Research Statement, September 2017
Amanda Pallais, Harvard University

My research areas are labor economics and the economics of education. I strive to understand the labor market performance and educational investment decisions of disadvantaged and socially excluded groups. My work has three main components.

The first concerns groups that have difficulty getting a foothold in the labor market: young workers, workers without social connections, and ethnic minorities. I find that employers' uncertainties and misperceptions about young and disadvantaged workers' abilities and qualifications can prevent them from being hired and can even harm their on-the-job performance. The second strand considers the gender earnings gap, exploring to what extent family responsibilities and marriage market concerns can limit women's labor market success. Finally, I focus on access to higher education. Students from disadvantaged families are less likely than are students from higher-income families to attend college, choose selective colleges, and ultimately graduate from college. I analyze the extent to which improvements in information, financial aid, and emerging online education can ameliorate socioeconomic status differences in access to and performance in higher education. Factors preventing workers from reaching their labor market potential not only lead to inequities, but they can also generate economic inefficiencies (and lost production) with workers' talents not being fully utilized and valuable human capital being misallocated.

My research approach primarily relies on the design, implementation, and analysis of large-scale randomized field experiments. Non-experimental observational studies face severe challenges in assessing the causal impacts of labor market and educational interventions because of the difficulty of fully controlling for many plausible confounding factors (omitted variables). Thus, field experiments that allow the researcher to generate exogenous variation in the interventions (explanatory variables) of interest are often necessary to determine which factors *cause* poor labor market performance. My experiments take place not in the lab but in the context of the participants' existing activities such as applying for a job or a college scholarship, thereby enhancing external validity and policy relevance. The experiments are large, typically with thousands of students or workers. They are tightly linked to economic theory and are explicitly designed to test the validity of different theories. Finding the right setting and institutional partner and creatively designing the treatments and data collection allow me to determine mechanisms, test theories, and rule out alternative explanations, which would not be possible using existing data and non-experimental observational methods.

I. Socially Excluded Groups

My research agenda started by focusing on young (and inexperienced) workers. Young workers have much higher unemployment rates than do older ones. Policies designed to help young workers often focus on training. However, my paper, **"Inefficient Hiring in Entry-Level Labor Markets" (*American Economic Review*, 2014)** suggests that young workers' poor labor market outcomes might not come solely from human capital deficiencies; instead, some workers may simply need help getting a foot in the door and an opportunity to demonstrate their abilities. New labor market entrants face an unfortunate catch-22: it's difficult for them to get a job without prior experience, but it's difficult for them to get experience without a prior job. I develop a model that formalizes this dilemma. Firms are uncertain about the abilities of inexperienced workers. Firms learn about workers' abilities by hiring them and observing their performance. This information is valuable because it allows a firm to employ more-productive workers. However, if employers cannot keep private the information they discover,

other firms will try to poach workers discovered to be more productive, bidding up these workers' wages. Thus, the initial employer will not capture the entire benefit from hiring a novice worker. Employers thereby may hire inefficiently few inexperienced workers, denying young labor market entrants the chance to showcase their abilities.

Without an experiment it is difficult to know whether young workers would benefit from being able to demonstrate their abilities and, even if so, whether any gains would merely come at the expense of more-experienced workers. I answer these questions by utilizing a large field experiment in an online labor market. In this marketplace, employers leave public ratings of their workers' job performance. Acting as an employer, I posted job openings and randomized which of the roughly 3,800 applicants I hired. Among the thousand workers I hired, I also randomized the amount of information provided in the (public) evaluations. Allowing workers to demonstrate their skills, through employment and more-informative evaluations, substantially improved their labor market outcomes, even without training. The gains for these workers did not come solely at the expense of others. Even under conservative assumptions, the benefit of hiring the workers and revealing their abilities to the market outweighed the cost, suggesting that inexperienced workers had been inefficiently underemployed beforehand.

Young workers are not the only ones who have difficulty getting a foot in the door of the labor market; those without social connections also have trouble. Most jobs are found through social connections, such as friends and relatives. Employers are more likely to hire referred applicants than observationally-equivalent applicants without referrals. In **"Why the Referential Treatment?: Evidence from Experiments on Referrals" (*Journal of Political Economy*, 2016)** with Emily Sands, I investigate why employers prefer referred workers. Is this simply nepotism or do referrals allow firms to identify a more-productive workforce? If the latter, providing unconnected workers with ways to demonstrate their abilities may help them as it did for inexperienced workers.

The question is difficult to answer with existing data. Because there is differential selection of referred and non-referred applicants into employment, comparing the performance of *hired* referred and non-referred workers does not reveal whether referrals provide information about applicants. We ran a series of experiments with more than 1,200 workers in an online labor market. We hired workers, asked for referrals, and hired all referred and non-referred applicants. We found that the referred applicants performed substantially better than non-referred applicants and had less turnover. Referrals contained information about workers' abilities that was not present on their resumes or in their initial job performance. Although this is, in part, because people work better when working with their referrers, we find that referred workers perform better than otherwise identical workers even when not hired through a referral and not working with their referrer.

Once workers from socially-excluded groups get a foot in the door with an employer, will managers recognize and utilize their talents? In **"Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores" (*Quarterly Journal of Economics,* 2017)** with Dylan Glover and William Pariente, I focus on ethnic minorities, a group that has traditionally low employment rates and wages. While a large literature has assessed discrimination against minorities at the hiring stage, we examine whether managers' biases can directly harm workers' performance after they are hired. The workers we study are cashiers who are quasi-randomly assigned to different managers (supervisors) on different days. Drawing on research from psychology, we measure manager bias towards minorities with an Implicit Association Test (IAT). When they work with more-biased managers, minority cashiers – but not majority cashiers – are absent more often, leave work earlier, scan items more slowly, and take more time between customers. This pattern appears to arise because biased managers are uncomfortable

with minorities and avoid them, leading minorities to exert less effort when supervised by biased managers. Worse minority performance under biased managers can lead to slower advancement within the firm and even feed back into lower minority hiring rates. We find that, on average, minority and majority workers perform equivalently, but minorities perform significantly better than do majority workers in the presence of unbiased managers. The findings are consistent with statistical discrimination in hiring where, in the presence of manager bias, the firm sets a higher hiring standard for minorities to offset the under-performance of minority workers when assigned to biased managers.

II. Gender Earnings Gap

Despite being more likely to graduate from high school and college than their male counterparts, women still earn significantly less. One common explanation is that family responsibilities lead women to take more-flexible jobs and women pay a compensating differential for this flexibility. However, it has been difficult to estimate compensating differentials for work arrangements because these amenities can be correlated with unobserved person- or job-specific characteristics**.** In **"Valuing Alternative Work Arrangements" (forthcoming, *American Economic Review*)** with Alexandre Mas, I estimate worker preferences through an experiment with more than 7,000 job applicants to phone survey positions. Each applicant chooses between two jobs, where both the work arrangement and wages are randomized, allowing us to estimate the distribution of willingness to pay for different work arrangements. Surprisingly, we find that the typical (median) worker does not value flexible scheduling, neither the ability to choose how much nor when to work. However, some workers do value this flexibility highly, allowing for the possibility of large compensating differentials. Workers value working from home and are willing to give up 20% of wages to avoid schedules that can be set by employers on short notice. Although women are more likely to be in work-from-home jobs and less likely to have irregular schedules than men, the differences in the incidence of the work arrangements we study are not large enough to explain a substantial fraction of gender wage gaps among hourly workers.

"'Acting Wife': Marriage Market Incentives and Labor Market Investments" (forthcoming, *American Economic Review*) with Leonardo Bursztyn and Thomas Fujiwara shows that women face tradeoffs between career and family even before marriage. Research from economics and psychology shows that men typically prefer female romantic partners who are less professionally ambitious and successful than they are. Thus, actions that help single women's careers may harm their marriage market prospects. We test whether single women respond to this tradeoff with two experiments on the entire entering cohort of an elite U.S. MBA program. In the first experiment, randomly-selected students were primed to believe that their responses to a placement questionnaire would be viewed only by a career counselor, while the remainder was primed to believe their classmates would also see their responses. When students thought only a career counselor would see their answers, single women, non-single women, and men all reported similar preferences. However, when they thought their classmates would see their answers, single women reported $18,000 lower desired yearly compensation and a willingness to travel seven fewer days per month and work four fewer hours per week. Men and non-single women did not report different preferences in front of their classmates. In a second experiment, students made choices over hypothetical jobs, knowing that a career counselor and a randomly-selected small group of classmates could see their answers. Single women made substantially less career-focused choices in front of single, male peers than in front of female or married male peers. Using observational data on grades, we find a similar pattern in the classroom. Unmarried women perform similarly to married women in class when their performance is kept private from their classmates (on exams and problem sets); yet, they have significantly lower participation grades.

III. Access to College and Upward Mobility for Low-Income Students

The final part of my agenda analyzes the barriers preventing financially disadvantaged students from investing in post-secondary education. Although college is a key stepping stone to upward mobility and economic security, students from low-income families are less likely to attend college, choose selective colleges, and persist in college than are higher-income students with the same grades and test scores. **"Small Differences that Matter: Mistakes in Applying to College" (*Journal of Labor Economics*, 2015)** finds that this is in part due to low-income students' lack of information about the college application process. Low-income students send fewer applications and send their applications to less-selective colleges than do their higher-income peers with the same grades and test scores. They do not follow recommended application strategies. However, when students taking the ACT college entrance exam were allowed to send their standardized test scores to four colleges free of charge instead of just three, low-income test-takers sent more applications. More than 70% of low-income students went from sending their scores to three colleges to four and many of these additional score reports went to more-selective colleges. Although the reduction in the cost of the fourth score report was small (from $6 to free), it signaled to low-income students that they should apply to more colleges and they complied. As a result, low-income students (but not higher-income ones) attended more-selective colleges. Back-of-the-envelope calculations imply that sending an additional score report increased the average low-income students' lifetime earnings by over $10,000.

The cost of college prevents low-income students from investing in their educations. **"Evaluating Post-Secondary Aid: Enrollment, Persistence, and Projected Completion Effects"** and **"Early Results from a Randomized Evaluation of Post-Secondary Aid" (NBER Working Papers No. 20800 and 23015)** with Joshua Angrist, David Autor, and Sally Hudson examine the effects of randomly providing 3,700 low-income students with full-tuition college scholarships. These scholarships provide aid for up to five years and can be used at any public Nebraska college through the Buffett Scholarship program. We find that the scholarships increase college enrollment, induce students to attend four-year instead of two-year colleges, and decrease dropout rates. Scholarships offered to groups with traditionally low college attendance – including non-white students and students with low academic achievement – generate the largest gains in enrollment and persistence. Scholarships increase the probability that the most at-risk students are enrolled in any college four years after high school graduation by 30 percentage points.

We are continuing work on this project and plan to soon release a paper analyzing the effect of the scholarships on college graduation. Our preliminary estimates suggest that the scholarships will increase college graduation but delay it for students who would have graduated without the aid. We have laid the groundwork to follow these students after graduation, using this experiment to explore the effects of college-going and improved college quality on longer-term outcomes including earnings and credit scores.

Many states have broad-based merit scholarship programs that award grants to students who meet certain GPA or test score thresholds. Our experimental results from the Buffett Scholarship program suggest that by targeting high-achieving students, many of these programs are giving scholarships to students for whom financial aid is less effective in improving college outcomes. However, in **"Taking a Chance on College: Is the Tennessee Education Lottery Scholarship Program a Winner?" (*Journal of Human Resources*, 2009)**, I find that these state merit scholarship programs can provide a strong incentive for students to work harder in high school. I analyze Tennessee's program which provided college financial aid to students who scored above a given threshold on the ACT. As a result of this

program, Tennessee high school students increased their standardized test scores and were substantially more likely to score above the cutoff.

Online education, including Massive Open Online Courses (MOOCs), has been heralded as potentially transformative in its power to increase access to higher education. Unlimited by physical capacity constraints, it can give more students access to the best professors. Economies of scale allow it to be offered more cheaply than in-person classes and it may be more attractive to students who can live anywhere and watch lectures on their own time. Yet, existing research has been relatively negative on online education, finding that students tend to perform slightly worse in online than in-person classes, implicitly treating these two types of education as substitutes. **"Can Online Delivery Increase Access to Education?" (forthcoming, *Journal of Labor Economics*)** with Joshua Goodman and Julia Melkers is the first paper we know of that analyzes whether online education substitutes for in-person programs or attracts new students. We study the nation's first degree program delivered entirely by MOOC-like courses: Georgia Tech's online master's of science in computer science. The program delivers rigorous courses at a much lower cost than the corresponding in-person program. (It has since been used as a model for other programs.) Using a regression discontinuity in admission, we find that the online program almost exclusively attracts students who wouldn't otherwise have obtained additional training. In particular, it attracts older U.S. students who are less likely to have computer science training and who attended less-expensive colleges. We conservatively estimate that by satisfying large, previously unmet demand for mid-career training, the Georgia Tech online program alone will boost annual production of American computer science master's degrees by 7 percent.

Online education and online labor markets provide an opportunity to learn more about education and labor markets more generally, and they are increasingly important in their own right. The nature of work and education is changing. Many more workers have non-standard work arrangements today than they did ten years ago. These arrangements tend to provide more flexibility but less stability and fewer benefits. Education is also facing a shift as more students take online classes and take classes outside of formal degree programs. Understanding what these changes mean for the future of education and work and how they will differentially affect disadvantaged populations is an important part of my ongoing research agenda.

IV. Summary and Future Work

I use large, randomized field experiments to better understand the labor market barriers facing inexperienced and unconnected workers, ethnic minorities, and women, as well as the educational barriers facing low-income students. Through careful intervention in real-world labor markets and higher-education settings, my research designs new ways to test economic theories and to measure the causal impacts of specific factors preventing socially-excluded groups from flourishing in the labor market. For example, to determine whether manager bias harms workers' job performance, I use tools from psychology to measure manager bias in a setting in which there is quasi-random variation in the manager on duty. To determine the effect of marriage market concerns on women's career investments, I experimentally vary whether single and partnered women make career decisions in the presence of (single, male) peers. And instead of inferring preferences for job amenities by estimating compensating differentials for difficult-to-measure amenities in observational data, I measure them directly by experimentally varying job applicants' real-life choices over different work arrangements and wages.

I plan to continue this direction in my future work. For example, my paper **"Labor Supply and the Value of Non-work Time: Experimental Estimates from the Field" (draft available)** with Alexandre Mas

experimentally estimates workers' value of leisure and other non-work activities relative to their market wages (labor productivity). This value of leisure parameter (capturing the ratio of value of time to productivity) is important in many applications in labor economics and macroeconomics, such as understanding wage dispersion and business cycle volatility and conducting cost-benefit analyses of programs that increase employment. This key parameter has typically been calibrated using difficult-to-assess observational approaches. We estimate it directly by giving job applicants choices over positions with different weekly hours and wages. Our approach allows us to determine workers' value of time in 5-hour increments from 5 to 50 hours of work per week. We estimate an intermediate value for this parameter relative to the existing literature. The implied pattern of labor supply elasticities – high at low hours of work but decreasing in weekly hours – can help rationalize why there are many full-time workers and many individuals who are out of the labor force entirely, but fewer part-time workers.

Representative of my planned future work is a new project with Leonardo Bursztyn, Zoë Cullen, and Thomas Fujiwara analyzing the effect of social dynamics in the workplace. Many high-powered work environments reportedly still have the flavor of an 'old boys' club' where women interact less with their (primarily male) coworkers and bosses, potentially contributing to women having poorer job performance, higher turnover rates, and lower salaries and promotion rates even conditional on performance. Understanding the effect of social dynamics within the firm on women's career trajectories is challenging due to the difficulty in quantifying these interactions and the lack of random variation in these interactions. Partnering with a large bank in Vietnam, we have data on emails between employees (including their content), time clock data (allowing us to determine who breaks together), performance, and promotions and bonuses. These data, combined with an employee survey on the quantity and nature of other interactions, allow us to characterize socialization patterns by gender and explore how these correlate both with performance and employment outcomes conditional on performance. We also have uncovered plausibly exogenous sources of variation in the ability of different workers to network, specifically in the ability to take smoke breaks with coworkers. Smoking, which is more prevalent among men than women, is common, but not ubiquitous. The ability to smoke with coworkers and bosses changes with the season (due to temperature) and the bank's smoking policy. We plan to analyze how changes in the relative ability of smokers and non-smokers to socialize with smoker and non-smoker bosses affect performance, employment outcomes, and gender gaps. We also have the opportunity to run experiments to assess the types of firm policies that can lead to both high performance and equitable outcomes.