Report on Amy Finkelstein
AEA Honors and Awards Committee

Amy Finkelstein is the leading scholar in Health Economics and one of the most accomplished applied micro-economists of her generation. Amy is distinguished by the ability to identify important and policy-relevant, yet tractable, research questions, a deep understanding of the institutions that are the focus of her research and creative, extremely careful and convincing empirical work. The methods she uses are sophisticated, informed by theory and appropriate for the questions under investigation. Her research agenda is centered on some of the most important and policy-relevant issues facing developed economies today. She has made important contributions to the literature on asymmetric information in health insurance and annuity markets and has provided convincing analyses of the effects of public intervention in health and long-term insurance markets. Her research serves as a model of how theory and empirics can be combined in creative ways in order to yield credible, novel and often unexpected insights into economic questions that will inform policy design.

Amy's research is focused on health insurance markets. Within this area, she has had three main interrelated lines of research: (a) Tests on the presence of asymmetric information in insurance markets; (b) Structural estimation and analysis of welfare implications of models with asymmetric information; (c) Effects of public intervention in health and long-term care insurance markets.

Amy's most influential research has investigated the presence of selection and asymmetric information in insurance markets. Her most widely cited paper in this strand of work is: "*Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market*," with Kathleen McGarry, which was published in the American Economic Review in 2006. The paper provides convincing evidence of the existence of multiple dimensions of private information in the long-term care insurance market. Traditional tests for the presence of adverse selection in insurance markets have focused on the positive correlation between insurance coverage amount and insurance utilization implied by theoretical models with one-dimensional unobserved heterogeneity. The results from this literature have been mixed, with some empirical studies finding strong support for adverse selection and others not. Amy's work suggests an explanation for these diverse findings: if individuals differ in more than one dimension, e.g., in their risk type as well as their risk preferences, then it is possible that the two types of private information offset each other to produce an equilibrium with no correlation between insurance coverage and utilization. While empirical researchers using standard tests may fail to find evidence for adverse selection in this case, the resulting equilibrium will still be inefficient. Empirically, they find that individuals who buy insurance are not only worse risks, but also more risk averse than those who do not, conditional on their underlying characteristics. It is this correlation between risk aversion and purchasing behavior that explains the lack of a first-order difference in

nursing home utilization among those who buy and do not buy long-term care insurance. The paper represents an important contribution to the literature on asymmetric information.

While the previous study focused on multiple dimensions of individual unobserved heterogeneity, the paper "*Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market*", with James Poterba, published in the Journal of Political Economy in 2004, makes the point that accounting for the multiple dimensions of insurance contracts is equally important when testing for adverse selection. The paper is motivated by the observation that annuity markets are in practice limited, despite the fact that standard lifecycle models suggest that the availability of such markets should raise welfare for most households by enabling them to purchase insurance against outliving their resources. One potential explanation for this paradox is adverse selection: only those who expect to live for many years buy annuities, thereby making a policy that is actuarially fair for annuity buyers unattractive from the perspective of an individual with the population-average mortality rate. To test this hypothesis, Amy and Jim examine data on the persons who purchased annuity policies from a large U.K. insurance company over two decades. The findings suggest that adverse selection plays a potentially important role in the behavior of annuitants, and that this feature of annuity markets may potentially explain their limited size. The results also indicate that selection can be subtler than simply a decision of whether or not to purchase an annuity – it can also involve decisions about the detailed characteristics of the annuity policy. Together with her AER paper on multiple dimensions of private information, this paper has established Amy as a leading expert on the role of asymmetric information in insurance markets.

A second line of Amy's research has used (semi)structural techniques to identify the relevant structural parameters in models of asymmetric information and conduct policy simulations. A representative example of this line of Amy's work is the paper: "*The Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market*", with Liran Einav and Paul Schrimpf, published in Econometrica in 2010. The paper focuses on the same market she analyzed in her JPE paper with Jim Poterba: the U.K annuity market. The authors develop a structural model of annuity contract choice that accounts for both differing mortality risk and different preferences across individuals. They estimate the parameters of the joint distribution of risk and preferences by relying on variation in the mortality rates and product choices in their data. Other parameters related to the utility function need to be calibrated. With the structural parameters in hand, the authors conduct simulations and find that asymmetric information results in a welfare loss that is approximately 2% of the annual premiums in this market. However, the general and important lesson that is learnt from these simulations is that government mandates, the canonical solution to adverse selection problems, can potentially have large negative effects on welfare. The key for understanding these effects lies in the multiple dimensions of unobserved heterogeneity explored in Amy's earlier work. If individuals differ in both their risk type and their preferences, then

mandates involve a trade-off between reducing the inefficiency generated by adverse selection and increasing inefficiency by eliminating self-selection. The paper demonstrates that the negative effects are empirically relevant.

In "*Estimating Welfare in Insurance Markets Using Variation in Prices*", published in the Quarterly Journal of Economics in 2010 with Liran Einav and Mark Cullen, the authors use plausibly exogenous price variation to estimate the demand for and cost of insurance. These estimates are then used to calculate the welfare loss associated with adverse selection. This is a nicely written paper that clearly explains how the welfare loss associated with mispricing in the presence of adverse selection can be estimated.

A third line of research has concentrated on examining the effects of public policy on various outcomes in the health and long-term care insurance sectors. There are many papers in this category, all characterized by the focus on policy-relevant questions and the ability to come up with plausibly exogenous sources of variation in the data that allow her to identify the causal relationships of interest.

One of the most important papers in this vein is her 2007 paper in the Quarterly Journal of Economics, "*The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare*." The paper examines the impact of the most dramatic change in health insurance coverage in American history, the introduction of Medicare in 1965, on the consumption of health care. The analysis focuses on the hospital sector, which was the largest component of health spending at the time of Medicare's introduction and of the subsequent growth in health spending. The empirical strategy exploits the fact that prior to the introduction of Medicare, the incidence of hospitalization insurance policies varied widely across regions. This implies that the fraction of elderly households who experienced substantial changes in their ability to pay for hospital-based care varied across regions. By exploiting these inter-regional differences along with an extremely rich database on hospital care during the years surrounding the enactment of Medicare, Amy provides clear evidence that in places where Medicare had an especially large effect in expanding insurance coverage, there was an especially large increase in health expenditures. The paper estimates that the impact of Medicare on hospital spending is over six times larger than what the evidence from individual-level changes in health insurance would have predicted. This disproportionately larger effect is attributed to the fact that market-wide changes in health insurance may have changed medical practices. Consistent with this explanation, the paper provides suggestive evidence that the introduction of Medicare was associated with substantial new hospital entry, increased adoption of cardiac technologies and increased spending on non-Medicare patients. Based on these estimates and some back-of-the-envelope calculations, Amy suggests that about half of the increase in real per capita health spending between 1950 and 1990 can be attributed to the spread of health insurance.

Another influential paper in this strand of work is her 2008 American Economic Review paper, "*The Interaction of Public and Private Insurance: Medicaid and the Long-Term*

Care Insurance Market", with Jeffrey Brown. The paper is motivated by the observation that the size of the Long-Term Care Insurance (LTC) market is surprisingly small given the potentially large late-life costs that nursing home care can impose on retirees. Rather than focusing on adverse selection as a potential explanation for the small size of this market, Amy and Jeff explore the role of public insurance programs in crowding out private insurance. In particular, the paper examines the interaction of the public Medicaid program with the private market for LTC insurance and estimates that Medicaid could explain the lack of private insurance purchases for about two-thirds of the wealth distribution, even if there were no other factors limiting the size of the market. The large crowd out effect stems from the fact that Medicaid is always the second payer when covered individuals need nursing home care; because of this, it imposes a very high implicit tax on the benefits associated with private LTC insurance policies. The extent to which this tax reduces the demand for private LTC insurance policies depends critically on the distribution of wealth and income among elderly households. The analysis in the paper demonstrates convincingly that the fraction of elderly households who face a high implicit tax is substantial. Hence, the paper offers a compelling demand-side explanation for the small size of the LTC insurance market. The paper provides a careful and theoretically informed analysis of an important public policy program and yields significant policy implications. Specifically, the findings imply that public policies designed to stimulate private insurance demand will be of limited efficacy as long as Medicaid continues to impose this large implicit tax.

Finally, Amy is involved in an exciting new project on the effects of Medicaid reform in Oregon. In 2008, Oregon expanded the coverage of its Medicaid program, but did not have enough resources to allow universal access. As a result, access to coverage was provided by lottery. Amy and her co-authors managed to collect baseline data on both insurance status and health conditions among those who would potentially be eligible for insurance coverage prior to the implementation of the reform, and examined the consequences of this randomized controlled trial. The early findings from their study suggest that the insurance expansion increased utilization of medical care, reduced out-of-pocket health costs and led to higher self-reported health status. A summary of these results was published in the New England Journal of Medicine in 2011 ("*The Effects of Medicaid Coverage: Learning from the Oregon Experiment*," with Kate Baicker). A longer discussion of preliminary findings can be found in the working paper "*The Oregon Health Insurance Experiment: Evidence from the First Year*", which is joint with various coauthors and forthcoming in the Quarterly Journal of Economics. The project is ongoing and will continue to track and study those who did and did not receive Medicaid coverage; accordingly, there is no final paper associated with this work yet. This work is likely to be viewed as a seminal contribution to both health economics and health policy discussions.