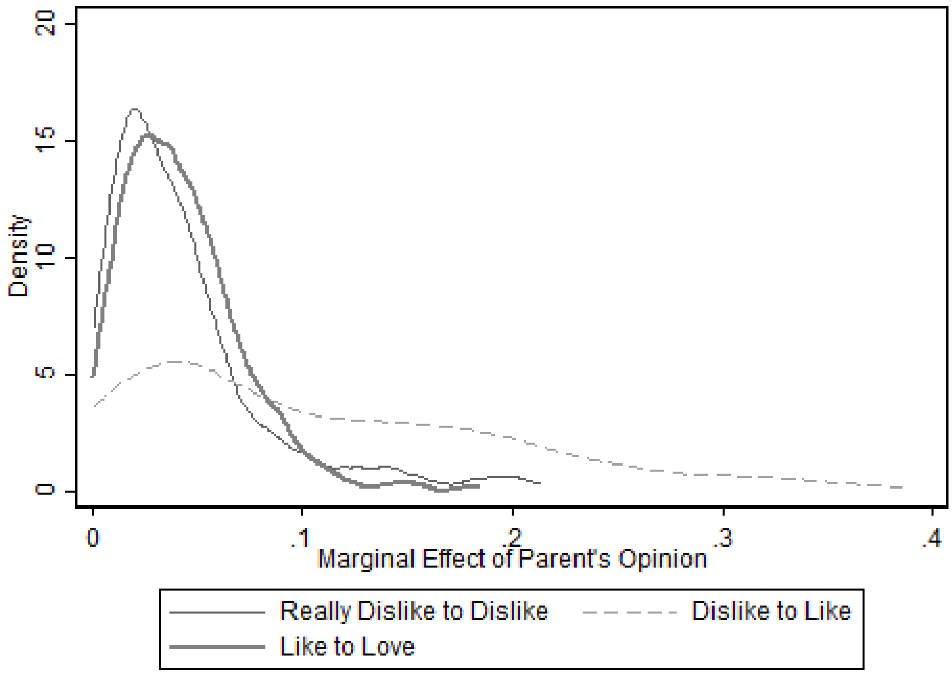

being very low for at least one margin. Black students are more likely than white students to have a very low response (by 12.4 percentage points) as are those with only a high school degree compared to having graduate education (by 6.7 percentage points) although these differences are only weakly significant. This result is consistent with collective choice between husbands and wives; Del Boca and Flinn (2012) show that a non-negligible portion of households engage in non-cooperative household allocation.

The incidence of very strong responses varies distinctly across the margins of parental preference, with the "Dislike to Like" marginal effect showing a much wider variance and many more strong responses. So while the almost-zero student responses are fairly rare, instances where parental opinion could shift the probability of choosing a college by 10 percentage points or more are fairly common (8.2%, 45.3%, and 3.5% by margin). It is important to students that parents not "Dislike" the college, and when the choice is along this margin the parent may have a large share of the bargaining power. But other margins provoke far fewer extreme reactions. There is not a strong demographic pattern for who has these strong responses for the lowest margin. For the middle margin, extreme responses are less likely among those thinking of themselves as academically average (compared to both below- and above-average, by about 30 percentage points), and Asians are more likely than whites to have a strong response at the top margin (9.9 percentage points).

Tables 3 and 4 show how the marginal effects overall, rather than just extreme values, vary by background characteristics. Some categories are combined to avoid small cell sizes. Each background variable type (age, gender, race, education, self-reported academic ability) is run as a separate regression to avoid collinearity. In each case the dependent variable is the marginal effect of moving from the lowest level of the given attribute to the highest.

Given role, respondent age is not a significant predictor of the response to college attributes. This finding supports the analysis in generalizing to the 16-17 year old age group that could not be included in the sample but are certainly of interest when considering college choice, as will be discussed in Section 4.4.

There is relatively little variation in the response to college attributes by background. There is only weak evidence that gender explains these marginal effects for students. For parents, women care less about the enjoyability of classes and the student's opinion than men do. There is also weak evidence that student men care more about social life than others. Racial differences are limited as well: black students care less about the enjoyability of classes and tuition than white students do, Asian students care more about their parents' opinions and less about tuition, and the relatively small Other race group cares much less about classes and social life. Asian parents care much more about earnings than white parents do.

There are surprisingly no significant differences across educational groups for parents. For students, higher education levels are associated with caring more about earnings and parental opinion.

For students, self-reported academic ability matters. Those who see themselves as above average care more about earnings and the enjoyability of classes than others. Reliance on parental opinion varies heavily, with the small group who see themelves as below average responding heav-

Table 3: Student Marginal Effects by Background

	Earnings	Classes	Soc. Life	Tuition	Parent's Opinion
Age	.001	-.004	-.001	-.001	-.002
	(.004)	(.004)	(.003)	(.004)	(.003)
Male (Ref.)					
Female	.006	-.012	-.020*	.021	-.014
	(.018)	(.019)	(.012)	(.016)	(.015)
Neither Gender	-.001	.074	-.080*	-.064	-.034
	(.071)	(.075)	(.048)	(.065)	(.059)
White (Ref.)					
Black	-.055	-.087**	-.032	.069**	-.011
	(.035)	(.036)	(.024)	(.032)	(.029)
Hispanic	-.027	.029	.020	.013	-.003
	(.035)	(.036)	(.024)	(.032)	(.029)
Asian	.028	-.031	-.017	.037*	.044**
	(.024)	(.025)	(.016)	(.022)	(.020)
Other	-.054	-.221***	-.094***	-.051	-.045
	(.048)	(.049)	(.032)	(.044)	(.039)
No College (Ref.)					
College but no BA	.049*	-.020	.005	-.016	.048**
	(.028)	(.030)	(.019)	(.026)	(.023)
BA	.078***	-.026	.004	-.034	.037
	(.029)	(.031)	(.020)	(.027)	(.024)
More than BA	.040	.006	-.011	-.007	.086**
	(.043)	(.046)	(.029)	(.040)	(.035)
Self-Reported Academic Ability					
Below Average	-.087	-.050	-.002	.082	.202***
	(.070)	(.074)	(.048)	(.065)	(.057)
Average (Ref.)					
Above Average	.077**	.075**	-.004	-.038	.066**
	(.034)	(.036)	(.023)	(.031)	(.028)
N	256	256	256	256	256

Students are those aged 18-25 responding in the student role. Dependent variable is the probability change of moving from college with the lowest value of the given attribute to the highest. Standard errors in parentheses. */**/*** indicates statistical significance at the 10%/5%/1% level.

Table 4: Parent Marginal Effects by Background

	Earnings	Classes	Soc. Life	Tuition	Student's Opinion
Age	.004	-.002	-.002	.002	.003
	(.004)	(.004)	(.003)	(.004)	(.005)
Male (Ref.)					
Female	.014	-.055***	-.010	.010	-.060***
	(.017)	(.019)	(.012)	(.019)	(.022)
Neither Gender	-.055	.157*	.048	.002	-.099
	(.074)	(.085)	(.051)	(.083)	(.094)
White (Ref.)					
Black	.009	-.004	.010	-.019	-.028
	(.027)	(.032)	(.019)	(.031)	(.035)
Hispanic	.007	-.016	.007	.012	.057
	(.029)	(.034)	(.020)	(.033)	(.038)
Asian	.121***	-.010	-.003	.014	-.016
	(.045)	(.053)	(.032)	(.051)	(.059)
Other	.077	.124**	.004	.004	-.056
	(.048)	(.057)	(.034)	(.055)	(.063)
No College (Ref.)					
College but no BA	-.004	-.004	-.001	-.020	-.034
	(.035)	(.042)	(.024)	(.039)	(.045)
BA	.025	-.004	.002	.015	-.068
	(.035)	(.041)	(.024)	(.039)	(.045)
More than BA	-.015	-.013	-.017	.011	-.044
	(.035)	(.041)	(.024)	(.039)	(.045)
Self-Reported Academic Ability					
Below Average	.017	-.099*	-.078**	-.002	-.048
	(.049)	(.057)	(.033)	(.055)	(.063)
Average (Ref.)					
Above Average	.014	.006	-.028*	-.001	-.034
	(.022)	(.026)	(.015)	(.025)	(.028)
N	228	228	228	228	228

Parents are those aged 39-47 responding in the parent role. Dependent variable is the probability change moving from college with the lowest value of the given attribute to the highest. Standard errors in parentheses. */**/*** indicates statistical significance at the 10%/5%/1% level, respectively.

Figure 3: Single-Agent and Collective Profile Rankings from a Single Randomly Chosen Student/Parent Pairing

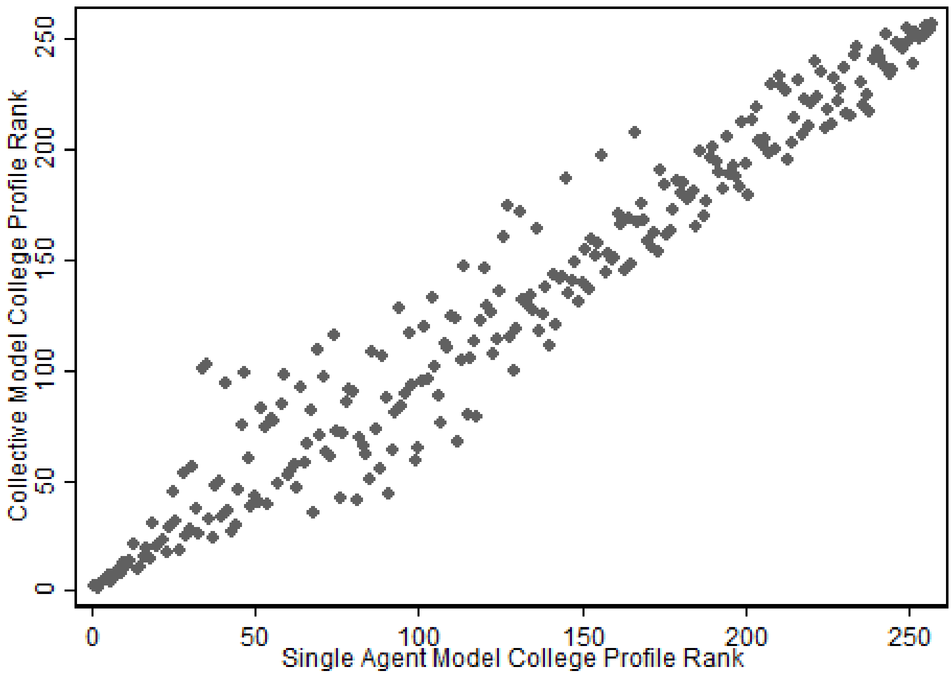

ily to parental opinion.

Based on these results, students and parents respond differently to college attributes. Background characteristics are only occasionally able to explain differences within these groups.

4.2 The Implications of Model Choice: Differences in Prediction

The results in Section 4.1 show differences between Student and Parent preferences, and show that an estimate of the bargaining parameter based on average coefficients is far from 0. The single-agent student-only model can be formally rejected for the average respondent. In this section and the next I evaluate the consequences of choosing a single-agent model over a collective model.

I compare the predictions of two separate models. The first is the collective model, as described in the previous section. The second represents a single-agent model of student choice in which parent's opinion is left out of both the estimation of student coefficients and in determining which college is chosen.

Since the single-agent model has fewer predictors, it will necessarily have worse fit in predicting individual choices. However, the difference may be small or large. As mentioned in Section 4.1, the collective interacted model has strong out-of-sample predictive power, assigning a probability of .777 to the option that is actually chosen. In contrast, the single-agent model assigns on average a probability of .604 to the option that is chosen. An out-of-sample predictive validity of .604 is still strong, but is meaningfully weaker than .777.

However, these estimates of fit reflect fit at the individual-choice level. The proper test of fit is at the collective level. To assess this, I compare the preference ordering of all colleges for each student under the two different models. In the single-agent model, this is a ranking of the student's indirect utility for each of the $4^4 = 256$ college profiles (omitting parental opinion as a predictor), plus the No college option.

This ranking is compared to a similar complete preference ordering as determined by the collective model, linking a single student to a single parent. These links are made by either pairing students and parents randomly or by Gale and Shapley (1962) matching, where students and parents both want to minimize a noise-augmented Mahalanobis (1936) distance between decision function coefficients. The Gale-Shapley matching allows parents to raise children with preferences similar to their own. Matching by other metrics, such as demographics, would generate correlations between the random-matching and the intentional matching versions shown here.

The matching process pairs a student to a parent, and compares single-agent versus collective preference rankings over all possible college attribute sets. In the collective model, the ranking is based on the household decision function as described in Equation 4. V^s and V^p are generated using estimates of the decision functions. μ is constructed as in Equation 8, using marginal effects so that the scale between students and parents is comparable.[11]

[11] Since the marginal effect of the other person's opinion varies based on the college profile, the marginal effect used is based on the other person's individual ranking of profiles. Profiles are broken into quintiles based on the probability of individual choice given the estimated decision function. The bottom, middle, and top quintiles use the Really

Figure 4: Distribution of μ Bargaining Parameter

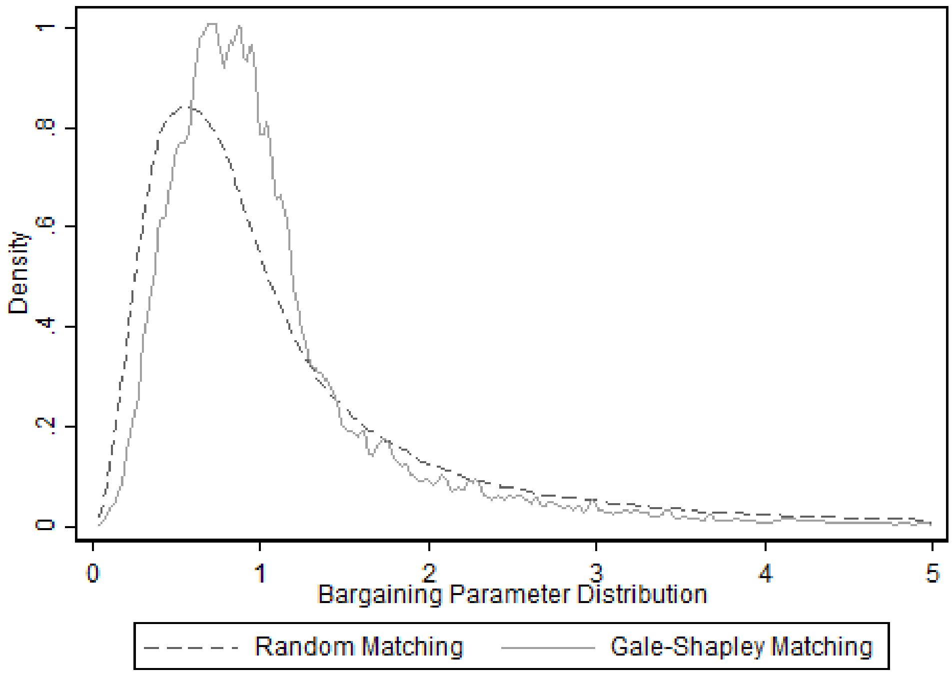

Figure 3 gives an example of one of these comparisons, choosing one student-parent pairing at random. The ranking generated by the single-agent model is along the x-axis, and the ranking generated by the collective model is along the y-axis. In this example, the two models are closely aligned and the rankings have a correlation of .971. I calculate the average correlation over all student-parent pairings. The student-parent pairing is performed 10,000 times using completely random matching, and another 10,000 times using noise-augmented Gale-Shapley matching.

Over all iterations, the average correlation between the single-agent and collective preference rankings when matching randomly has an average of .881 with a standard deviation of .006. When using positive matching, the average is .917 with a standard deviation of .004. Both are close to normally distributed.

The rankings from the single-agent and collective models are fairly closely related, as we might

dislike to Dislike, Dislike to Like, and Like to Love marginal effects, respectively. The second and fourth quintiles use a weighted average, linearly phasing out one marginal effect for the next.

expect given that student preferences are determinants of both, and since parental coefficients are not wildly different from those of students. However, the rankings are not perfectly matched, and the differences between the single-agent and collective rankings are still indicative of a difference between the models.

The matching process also gives a better picture of the distribution of the μ bargaining parameter, which naturally varies depending on which student is matched with which parent, and also depending on the college profile in question. Figure 4 shows the distribution of μ over all college profiles for the first 200 student-parent matchings. Very large values of μ above 5, generated by tiny parental responses to student opinion, are cut off in the image.

The majority (about 60%) of the weight of the μ distribution is below 1. Student opinion matters more than parental opinion in a majority of cases, but certainly not universally. The median μ is .860 using random matching, and .889 using Gale-Shapley. The weight of the decision falls on students, but parental opinion is also clearly important. About .4% of the matches have a μ outside the bounds of .1 to 10, which would indicate heavily one-sided decision-making.

4.3 The Implications of Model Choice: Differences in Inference

The single-agent and collective models can also be compared in terms of the inference that can be drawn from them. Commonly, it is assumed that data on actual student behavior, which is really generated by household preferences, reveals information specifically about the student's utility function and elasticities.

Policy often must be targeted at a particular category of individuals. Informing such policy with individual utility functions inferred from behavioral data actually generated by a household decision function can lead to policy that is suboptimal or completely ineffectual due to being incorrectly targeted, if the other individual is responsible for responsiveness to the given input.

College policies that must be targeted include outreach programs, which decide whether they are attempting to speak directly to students, parents, or both. There is growing interest in a partic-

Table 5: Marginal Effects of Dual vs. Single Agent Models

	Collective Students [SD]		Single-Agent Students [SD]		Difference (SE)	
Earnings at Age 25						
$48,000 to $54,000	.081	[.041]	.095	[.036]	.014***	(.002)
$54,000 to $60,000	.085	[.063]	.097	[.061]	.012***	(.004)
$60,000 to $66,000	.082	[.06]	.102	[.059]	.02***	(.004)
Enjoyability of Classes						
Not a lot to *A little*	.153	[.062]	.171	[.063]	.018***	(.004)
A little to *Fairly*	.067	[.059]	.08	[.065]	.013***	(.004)
Fairly to *Very*	.128	[.089]	.157	[.093]	.029***	(.006)
Social Life						
Poor to *Okay*	.12	[.055]	.148	[.047]	.028***	(.003)
Okay to *Good*	.047	[.033]	.057	[.034]	.01***	(.002)
Good to *Great*	.028	[.038]	.04	[.038]	.012***	(.003)
Tuition						
$5,000 to $10,000	-.081	[.066]	-.119	[.074]	-.038***	(.005)
$10,000 to $15,000	-.066	[.054]	-.072	[.047]	-.006**	(.003)
$15,000 to $20,000	-.064	[.04]	-.073	[.035]	-.009***	(.002)

Students are respondents aged 18-25 responding as students. Standard deviations are shown in brackets, and the standard error of the difference estimate is shown in parentheses. */**/*** indicates statistical significance at the 10%/5%/1% level.

ular flavor of low-cost outreach program, which attempts to improve educational choice, or learn about it, by providing information directly to students or parents (Jensen, 2010; Hoxby and Turner, 2013; Hastings et al., 2015; Bleemer and Zafar, 2015, e.g.). These interventions must be properly targeted, deciding who to educate, in order to function effectively.

I evaluate how modeling affects inference by generating household-choice data using the Gale-Shapley matching method from the previous section. I generate ten sets of student-parent pairings. For each pairing, I randomly generate twelve choice tasks of the same type as were given to the original sample, and have the household choose using their previously estimated decision functions. I use this data to re-estimate the model as though the choices were made by the student alone, and compare the estimates to the "true" student coefficients used to generate the data.

Table 5 shows student marginal effects from the collective model, used as the true parameters to generate simulated choice data. These are compared to marginal effects estimated using a single-agent student model with the simulated choice data, intending to recover the true parameters.

In general, the collective-model student parameters show smaller responses to changes in college attributes than does the model estimated as though the household is a single agent, which makes sense since the single-agent student coefficients also incorporate the responses actually made by parents, which are of the same sign. The difference is most pronounced for the response to the quality of social life. There are also significant differences moving from low to middling levels of earnings and tuition.

These differences point to the potential for being misinformed by studies that treat household-generated behavioral data as though it represents student preferences only. For example, consider a reduced-form study that uses an instrumental variable to identify the influence of the college wage return on observed college attendance. This study will pick up the total effect of earnings on the college decision, including how earnings factor into student preference and how they factor into parental preference. A policy based on this IV study that targets students only and informs them about the college wage return would not take advantage of the part of the behavioral response that operates through parental preference, and so the intervention would have a smaller effect than the IV study indicated it should.

4.4 Generalizing From the Sample

While one goal of this paper is to offer a methodological contribution to the literature on collective choice, Sections 4.1-4.3 aim to present useful, generalizable, results about college choice. But in the context of this paper generalizability is threatened in three ways.

The first concern is the use of an online survey. An online survey sample may not be representative of the general population. However, internet-based surveys, including conjoint analyses, have a good record of providing data that are similar to those collected offline, when adjusted for demographic differences (Melles et al., 2000; Sethuraman et al., 2005). In this study, collegeeducated respondents are overrepresented in the sample relative to the population, and there is no explicit control to correct for this. The sample is more representative of the groups who would actually be likely to be considering college options, rather than the general population.

The second concern relates to the fact that subjects identified as parents or students by age group are assumed to be representative of actual parents and students. The age-assigned Student and Parent groups may not be representative of actual students and parents. Two features of analysis address this issue. First, the general relative importance for students of consumption value, cost, and earnings in college choice mimics results from studies that use data on actual students (e.g. Wiswall and Zafar, 2015; Dynarski, 2002). Second, I find in Section 4.1 that respondents taking the student or parent roles give similar answers regardless of age. Extrapolating from this result, we may expect that respondents who are younger still (16-17, the age of students currently making college decisions) would also be likely to give similar answers.

The third concern is that people, while alone, may not be able to accurately predict how they would respond when choosing alongside an actual bargaining partner. While there is support for the assumption that people make similar (although not identical) choices in stated choice conjoint experiments and in real life (Orme, 2006; Wiswall and Zafar, 2016), to my knowledge there is no work on whether this extends to bargaining problems. Further work contrasting choice in paired and separated conditions is called for, but this remains as an assumption for this paper.

5 Discussion

In this paper, I use a conjoint choice experiment in a novel way to identify a collective household function for choosing between colleges as a function of their attributes.

Consistent with previous literature, decisions respond very strongly to consumption value (Alstadsæter, 2011; Wiswall and Zafar, 2015; Belfield et al., 2016) and tuition (Dynarski, 2002) relative to salary. However, students place significantly more weight on certain college attributes than do parents, in particular future earnings and classroom consumption value.

On average, students have a stronger bargaining position in the choice than parents, but not full autonomy. Compared to a single-agent model, a collective choice model leads to different (and more accurate) predictions of behavior in choice tasks, and better inference of student decision functions.

There are several caveats to the analysis that invite further study.

First, while I add a second agent to the choice model, I implicitly assume a single parent or a pair of homogeneous parents, rather than modeling two parents separately.

Second, the hypothetical nature of the choice tasks removes subjects from the real task of choosing a college while sitting next to an actual bargaining partner. Future work could use the same method with actual paired student-parent data to sidestep part of the issue.

This weakness in the use of a hypothetical conjoint choice method to study intra-household choices must be set against the potential benefits. In particular, this approach allows for a simple way of identifying collective choice models, which often requires very detailed in-house consumption data or convenient natural experiment. Conjoint choice also allows me avoid the issue of choice set construction in college choice, and to randomly vary attributes like social life for which it is nearly impossible to find real-world exogenous shocks, or even to measure. Conjoint analysis is typically able to predict real-world market shares well (Orme, 2006).

The differences in prediction and inference between the single-agent and collective models are found to be meaningful but modest in this paper, rather than dramatic. But even modest improvements in prediction and inference with the collective model suggest its use, and incorporation into a general model of choice when answering questions about education. Differences in inference about the student choice function are particularly noteworthy. Because many studies aim to understand student choice functions, these studies should keep the collective model in mind. The growing number of studies focusing on the influence of uncertainty and risk on educational choice, for example, often implicitly rely on an assumption that the uncertainty in the system is faced solely by a student. This single-agent model will lead to a misunderstanding of the influence of uncertainty on choice. Attanasio and Kaufmann (2014) and Giustinelli (2016) both incorporate collective choice-making into their studies of uncertainty in educational choice, and can act as a guide for moving forward for this line of work.

Policy, as well, can be improved by better understanding the distinction between student and household choice functions. In line with policy prescriptions from the wider husband/wife household choice literature, policy makers could gain leverage by properly considering policy targeting. In addition to college-specific policy that reaches out to potential students in an attempt to entice them away from competitors (Jacob et al., 2017), there is a wide range of potential and existing informational and persuasive outreach policies that attempt to sway student and parent choice about whether or not to go to college, where to go, or what to major in (Swail and Perna, 2002; Domina, 2009; Jensen, 2010; Hoxby and Turner, 2013; College Affordability and Transparency Center, 2014; Hastings et al., 2015; Bleemer and Zafar, 2015).

These outreach policies necessarily pick targets and must attempt to understand what is important to their target group. Policy makers can maximize their effectiveness by targeting themselves properly, both understanding who holds the weight of bargaining power, and who is most interested in the attribute they want to relay information about. Should an attempt to advertise a new aid grant reach out to students or parents? Should a program with the ears of students focus on tuition or "fitting in"? Understanding the difference between what the household values and what their target group values will make these policies more effective. This will mean targeting information about particular attributes towards those who value them the most.

References

Alstadsæter, Annette. 2011. Measuring the Consumption Value of Higher Education. *CESifo Economic Studies* 57(3):458.

Altonji, Joseph G., Erica Blom, and Costas Meghir. 2012. Heterogeneity in Human Capital Investments: High School Curriculum, College Major, and Careers. *Annual Review of Economics* 4(1):185–223.

Andrews, Rick L., Asim Ansari, and Imran S. Currim. 2002. Hierarchical Bayes Versus Finite Mixture Conjoint Analysis Models: A Comparison of Fit, Prediction, and Partworth Recovery. *Journal of Marketing Research* 39:87–98.

Aribarg, Anocha, Neeraj Arora, and H. Onur Bodur. 2002. Understanding the Role of Preference Revision and Concession in Group Decisions. *Journal of Marketing Research* 39(3):336–349.

Aribarg, Anocha, Neeraj Arora, and Moon Young Kang. 2010. Predicting Joint Choice using Individual Data. *Marketing Science* 29(1):139–157.

Attanasio, Orazio P., and Katja M. Kaufmann. 2014. Education Choices and Returns to Schooling: Mothers' and Youths' Subjective Expectations and their Role by Gender. *Journal of Development Economics* 109:203–216.

Belfield, Chris, Teodora Boneva, Christopher Rauh, and Jonathan Shaw. 2016. Money or Fun? Why Students Want to Pursue Further Education. IZA Discussion Paper No. 10136.

Bleemer, Zachary, and Basit Zafar. 2015. Intended College Attendance: Evidence from an Experiment on College Returns and Cost. IZA Discussion Paper No. 9445.

Carneiro, Pedro, and James J. Heckman. 2002. The Evidence on Credit Constraints in Post-Secondary Schooling. *The Economic Journal* 112(482):705–734.

Chiappori, Pierre-Andre. 1992. Collective Labor Supply and Welfare. *Journal of Political Economy* 100(3):437–467.

College Affordability and Transparency Center. 2014. College Scorecard.

Cunningham, Alisa F., and Deborah A. Santiago. 2008. Student Aversion to Borrowing: Who Borrows and Who Doesn't? Tech. Rep., Institute for Higher Education Policy, Washington, D.C.

Dauphin, Anyck, Abdel Rahmen El Lahga, Bernard Fortin, and Guy Lacroix. 2011. Are Children Decision-Makers within the Household? *Economic Journal* 121(553):871–903.

Del Boca, Daniela, and Christopher Flinn. 2012. Endogenous Household Interaction. *Journal of Econometrics* 166(1):49–65.

Domina, Thurston. 2009. What Works in College Outreach: Assessing Targeted and Schoolwide Interventions for Disadvantaged Students. *Educational Evaluation and Policy Analysis* 31(2): 127–152.

Duflo, Esther. 2003. Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *The World Bank Economic Review* 17(1):1–25.

Dynarski, Susan. 2002. The Behavioral and Distributional Implications of Aid for College. *American Economic Review (Papers and Proceedings)* 92(2):279–285.

Ermisch, John. 2003. *An Economic Analysis of the Family*. Princeton, NJ: Princeton University Press.

Fortin, Bernard, and Guy Lacroix. 1997. A Test of the Unitary and Collective Models of Household Labour Supply. *The Economic Journal* 107(443):1–26.

Gale, David, and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69(1):9–15.

Giustinelli, Pamela. 2016. Group Decision Making with Uncertain Outcomes: Unpacking Child-Parent Choices of High School Tracks. *International Economic Review* 57(2):573–602.

Giustinelli, Pamela, and Charles F. Manski. 2017. Survey Measures of Family Decision Processes for Econometric Analysis of Schooling Decisions. *Economic Inquiry* (Forthcoming).

Hainmueller, Jens, Daniel J. Hopkins, and Teppei Yamamoto. 2014. Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments. *Political Analysis* 22:1–30.

Hastings, Justine S., Christopher A. Neilson, Anely Ramirez, and Seth D. Zimmerman. 2017. (Un)informed College and Major Choice: Evidence from Linked Survey and Administrative Data. *Economics of Education Review* (Forthcoming).

Hastings, Justine S., Christopher A. Neilson, and Seth D. Zimmerman. 2015. The Effects of Earnings Disclosure on College Enrollment Decisions. NBER Working Paper No. 21300.

Hoxby, Caroline, and Sarah Turner. 2013. Expanding College Opportunities for High-Achieving, Low Income Students. SIEPR Discussion Paper No. 12.

Hoxby, Caroline M., and Christopher Avery. 2013. The Missing ”One-Offs”: The Hidden Supply of High-Achieving, Low Income Students. *Brookings Papers on Economic Activity* 2013(1): 1–65.

Jacob, Brian A., Brian McCall, and Kevin M. Stange. 2017. College as Country Club: Do Colleges Cater to Students' Preferences for Consumption? *Journal of Labor Economics* (Forthcoming).

Jensen, Robert. 2010. The (Perceived) Returns to Education and the Demand for Schooling. *The Quarterly Journal of Economics* 125(2):515–548.

Kalenkoski, Charlene. 2008. Parent-child Bargaining, Parental Transfers, and the Post-secondary Education Decision. *Applied Economics* 40(4):413–436.

Krishnamurthi, Lakshman. 1988. Conjoint Models of Family Decision Making. *International Journal of Research in Marketing* 5(1982):185–198.

Lenk, Peter J., Wayne S. DeSarbo, Paul E. Green, and Martin R. Young. 1996. Hierarchical Bayes Conjoint Analysis: Recovery of Partworth Heterogeneity from Reduced Experimental Designs. *Marketing Science* 15(2):173–191.

Lundberg, Shelly, Jennifer L. Romich, and Kwok Ping Tsang. 2009. Decision-making by Children. *Review of Economics of the Household* 7(1):1–30.

Lundberg, Shelly J., Robert A. Pollak, and Terence J. Wales. 1997. Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *Journal of Human Resources* 32(3):463–480.

Mahalanobis, Prasanta C. 1936. On the Generalised Distance in Statistics. In *Proceedings of the national institute of sciences (calcutta)*, 49–55. Calcutta, India.

Manski, Charles F., Kenneth I. Wolpin, and Elke U. Weber. 1999. Analysis of Choice Expectations in Incomplete Scenarios. In *Elicitation of preferences*, ed. Baruch Fischhoff and Charles F. Manski, 49–72. Springer Netherlands.

Melles, Torston, Ralf Laumann, and Heinz Holling. 2000. Validity and Reliability of Online Conjoint Analysis. In *Proceedings of the sawtooth software conference*, 31–40. Sequim, WA.

Orme, Brian K. 2006. *Getting Started with Conjoint Analysis*. Madison, WI: Research Publishers.

Raghavarao, Damaraju, James B Wiley, and Pallavi Chitturi. 2011. *Choice-based Conjoint Analysis*. Boca Raton, FL: CRC Press.

Rossi, Peter E., Greg M. Allenby, and Robert Edward McCulloch. 2005. *Bayesian Statistics and Marketing*. Hoboken, NJ: J. Wiley & Sons.

Sawtooth Software. 2009. The CBC / HB System for Hierarchical Bayes Estimation Version 5.0. Tech. Rep. 360, Sawtooth Software.

Sermas, Ryan. 2014. Package 'ChoiceModelR'. Tech. Rep., Decision Analyst, Inc., Arlington, TX.

Sethuraman, Raj, Roger A. Kerin, and William L. Cron. 2005. A Field Study Comparing Online and Offline Data Collection Methods for Identifying Product Attribute Preferences using Conjoint Analysis. *Journal of Business Research* 58(5):602–610.

Stinebrickner, Ralph, and Todd R. Stinebrickner. 2014. A Major in Science? Initial Beliefs and Final Outcomes for College Major and Dropout. *Review of Economic Studies* 81(1):426–472.

Swail, Watson Scott, and Laura W. Perna. 2002. Pre-College Outreach Programs: A National Perspective. In *Increasing access to college: Extending possibilities for all students*, 15–34. New York City, NY: State University of New York Press.

Thomas, Duncan. 1990. Intra-Household Resource Allocation: An Inferential Approach. *The Journal of Human Resources* 25(4):635–664.

Wang, Jing, Anocha Aribarg, and Yves F Atchadé. 2013. Modeling Choice Interdependence in a Social Network. *Marketing Science* 32(6):977–997.

Wiswall, Matthew, and Basit Zafar. 2015. Determinants of College Major Choice: Identification using an Information Experiment. *Review of Economic Studies* 82(2):791–824.

———. 2016. Preference for the Workplace, Human Capital, and Gender. NBER Working Paper No. 22173.

Zafar, Basit. 2012. Double Majors: One for Me, One for the Parents? *Economic Inquiry* 50(2): 287–308.

Appendix A Experimental Instrument

Respondents in the sample pool are shown the opportunity to take the survey. Those who opt in are shown a standard disclaimer provided by the institutional review board before beginning. The experiment is presented as a series of slides. On each slide the participant is given a set of instructions or a choice task. In this section I list the exact wording used on each slide.

Slide 1

In this survey, you will choose between hypothetical colleges for yourself or for your child.

For each question, you will be shown the profiles of four colleges. The profiles will list a few things about each college. Assume that for everything not mentioned, the colleges are identical. When choosing a college, you will have to make trade-offs. One college may be better than another at one thing (perhaps it has lower tuition) but worse at another (perhaps the classes are not enjoyable). Try to think about which mix of attributes you would prefer most. It is important that you try to choose the college that you would actually select if presented with the decision in real life.

Slide 2

[An example of a hypothetical college is shown here with the below-listed attributes]

Here is an example of a college. In the college described above, if you attend College A:

- You will earn a salary of $60,000 per year at age 25
- Your parents would like for you to go to this school
- You would not enjoy the classes at all
- You would rate the social life as poor
- You would pay $15,000 in tuition per year (or take on loans to cover what you can't afford)

Slide 3

[This slide is shown first to those who are randomized into the "perform choice tasks as student first" version. Otherwise, the parental version (here slide 11) is shown first.]

Imagine that you are just finishing high school. You can choose which college to go to and you are trying to decide which college to choose (or no college at all). Assume that for all attributes not listed, the colleges are exactly the same. If the tuition is too high for you and/or your parents to afford, assume the rest can be covered by student loans.

Slides 4-10

[These seven slides show the choice tasks where the respondent plays the role of a student.]

Slide 11

Imagine that your child is just finishing high school. You can choose which college to send them to and you are trying to decide which college to choose (or no college at all). Assume that for all attributes not listed, the colleges are exactly the same. If the tuition is too high for you and/or your child to afford, assume the rest can be covered by student loans.

Slides 12-17

[These six slides show the choice tasks where the respondent plays the role of a parent.]

Slide 18

[This slide shows questions about respondent background, including gender, age, and educational attainment.]

Appendix B Additional Results

Table 6: Coefficient Means in Interacted Linear Model

Attribute	Students	[SD]	Parents	[SD]
Earnings at Age 25	.173	[.187]	.114	[.157]
Enjoyability of Classes	1.243	[1.016]	.696	[.715]
Social Life	.084	[.118]	.146	[.203]
Tuition	-.277	[.259]	-.236	[.331]
Parent's/Child's Opinion	1.158	[1.160]	2.148	[1.863]
Earnings × Classes	.672	[.639]	.301	[.726]
Earnings × Social	.339	[.638]	.253	[.650]
Earnings × Tuition	-.215	[.639]	-.036	[.738]
Earnings × Opinion	.383	[.805]	.266	[.879]
Classes × Social	.511	[.833]	.195	[.747]
Classes × Tuition	-.324	[.445]	-.289	[.548]
Classes × Other	.352	[.722]	.724	[.837]
Social × Tuition	-.165	[.619]	-.097	[.666]
Social × Tuition	.197	[.590]	.158	[.631]
Tuition × Other	-.282	[.613]	-.39	[.753]

Students are respondents aged 18-25 responding as students; Parents are respondents aged 39-47 responding as parents. Standard deviations of coefficient distributions are shown in square brackets.