

An Economic View of Culture, Identity, and Family Change

Essay prepared for the Working Group on Explaining Family Change, NICHD

May 2006

Rachel Kranton

Introduction

Economists often face the following puzzle. We observe that different people, facing the same set of economic circumstances, make different choices. For example, people may have the same economic incentives (as far as we can observe them) to get and stay married, but some people do not marry, some divorce, and others stay married for years. An easy explanation for such differences in behavior is that people have different “tastes.” Some people just like to be married more than other people. The differences in choices per se are not so puzzling if different choices can be understood as simply arising from differences in people’s tastes.

But what is puzzling is that these differences in choices are often correlated with race, ethnicity, gender or other sociological distinctions. For example, black and white women in identical positions in terms of income, education, wages, etc., make different fertility decisions (on average). They marry at different ages and engage in pre-marital sex at different rates. How then to understand these differences?

Of course, there is yet another easy way to explain these differences. Perhaps we do not observe all the relevant economic variables. For example, we might observe an individual’s level of education, but we do not observe the quality of education. If blacks and whites have systematically different qualities of education, and quality of education is what really matters for decision-making, then we have a problem with our data, but not a conceptual problem. The differences we see can just be explained by unobserved differences in economic characteristics and incentives.

Many economists do not find this explanation satisfactory. The findings of subgroup variation are so robust, in family decisions and many other realms (such as smoking, saving, and schooling), they beg for a better explanation. Many believe there is something else at work other than economic incentives and individual tastes. There might be some other factors involved in decision-making, motivations related to “culture” or “social norms” or “identity.”

This essay describes how economists understand these notions. It then suggests ways such economic theories can inform further theoretical and empirical work on family change.

In recent years, economists have made progress in studying social norms and incorporating social norms into economic analysis. There are two basic understandings of social norms. The first is that social norms are prevailing patterns of behavior. The task, then, is to understand how different patterns emerge. Work in this area describes social norms as emerging from interaction of individuals, each acting in his or her own self-interest. A social norm – a pattern of behavior – is, then, one *equilibrium* of this interaction. Different social norms, i.e., different equilibria, can emerge from the same individuals and the same set of economic circumstances.¹ A central message of this work is that social norms are the product of interaction and can involve behaviors that ultimately are detrimental to individuals and societies.

¹ Classic contributions include George Akerlof, “The Economics of Caste and of the Rat Race and Other Woeful Tales,” *Quarterly Journal of Economics*, Vol. 90, No. 4. (November 1976), pp. 599-617. Jean Tirole, “A Theory of Collective Reputations (With Applications to the Persistence of Corruption and to Firm Quality),” *Review of Economic Studies*, Vol. 63 (1), 1996, pp.1-22, and Peyton Young, *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton University Press, 1998).

The second understanding of social norms is that norms are prescriptions or imperatives for behavior. Different people *should* act in particular ways. People internalize these norms, and hence the norms guide individual decisions. A primary message of this work is that there can be tradeoffs between earning economic returns and abiding by the norms. A second message is that individuals – parents, priests, policy makers etc., can act to try and change the norms and rules that govern behavior.

These theories both raise challenges and suggest new directions for research. Both theories suggest that social norms are endogenous to a particular community or environment. Hence, it will be difficult to identify the impact of norms on individual behavior using traditional data analysis. But there may be ways to explore the connection between norms and behavior by combining traditional data analysis with an understanding of norms gained through historical research, participant observation, or experiments.

The essay has three sections. First, I will give an overview of a standard economic model of decision-making. Second, I will discuss how economists have modeled “culture/social norms.” I will discuss the two traditions mentioned above – norms as deriving from interactions among individuals and norms as imperatives for behavior. The latter derives from work that incorporates culture and identity into economic models of decision-making from my on-going collaboration with George Akerlof.² Finally, the essay will consider new avenues for research, including further theoretical work.

² George Akerlof and Rachel Kranton, “Economics and Identity,” *Quarterly Journal of Economics*, 115(3), August 2000, pp. 715-733. See also Akerlof and Kranton, “Identity and Schooling: Some Lessons for the Economics of Education,” *Journal of Economic Literature*, 40(4), December 2002, pp.1167-1201.

1. Standard Economic Modeling

This section describes a standard economic model of decision making. I will first discuss how economists describe individuals' motivations for behavior. I will then describe how we place individual decision making in a wider context of the individuals' economic and social environment. Along the way, I will also describe why economists find such a method useful. The section illustrates these points with applications to the question of union formation.

Economists by and large focus on individual decision-making. The individual is the primary actor and makes decisions given his or her social and economic circumstances. Economists generally describe the motivations of an individual with a *utility function*. The utility function specifies what gives an individual *utility*, which we can understand as an individual's benefits or pleasure from taking certain actions or consuming certain goods. The utility function describes an individual's *preferences* (or tastes) over different types of actions and goods. Preferences are often narrowly defined; individuals care only about "economic" or pecuniary outcomes. For example, in a description of preferences for different consumer goods, the consumer would consider the quality and different technical attributes of product.

An expansive view of preferences, often attributed to the work of Gary Becker, allows us to use utility functions to study a wider range of decisions, such as the decision to marry.³ A utility function would describe an individual's gains from marriage or other

³ See Gary Becker, *Treatise on the Family*. Cambridge: Harvard University Press, 1993.

form of union. Different people would then have *preferences* to live as a married rather than unmarried couple, for example.

In general, we can describe a person's utility as deriving from her own actions and the actions of others. We can represent the utility of person j as

$$W_j = W_j(a_j, a_{-j})$$

where a_j are j 's actions, a_{-j} are others' actions. The inclusion of a_{-j} captures the possibility that other people can affect j 's utility. For example, a_j is the amount of food j consumes, and a_{-j} is the amount of pollution in the air caused by the neighboring factory. For marriage, for example, a_j is the decision of j to marry or cohabit, and a_{-j} is other's decisions to marry or cohabit.

An economic model – whether to purchase a product or whether to marry or cohabit – would embed this utility function into a wider economic and social context. There are three basic ingredients to an economic model. First, we would have the description of the motives of individuals, the utility functions. Second, we would have a description of different individuals' economic circumstances – the prices they face, their income, etc. Third, we would have a description of how individuals interact with one another. Economists use mathematics as a way to describe each of these ingredients. We then can solve for the outcome – or solution – also using mathematics. The solution describes the decision that people make. And it indicates how the decision depends on each of the three ingredients of the model – an individual's motivation, her economic circumstances, and how individuals interact with one another.

Economists find this method useful for at least two reasons. First, the imperative to write a model forces the researcher to be precise in his or her description of the

decision at hand, and the circumstances surrounding the decision. Second, this method allows us to see how decisions would change when any one of the underlying circumstances change.

For example, an economic model of union formation would specify: (1) a utility function that describes what might motivate individuals to marry – financial security, companionship, childbearing, etc., (2) individuals’ economic circumstances – income, age, health, “marriage taxes,” legal privileges of married couples, etc., and (3) how individuals’ interact – how they find partners, how they divide household chores and resources, etc. The outcome of the model would be the decision to cohabit or to marry, given the underlying economic conditions and interaction structures. The model can then tell us how the decision to cohabit or marry would be different for people at different income levels, under different tax regimes, or for new technology for finding partners, and so on.

Thus far, social norms and culture have not played a role in the analysis. Importantly, economists generally have viewed the preferences in a utility function as individualistic. Preferences are just characteristics of individuals and display no particular pattern within society. People simply have different tastes or opinions than others. “There’s no accounting for taste” one might say. In the marriage examples, some people just enjoy being married more than others, some care more about financial security, some just care more about childbearing in wedlock.

The supposition of individualistic preferences is at the heart of the puzzle facing economists. If preferences are individualistic, how can we account for systematic differences in people’s decisions – differences that do not derive from differences in

economic circumstances or the interaction structure? As discussed above, many researchers would say that “social norms” or “culture” is at play. The next section of this essay discusses how economists have understood social norms, or culture.

2. Social Norms in Economic Theory

I describe here the two generations of economic thinking about social norms. The first explained group patterns of behavior as deriving from the interaction of individuals, the third ingredient in a model, described above.⁴ In this work, the interaction of individuals is often described as a *game* (as in game theory). The solution to the game, the *equilibrium*, is then a description of a prevailing social norm. The second generation explains group patterns as arising from different preferences and utility, the first ingredient in a model described above.⁵ In particular, this work argues that preferences that relate to person’s social identity can give rise to subgroup variation within a population.

Let me describe each in turn.

Social Norms, or Culture, as “Equilibria” of a Game

Economists have developed two basic game-theoretic models to describe social norms, or culture. As in all economic models, these models describe each of the three basic ingredients: the motivation of individuals (the utility functions), the underlying

⁴ For one of the first contributions, see Akerlof, “The Economics of Caste and of the Rat Race and Other Woeful Tales,” *Quarterly Journal of Economics*, 1976.

⁵ See Akerlof and Kranton, “Economics and Identity,” *Quarterly Journal of Economics*, 2000. For another example, see Alberto Bisin & Thierry Verdier, “Beyond the Melting Pot: Cultural Transmission, Marriage, and the Evolution of Ethnic and Religious Traits,” *The Quarterly Journal of Economics*, 115(3), pages 955-988, August 2000.

economic environment, and how individuals interact. The interaction is particularly important here. The interaction structure tells us how other people's actions (a_j in the utility function above) affect an individual's utility and thereby the individual's own incentive to take certain actions. Through this feedback loop, individuals' actions ultimately affect all others' incentives and hence all other individuals' decision. A solution, or "equilibrium," arises when the feedbacks are internally consistent: everyone is acting to maximize her own utility, doing the best she can given what all other individuals are doing. (This concept should become clearer in examples below.)

Two interaction structures have dominated economists' modeling of social norms and culture. The first is a *coordination game*. The second is a *repeated game*.

Coordination Games

A coordination game describes a situation where individual utility depends on whether or not two people have coordinated on the same action. When two individuals meet, for example, they each will have higher utility if they both speak the same language. The question then becomes, which action will people choose? This basic coordination problem is embedded in an interaction structure. Individuals are described as meeting other people, usually at random. For example, suppose individuals can (and must) choose what language to speak, English or Spanish. Individuals in the population then randomly meet each other. Individual payoffs are higher when two individuals who meet speak the same language rather than different languages. It is important in these games that individuals must choose their action (e.g., which language) before knowing who their partners will be or what action their partners have chosen. The choice is

viewed as a permanent choice, and individuals then choose an action anticipating what other individuals' choices will be. An "equilibrium" describes the ultimate choices of individuals in the population. In the language example, when the cost of acquiring the language are the same for everyone in the population, there would be two equilibria – one where everyone speaks English and one where everyone speaks Spanish. The different equilibria then are said to describe two different social norms. There is a Spanish social norm, or an English social norm.

When applying this type of model to marriage, we could construct a simple coordination game as follows. Suppose individuals decide whether to marry or not. Then couples randomly meet. If a married couple meets another married couple, their utility (from socializing say) is higher than when a married couple meets a cohabiting couple (or meets single individuals). There would then be two equilibria: one where everyone gets married, and one where no-one gets married. There are then two possible social norms.

This type of modeling gives us some insight into social norms when coordination is particularly valuable to economic payoffs. The model has been applied to such conventions as driving (right or left side of the road), and language and assimilation. The modeling, however, has two major shortcomings. First, the typical model does not allow individuals to choose whom they interact with. It is easy to imagine that English speakers would seek out other English speakers, and married couples would seek out other married couples. Allowing for this type of sorting could destroy the power of the model – meaning that equilibria may not exist and we may gain no insight into the prevalence of one decision over the other. Second, like all game theoretic models, there

is no way to see why we would end up in one of the equilibria rather than another. In order to move from one equilibrium to another, something must happen “outside the model.” There must be some change in policy or change in the composition of the population that would act to coordinate the population on another choice. For example, making English the only language allowed in schools would change the individual payoffs to speaking Spanish and perhaps eliminate the equilibrium where everyone speaks Spanish. Similarly a policy that increases the economic benefits of marriage would eliminate the equilibrium where no-one marries.

Repeated Games

The second way economists have understood social norms is as an equilibrium of a repeated game. In a repeated game, a group of individuals interact today, and again in the future, and again far into the future. In this setting, an individual’s actions (a_j) and others’ actions (a_{-j}) not only affect utility today, but also utility in the future. In particular, one person’s actions today can lead others to take actions in the future that harm the first person. For example, consider companies that are competing for customers. They compete in the same market again and again. For one company, setting a low price today can attract a new set of customers. But a competitor can retaliate by setting a lower price in the future. This type of price war can ultimately destroy all of their profits. The question then becomes what price will companies set? There are many equilibria. They could all set low prices, for example. Or they could all set high prices, where each company does not lower its price to attract new consumers because that action could start a price war.

A repeated game analysis could easily be used to study marriage and divorce. Here is an extremely simple model. Suppose that individuals are motivated to marry for companionship. Individuals will want to stay married to someone who provides companionship, but divorce someone who does not provide companionship. Individuals also care about having friends. The question we ask is: what is the equilibrium behavior? Is it possible that individuals stay married even though the partner does not provide the desired companionship? The answer is yes – there is an equilibrium where people always stay married. Consider the following behaviors. (1) No-one socializes with someone who is divorced, and (2) no-one socializes with anyone who socializes with someone who is divorced, and so on. Effectively, divorced people are shunned. If individuals care enough about having friends, these behaviors are self-sustaining. No-one will socialize with a divorced person because they in turn will lose their friends. Through this repeated string of “punishments,” we can sustain an equilibrium where no-one seeks a divorce. We call this equilibrium a social norm. As in many games, there can be several equilibria, and the same circumstances could lead to different social norms (e.g., a norm with divorce and a norm without divorce).

Like the coordination game, this repeated game modeling of interaction gives us some insight into social norms. In particular, this model tells us how people can sustain behaviors that are ultimately highly detrimental. For example, this model applies well to the persistence of practices such as female genital cutting. This practice is very harmful to girls and women, yet it continues. The logic of the model gives us some insight into why. A woman who is not cut does not find a marriage partner. She is shunned. She is shunned because anyone who marries her will also be shunned, and so on. This model

may also apply to understanding what the “appropriate” age to marry. A woman who does not marry by a certain age is deemed an “old maid,” and her social life is circumscribed.

While this model lends some insight into social norms, it also has shortcomings. First, it relies on a long string of punishments to sustain the equilibrium behavior. As soon as one person decides not to punish someone who deviates from the prescribed behavior, the equilibrium falls apart. Second, we have no way of understanding why different equilibria emerge. We must again appeal to features “outside the model,” such as historical precedent. But we cannot apply such explanations systematically. We again must appeal to something outside the model, in particular if we want to know how to exit from a particular equilibrium. For example, in a repeated game setting, individuals acting on their own cannot end the practice of genital cutting. They also cannot act on their own to change the age at which one is considered “too old” for marriage. A coordinated effort and perhaps government intervention would be required.

Cultural Change and Multiple Equilibria

Can the notion of multiple equilibria help us understand cultural change? That is, can we understand cultural change as a shift from one equilibrium to another in a game? I would say no. As discussed above, we have no theory as to why one equilibrium emerges over another. Any equilibrium is as likely as any other equilibrium. To understand shifts in equilibrium, we must appeal to our knowledge of the environment outside the game. Such an analysis would depend very much on the specific setting.

From such analyses, we could not conclude much in general about shifts from one equilibrium to another.

A different approach to understanding cultural change would be to see cultural change as a change in the game people are playing. A change in the underlying interaction structure, or in the utility functions, could lead people to change their behavior and a different pattern of behavior would emerge.⁶

Identity and Economics

The second generation of work on social norms posits that people's preferences themselves may be systematically different in different groups. That is, preferences themselves can be a way of modeling culture and "social norms." The basic premise is that preferences may be systematically different in different groups. Preferences may be socially defined. In particular, in the work with George Akerlof we build models where people have a notion of their *identity*, and this identity affects their preferences.

Our understanding, and modeling, of identity draws on the vast literature on social identity outside of economics. We could say that we translate the notion of identity into economic analysis. Our modeling begins with the notion of social difference. We posit different social categories, such as man, woman, black, white. People think of themselves and others in terms of these categories. Associated with each category are appropriate and inappropriate behaviors, i.e., norms. E.g., there is a long list of things that

⁶ This suggestion, of course, just pushes the question one step back. What now accounts for a change in the game? Perhaps there is a meta-game where individuals choose which games to play. I am not sure that pursuing this question yet one more step back would ultimately be very fruitful, as in there is no obvious end.

men or women respectively should or should not do. We then posit a utility function where an individual's utility depends on the extent to which he or she follows the norms.

With an identity model, we can see tradeoffs between norms and economic incentives. In our modeling, if an individual thinks of herself as a “woman,” she may act in ways appropriate to being a “woman,” even when those actions reduce her income. She may be offended by others' actions that violate norms for behavior. And she may take actions to distance myself from them or otherwise act against them, even when those actions reduce her income.

It is thus the first ingredient in our model – our utility function – that distinguishes our model from other economic modeling. We propose that individual's utility depend on own actions, a_j , and others actions, a_{-j} , as in a standard economic model. In addition, we propose that the utility from different actions depends on the set of social categories (which we denote C), the norms (N) that give the appropriate behavior and ideal attributes for people in these categories, and the assignment of oneself and others into the categories. Hence, we have utility function

$$W_j = W_j(a_j, a_{-j}, I_j)$$

where I_j is j 's self-image, or identity, and it in turn depends on actions, categories, and norms:

$$I_j = I_j(a_j, a_{-j}, c_j, \varepsilon_j, N)$$

Identity depends on status of assigned category, given by $I_j()$. The match between own attributes, ε_j , and the ideal of assigned category (given by N), and especially the extent to which own and others' behavior matches norms.

In the simplest case, an individual j chooses actions to maximize this utility, taking as given c_j , ε_j , and N and the actions of others. We use the verb “choose” advisedly. We do not presume one way or another that people are aware of their own motivations, as in standard utility theory which is agnostic as to whether an individual shopper is aware or not of the reasons for her choices.

Beyond actions, to some extent an individual may also choose the category assignment c_j . Social categories may be more or less ascriptive, and in general, the individual is likely to have some choice over identity, as indeed people may even have some choice over their gender. Again, this “choice” may be more or less conscious.

Individual actions may also affect the norms N , the set of social categories, C , as well as the status of different categories reflected in $I_j(\cdot)$. With respect to gender, for example, status differences between men and women have diminished over time, and prescribed behavior and physical ideals have changed. Gender categories themselves have become varied and complex. There may be no universal agreement about social categories and prescriptions. Indeed, they are the subject of much debate and controversy.

Let me apply this utility function to problem of union formation and dissolution. We can build a simple model that includes identity and norms for behavior. The model has the same three basic ingredients as standard models – descriptions of motivation, economic circumstances, and interaction structure. Our description of motivation will include the identity elements outlined above. In our model, we would have social categories, such as black and white and man and woman. We would have norms for each category, such as whether a “woman” should be married and have children, or whether

children should be born in wedlock, or whether a “man,” in order to marry, must be able to support his wife and future children, or whether a “black” should or should not marry a white. These norms, or the strength of the imperative of these norms, could vary by race. Our description of economic circumstances would include the standard variables such as education, health, age, the legal environment, etc. And our description of the interaction would include how men and women meet, and how they bargain over household resources

Such a model would easily generate subgroup variation in marriage rates and out of wedlock birth rates. Blacks and whites make different decisions, according to their different norms. Like standard models, it would easily show how changes in the economic circumstances will change these patterns, and how changes in the interaction structure will change the patterns. Unlike standard models, this identity model would show how changes in norms affect the patterns.

Let me briefly illustrate the difference between this view of norms and the previous modeling. Consider the norm for an “appropriate” age for marriage and childbearing. In a repeated game model, women decide to marry by a certain age because they fear social ostracism if they do not. In an identity model, women have internalized the norm. They feel bad about themselves if they do not marry by a certain age, or have children by a certain age. They feel personally unsatisfied.

This identity model provides a general theory for how culture and social norms affect economic decisions and outcomes. It also shows us how norms can change. Norms themselves are the product of individual actions. Some individuals may have incentive to change societal norms. Advertising is an obvious example of how

companies try and change norms to increase consumption of their products (think of gender and cigarette advertising). Yet other actors also may change norms, including politicians, policy makers, and activists. While we have yet to develop a coherent theory of change in social norms, this model is a start and we describe other directions for future research in the next section.

3. New Research Directions

The section outlines directions for new research on culture and social norms and family change. I will discuss how economic theories might guide empirical work, and I will also discuss new directions for theoretical work.

Empirical Work

Testing economic models can be difficult. The first challenge is obtaining measures of a model's three basic ingredients: individuals' preferences, their economic circumstances, and the interaction structure. Individuals' economic circumstances – age, income, education, etc., – are often well-measured and found in many data sets. It is very difficult, however, to find measures of individual preferences. Some surveys (such as the General Social Survey) ask people about their preferences, and this is an obvious place to begin. I will discuss below how researchers might be able to infer preferences from observing people's behavior. As for interaction structures, they are also extremely difficult to find in standard data sets. While we might have some information on friendship networks (as in the National Longitudinal Study of Adolescent Health) or individual neighbors or neighborhoods, we typically do not see how often and under what

circumstances people interact with one another. Again, there is much room for work here on developing methods from observation, which I will discuss below.

Of course, as in all empirical work, the second challenge is identification. Identification is a particular challenge here since, as discussed, it is hard to observe and measure culture and norms, and, second, culture and norms are likely to be endogenous to a social and economic setting. Indeed, the economic theory suggests that norms emerge from the interaction of individuals, and individuals can act to change norms.

So how to proceed? Of course the first suggestion is to find, as much as possible, a “natural experiment.” One possibility is to locate two communities where the norms are different but everything else is the same (or the researcher can control for the differences). To find such an experiment, the researcher might rely on historical or other evidence of difference in norms or changes in norms. For example, a researcher might be able to exploit some time variation in the speed at which norms change. Take urban African American communities on the East Coast and urban African American communities on the West Coast. When new norms emerge on one West Coast (as with the Black Panther movement in the 1960’s, for example), it might take some time for the norm to be adopted on the East Coast. We can then see if the change in norms is associated with changes in behavior. Other examples of this sort would be two high schools with similar demographics but with apparently different norms governing sexual activity or school achievement.

Short of such “natural experiments,” it might be possible to see how norms affect individual behavior in a smaller setting. A second suggestion is to combine survey research with qualitative interviews and participant observation. There may be many

indicators of identity and norms that cannot be captured in large data sets administered at a distance. But a researcher who spends time within a particular environment or community may be able to pick up on clues that indicate a person's identity and associated norms for behavior. Such a researcher would also be able to observe and measure individual interactions. To take another example from high schools, researchers who spend time in high schools are able to observe, and record, which students belong to different groups from the way students dress, where they congregate in the school, and so on. This researcher would also be able to observe how students interact. Do particular students interact frequently or infrequently, in school or out of school, etc.? The researcher could then classify the students in terms of groups, and classify their interactions, and use this classification in combination with other data to determine how groups and identity influence achievement and other outcomes.⁷

A third suggestion is to combine survey research with experiments. There is a large body of social psychological experiments on identity and group membership.⁸ Certain experiments try and elicit a person's identification with particular groups (ethnic identity, gender identity, etc.). It might be possible to combine such questions or such experiments with surveys containing more traditional data in order to study the extent to which norms and identity are associated with different individual characteristics and choices.

Theoretical Research

⁷ For an extensive discussion see pgs. 1177-1178 in Akerlof and Kranton, "Identity and Schooling," *Journal of Economic Literature*, December 2002.

⁸ For an excellent review see S. Alexander Haslam, *Psychology in Organizations*. (London: Sage Publications, 2004).

Theoretical research in economics on social norms and culture is, relatively, in its infancy. There is much work to be done. Three avenues of research could be particularly informative for understanding culture and family change.

First, we could do much more to understand cultural changes and shifts in norms. I have already explained that, unfortunately, the game theoretic approach to norms is not likely to help us much here. The identity approach is likely to be more fruitful, as there is room for individual actions to change norms. We would need to construct a theory as to how individual actions to change norms – through activism, policy, etc. – ultimately reach large numbers of people and affect their own understanding of proper and appropriate behavior. The latter part of the 20th Century saw a great change in norms governing gender and race. Many would credit the civil rights movement and the women's movement for these changes, along with legislative reforms that one can view as products of these social movements. There is empirical evidence that women, for example, have adopted the new norms concerning work, and their labor force participation decisions reflect these new norms.⁹ Of course, there is much work in the social sciences on these dramatic changes. Economists have yet to develop a complementary theory. What would the value-added be of such an effort? Again, like all economic modeling, such a theory would allow us to see how outcomes would change if circumstances were different. We would then be able to see how cultural change would or would not emerge under different economic conditions or the different ways in which people interact and communicate.

⁹ See Claudia Goldin "The Quiet Revolution that Transformed Women's Employment, Education, and Family" 2006 Ely Lecture, American Economic Association Meetings, Boston MA (Jan. 2006)

A second avenue for research would be further exploration of social norms and group dynamics and decision-making. As suggested above, there is a large literature in social psychology on group dynamics. Norms, of course, also govern how a group makes a decision. Thinking of the family as a group would open a new avenue for research on family decisions. Currently, economists generally understand family decision making as a negotiated decision between husband and wife. Yet, there are many other parties who are affected by decisions, including children and extended family members. It is likely that social norms influence which parties' needs are taken into account and the extent to which different parties have a say in the decision – are the grandparents' needs or children's needs paramount, for example? And how does this complicated group arrive at a decision? These are wide open questions, yet have significant implications for family change. A research effort combining the social psychology of groups and the identity models of norms could make great advances in this domain.

A third direction would be integrating network theory with identity theory. Sociologists have long studied social networks and argue that individual decisions are influenced by social structure. Economists are now beginning to study the influence of networks on economic outcomes. We are using the standard economic methodology, using models with the same three basic ingredients. The innovation, here, is the interaction structure. The interaction structure is conceived as a network. We ask how individuals, with given preferences and economic circumstances, would make different decisions depending on their positions in a social network. Yet, we also learn from sociologists and anthropologists that individual preferences themselves may be influenced by social connections and network structure. That is, we should not

necessarily view the interaction structure as independent from the preferences. A person's network can shape their preferences. To use the language of identity theory, a person's network can shape their understanding of norms, N , the social categories, C , and their own self-categorization, c_j . There is much to be gained in understanding into social norms and culture by integrating network theory with identity theory. We will be able to see how networks influence the diffusion of norms, and how norms themselves can structure network relations.