

Human Capital and Economic Opportunity Global Working Group FINet Chicago 2012 Conference Report

Introduction

Research presented at the Family Economics and Human Capital in the Family Workshop sheds new light on issues of marriage patterns, women's labor force participation, parental investments in children, child poverty, and child-parent decision-making. Among the diverse insights in work presented at the conference were the startling fact that changes in cohort size alone appear to explain a large fraction of the variation in marriage rates over time, that poverty may be starkly higher at the individual level than it is when measured at the household level only, and that parents may underinvest in children not only because they are time constrained, but also because they underestimate the value of such investments. Pierre-Andre Chiappori, the group leader of the Family Inequality Network, mentioned in introductory remarks that a major focus of the network is promoting cutting edge research by young scholars, ensuring that these important issues of family welfare continue to be explored by the next generation of researchers. The papers shared at the conference provide a rich basis for such future work.

Analyzing Marriage Patterns and Decision-Making

The decision to marry is often the first step in forming a family unit, precipitating a host of economic decisions including how to contract over and divide resources, whether to have children and how much to invest in them, and even when to participate in the labor force and how much labor to supply. Research presented at the conference provided insight into why people marry, how they divide assets and utility within a marriage, and how the marriage market may interact with pre- and post-marriage decisions.

Marriage rates have been falling in the United States for several decades, but a historical examination of data back to 1900 reveals substantial variation in marriage rates prior to this decline. Research presented by Maurizio Mazzocco indicates that these changes in marriage rates in the US throughout time are tightly linked to changes in cohort size. Why? If men can marry when either old or young, but women can only marry when young (empirically, older women have low marriage rates, and women are much less likely than men to remarry following divorce), changes in cohort size are, effectively, changes in the gender ratio for the current single population, which can drive changes in marriage rates. Mazzocco, with co-author Mary Ann Bronson, offers a search model of the marriage market, where individuals will only marry when the gains from the current match exceed the expected value from continuing to search. If the current cohort size is large relative to the previous cohort size, then the young women are facing fewer potential male partners, since the pool of eligible men, made up of last cohort's unmarried older men and this cohort's young men, is smaller. Additionally, the current cohort of men faces a larger pool of potential spouses, and thus has an incentive to search for longer before choosing a mate. Taken together, these factors will result in a negative correlation between the marriage rate and cohort size. When cohort size increases, fewer people marry. When it decreases, women experience a relative increase in available men, and men face a tighter marriage market, and more people marry. Indeed, Bronson and Mazzocco demonstrated that in data since the 1900s, the

FINet Chicago 2012 Conference Report, continued

cohort size is tightly linked to marriage rates (as measured by the share of individuals never married by age 30, 35, or 40), for both black and white individuals. This correlation breaks down somewhat at the end of the 20th century, but is restored if cohabitation as well as marriage is considered, since cohabitation increasingly has substituted for marriage.

Existing theories in changes in marriage rates over time, such as birth control, household division of labor, incarceration, and government assistance, only apply to certain time periods and certain demographic groups, but fail to match the pattern in marriage rates throughout the entire century. Cohort size remains a strong predictor of marriage rates when variation between states is considered, as well as when an instrument for cohort size, World War II deployments, is considered.

Research presented by Alessandra Voena examines potential welfare gains from prenuptial agreements. Such agreements are uncommon in the United States, perhaps due to high upfront costs and low enforcement, but Mahar (2003) finds that 20% of US couples have either signed a prenuptial agreement or would like to sign one. Voena and coauthor Denrick Bayot examine the role of explicit contracting in marriage in an environment where a rudimentary prenuptial agreement is easy and cheap to form. In Italy, all couples can decide whether to enter into a communal property or separation of property agreement at the onset of marriage. As in the U.S., the default choice is communal property, but unlike in the U.S., the ease of contracting results in 64% of newlyweds choosing separation of property. Legal scholars advise that separation is suitable in household in which both spouses contribute to income, but if one spouse works and the other specializes in household duties, then the couple should consider community property. Using administrative data from 2000 to 2009, Bayot and Voena find systematic differences between couples that choose separation versus community property, both at the time of marriage and at the time of divorce. Couples where the woman works at home are more likely to choose community property, whereas women who work outside the home are more likely to be in couples choosing separation. Similarly, more educated women choose separation, while women who marry young are more likely to choose community property. Couples in which one member is self-employed are also more likely to choose separation, perhaps to protect the other partner from financial liabilities resulting from the business. Among divorcing households, mothers in community property agreements are more likely to get custody of the children and to be granted alimony. Separation of property is more common among couples that divorce than those that stay together, either indicating selection or ex-post impacts of the choice to separate property.

Voena presented a model where spouses face a joint budget constraint, but, under separation of property, they maintain their own assets upon divorce, whereas under community property, they split property equally upon divorce. Because even with separation of property spouses could always choose to share property the way they would under community property, there has to be a reason for some couples to prefer community property. Candidate mechanisms presented by Voena include: communal property being the default value (but not actually preferred), borrowing constraints that limit transfers from one spouse to the other, and lack of ability to commit to share property otherwise. A preliminary calibration of the model indicates that the choice to specify separation of property at the start of marriage offers large welfare gains, and in fact there may be additional gains to changing the default choice to separation. In addition, the model indicates that the higher divorce rates for couples who choose separation of property is due to selection, and not to a causal impact of prenuptial agreements on marital stability. These findings suggest that

FINet Chicago 2012 Conference Report, continued

lowering the cost of prenuptial agreements in the U.S. and other countries might increase the wellbeing of married couples.

Education choices affect one's marriage market opportunities and post-marriage labor supply, and incentives to participate in the labor market after marriage will impact pre-marriage investments, as well as matching decisions. Yet, much research evaluating the impact of policies that affect educational investments or labor supply decisions look at them in isolation, without acknowledging ancillary effects throughout the lifecycle. Research presented by Monica Costas Dias tackles this issue. Dias, together with co-authors Pierre-Andre Chiappori and Costas Meghir, generates an integrated model of pre-marriage education choice, the marriage market, and labor supply, which may be useful in evaluating policies such as education subsidies (which increase pre-marriage educational attainment) or earned-income tax credits (which affect post-matching labor supply) over the entire lifecycle.

In this model, individuals first make their educational choices, then they match with a marriage partner without knowing what outcomes will realize for each partner (meaning that there exists a risk that will be efficiently shared between the spouses), and finally they make decisions about labor supply and consumption in a collective household framework. The model predicts positive assortative matching, with a sharing rule that determines the division of the marital surplus. This anticipated surplus from different decisions in turn determines the size of educational investments undertaken, and who chooses to enter the marriage market. The equilibrium sharing rule must ensure that the same number of men as women choose to enter the marriage market, or receive utility gains from marriage versus remaining single.

This model can be used to simulate the impacts of policy interventions that change incentives to make educational investments or participate in the labor force.

Macro and Micro Impacts of Changes in Women's Labor Force Participation

The 2008 to 2009 recession has been called a "mancession" due to the male unemployment rate exceeding the female unemployment rate. However, research presented by Stefania Albanesi points out that this trend actually started in the 1980s. Prior to 1980, the female unemployment rate tended to be higher than the male unemployment rate. After the 1980s, this gap disappeared, with female unemployment and male unemployment tending to be equal, except in recessions when male unemployment is higher. What accounts for this trend reversal? Albanesi's research points to greater female labor force participation and labor force attachment as the driver of this reversal. The rate of unemployment is calculated as the number of unemployed over the number employed plus unemployed workers. Non-participants are not counted among the unemployed. So, even though greater labor force attachment—women who become unemployed looking for new work, rather than leaving the labor force—increases the number of unemployed women, as long as labor force participation has also increased, the number of employed over unemployed will rise, and unemployment will fall.

This greater labor force participation and attachment (and at the same time slightly lower participation and attachment on the part of men) explains the convergence in unemployment rates better than convergence in age or skill distribution. Moreover, changes in industry composition cannot explain the shift.

FINet Chicago 2012 Conference Report, continued

What industry composition does account for, though, is the gap between male and female unemployment in recent recessions. Men tend to work in more cyclical industries—without this industry difference the gap in recession unemployment would disappear. A recent flattening of female labor force participation also helps explain another recession trend: jobless recoveries. In recessions in the 70s, 80s, and early 90s, labor force participation tended to increase as unemployment recovered from recessions. This was largely driven by women continuing to join the labor force as employment picked up. In later recessions, though, women's overall labor force participation has started to flatten out, decoupling participation growth from recession recovery.

Research presented by Valerie Lechene shows that the overall greater labor force participation by women since the 1980s might also have micro impacts on how families spend their time, divide tasks, and consume. In the UK, the share of the food budget being spent on processed food increased from 40% to 60% between 1980 and 2000. One theory could be that the price of processed food has fallen, making it cheaper relative to home cooked food. But, although the price of processed food has decreased, the fall in ingredients for home-cooked food has actually been greater, making processed food relatively more expensive in later years. But then what accounts for the increase in consumption? Prices at the grocery store aren't the only cost that might matter in food choices. Processed food is immediately consumable, whereas ingredients have to be combined with time to make a meal, and that time has been getting more expensive as women's opportunity cost of time has risen. Looking instead at the shadow price of home-cooked food yields the opposite conclusion: once the increase in women's labor force participation, hours, and wages is taken into account, home-cooked food's price has risen sharply relative to processed food, accounting for the changing consumption patterns.

Beliefs About and Impacts of Parental Investments

A primary function of the family unit is the production of children, with parental investment of both time and resources being one of the most important inputs. How much investment parents put into children has lasting consequences for the child's human capital development and future earning capacity. Research discussed at the workshop explored what accounts for the substantial variation in this investment, and what areas policy-makers should consider in trying to encourage human capital development through optimal parental investments.

Research presented by Limor Golan looks to estimate the returns to time spent with children. Golan's work, along with coauthors George-Levi Gayle and Mehmet Soyta, shows that there is lower time-investment in children among black mothers than white mothers. Some of this can be accounted for by differential marital status, because married women invest more time in children, and more white women are married. To better understand the drivers of this investment choice, Golan and co-authors developed a dynastic, altruistic investment model, where parents invest time in kids because they partially incorporate their children's utility into their own decision-making. They use this model to examine optimal parental investment decisions, controlling for costs and returns to parental investments. The costs include such things as the opportunity cost of time based on labor market returns. The returns take into account the educational attainment of children, which affect their lifetime earnings and their marriage market outcomes. Parents in this model choose their own labor supply, fertility, and investment in children through time spent with kids. Estimating the model empirically with data from the PSID, they find that mother's and father's time investments are complementary, with mother's time increasing probability of

FINet Chicago 2012 Conference Report, continued

attending college, while father's investments appear to relate to finishing high school. They also find that the marginal return to mother's time is higher for black families than white families. Moreover, the opportunity cost of time is lower for black mothers in terms of wage rates, and conditional on education, earnings of black individuals are lower than earnings of white individuals and they have weaker patterns of assortative mating. This indicates that the reason for lower investment from black parents is not differing productivity, nor is it differing opportunity costs, but rather a binding time constraint for working black mothers, especially single mothers, who may face drastically lower budget sets and therefore require more hours of work to provide material necessities.

Much research has focused on understanding the human capital development formula, while another strand of literature examines the determinants of parental investment in children. Research presented by Flavio Cunha brings these two strands together, to ask what role lack of knowledge of the technology of childhood development could play in apparent underinvestment by low-income mothers. Cunha notes that there is large divergence in skill by age five between children from the bottom and top income quartiles. Moreover, mothers at the top income quartile invest more hours per day in their children. This difference may be due to differing preferences and constraints (e.g., poorer mothers needing to work more and having less time to spend with children), but it also could be due to parents misjudging the importance of their time as an input into the production function. If parental beliefs about return to investment diverge, then investments would diverge even under similar preferences and constraints. Cunha and coauthors Jennifer Culhane and Irma Elo have compiled new data to address this question, surveying 200 Medicaid-eligible mothers to understand subjective beliefs on skill formation relative to objective measures. Parents' response to scenarios indicating varying child weight/height and varying parental time investment demonstrated an understanding that both health and parental investment are positively correlated with child development, and also a high perceived degree of complementarity between child health and investment. Compared to estimation of the importance of time investments from CNLSY, though, mothers in this sample consistently underestimate the importance of parental time investment. These underestimates are less severe for "less disadvantaged" mothers among the disadvantaged population. This research indicates that parental beliefs about the production function, and not just preferences or constraints, may be driving low parental investment among disadvantaged parents, and therefore informational interventions on the importance of parental time may be effective.

Going Beyond Household Poverty Measures

Parental investment in children includes resources as well as time, and these resource allocation decisions have consequences for child welfare. But research presented by Arthur Lewbel and Krishna Pendakur suggests that traditional measures of household-well-being may systematically underestimate child poverty, by assuming all household-members share resources equally.

In general, poverty is measured as a function of total household income, relative to the number of members in the household, and assuming some "returns to scale" in household size. For example, in the US, a single person living alone is considered to be in poverty if he or she earns less than \$11,170 per year, a family of two less than \$15,130 per year, a family of 4 less than \$19,090, and so on. The traditional approach to these questions relies on the notion of an "equivalence scale," defined from the respective levels of income that would be needed by families of different

FINet Chicago 2012 Conference Report, continued

compositions to reach “the same level of utility.” The shortcomings of this definition, however, are well known. From a theoretical perspective, comparing “utility” of different families does not make much sense (actually, the mere notion of family utility is debatable). Moreover, equivalence scales are not in general identifiable empirically. Lastly, from a more applied perspective, the crucial notion of intra-household allocation is disregarded. What if the household head consumes more than other members of the family? What if a four-person household has total income of \$20,000, but the parents consume almost all of that income, leaving the children underfed and neglected?

A recent development of this literature, due to Browning, Chiappori and Lewbel, has been the notion of “indifference scales”; these are defined from the income a single person would need to be as well off as that individual would be in a given family context. The concept of indifference scales only refers to individual utility, and does not require interpersonal comparisons; moreover, it emphasizes the crucial role of intrahousehold allocation. In their initial paper, Browning, Chiappori and Lewbel showed that indifference scales can be non-parametrically identified from consumption data. However, such an identification requires price variations and the independent observation of singles’ demand functions. The work presented by Lewbel and Pendakur extends this approach in two directions. First, they relax the need for price variations; in their context, indifference scales can be identified from Engel curves. Second, they only require demand functions to be observed for families of different composition. The crucial assumption is independence of base; i.e., the shares of income going to each family member should be independent of income, at least for low-income levels (although they may of course depend on family composition). This assumption is well supported by the data; for instance, the ‘revealed preferences’ approach of Cherchye, De Rock, Lewbel and Vermeulen (2012) finds the wife’s resource share is between .4 and .6 and does not appear to vary with total income. However, the resource shares do appear to vary systematically with relative wages—the higher the wife’s wage is compared to the husband’s wage, the greater share of resources she appears to get. A second assumption is needed to obtain non-parametric identification; namely, preferences have to be “similar” across different individuals. Then the observation of demand for at least one “assignable” good, such as clothing specifically for men, women, and kids, is sufficient to recover the resource share.

Dunbar, Lewbel, and Pendakur (2012) use this method to examine poverty in Malawi. Using survey data that includes households with different numbers of kids, they find that the man’s share appears to remain fixed as the number of household members increases—that is, they can’t reject that the father consumes 40% of all household resources constantly, no matter how many children are present. As a result, the share consumed by the mother and each child shrinks as the number of children increases, perhaps leaving mothers and children in larger families more vulnerable to poverty than total household income would suggest. This measurement strategy goes beyond the question of whether the household is living in poverty and begins to address how many men, women, and children are living below the poverty threshold. Using this measure in Malawi, for example, might result in a finding of fewer men living in poverty than suggested by household statistics, but potentially many more women and children. Identifying these individuals in need is a first step toward public policy that addresses those needs.

Children as Agents

The research looking at parents' investment in children and resulting child welfare tends to treat children as passive recipients of parental benevolence. Several papers presented at the workshop go beyond this framing, to examine the interactions between parents and children as multi-party games.

Research presented by Pamela Giustinelli empirically examined the decision framework for Italian teenagers deciding on a high school track. This decision is complicated not only by the dual decision-makers, parents and children, who may have competing preferences, but also by differing beliefs regarding the child's future desires and likely outcomes. In this setting, the choice of a high school track (academic versus vocational) determines future career and education options. Giustinelli's unique data set of 1,029 9th grade students and their parents, by separately eliciting parents' and children's beliefs about future outcomes and preferences over tracks, and relating them to the final observed decision, provides insight into the decision rules used.

Through retrospective surveys on the decision-making process, Giustinelli finds that about half of the parent-child pairs were already aligned at the time of decision-making, and for those that were not aligned, the child's preferences appear to dominate more frequently. The data also includes the self-reported decision rule used by the family. When the child reports deciding unilaterally, then the final decision is more likely to align with the child's stated preferences, regardless of their parents' preferences. However, when the child reports either deciding together with the parents, or listening to the parents' input and then making a decision, the results are more mixed. This shows that the degree to which the parents' beliefs and preferences are incorporated, or the parental weight in decision-making, is indeed more strong when the child reports joint, rather than unilateral decision-making. This work indicates that children cannot be treated as "blank slates" for parental decisions, but must be considered as potential decision-makers within the household. This is true even before children become fully autonomous agents—Giustinelli's work allows for an "in-between" period of adolescent development where parental input and beliefs are still incorporated into decision-making, but the child's own preferences and beliefs play an important (and sometimes dominant) role.

In the presence of both subjective beliefs about the future and divergent preferences, the most insistent decision-maker may not always be the "best" decision-maker. For instance, children might prefer the less-challenging "vocational" track because they are driven by short-sighted utility concerns, in which case the parent might be the "better" decision-maker. Alternatively, the child could have better insight into their potential future preferences, and therefore may be more likely to make the right decision. Future research might therefore examine how policy intervention should aim to redistribute decision-making power, if at all, or improve the decision-making process for choosing binding educational tracks.

Research presented by Marco Cosconati examines the impact of differing parenting styles on children's human capital accumulation, including cognitive and non-cognitive skills. By generating a model of parental behavior as a dynamic game where parents and children have different preferences, Cosconati aims to shed light on whether it is better to be the strict "bad cop" parent, or lenient "good cop" parent.

FINet Chicago 2012 Conference Report, continued

The child in this model cares about leisure and adult human capital. The parent cares about adult human capital and experiences a cost of monitoring the child's behavior. The expected outcome of this interaction is an important social issue if the socially optimal parental monitoring is higher than realized monitoring, due to the high private cost born by the parent of monitoring. But is stricter parenting always better? By capturing the dynamic impacts of strict parenting, in terms of its effects on the child's subsequent behavior, Cosconati's model captures the potential for costs and benefits of strict parenting, and thus variance in optimal parenting behavior. For example, the model predicts that being an extremely strict parent may create "backlash" by giving the child a low incentive to exert effort, an effect that would not be possible in a non-dynamic model. Additionally, the optimal parenting style will depend on the amount of self-control already being exercised by the child, and naturally this may also change with time as the child acquires greater human capital. Using data from the NLSY 97, which asks youth about who makes decisions regarding different discretionary activities, Cosconati finds his model matches several features of the data, such as that parents' strictness falls with age, as self-control increases. However, the model predicts more strictness than is seen in the data. This model has implications for the study of government policies aimed at promoting certain parental behavior, or decreasing the scope for moral hazard by reducing monitoring costs in parent-child interactions.

Research presented by Juan Pantano examined a potential origin of birth order effects, a commonly documented phenomenon whereby children born later have lower outcomes, such as schooling attainment, than early-, and especially first-, born children. Birth order effects are robust to controlling for initial ability and family fixed effects, and are present even in the sub-sample of intact families. Pantano presented one potential mechanism that would create birth order effects—parents taking actions to enforce rules on children both for the benefit of the current child and for the effect on their reputation for "toughness" with future children. As birth order increases, the latter effect is diluted, and so the cost of being a "tough" parent may gradually become too great relative to the benefit for the remaining children. As the benefit of retaining the reputation for toughness declines, parents become more lenient, potentially leading to lower educational outcomes for later children.

Pantano's model involves two types of parents—those that are naturally tough, and those that are lenient but may wish to mimic tough types to discourage bad behavior. But although these parents are rewarded for toughness through their child's school performance, they receive direct disutility from punishing their children. The children, on the other hand, receive some utility from their outcomes, but disutility from putting in effort—they also, naturally, dislike being punished. If a parent punishes children when low outcomes are seen, they develop a reputation for toughness, that carries over to future children, and is anticipated when these children choose how much effort to put in. This "reputation bonus" to punishment decreases from first-born to last-born. Therefore, the model predicts that punishment for the same level of underperformance should decline with birth order; and Pantano shows this effect is seen empirically in data from the NLSY. This model presents the possibility that some portion of birth-order effects may be due to the declining need with each subsequent child for parents to establish a reputation for toughness, relative to the cost of punishing children. By decreasing the cost of punishing bad behavior or increasing the ability to monitor a child's effort, it may be possible to reduce this component of the birth order effect.