Summary of Research

Adam Szeidl

July 2009

My research is in applied theory. My objective has been to develop models of individual behavior that incorporate realistic constraints (such as transactions costs or social networks) and study their aggregate and welfare implications both theoretically and empirically. This summary groups my work into four topics. *(1) Networks.* I combine models, experiments and data to study social capital, social networks and their interactions with markets. *(2) Consumption and household finance.* I study the implications of adjustment costs and borrowing constraints on risk preferences and consumption. *(3) International trade.* I explore the effects of contractual and physical transaction costs on trade and productivity in international markets. (4) *Auctions and markets.* I study the foundations for market efficiency using models of auctions with finitely many participants, and when some bidders suffer from behavioral biases.

I. Networks

Neoclassical economics often assumes that agents can write formal contracts that are costlessly enforced by the legal system. However, in many economies, low-cost contract enforcement is imperfect, and agents often rely on informal arrangements that take place in the social network. In **Trust and Social Collateral** (Quarterly Journal of Economics, 2009) Dean Karlan, Markus Mobius, Tanya Rosenblat and I build a theory of trust based on informal contract enforcement in social networks. In this model, network connections between individuals can be used as social collateral to secure informal borrowing. We define network-based trust as the highest amount one agent can borrow from another agent, and express this quantity as the sum of the weakest links on all disjoint paths connecting borrower and lender, also known in mathematics as the maximum network flow. We then use this reduced form in three applications related to the theory and measurement of social capital. (1) We predict that dense networks generate bonding social capital that allows transacting valuable assets, while loose networks create bridging social capital that improves access to cheap favors like information. (2) For job recommendation networks, we show that strong ties between employers and trusted recommenders reduce asymmetric information about the quality of job candidates. (3) Using data from Peru, we show empirically that network-based trust predicts informal borrowing, and we structurally estimate and test our model. A broader contribution is a framework for informal contracting in networks which we use in several other papers that I discuss below.

The previous paper provides a theoretical framework for informal borrowing in networks. But how valuable are networks for loan transactions in practice? Sociologists often argue that allocations are primarily determined by the social context, with little room left for price incentives. In contrast, neoclassical economics emphasizes prices as the main determinant of allocations. In the working paper **Measuring Trust in Peruvian Shantytowns**, Karlan, Mobius, Rosenblat and I use a field experiment to measure the relative importance of networks and prices for borrowing

in shantytown communities in Lima, Peru. The experiment randomizes the interest rate on loans provided by a microfinance agency, as a function of the social distance between the borrower and the cosigner. This design effectively varies the relative price (interest rate differential) of having a direct friend vs an indirect friend as a cosigner. After loans are processed, a second randomization relieves some cosigners from their responsibility. Data from these experiments yield three main results. (1) As emphasized by sociologists, connections are highly valuable: having a friend cosigner is equivalent to about 18% of the face value of a 6 month loan. (2) While networks are important, agents do respond to price incentives and switch to a non-friend cosigner when the interest differential is large. (3) Relieving responsibility of the cosigner reduces repayment for direct friends but has no effect otherwise, suggesting that different social mechanisms operate between friends and strangers: Non-friends cosign known high types, while friends also accept low types because of social collateral or altruism. The main lesson is that social connections can be exploited for substantial economic gain.

The importance of social connections for informal transactions may help explain why many developing villages obtain good consumption insurance even though formal financial and insurance markets are missing (Townsend, 1994). In **Consumption Risk-sharing in Social Networks**, Attila Ambrus, Mobius and I explore this idea using the framework developed in Trust and Social Collateral. We build a model where agents use their network connections to enforce risk-sharing contracts, and develop two main results. (1) Expansive networks, where every group of agents have a large number of links with the rest of the community relative to the size of the group, facilitate better risk-sharing. In particular, "two-dimensional" village networks organized by geography are sufficiently expansive to allow very good but imperfect risk-sharing. (2) In second-best arrangements, agents organize in endogenous connected "risk-sharing islands" in the network, where shocks are shared fully within but imperfectly across islands. As a result, risk-sharing in second-best arrangements is local: socially closer agents insure each other more. These findings help explain the empirical results of Townsend (1994) and others, and yield new testable implications. A more general lesson is that expansiveness facilitates network-based transactions by allowing groups of agent to access other parts of the network.

The results of these papers suggest that variation in network structure can have important effects on economic outcomes. But what determines network structure? In **Community Size and Network Closure** (American Economic Review Papers and Proceedings, 2007), Hunt Allcott, Karlan, Mobius, Rosenblat and I argue that community size should be an important determinant of network closure, i.e., the degree to which the social network is interconnected. The intuition is straightforward: in a smaller community, the pool of potential friends is limited, which makes it more likely that the friend of one agent is also a friend of another. We document evidence consistent with this hypothesis in data on the social networks of students in the National Longitudinal Study of Adolescent Health. Using the measure of network closure we developed in Trust and Social Collateral, we show that in U.S. middle and high schools, increases in grade size predict lower network density (i.e., lower network closure). We also show that this effect is negatively related

to GPA and other outcomes. Our results suggest that community size may influence outcomes through a previously unexplored channel, through variation in the structure of the social network.

While there is significant variation in network structure across different environments, there are also some common regularities. Two features documented for many economic and social networks are: (1) a core-periphery structure; (2) positive correlation between network centrality and payoffs. In **Core and Periphery in Networks** (Journal of Economic Theory, 2008), Daniel Hojman and I develop a model of endogenous network formation to explain these features. In our model, agents benefit from access to others in the network, but these benefits exhibit decreasing returns to scale and decay with network distance. These simple assumptions can be microfounded in a model of information diffusion in the network, but are likely to hold in more general environments as well. We show that these assumptions have strong implications: there is a unique equilibrium network architecture, which we call the periphery-sponsored star. In this equilibrium one agent, the center, maintains no links and earns a high payoff, while all other agents maintain a single link to the center and earn lower payoffs. Hence our model generates both of the common features of social networks described above in a stylized way. We also develop conditions under which the star architecture and payoff inequality are preserved even if agents are allowed to make transfers and bargain over the formation of links.

In earlier work, I have explored the connection between network structure and the evolution of social conventions. Kandori, Mailath and Rob (1993) and Young (1993) develop theories of social conventions by modeling a society where agents play a coordination game using myopic best responses with infrequent mistakes. They show that in this environment the risk dominant equilibrium emerges as the long-term stable "convention" because it is more robust: it is a best response to a wider range of possible opponent strategy profiles. Much of this research assumes that agents interact with everybody else in the community. In **Endogenous Networks, Social Games and Evolution** (Games and Economic Behavior, 2006), Hojman and I endogenize the interaction structure in this model, by allowing agents to choose their friends as well as the convention they follow. In our model, scale effects introduce a new trade-off: It may be optimal to leave a large group of agents coordinating on the risk-dominant equilibrium, and instead connect to a smaller group playing the efficient action. This trade-off modifies the equilibrium selection result of the earlier literature. We show that the unique stable convention is the payoff dominant equilibrium if it is not very inferior on the risk dimension, and is the risk dominant equilibrium if it is not too inefficient. More broadly, we identify a general pattern of long term equilibrium selection, namely that flexibility favors payoff dominance: In environments where agents find it easier to choose both the number and type of partners they interact with, the efficient equilibrium is more likely to be selected. This finding suggests that more mobile societies may select more efficient social conventions.

These results about stochastic stability assume that agents experiment with alternative conventions with similar probability when they play either strategy. Bergin and Lipman (1996) criticized

the robustness of this approach by showing that when some experiments are more likely, any equilibrium can be selected in the long term. In **Contagion and State-Dependent Mutation** (B.E. Journals, Advances in Theoretical Economics, 2003), In Ho Lee, Akos Valentinyi and I argue that this conclusion depends on the interaction structure. We show that for arbitrary mutation probabilities, in a large enough two-dimensional network the risk dominant equilibrium is uniquely selected. This result is driven by two forces. First, the dynamics of the model exhibits "local conventions," in which a group of agents coordinating on the risk-dominant action find it suboptimal to deviate. Second, at the boundaries of local conventions, relatively few experiments can result in "contagion" where the risk-dominant strategy spreads further, until the system settles in a larger local convention. The dynamics of persistent local conventions alternating with infrequent contagion lead to the risk dominant equilibrium in the long term. The results of these two papers emphasize the importance of the interaction structure for the selection of long-term equilibrium "conventions" in evolutionary models.

II. Consumption and household finance

My work in this area studies the implications of adjustment costs and borrowing constraints for consumption and investment. Many consumption goods like housing or cars involve adjustment costs and hence may not be adjusted in response to small shocks. In **Consumption Commitments and Risk Preferences** (Quarterly Journal of Economics, 2007), Raj Chetty and I study the effect of such consumption commitments for choice under uncertainty. We first document using data from the PSID and the CEX that the average household spends at least 50% of its budget on commitment goods. We then characterize risk preferences in a model with two consumption goods, one of which involves commitments. We show that commitments affect risk preferences in two ways: (1) they amplify risk aversion with respect to moderate-stake shocks; (2) they create a motive to take large-payoff gambles. This model helps resolve two basic puzzles in expected utility theory: the discrepancy between moderate-stake and large-stake risk aversion, and lottery playing by insurance buyers. The model has implications for the design of social insurance: amplified moderate-stake risk aversion increases the optimal rate of unemployment benefits. The model also helps explain the added-worker effect in labor economics, because commitments increase the short-term wealth elasticity of labor supply.

Adjustment costs affect risk preferences because preexisting commitments act as a reference point by altering indirect utility over total consumption. In **Consumption Commitments: A Foundation for Reference-Dependent Preferences and Habit Formation**, Chetty and I use this idea to develop a new theory of endogenous reference points based on adjustment costs, which we allow to include the mental cost of changing consumption plans. The endogenous evolution of the reference point in our model captures several features of existing theories in which the evolution of the reference point is specified exogenously. In particular, our model predicts that reference points (i) depend on past consumption levels, (ii) reflect recent expectations, and (iii) diminish in importance when agents experience large shocks. We analyze the dynamics of the commitment-based reference point in an economy with many consumers, and show that when

the ratio of idiosyncratic to aggregate risk is large, the model is isomorphic to standard habit formation specifications, in which the reference point is a weighted average of past consumption. This result provides foundations for reduced form habits used in applied work. We illustrate the implications of endogenizing the reference point in three applications. (1) Consumption is excessively smooth during normal times like in habit models; but it responds quickly in extreme events when agents abandon commitments. Hence VAR estimates can break down in crises due to changing correlations. (2) The welfare costs of large shocks are smaller with commitments than with habit, because agents abandon consumption plans, showing that while observationally similar, the two models have different policy implications. (3) Like in the Lucas critique, changes in the environment change the dynamics of commitments, but not reduced form habit. For example, insuring idiosyncratic shocks increases welfare as in standard models, but can also raise the welfare cost of aggregate shocks by making reference points more persistent. A broader lesson from this paper is that planning costs can be used to develop explicitly dynamic models of reference points.

Do consumption commitments act as a state variable that affects household choices in practice? In **The Effect of Housing on Portfolio Choice**, Chetty and I focus on housing, and test whether it affects choice under uncertainty using data on portfolios. A recent theoretical literature predicts that housing should have substantial effects on financial markets both because of the commitment effect and because of home price risk. We estimate the causal effect of changes in mortgages and home equity on portfolio allocations. We use two instruments – average house prices in an individual's state in the current year and in the year he purchased his home – to generate variation in home equity and mortgages that is plausibly orthogonal to unobserved determinants of portfolios. Using data from the Survey of Income and Program Participation, we estimate that a $10,000 increase in a individual's mortgage (holding fixed total wealth) reduces the share of liquid wealth he holds in stocks by 6.5%. This effect is driven by a combination of extensive margin (stock market participation) and intensive margin effects. Auxiliary evidence suggests that housing induces individuals to hold more conservative portfolios primarily because of consumption and mortgage commitments rather than because of house price risk. The impact of housing on portfolio choice is substantial: our estimates imply that if their mortgage debt were forgiven, households would increase the amount they invest in stocks by 40%. A contribution of this paper is to bring new evidence to an area mostly dominated by theoretical models.

Besides missing insurance markets and adjustment costs in consumption, households often face a third market imperfection: credit constraints in borrowing. In **Invariant Distribution in Buffer-Stock Savings and Stochastic Growth Models**, I explore the effects of such borrowing constraints at the household level on the dynamics of aggregate consumption. I first establish a result in probability theory: I derive a sufficient condition for the existence of a stable invariant distribution for a class of unbounded Markov processes. I then use this condition to settle a conjecture of Christopher Carroll (1997) by characterizing the conditions when consumption has a stable invariant distribution in the buffer stock savings model. This result helps characterize consumption behavior with borrowing constraints. When the model has a stable invariant distribution, buffer

stock behavior obtains: consumption growth is predictable and excessively sensitive to temporary shocks. When there is no stable invariant distribution, eventually consumption becomes a random walk (as in Hall, 1978), and follows the standard permanent income model: innovations are unpredictable and not excessively sensitive. For applied purposes, my conditions are also useful in helping to identify whether the "right" numerical model should be stationary or integrated to match long term behavior. In a second application of my theoretical results, I show that in the Brock-Mirman growth model, a stable invariant distribution exists for a large class of production functions, even if shocks and the capital stock can be unbounded.

III. International trade

My research in international trade develops and estimates models of firm behavior in the presence of transactions costs, to study aggregate implications in international markets. A body of recent work suggests that access to imported inputs improves productivity. To understand the mechanism responsible for this effect, in **Imported Inputs and Productivity**, (American Economic Review, revise and resubmit), Miklos Koren, Laszlo Halpern and I develop and estimate a structural model of importers using product level data for all Hungarian manufacturing firms during 1992-2003. We obtain three main findings. (1) We confirm that imported inputs have large productivity effects in these data: increasing the share of imported goods from 0 to 100 percent increases productivity by 11 percent. (2) We show that about 60% of this gain is due to imperfect substitution, i.e., the idea that combining different inputs is "more than the sum of the parts." This is consistent with an old view in development economics about the importance of complementarities along a production chain for productivity. (3) Tariff cuts have a highly non-linear effect on productivity, due to firm entry into import markets for new varieties. We show that this non-linearity can rationalize differences between estimated tariff effects in different studies. We also use our structural model for counterfactual policy analysis, and to study the different implications of the quality and complementarity mechanisms. A broad lesson from this paper is that firm level analysis helps understand cross-country differences in macro elasticities, because firm responses that vary depending on size, importing status and other characteristics create nonlinearities.

To economize on transaction costs, multinational firms often conduct different stages of production in different countries. This activity plays a large role in the world economy: e.g., the gross product of foreign affiliates accounted for 11% of world GDP in 2001. In **Optimal Integration Strategies for the Multinational Firm** (Journal of International Economics, 2006) Gene Grossman, Elhanan Helpman and I examine integration strategies of firms that can produce intermediate inputs and conduct assembly operations in one or more of three locations. With many countries and stages of production, the standard categories of horizontal and vertical FDI no longer describe all possible organizational forms: more complex integration strategies are possible. We study the equilibrium choices of firms that differ in productivity, focusing on the role of industry fixed costs, the cost of transporting intermediate and final goods, and the regional composition of the consumer market in determining integration strategies. In the process, we identify three distinct "complementarities" in allocating different stages of production to foreign countries. In

our model with heterogenous firms, these complementarities are reflected by the fraction of firms that choose different integration strategies and can be subjected to empirical analysis. A general implication is that horizontal and vertical FDI can simultaneously exist side-by-side in the same firm.

In **Complementarities between Outsourcing and Foreign Sourcing** (AER Papers and Proceedings, 2005), Grossman, Helpman and I extend the above model to incorporate contracting frictions as well. Firms minimize the transaction costs of incomplete contracts by optimally choosing the ownership structure of their affiliates. We assume that each firm must acquire intermediate inputs and assemble final products, and can perform these activities internally or externally, at home or abroad. We identify conditions under which outsourcing and foreign sourcing are positively correlated across industries, and show that this correlation comes from two complementarities. The results in these two papers help explain the broad range of corporate strategies found in the data.

IV. Auctions and markets

A common theme of my research is to develop models of exchange when markets function imperfectly. In **Existence of Equilibrium in Large Double Auctions** (Journal of Economic Theory, 2007), Drew Fudenberg, Mobius and I address this subject by characterizing the degree of inefficiency in spot "double auction" markets. Departing from the assumption of price-taking behavior in Walrasian markets, we consider an environment where a finite number of buyers and sellers can submit bids based on their correlated private valuations, and trade at the endogenous price where the auction clears. In this market, participants act strategically by taking into account the effect of their bid on the market price. We show the existence of a pure strategy, symmetric, increasing equilibrium when there are many participants, and prove that this equilibrium is arbitrarily close to fully revealing as the market size grows. These results can be interpreted as strategic foundations for approximate price-taking and efficiency in thick spot markets, while leaving open the possibility that thin markets function imperfectly.

In **Fishing for Fools**, Ulrike Malmendier and I explore the efficiency of thin markets by asking how allocations and prices are affected when some bidders suffer from behavioral biases. We show that in one-sided auctions, efficiency breaks down: even a small share of overbidding behavioral agents can have a large effect, because the auction format "fishes" for overbidders. In these auctions, behavioral buyers disproportionately increase prices, and may also create large welfare losses by crowding out the demand of rationals. In ongoing work, we develop three applications of this idea. (1) Profit maximizing sellers optimally choose auctions even when fixed-price markets are more efficient, because of the amplified profits generated by fools. This helps explain the popularity of auctions in practice. (2) The effect of fools is most powerful in markets such as housing where transactions are separated in space and time, because in the resulting thin local markets they are more likely to become marginal. (3) In such markets, the fishing mechanism can create bubbles when there is an inflow of overoptimistic agents, and declining prices when these agents purchase and exit. This prediction helps explain fluctuations in home prices and the afternoon effect in

wine and art auctions. Evidence from eBay supports the existence of overbidding and confirms its amplified effect on allocation and profits. A more general lesson is that agents with behavioral biases can have amplified effects by endogenously selecting themselves into particular activities such as trading.

References

[1] Bergin, J. and Lipman, B. (1996): "Evolution with State Dependent Mutation," *Econometrica* 64, 943–956.

[2] Carroll, C. D. (1997), "Buffer stock saving and the life-cycle/permanent income hypothesis," *Quarterly Journal of Economics* 62, 1–56.

[3] Hall, R. E. (1978), "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86: 971-987.

[4] Kandori, M., Mailath, G. and Rob, R. (1993): "Learning, mutation and long run equilibria in games," *Econometrica* 61, 29–56.

[5] Townsend, R. (1994): "Risk and Insurance in Village India," *Econometrica* 62, 539–591.

[6] Young, H. P. (1993): "The evolution of conventions," *Econometrica* 61, 57–84.