

Research Statement

Annika M. Mueller

Institute for Quantitative Social Science, Harvard University

I am a Post-Doctoral Fellow at the IQSS at Harvard University with an interest in contributing to a deeper understanding of individual decision-making. Human behavior is often characterized by deviations from perfect rationality and influenced by numerous factors that cloud the researcher's view of underlying causalities. I employ field, lab, and lab-in-the-field experiments, develop survey tools, and analyze large panel data sets to better understand economic decision-making and its psychological underpinnings, generating insights to inform theory as well as real world decision makers in the areas of public policy and management.

The majority of my current research applies field and lab-in-the-field experiments to test insights from classical and behavioral economics in the contexts of economic development (Sub-Saharan Africa) and the design of innovation contests (USA/internationally).

This research statement is organized as follows: The first section discusses my work in the area of development economics/public policy, with a focus on my dissertation papers. The second section discusses my work in the area of contests/management from my post-doctoral work. Both sections include plans for future research in the respective areas.

Experiments in the Area of Development Economics

Current Work:

The two chapters of my dissertation, which together form the basis of my first job market paper, are derived from a series of experiments in Malawi designed to investigate social norms:

The first chapter of my dissertation explores the role of culture as a co-determinant of economic growth by investigating the prevalence of six widely discussed norms of distributive justice in poor, rural communities of Malawi. I do this using a lab-in-the-field experiment in the form of a one-shot two-person dictator game with a production phase. More precisely, subjects had the option of generating social surplus through sorting beans. The social surplus generated by two matched players would then be shared by the dictator, who would be randomly determined among them, after the sorting decisions had been carried out by the participants. Alternatively, subjects had the choice to return all or part of the unsorted beans for money that would not go into the common pot but was theirs for certain. Subjects faced different rates of return for sorting, had different amounts of beans available for sorting as a proxy for differing opportunities to exert effort, and faced random income shocks in some treatments after carrying out their sorting decisions.

I find that decisions are guided by self-interest balanced with a multiplicity of norms. Interestingly, and despite a large influence of strict egalitarianism on sharing behavior, subjects in my sample react as intuitively expected to all experimental variations. Specifically, subjects reward own and others' effort, take relative rates of return into account, and compensate themselves and others for income shocks and limited opportunities to exert effort.

The knowledge gained contributes to the debate of culture as a determinant of economic growth in important ways. Arguably, the link between norms of distributive justice and redistributive outcomes is closer in environments with a relative lack of formal institutions. Therefore, in such environments, norms set incentives for effort-based production in a more immediate way, co-determining whether and which productive activities are encouraged in society. In this context, egalitarian norms have often been pointed out to inhibit the emergence of entrepreneurial activity in rural areas of developing countries. Despite the consequent importance of the relative prevalence of norms for rural development, previous evidence about sharing behavior for developing countries has generally not allowed for the distinction between various norms of distributive justice. Norms map a set of variables which can be identified in the real world, as well as in experimental settings, into an optimal distribution of effort-generated income. Thus, identifying underlying norms of distributive justice is of crucial importance for gaining a deeper understanding of the impact of cultural values on output, and ultimately development, since it allows for a better assessment of the generalizability/external validity of results. This chapter is an improvement over the current literature because studies that *did* focus on norms of distributive justice usually did not focus on effort-generated income; those that *did* have tended to employ a design that lets subjects invest part of a windfall gain obtained in the course of the experiment, which is likely an unsatisfactory proxy for effort-generated income. For the reasons outlined above, these may be serious omissions in the context of development.

Predicting the expected impact of values on behavior in the medium to long run remains difficult nevertheless, since sharing norms may not be independent of the current social environment of those who subscribe to them. Specifically, sharing norms may either be fully internalized or domain-specific to a particular informational environment, two options between which I discriminate in the second chapter of my dissertation. In this chapter, which together with the first chapter forms the basis of my first job market paper, I extend the experimental framework of the first chapter using imperfect information regarding the size of the pie and share taken by the dictator on the side of the receiver. The purpose is to investigate the origins of the high prevalence of equal sharing of effort-generated income that I found among the subjects of my first chapter, and, more generally, to understand how subjects balance their behavior between self-interest and norm-adherence if their degree of norm-adherence can no longer be fully judged by others.

I find that dictators act more selfishly under incomplete information, i.e. deviate further from the various norms of distributive justice in situations where their social image cannot be damaged by such action. This finding is significant for two reasons, which help us gain perspective on the findings of the first chapter: a) when norms are domain-specific to a particular informational environment, changes in community structures arguably lead to more abrupt changes in sharing behavior, and b) community structures are rapidly changing in the developing world (for example through internal migration, out-migration and remittances, but also increased financial infrastructure such as the introduction of bank accounts).

The third chapter of my dissertation, which is co-authored with my colleagues at the University of Michigan, the IPA, and the World Bank, draws lessons from a randomized field experiment in rural Malawi that temporarily increased access to markets for previously isolated areas through regular subsidized transport. The cost of being isolated is a growing subject of research in development economics. Rural road investments have been pointed out to be critical for households to overcome spatial poverty traps. Over the past decades, many such investments in Sub-Saharan Africa seem to have been built on donors' and governments' assumption that the provision of roads that are passable for motorized vehicles will automatically lead to market provision of such transport and thus, poverty reduction and income generation. There is, however, a lack of rigorous evidence on these assumptions; namely the lack of randomized field studies evaluating the profitability of routes.

The randomized experiment summarized in this paper is the first of this kind in Sub-Saharan Africa and illustrates that improving the condition of roads does not necessarily generate transport provision at an affordable price for villagers. In other words, based on take-up, this experiment demonstrates that a bus provider cannot, at any price, break-even on these routes. This result is strengthened by the fact that subsidized transport was provided on a regular basis over a period of six months. This contrasts sharply to an often encountered situation in real life in which a bus provider often does not possess the required funds to scout out a new area for more than a few days. Potential customers cannot predict the supply of transport well in such a case and hence (depending on the specifics of the area and its population), observed demand during such a scouting period may seem lower than actual demand would be if regular supply of transport was available, leading the bus provider to underestimate potential demand for transport service and not provide services to potentially profitable routes. Given that in our study a regular bus schedule was communicated to all potential customers and this schedule was honored by the transport provider, we can conclude that in our case true demand was revealed and yet a bus provider cannot make a profit on these routes.

This explains why many rural roads in relatively good condition are currently not being used by motorized vehicles and this has important public policy implications. It helps policy makers design solutions that actually increase access to markets through a) an increased understanding of why, despite massive investments in infrastructure, the expected outcomes on access to markets have often not materialized, and hence, b) an increased understanding about which data to collect in order to be able to assess whether rural road investments may crowd out investments in other sectors that may have a greater impact on economic and social development.

Future work:

A related paper I am currently working on with my co-authors is based on the same field experiment and looks at income, health, and time-allocation effects of the transport subsidy (including spill-over effects) over a period of six months.

I view all of the above-mentioned papers as work in progress that I intend to complete and submit for publication within the next eight months, and that may lead to further related projects. I have been to Malawi for a total of 1.5 years and have access to a large logistical network. I intend to harness this research capital, as well as my personal contacts with government officials and villagers, during the next years of my career for several additional projects.

Specifically, I intend to conduct another lab-in-the-field experiment in Malawi in Spring 2013 in which I will look at how stable spiteful preferences are, and am currently also going through the IRB approval process for a methods paper, the idea for which stems out of my experiences during past piloting.

Experiments in the Area of the Economics of Contests

Current work:

Two papers coming out of my first year at the Harvard-NASA tournament lab deal with creative workers in an open crowdsourcing contest setting.

The first, which is also my second job market paper, is a field experiment that investigates the impact of various levels and combinations of cash, job market signaling, and peer signaling incentives on participation, effort, and performance of creative workers in a real innovation tournament on an open crowdsourcing platform. Over the last decade there has been an increased focus on the optimal application and academic study of non-monetary incentives, including signaling. But while contests feature a multitude of such incentives, no previous field experiment, to our knowledge, has disentangled the effects of signaling (in the form of the revelation of the rank-order of contestants) from monetary incentives on the composition of entrants in public contests.

Theoretically, the outcome of a tournament can reveal information about the abilities of the entrants. This revelation can bring status and prestige to the winners (e.g. in the case of the X-Prize) and thus act as a reward. However, for entrants who do poorly, the revelation can be damaging. In effect, publicizing the tournament outcome creates personalized entry fees for entrants so that those with lower ability opt out. All in all, our model shows that the ability-dependent incentives created by the public nature of some tournaments imply that entrants in such tournaments will have higher ability and exert more effort than those in an equivalent private tournament.

We test the implications of the theoretical model with data from a field-experiment that features a real-world, non-trivial, prestigious innovation task by randomizing coders across contests which differ only with respect to incentives offered. Specifically, we offer various combinations of three types of prizes on two levels: high and low cash prizes, public announcement of the results to their peers or not, and an optional announcement of results to potential employers in some treatments.

All participants were part of the TopCoder community of programmers and competed over 4 days to produce the highest performing algorithm to solve an abstract computer science problem. Cash prizes were offered for the top five solutions in each treatment cell, totaling \$500 and \$1500 in the low and high cash treatments, respectively. The results were made public in the public announcement treatment by publishing a “glory page” on the TopCoder website listing the rank-order results of the top five participants in the respective treatment cells. The winners in the optional job market signal treatment were offered a reference letter including their rank-order in the tournament that could optionally be sent to Google and the NASA Jet Propulsion Laboratory recruitment headquarters. Both organizations participated in organizing the tournament and are popular potential employers for members of the TopCoder community.

Our data include a high-quality objective skill measure which enables us to understand how individuals of different skill levels react to different incentive mixes. In addition, survey data allow us to investigate how effects differ for people driven by different self-reported motives for participation. Our results are supportive of the theoretical implications. Publically announcing results increases participation of high ability entrants by about 30% over when there is no announcement. We can also verify our hypotheses regarding signaling and the opting-out of low ability entrants and identify situations in which cash as additional incentive can be used to compensate such entrants for the expected cost of sending a negative signal about their ability. Given that tournaments are a historically important and increasingly popular mechanism to elicit innovation; our findings highlight a low cost mechanism by which to control entry.

The second paper looks into sorting and participation dynamics of differently motivated software developers in different types of contests on an open crowdsourcing platform over a three-year period. To examine these dynamics we have put together an extensive panel data set on contest characteristics/incentives and participation information, coders’ motives for participation at the time of joining, current motives, and information about why these have changed from when they joined the platform (if they did). One feature of the data that is particularly novel is that the classification of contests regarding subjective incentives (such as the status gained from winning a particular contest) that we use stems from a

meta-contest we launched among platform members for this purpose. This feature allows us to explicitly consider a multitude of non-pecuniary besides pecuniary motives and incentives for driving observed behavior.

We confirm that observed behaviors are consistent with an evolution of motives. Consistent with there being a “life cycle” of non-monetary motives, we observe sorting patterns that initially intensify in accordance with early work motives and eventually wane. In many cases, workers eventually stop participating. This study is significant for multiple reasons. Despite wide recognition in the (open source) innovation literature that the appeal of non-pecuniary incentives appears to be especially conspicuous for creative workers, model-based analyses and considerations of institutional designs that focus on individuals’ motives are static. Additionally, compared to literature in other fields on dynamics similar to those we consider (pertaining to individuals’ career paths), we have a relative wealth of data on non-pecuniary motives and incentives, as well as on how they evolve. Our findings highlight that institutions designed to harness non-monetary motivates may need to carefully contend with the possibility of muted incentives or accelerated attrition; “staging” or life cycles might therefore be inherent factors in managing crowdsourcing platforms.

Future work:

I have two other works in progress. The first is a lab experiment on a modified trust game that partially untangles the multitude of facets of trusting behavior as measured in a traditional trust game. The second is a project with colleagues from the Peking University HSBC Business School on a follow-up study to their paper on gender signaling in all-pay auction experiments.

There are also plans with my co-authors at Harvard and LBS to collaborate on future projects. One question which particularly interests us is the so-called “alliance formation puzzle” in contests, which we hope to investigate using a field experiment.

To summarize, at this point in my career, my primary interests are in experimental economics applied to the field of economic development and management. I have pursued these interests to date by conducting field experiments in two very different contexts, and lab-in-the-field experiments on social norms. In the immediate future I would like to go further with lab and lab-in-the-field experiments in the realm of social norms. In the long term I see myself working on expanding my expertise in these areas by a mixture of experimental and empirical work.