

Giorgio Topa
PERSONAL RESEARCH STATEMENT

My research to date has focused on various areas in applied microeconomics, ranging from modeling social interactions to studying the transmission of cultural or cognitive traits, the impact of microcredit in developing economies, or the nature of discrimination in the labor market. These topics span several traditional fields, including labor economics, local public finance, and development economics. In addition, my research is quite interdisciplinary, being closely connected to several areas and ideas in sociology, political science, geography, and anthropology. A common feature in my research is an emphasis on devoting attention to both the theoretical and the empirical aspects of a given question.

Research in applied microeconomics often displays a gap between theoretical modeling and empirical analysis. I place great value on maintaining a tight connection between theory and estimation, be it by conducting a fully structural estimation of the ‘deep’ parameters underlying a given model, or by looking for precise observable implications of a model that can be tested in the data using more flexible tools. In my research, I also investigate methodological issues that may arise from a specific application: for example, the study of local interaction and cultural transmission models has led me to think about general estimation procedures for models in which multiple equilibria are possible.

In what follows, I describe my past and current research in each area, and outline my agenda for future work.

1. Local Interactions

The bulk of my research has focused on the empirical evaluation of local interaction models. These are models in which agents’ decisions and outcomes are affected by other agents’ behavior directly, rather than being mediated by markets. The underlying idea is that agents do not exist as isolated atoms, but rather are embedded within networks of friends, relatives, acquaintances, neighbors, colleagues, that jointly provide cultural norms, economic opportunities, information flows, social sanctions and so on. A cursory list of applications in economics includes neighborhood effects and the spread of social pathologies; peer effects in education; informal hiring networks in job search; endogenous growth and income inequality; cultural fads; the adoption of new crops or other technological innovations; network externalities in consumption; localization decisions by households and firms. A common assumption in these models is that individuals interact *locally*, with a set of neighbors defined by a social or economic distance metric. My objectives in this area have been on the one hand to provide estimates of the presence and magnitude of spillover effects in specific applications, and

on the other hand to investigate several methodological issues, such as identification and aggregation, that arise when one brings such models to data.

My paper (5) presents a model of local interactions in the context of urban unemployment. Agents exchange information about job opportunities within their social networks. In particular, useful tips or referrals are transmitted by currently employed agents to their unemployed contacts, in the expectation of receiving similar tips when unemployed. Such information exchanges may be viewed as informal mutual insurance arrangements that are sustainable even in the presence of limited commitment.

Physical distance between agents is assumed to be a good proxy for social distance. The model then generates positive spatial correlations of unemployment across nearby Census tracts (with the correlation becoming weaker as distance increases). The structural parameters of the model are estimated by matching several empirical moments of the spatial distribution of unemployment in the city of Chicago with their simulated counterparts generated by the model.

A crucial issue in this empirical exercise is identification. The same spatial patterns in unemployment could be generated by positive sorting of households across locations, by correlated unobservables, or by other kinds of inter-tract linkages (such as crime spillovers), rather than by social interactions (these issues are best described in Manski's work on the identification of social effects). In the paper I try to isolate social spillovers as much as possible by looking at whether the estimated spillovers are weaker across local community boundaries that have been identified by local residents, as well as across physically adjacent tracts that have very different ethnic compositions.

The estimation results support the model with local interactions. The estimated spillovers are positive and significant, and vary with observable characteristics in a way that is consistent with the existing evidence from the literature on informal hiring channels. Furthermore, they do in fact become weaker across ethnic and local community boundaries, suggesting that they may indeed be related to social interactions as opposed to other possible explanations. This paper is becoming widely cited as one of the few existing empirical applications of local interaction models, in a growing literature on endogenous social effects that spans a variety of traditional fields: labor, local public finance, public policy, micro theory, growth with externalities, spatial econometrics.

My subsequent research in this area explores some important issues related to the estimation of local interaction models. Paper (4), joint with Tim Conley (GSB, Chicago), presents several socio-economic distance metrics that attempt to capture different dimensions along which social networks exist, and studies urban unemployment spatial patterns according to these metrics. Evidence from a rich sociological literature on networks and local communities is drawn upon to justify the choices of metrics employed in the paper. The wider goal here is to provide a flexible econometric framework that enables the researcher to consider any likely determinant of "social locations" or "social proximity" that is appropriate to a specific application, and to carry out an empirical

evaluation of models where social and economic distance matters (this is related to theoretical work by Akerlof and Manski on social interactions and endogenous social effects, resp.).

Given each distance metric, we present non-parametric estimates of spatial Auto-Correlation Functions (ACFs) both for unemployment rates across Census tracts, and for residuals from a regression of unemployment on a set of tract characteristics. ACFs are reported as a function of a single distance metric, as well as for combinations of metrics in order to study possible interactions between metrics. The results indicate a strong positive correlation of unemployment for all metrics employed, decreasing with distance. This spatial correlation is greatly reduced once we control for observable tract characteristics, in particular racial and ethnic composition. Among the different metrics, ethnic distance seems to dominate.

Paper (2) (also joint with Conley) focuses on identification issues for local interaction models, under different assumptions about what information is available on agents' locations in the abstract space of social relations. In particular, three situations are studied. The first is one of perfect information about individual agents' locations. In this case we show that global identification is attained for the model parameters. We then consider a situation in which information about individual locations is available only up to a certain level of geographic detail (e.g., Census tract or zip code). This is a very common scenario when one uses publicly available datasets such as the Census microdata on households and firms. Here we provide simulations that indicate that local identification is attained. Finally, we study the case in which individual locations are only measured with error: this is again quite likely given that social distances between actors are very difficult to observe precisely. This case is very hard to analyze, but we provide an example in which lack of identification arises almost surely w.r.t. the probability distribution of the distortion map. We then propose an estimation procedure that can yield useful estimates for the likely range of the interaction parameters even when point identification fails.

In papers (9) and (16), Conley and I perform an estimation and calibration exercise to study the dynamic properties of a local interaction model of urban unemployment, defined at the level of individual agents. In particular, we estimate the structural parameters for an extension of the information exchange model in (5), matching moments of the cross-sectional distribution of unemployment (we use Census data for the Los Angeles metropolitan area, as well as Current Population Survey data at the level of individuals). We can then answer questions regarding persistence over time of unemployment clusters, the rise and fall of neighborhoods, and the propagation of local adverse shocks (such as a plant closing down, or a rise in localized crime) both in time and in space. An interesting by-product of this approach is that it provides additional tests for the “goodness of fit” of a local interaction model, by generating implications for individual unemployment duration data.

Paper (9) focuses on the estimation of the basic framework, and compares the performance of the full model with local interactions with that of a restricted model in which the local interactions channel has been shut off, and the only source of correlation in outcomes is the sorting on the basis of observable characteristics. The model with interactions seems to perform much better not only in terms of matching the observed spatial patterns of unemployment in LA, but also in terms of generating a distribution of in-progress unemployment spells close to the empirical one.

Paper (16) is in progress, and focuses on the dynamic properties of the model – in particular, the speed of convergence to the stationary distribution and Impulse Response Functions to local shocks. The model seems to converge to the stationary distribution roughly within two years, which is consistent with our assumption about fixed agent locations. Interestingly, a negative local labor demand shock seems to travel very little in space (only a few kilometers), and to be absorbed fairly quickly in time (within one calendar year). Papers (2), (4), (9) and (16) constitute different portions of an ongoing research project that is being jointly supported by three separate NSF programs (through grant SES-0080826): Economics; Geography and Regional Science; Methodology, Measurement and Statistics.

Paper (19) (joint with Frank Vella (IUE, Florence, Italy)) is still at a very early stage. Here the plan is to use cross-sectional data on the spatial distribution of income and unemployment across all metropolitan areas in Australia to distinguish social interaction effects from locational decisions, following a strategy suggested by Glaeser and Scheinkman in a recent paper. They decompose the variance of a given outcome variable into a portion driven by local interactions and a portion driven by endogenous selection of households into locations. A crucial finding is that the term related to local interactions does not vary with city size, whereas the sorting related term is increasing in city size. This result provides an identification strategy to distinguish local interactions from sorting. We plan to pursue this strategy by comparing variances of outcome variables across sets of cities in different Australian regions.

I plan to build on the research pursued so far in several directions (most of this work will be done jointly with Conley). Paper (9) provides a framework that can be extended in several directions. The current setup takes as given the labor demand side: in particular, the probability of finding a job is independent of the aggregate stocks of employed vs. unemployed agents. We plan to modify this by explicitly considering an aggregate matching function that depends on vacancies and unemployment. A second extension that we wish to pursue is to allow agents' individual characteristics to be themselves a function of their employment histories. For example, the longer one's unemployment spell, the more one's human capital deteriorates. Building this feature into the model should improve the goodness of fit with respect to the distribution of individual spells, as well as increase the persistence of unemployment in a given geographic area, with neighborhoods becoming "stuck" in a high-unemployment situation.

We also want to explore the policy implications of these models by studying the effects of different policies aimed at improving the flow of information about jobs in poorer neighborhoods. In fact, the presence of local spillovers and contagion effects implies that there exists a “social multiplier” effect: each new job has a positive impact not only on the person becoming employed, but also on her social contacts since it raises their expected employment and wages. Finally, we intend to introduce residential mobility into the model (using market clearing for each piece of land as an equilibrating mechanism) to quantify, for example, the extent of racial residential segregation and simulate the effects of different housing policies.

I also intend to study the effects of information networks in search in the context of marriage markets in developing economies. In paper (13), joint with Andrew Foster (Brown University), I plan to analyze and estimate a model in which grooms entering the marriage market can choose whether to search for a partner using undirected search, or rather to use referrals from brides that are already present in the village. We intend to use data on marriages in a set of villages in Bangladesh over 25 years: in particular, the data contain information about the village of origin of the bride, and the current village of residence of the couple (usually the village of origin of the groom). For a given village, the presence of undue clustering in the place of origin of brides that have migrated there can be used as evidence of the use of referrals in the search process. Thus, the spatial patterns in the data can be used to estimate the parameters of the marriage search model. Since the quality of the marriage match affects the level of investment in human capital of the offspring and thus the long run prospects for development and growth, it is particularly important to better understand the process through which marriage unions are formed within and across villages.

In paper (14), joint with Pat Bayer (Yale) and Steve Ross (U. Conn.), we plan to use a very similar idea to that in (13) to detect the presence of informal hiring networks in job search. The basic idea is to examine whether for individuals working at the same location or firm there is undue clustering of residential locations, suggesting the presence of informal hiring networks. Suppose for the moment that space is homogeneous. A standard spatial search model, where search costs are increasing in physical distance to jobs, would predict that the place of residence of workers at a given firm will be uniformly distributed within a circle with the firm at its center. The radius of this basin will be dictated by the shape of the search costs and the level of wages. On the other hand, suppose that workers use informal hiring channels in their search, such as referrals from current employees at a given firm, and that agents interact more frequently with social contacts who reside physically close. Then workers at a particular firm will be no longer uniformly distributed around the firm, but rather they will be clustered together in a specific sub-area of the disc centered around the firm. We plan to verify this conjecture by using individual level data on place of residence and place of work in Boston (geocoded down to the Census block level), supplemented by very

detailed neighborhood level information.

In the longer term, I would like to apply these ideas to other contexts, such as the diffusion and dynamics over time of “social pathologies” (like crime, alcohol and drug consumption, smoking, teenage childbearing); the economic growth implications of the existence of network and agglomeration externalities among firms; the diffusion of new technologies; the dynamics of ethnic conflict (more on this below).

2. Cultural Transmission and Endogenous Preferences

This area of my research (joint with Alberto Bisin (New York University) and Thierry Verdier (DELTA, CEPR)) studies economies in which different cultural or preference traits are endogenously formed and transmitted across generations. The general objectives of this project are to better understand cultural transmission mechanisms and study the dynamics of the distribution of traits in a population.

In paper (10), we analyze and structurally estimate a marriage and socialization model of the transmission of religious traits, using data from the General Social Survey. We focus on religion for several reasons: *(i)* religious traits are fairly well defined and measured; *(ii)* they represent a cultural trait that most families are keen to transmit to future generations; *(iii)* parental incentives to transmit their own religious trait to their offspring are not confounded by related economic opportunities considerations; and *(iv)* there exists a rich sociological literature that has studied religious minorities in the US and, based on linear extrapolations of current trends, has made predictions on the long run distribution of religious traits (such as ‘the triple melting pot’ or ‘the vanishing of American Jews’): we wish to emphasize how casting agents’ decisions in a choice-theoretic framework delivers starkly different implications.

We assume that parents have a weak preference for transmitting their own trait to their children. Young adults are socialized to a certain trait both within the family and through societal influences. We also assume that the socialization technology within the family is more effective if both parents share the same trait as opposed to having different traits. Agents choose how much effort to spend in searching for a partner with the same trait in the marriage market, and the level of effort to socialize their children to their own trait. Given the assumptions, homogamous marriages are valued more than heterogamous ones, and agents choose a positive level of effort both in their marriage search and in socialization. Moreover, the effort spent will be higher the more a religious group is a minority in the population.

The data support our model, both in terms of its “reduced form” implications, and more formally once we structurally estimate it. We perform a non-nested test comparing our model to several alternatives, in which for example socialization is exogenous or the value of a marriage is intrinsic and not necessarily related to the desire to socialize one’s offspring: the statistical tests reject these alternatives. We also

explore the possible effects of agent mobility and unobserved heterogeneity in religious preferences. The estimation results are also fairly robust to different definitions of the religious groups.

Given the structural estimates of the model, we simulate the population dynamics of the distribution of religious traits. The dynamics are highly non-linear (in stark contrast to linear extrapolations from current trends), since agents optimally adjust their marriage and socialization choices as the share of their own group in the population varies. Accordingly, our simulations support neither the triple melting pot hypothesis, nor the vanishing of American Jews (by intermarriage and assimilation).

Paper (11) analyzes a simple model of cultural transmission along the lines of (10), in the context of strategic interactions such as those modeled in the Prisoner’s Dilemma. In our model, agents are characterized by a ‘cooperative’ or a ‘non-cooperative’ preference trait, and rationally take costly actions to socialize their children to their own trait. We then ask whether ‘cooperative’ agents survive in the long run, when they interact strategically with agents who do not cooperate. In contrast to the evolutionary game theory literature, we find that our cultural selection mechanism generally leads to long run polymorphic populations, in which cooperation can be sustained. This is because parents of the minority cultural trait have higher incentives to spend resources to socialize their offspring, than parents in majority groups. This result holds both under complete and incomplete information in the matching process.

Paper (12) studies the dynamics of a population with N distinct preference traits in a general model of cultural transmission and socialization. The focus here is on identifying conditions under which heterogeneous distributions are (locally) stable, as well as studying their basins of attraction. We show that homogeneous stationary distributions (in which a single trait survives in the long run) are locally unstable. Then, for any $1 < k \leq N$, we derive conditions that ensure that the unique (locally stable) stationary distribution exhibits k traits (those with the highest “intolerance” towards other traits). Further, the less heterogeneity there is across “intolerance” levels of each group towards the other groups, the more long-run distributions become diverse. Finally, we perform simulations to study the transition paths to the stationary distributions, to characterize their basins of attraction, and to study the robustness of our analytical results to the relaxation of several assumptions on the symmetry of inter-group preferences and on the shape of socialization costs. This is still work in progress.

Future work in this project will focus on the following areas. An important extension of (10) is to allow assortative matching in marriage to take place not only along religious lines, but also along race, education, and other individual characteristics (a preliminary estimation suggests that our results are robust to the introduction of a taste for marrying an educated spouse). Another desired extension is to endogenize fertility, since this crucially affects the long-run population distributions of traits.

Papers (20) and (18) are both at a very early stage. In (20), we wish to take a closer look at the intergenerational transmission of religious traits within American Jewish communities. In (18), we study the early rise of Christianity during the Roman Empire. Several historians and sociologists have examined this phenomenon, most prominently Rodney Stark in his 1997 book. We develop a model in which agents are more likely to convert if their social contacts have already converted, and where parents take conscious socialization actions to transmit their religious values to their offspring. Thus this paper combines a social network component with a cultural transmission one. We then try to determine how strong each separate channel needs to be in order to explain the growth patterns observed in the data.

3. The Econometrics of Multiple Equilibria

This research project has stemmed in part from my work on local interactions, and in part from the work on cultural transmission models. In fact, multiple equilibria arise quite easily in both contexts. This possibility poses severe econometric and computational problems for estimation.

In paper (15), joint with Alberto Bisin and Andrea Moro (Minnesota), we lay out a general theoretical framework for model with multiple equilibria, and propose a computationally feasible estimation strategy. Formally, we define an equilibrium mapping $\pi(\theta, u)$ from a vector of parameters θ and random disturbances u into a vector of equilibrium choices and prices π . In general, this is a correspondence (to allow for multiple equilibria). The econometrician observes a vector y of variables, which is a noisy transformation of the equilibrium vector π . We assume that the same equilibrium is realized for all units in the economy (although this is relaxed later), and that the parameters θ are identified. The question concerns how to consistently estimate θ in a computationally feasible way.

A brute force estimation strategy is to compute, for every choice of θ , the sample log likelihood $l_N(y|\theta)$ and maximize it w.r.t. θ . This likelihood is very complex, since it involves computing all possible equilibria for every choice of θ , as well as integrating over correspondences. This is not only computationally difficult to implement, but it is also prone to failure since numerical approximations may miss one or more equilibria.

Therefore, we suggest an alternative procedure that is much easier to implement. It exploits a decomposition of the original likelihood function $L(Y|\theta)$ into the product $L(Y|\pi, \theta) \cdot L(\pi|\theta)$. In other words, the equilibrium choices π are treated as “nuisance parameters”, and are estimated directly alongside θ in the first, unconstrained, step. In the second step, the model parameters θ are re-estimated holding π fixed at $\hat{\pi}$ and imposing the equilibrium restrictions. This approach vastly reduces the computational burden. Finally, we extend our proposed estimation strategy to the case in which different “regions” of the economy (be they sectors, cities, or the same economy at

different time periods) may be in different equilibria.

The paper is still in progress: we plan to investigate the asymptotic properties of our proposed estimators in terms of consistency at least. We also intend to carry out montecarlo simulations to evaluate the small sample behavior of our estimators, as well as their computational burdens. Simple economic examples will be used as applications for these simulations.

In the longer term, we plan to apply our estimation strategy to several concrete contexts. One area in which multiple equilibria arise quite naturally is that of local interaction models for various socio-economic outcomes (such as smoking or crime). We are also quite interested in job search models with endogenous search intensity, or models in which very different labor market institutions may arise endogenously (in terms of contract flexibility, employment protection, etc). Finally, a third natural and potentially fruitful application is to financial crises (bank runs, attacks on currency pegs, market crashes).

4. Ethnic and Religious Conflict

Ethnic and religious conflict have increasingly become prominent in recent years in many parts of the world. Such conflicts manifest themselves at very different scales (ranging from urban tensions or riots, to low-level civil wars, to ethnic cleansing and genocide). Unfortunately, examples abound: riots in Los Angeles in the 1990's, religious violence between Hindus and Muslims in India, the conflict in East Timor, ethnic cleansing in the former Yugoslavia, genocide in Rwanda. At the positive level, one would like to understand the mechanisms that bring about such phenomena, and the conditions that may favor or hinder their occurrence. More at the normative level, it is important to determine what policies might minimize the probability and/or the extent of ethnic and religious conflicts.

One of the most interesting features of ethnic and religious conflict is the extreme variance in outcomes, both across different places at one point in time, and in the same place over time. Cities with similar ethnic or religious mixes and similar potentials for violence experience very different histories of conflict. Similarly, it is quite common to observe long periods of peaceful coexistence between two groups in the same city or country, followed by sudden spells of very intense violence. This pattern is often repeated cyclically over generations.

I strongly believe that local interaction models on the one hand, and cultural transmission models on the other, can explain these patterns quite naturally and help shape policies to contain the occurrence of conflict. Models in which agents are embedded in an explicit network structure and interact locally with their relatives or neighbors are characterized by high correlations in outcomes within areas, and extreme variance in outcomes across areas. Because agents are affected by their nearest neighbors' ac-

tions, small initial differences can lead to dramatically different long-run outcomes. As a mental experiment, imagine two cities (A and B) with the same ethnic mix and with similar observed characteristics. A small initial “negative” shock (such as a particularly inflammatory speech by the leader of one ethnic community, or a beating involving members of different groups) can induce widespread violence in city A because agents’ behavior affects their neighbors and is in turn affected by them (this is a “social multiplier” effect), while city B remains peaceful.

Models of intergenerational cultural transmission may also help explain why cycles of violence and peaceful coexistence can be repeated over time. Suppose the relevant “cultural” traits in this context are tolerance on the one hand and fanaticism or extreme intolerance on the other. The model of cultural transmission described earlier can (under certain parameterizations) exhibit repeated cycles of temporary dominance, in which one trait suddenly grows very fast and becomes the dominant trait for a while, only to fall back and be replaced by the other trait after a few generations.

This project necessarily has an interdisciplinary dimension: I plan to interact with political scientists, sociologists and anthropologists to develop this specific research agenda.

5. Microcredit in Developing Economies

Paper (3) (with Ashok Rai (CID, Harvard) and Sajeda Amin (Population Council, NY)) empirically evaluates the targeting of microcredit programs in Bangladesh. Microcredit institutions (such as the Grameen Bank) have attracted a lot of attention as a supposedly successful anti-poverty intervention. But while there has been quite a lot of research into the impact of these programs, there has been relatively little research into their targeting success. We try to fill this gap by using a unique panel data set on two villages in Bangladesh, which allows us to measure poverty and vulnerability conditions of each household *before* they joined a microcredit program.

The main contribution of the paper is to bring together a development literature on targeting with a rich literature on general equilibrium risk sharing models. Following several contributions by Townsend, we define vulnerability as the inability to fully smooth consumption in the face of idiosyncratic risk. One of the advantages of this approach is that it delivers testable implications without having to be explicit about the specific formal or informal smoothing devices that may be available to village households. In practice, we estimate vulnerability for each household as the excess sensitivity of changes in household consumption to changes in household income, after controlling for changes in aggregate consumption, proxies for preference shocks, and other household characteristics. We then check whether poorer and more vulnerable households are more likely to join microcredit programs.

Our results indicate that, while microcredit is quite successful at reaching poor

households, it is not as successful at reaching vulnerable ones. More importantly, we find no evidence that microcredit reaches those who would benefit the most (in terms of welfare gains) from the availability of credit, namely the vulnerable poor. Of course, these households may also pose greater risks to credit providers, which points to a possible limitation of microcredit programs as an anti-poverty intervention. In the paper, we also contrast our vulnerability measure with more commonly used ones, based on variability of consumption, and argue that the latter are inadequate proxies for vulnerability. This paper is already being widely cited both by practitioners and researchers in development economics.

In future work in this area, we plan to expand our analysis in the following directions. The current dataset contains information about pairwise distances between households, with respect to several metrics (purely geographic, kinship, local neighborhoods). Informal smoothing devices should be relatively more important among households who are “close” in terms of their social distances. We intend to test whether our vulnerability measures are spatially correlated with respect to these distance metrics (as would be expected if these informal smoothing devices affect one’s vulnerability, and if our measures correctly capture vulnerability).

We also intend to use a second round of data collection (currently at the planning stage) to investigate the impact of microcredit on the poverty and vulnerability conditions of its members, and on several health and education measures. By collecting data on default rates, we can also investigate whether the households that were more vulnerable when they first joined microcredit were in fact more likely to default on their loans. This might help explain why such vulnerable households were less likely to become members in the first place.

6. Discrimination in the Labor Market

Paper (17) is work in progress with J.P. Benoit (NYU) and Raquel Fernandez (NYU). In it, we perform an empirical evaluation of a taste-based discrimination model that delivers testable implications for the variability of wages. Employers are assumed to care for profits but also have a taste for discrimination against black workers. We also assume that worker ability is measured with error, and that ability itself may have a varying impact on profits.

Then, in situations where the ability signal is very noisy and/or productivity is not affected much by individual ability, employers will be better able to indulge their taste for discrimination and thus will hire fewer black workers. At the same time, these will also be situations in which the variability of earnings, for a given job, will be lower. Therefore, the model implies the following testable implication: the variability of wages for workers in a given job and with similar characteristics will be positively correlated with the percentage of black workers in those jobs, across firms or across industries.

Further, if discrimination also results in a larger black/white wage gap, the variability of wages in a given job will be negatively correlated with the wage gap.

We evaluate these predictions using March CPS individual data on earnings, industry, occupation, location, as well as several individual characteristics. Preliminary estimation results indicate that indeed there is a significantly positive correlation between variability of wages and percentage blacks across industries, for a specific occupation, as well as a significantly negative correlation between wage variability and the black/white wage gap. These results hold even after controlling for local market conditions, gender, education, experience and union membership status.

We plan to conduct a series of robustness checks with respect to the empirical results, as well as work out the implications of possible competing explanations for these patterns, in order to test our model against some alternative hypotheses. In the long term, we would like to repeat the empirical analysis using more detailed firm level data from linked worker establishment datasets currently under construction at the Census Bureau. This would enable us to better control for possible biases resulting from unobserved heterogeneity in worker and job characteristics.

7. Other: Environmental Innovations

During my Ph.D. program in Chicago I have published two papers on environmental economics, with Carlo Carraro (University of Venice, Italy). Paper (6) analyzes a model of R&D and technological innovation to study the properties of an optimal fiscal policy instrument to induce environmental innovations. This is done under asymmetric information on the part of the regulator with respect to the innovation costs incurred by firms. The main result is that optimal adoption patterns can be implemented by a combination of emission taxes and innovation subsidies. Paper (7) extends the analysis to the case in which polluting emissions spill over across countries, thus giving rise to international coordination issues.