Comments and Discussion

COMMENT BY

CONRAD MILLER Will Dobbie and Crystal Yang characterize the economic costs of pretrial detention with a focus on its consequences for the labor market and economic insecurity. They discuss both the micro effects of pretrial detention on the detained and the macro effects on the broader community. At the micro level they summarize findings from their seminal paper, Dobbie, Goldin, and Yang (2018), where they find that pretrial detention increases the chances of conviction and decreases formal employment rates, earnings, and public benefits receipt using data from Philadelphia and Miami. At the macro level they present new findings on the county-level correlates of pretrial detention rates. They document that counties with larger increases in pretrial detention rates experience larger reductions in employment rates and increases in poverty rates.

The county-level results are particularly provocative. For example, the authors find that a 10 percentage point increase in county pretrial detention rates between 2000 and 2009 is associated with a 2 percentage point decrease in county employment rates for prime working-age adults.[1] The relationship is particularly strong and negative for Black employment rates. As Dobbie and Yang readily acknowledge, the aggregate relationship between pretrial detention rates and economic security is only suggestive of a causal relationship and likely reflects many confounding factors. One potential factor is that bail judges and magistrates set more generous bail conditions when defendants are gainfully employed. Another potential factor is that criminal justice jurisdictions with high pretrial detention rates may

1. County employment rates are calculated for 2000 and 2010, and exclude institutionalized individuals.

also be punitive along other dimensions, including conviction rates and confinement rates.

However, for the remainder of this discussion, I will assume that there is in fact a causal relationship between pretrial detention rates and employment rates and that the magnitude is economically important. My discussion will focus on what this causal relationship (should it exist) tells us about how employers consider applicants' criminal records in hiring decisions. Dobbie, Goldin, and Yang (2018) argue that pretrial detention lowers employment rates at least in part by increasing the chances that a defendant is convicted and that having a criminal conviction makes it harder to find a job. I will argue that the aggregate pattern suggests that employers are not primarily concerned about a job seeker's conviction because it signals something about worker productivity. If it were employers' primary concern, it is not clear why an exogenous increase in pretrial detention rates, and hence conviction rates, would significantly decrease aggregate employment. Instead, the aggregate pattern suggests that employers care about conviction status above and beyond its signaling content and that employers screen on conviction status directly, likely because convictions increase the perceived risk of negligent hiring lawsuits (Cavico, Mujtaba, and Muffler 2014; Lageson, Vuolo, and Uggen 2015) or restrict the set of tasks an employee can legally perform (Jacobs 2015).

Why does pretrial detention affect a defendant's labor market outcomes? In the short run, pretrial detention may lead to immediate job loss or disrupt educational attainment, housing, or family stability. As Dobbie, Goldin, and Yang (2018) argue, pretrial detention has longer-run effects on defendants' job prospects by increasing their chances of conviction, likely by weakening their bargaining position.[2] This conviction record follows the individual into the labor market.

These findings are consistent with an accumulating body of evidence that shows that having a prior conviction substantially harms an individual's job market prospects, regardless of the nature or length of the associated sentence (Pager 2003). Recent estimates indicate that about 25 percent of US adults have an arrest or conviction record (Jacobs 2015), and 13 percent of adult males have a felony conviction (Shannon and others 2017). We know that many employers use information on convictions in the hiring process; a recent survey of human resource professionals found that over 70 percent of firms conduct background checks for new hires (Holzer,

2. Heaton, Mayson, and Stevenson (2017) and Gupta, Hansman, and Frenchman (2016) also document that pretrial detention increases conviction rates.

Raphael, and Stoll 2004; Society for Human Resource Management 2018). While arrest records are more difficult to obtain in some states, court records on convictions are widely available (Bushway and Kalra 2021).

Many employers report that they are unwilling or reluctant to hire workers with criminal records (Holzer, Raphael, and Stoll 2004). A series of audit and correspondence studies provide compelling evidence that having a criminal record substantially reduces callback rates (Pager 2003; Agan and Starr 2018). This is true even for relatively minor offenses (Uggen and others 2014).

Mueller-Smith and Schnepel (2021) offer another compelling example. They study the criminal justice practice of diversion, where instead of getting a conviction on their record, defendants can avoid a conviction by successfully completing a probation sentence. In the setting they study, the marginal defendants who have their case diverted don't actually get a reduced sentence, they are just less likely to have a conviction on their record. Despite this seemingly artificial difference, the authors find large labor market gains to case diversion.

Most recently, Agan, Doleac, and Harvey (2021) study misdemeanor arrests and whether the prosecutor assigned to a defendant's case decides to file a court charge or declines to pursue the charge. Getting charged creates a criminal record and increases the chances of a conviction. But the authors study minor offenses where even defendants that get charged do not face significant punishment. They find that getting charged increases recidivism, as well as the chances that a defendant has a criminal record in the state repository. Again, a likely explanation for the increase in recidivism is that the presence of a criminal record worsens an individual's labor market prospects.

If we take for granted that pretrial detention affects a defendant's labor market outcomes primarily by increasing the chances of a criminal conviction, that leads to a natural follow-up question: Why do employers screen for a conviction in the hiring process? There are at least three reasons.

First, a natural labor economics view is that a conviction record can affect labor demand by serving as a negative signal of an individual's productivity (the productivity view). For example, employers may view applicants with criminal records as untrustworthy. Note that, for a prior conviction to be an informative signal, it must predict productivity above and beyond information on job applicants that is already readily available to employers.

A second view (the stigma view) is that a conviction record is associated with social stigma and that employers, employees, or customers would prefer not to employ, work with, or interact with individuals with convictions, regardless of on-the-job productivity. This is akin to the "taste-based" discrimination in Becker (1957).

A third view (the legal costs view) is that employers prefer not to hire individuals with criminal records because of the legal restrictions or costs associated with a criminal conviction. Laws prevent those with certain convictions from working in some occupations. Employing workers with a prior conviction may increase an employer's vulnerability to a negligent hiring lawsuit. If an employee harms a coworker or customer, the employer may be held liable for damage if it can be shown that the employer was negligent in hiring that worker in the first place (Cavico, Mujtaba, and Muffler 2014). A prior conviction can be used as evidence for such negligence.

Survey evidence provides support for all three views (Society for Human Resource Management 2012; Lageson, Vuolo, and Uggen 2015). However, it is difficult to infer from employer survey responses alone what considerations drive hiring behavior (Pager and Quillian 2005). Determining which views are relevant is important because different mechanisms suggest different policy solutions. To the extent that employers are focused on the legal costs and restrictions of an employee's conviction status, policies that shape those costs and restrictions will influence labor demand for job seekers with conviction records.

I will argue that the notion that pretrial detention reduces employment rates in the aggregate is difficult to reconcile with the productivity view and most consistent with the legal costs view for why employers care about conviction records.

Suppose that local courts increase conviction rates while holding criminal conduct fixed, which is arguably the first-order effect of increasing pretrial detention rates. How would we expect this to affect the functioning of the labor market and how employers infer job seeker productivity in particular? We can view this policy as increasing the set of information about job seekers that is available to employers.[3] That's because conviction records are typically more readily available than arrest records (Jacobs 2015; Bushway and Kalra 2021). When conviction rates are low, there are more job seekers

3. Alternatively, we can view this policy as changing the categorization of job seekers rather than increasing the set of available information per se. For example, suppose the way the criminal court system works is that cases are ranked by the severity of the underlying criminal conduct and courts vary in the threshold they set for determining conviction. Suppose further that employers can observe who is convicted and some case details but not underlying criminal conduct. Then lowering the threshold for conviction increases the number of applicants for whom case details are available but may reduce the strength of the signal that a conviction provides.

with criminal records that some employers do not have access to. When conviction rates are high, those employers will have more information on arrests for the marginal convictions.

If we interpret an increase in conviction rates as increasing the set of information available to employers, it is not clear why this change would reduce employment rates in the aggregate. On the one hand, we would expect job seekers with those marginal convictions to have more trouble finding a job. On the other hand, when conviction rates are low, employers may try to infer the criminal history of job seekers from other available information, including applicants' job histories and personal characteristics. In the absence of objective information about who has a criminal record, employers will depend on their own subjective assessments of who is likely to have a criminal record. This behavior disadvantages job seekers who are stereotyped as likely to have a criminal record. Hence, an increase in conviction rates will help job seekers with no criminal history who may nonetheless be stereotyped as likely to have a criminal history when conviction rates are low.

There is an analogy here to the literature on ban-the-box policies, which prevent some employers from asking job applicants about their criminal history at the initial screening stage (Raphael 2021). Prominent research has argued that ban-the-box policies widen Black-white inequality in labor market outcomes by making it more difficult for employers to distinguish between Black applicants with and without criminal records (Agan and Starr 2018; Doleac and Hansen 2020). The argument goes that, in the absence of direct information on criminal history, employers may statistically discriminate against Black applicants. In the ban-the-box case, research indicates that a reduction in the criminal history information available to employers worsens Black labor market outcomes in particular. In this paper we have a change that arguably makes criminal history information more available, yet employment rates are reduced, particularly for Black adults.

Dobbie and Yang provide another clue that employers care about convictions per se rather than the information they convey about worker productivity. To identify the causal effect of pretrial detention Dobbie, Goldin, and Yang (2018) use what's known as a judge design—they take advantage of the fact that cases are essentially randomly assigned to bail judges or magistrates, and those judges and magistrates vary systematically in their tendency to detain defendants or subject them to monetary bail. They compare the outcomes of defendants who are assigned to low detention rate judges to otherwise similar defendants who are assigned to high detention rate judges. An important feature of this research design is that, for the marginal convictions the authors study, conviction actually conveys no information about the defendant, including the defendant's productivity as a worker. Hence, if employers (a) only care about an applicant's conviction record to the extent that it conveys information about productivity and (b) can identify marginal convictions, then we may expect Dobbie, Goldin, and Yang (2018) to find no effect of pretrial detention on employment and earnings. The fact that they find large and negative effects suggests that either (a) or (b) does not hold.

In practice, employers cannot identify marginal convictions, at least not perfectly. (In fact, neither can the authors.) But when employers can access arrest records, prior work suggests that, conditional on initial charges, conviction is sufficiently arbitrary that it is not clear how much information it provides about worker productivity. The fact that labor market outcomes are so responsive to marginal convictions strongly suggests that employers care about convictions above and beyond their informational content.

While the findings Dobbie and Yang present are difficult to reconcile with the idea that employers use convictions in screening for their informational content, they are consistent with the view that convictions per se increase the legal costs of employing someone.[4]

Convictions restrict the set of occupations that individuals can legally work in (Jacobs 2015). For example, job seekers convicted of sex offenses may be banned from working with children. An increase in the conviction rates increases the set of job seekers subject to these legal restrictions.

A conviction on an employee's record makes employers more vulnerable to a negligent hiring lawsuit if that employee harms a fellow employee or a customer (Cavico, Mujtaba, and Muffler 2014). This increases the perceived risks associated with hiring a worker with a prior conviction. In both cases, it is easy to see why an increase in conviction rates would deteriorate labor market prospects for job seekers with criminal records and potentially decrease employment rates overall.

In summary, I interpret the evidence that Dobbie and Yang provide, in combination with prior work, as suggesting that employers care about a job seeker's conviction record per se above and beyond what it conveys about that job seeker's productivity on the job. If employers were primarily interested in the signaling content of a conviction record, it is not clear why an exogenous increase in pretrial detention rates, and hence conviction rates,

4. Under the stigma view, increasing pretrial detention increases the size of the stigmatized population. Whether we would expect this increase to reduce aggregate employment depends on whether there are sufficient nondiscriminatory employers (or employers serving nondiscriminatory customers) to absorb the increased number of stigmatized job seekers.

would meaningfully decrease aggregate employment. Instead, their findings support the view that employers are primarily concerned with the legal costs associated with a conviction record. This interpretation is subject to the caveat that the causal relationship between pretrial detention rates and aggregate employment may in fact be negligible or nonexistent. Dobbie and Yang make a convincing case that this aggregate relationship warrants further investigation.

REFERENCES FOR THE MILLER COMMENT

Agan, Amanda Y., Jennifer L. Doleac, and Anna Harvey. 2021. "Misdemeanor Prosecution." Working Paper 28600. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w28600.

Agan, Amanda Y., and Sonja B. Starr. 2018. "Ban the Box, Criminal Records, and Statistical Discrimination: A Field Experiment." *Quarterly Journal of Economics* 133, no. 1: 191–235.

Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago: University of Chicago Press.

Bushway, Shawn D., and Nidhi Kalra. 2021. "A Policy Review of Employers' Open Access to Conviction Records." *Annual Review of Criminology* 4:165–89.

Cavico, Frank J., Bahaudin G. Mujtaba, and Stephen C. Muffler. 2014. "Criminal Background Checks in Employment: An Unfolding Legal Quandary for Employers." *Journal of Law and Criminal Justice* 2, no. 1: 41–103.

Dobbie, Will, Jacob Goldin, and Crystal S. Yang. 2018. "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges." *American Economic Review* 108, no. 2: 201–40.

Doleac, Jennifer L., and Benjamin Hansen. 2020. "The Unintended Consequences of 'Ban the Box': Statistical Discrimination and Employment Outcomes When Criminal Histories Are Hidden." *Journal of Labor Economics* 38, no. 2: 321–74.

Gupta, Arpit, Christopher Hansman, and Ethan Frenchman. 2016. "The Heavy Costs of High Bail: Evidence from Judge Randomization." *Journal of Legal Studies* 45, no. 2: 471–505.

Heaton, Paul S., Sandra Gabriel Mayson, and Megan Stevenson. 2017. "The Downstream Consequences of Misdemeanor Pretrial Detention." *Stanford Law Review* 69, no. 3: 711.

Holzer, Harry J., Steven Raphael, and Michael A. Stoll. 2004. "Will Employers Hire Former Offenders? Employer Preferences, Background Checks, and Their Determinants." In *Imprisoning America: The Social Effects of Mass Incarceration*, edited by Mary Pattillo, David Weiman, and Bruce Western. New York: Russell Sage Foundation.

Jacobs, James B. 2015. *The Eternal Criminal Record*. Cambridge, Mass.: Harvard University Press.

Lageson, Sarah Esther, Mike Vuolo, and Christopher Uggen. 2015. "Legal Ambiguity in Managerial Assessments of Criminal Records." *Law and Social Inquiry* 40, no. 1: 175–204.

Mueller-Smith, Michael, and Kevin T. Schnepel. 2021. "Diversion in the Criminal Justice System." *Review of Economic Studies* 88, no. 2: 883–936.

Pager, Devah. 2003. "The Mark of a Criminal Record." *American Journal of Sociology* 108, no. 5: 937–75.

Pager, Devah, and Lincoln Quillian. 2005. "Walking the Talk? What Employers Say Versus What They Do." *American Sociological Review* 70, no. 3: 355–80.

Raphael, Steven. 2021. "The Intended and Unintended Consequences of Ban the Box." *Annual Review of Criminology* 4:191–207.

Shannon, Sarah K. S., Christopher Uggen, Jason Schnittker, Melissa Thompson, Sara Wakefield, and Michael Massoglia. 2017. "The Growth, Scope, and Spatial Distribution of People with Felony Records in the United States, 1948–2010." *Demography* 54, no. 5: 1795–818.

Society for Human Resource Management. 2012. *Background Checking—The Use of Criminal Background Checks in Hiring Decisions*. https://www.shrm.org/hr-today/trends-and-forecasting/research-and-surveys/pages/criminalbackgroundcheck.aspx.

Society for Human Resource Management. 2018. *Workers with Criminal Records: A Survey by the Society for Human Resource Management and the Charles Koch Institute*. https://www.shrm.org/hr-today/trends-and-forecasting/research-and-surveys/pages/second-chances.aspx.

Uggen, Christopher, Mike Vuolo, Sarah Lageson, Ebony Ruhland, and Hilary K. Whitham. 2014. "The Edge of Stigma: An Experimental Audit of the Effects of Low-Level Criminal Records on Employment." *Criminology* 52, no. 4: 627–54.

COMMENT BY

JUSTIN WOLFERS This paper by Will Dobbie and Crystal Yang is best understood as trying to draw policy implications from an important prior study they conducted with Jacob Goldin. That earlier study—Dobbie, Goldin, and Yang (2018)—found that detaining a defendant before their trial causes them to subsequently experience worse labor market outcomes. This conclusion follows from a clever quasi-experimental design which compared the long-run labor market outcomes of those defendants who were randomly assigned to lenient bail judges (who rarely require pretrial detention) with the outcomes of those defendants who were randomly assigned to less lenient judges (who were more likely to require pretrial detention).

The paper by Dobbie, Goldin, and Yang (2018) was a hit in the applied micro literature at least partly because of its methodological sophistication.

It was an early and convincing application of a “judge fixed effects” or “judge leniency” design, and it represents a cutting-edge application of the quasi-experimental methods that are often used in program evaluation. The present paper, in which the authors draw policy implications from that earlier analysis, can be considered a case study of the difficulty in drawing macro policy conclusions from even very well-identified micro econometric studies.

THE HAZARDS IN DRAWING POLICY RECOMMENDATIONS FROM WELL-IDENTIFIED CAUSAL ESTIMATES There are (at least) three sets of concerns that naturally arise when trying to draw policy conclusions from any empirical study. The first is the question of internal validity, which in this case means asking whether this research design yields a reliable estimate of the local average treatment effect for those marginal defendants affected by this natural experiment. This question tends to dominate debate within the academic applied micro research community, and indeed the “causality police” have been called to explore whether a judge fixed effects design really will yield internally valid inference about a local average treatment effect.[1]

Second is the question of external validity, which asks whether the findings from this natural experiment can be generalized to other settings and populations. Applied microeconomists have become much more interested in differences in treatment effects (treatment effect heterogeneity) in recent years. Dobbie and Yang are admirably clear that their earlier analysis in Dobbie, Goldin, and Yang (2018) only estimated the effect of pretrial detention on those “‘marginal’ defendants” (203) who might be released by a lenient judge but not by their stricter colleague. Yet the policy proposals they evaluate—such as the elimination of cash bail—are much more dramatic than these marginal changes, as they would almost eliminate pretrial detention. This is a problem because their prior study evaluated the consequences of allowing pretrial release of those defendants who at least some judges would have recommended be released, but in the present paper they

1. Frandsen, Lefgren, and Leslie (2019) provide a useful discussion of just how subtle the identification assumptions are. In brief, identification is straightforward if a strict judge would detain all those defendants that a lenient judge would, plus a few more. But if judges vary not only in leniency (that is, how many defendants they would detain) but also in who they judge important to detain, then the people who are detained by a strict judge are different from those detained by a lenient judge. (Formally, if the quasi-random judge assignments are instrumental variables, this is a violation of the usual monotonicity assumption.) In this case, the judge fixed effects design will yield estimates of the local average treatment effect that are confounded by differences in the average treatment effect across the groups of people that each judge would detain.

are extrapolating those findings to a population that includes defendants for whom all judges might otherwise have opposed pretrial release. It seems likely that this extrapolation is too optimistic about the effects of policy interventions that would largely eliminate pretrial detention.

Third, for a study to be relevant to policy, it must account for the full range of effects and not just the direct effects on the policy's direct beneficiaries. Dobbie and Yang note that their earlier study only identified the causal effect on an individual of that person receiving pretrial detention—which they call the direct effect—while there may also be unmeasured spillover effects to consider. As Baird and others (2014) note, "the impact of a program only on its beneficiaries becomes an unsatisfying answer to the real policy impact" (1). This problem of unmeasured spillover effects is sometimes also described as a problem of construct validity, as the measurements that were the focus of Dobbie, Goldin, and Yang (2018)—in this case only of the direct beneficiaries of a policy—are not the measurements that a policymaker would seek to rely on. None of this is news to Dobbie and Yang, and the present paper can be read as an attempt to address this set of concerns. The task they set themselves is to supplement their prior analysis with research that takes account of the relevant spillover effects.

But what are the most important spillovers in the present context? Dobbie and Yang highlight a set of socially mediated spillovers, arguing that "pretrial detention is likely to generate spillover effects . . . given its potential impact on families and communities." They argue that these socially mediated spillovers are likely negative, "as the costs of paying money bail and other related court fees and fines often fall on other family and community members of detained individuals." And so it follows that these "harms" amplify the direct effects. Thus, the authors could argue that their prior estimates of the direct effect of pretrial detention provide a lower bound for the aggregate effects. (Much of their text reads as if this is the implicit hypothesis.)

But as I'll argue below, there are also important market-mediated spillovers which likely attenuate the direct effect. Indeed, in standard models of the labor market, these market-mediated spillovers may even completely offset the direct effect. The idea is simply that if pretrial detention doesn't cause a shift in aggregate labor demand, then a job that a former detainee doesn't get is a job that goes to someone else. And so focusing on these market-mediated spillovers might lead one to argue that well-identified estimates of the direct effect are instead an upper bound on the aggregate effects.

When competing conjectures lead one to believe that Dobbie, Goldin, and Yang (2018) is either an upper bound or a lower bound of the aggregate effect of pretrial detention, it becomes clear that the prior study does not speak particularly clearly about the relevant policy issues.

It's worth pausing for a moment on this point, because this problem applies not just to Dobbie, Goldin, and Yang's (2018) careful study, but to literally hundreds of very careful and credible applied micro studies. After all, the vast majority of natural experiment papers focus on a causal estimate on individual people, families, businesses, or some relatively small subpopulation and so fail to account for spillover and general equilibrium effects. The difficulties in the present paper are simply a proxy for the broader question of how to get well-identified quasi-experimental methods to speak to policy questions. The current paper should be read as an attempt to solve this conundrum.

There are two ways forward. One strategy is to take seriously prior measures of the direct effect in Dobbie, Goldin, and Yang (2018) and then supplement them with further analysis that aims to directly measure, model, or otherwise bound these spillovers. This strategy builds on the existing evidence about direct effects and so is particularly appropriate when that prior evidence is of high quality. Chodorow-Reich (2019) is an example of this approach, as he describes conditions under which well-identified prior measures of the effect of state fiscal shocks on local state economies provide a lower bound for a particular national multiplier.

The alternative approach is to measure the aggregate effect of changing rates of pretrial detention, effectively estimating the combined consequences of both the direct and spillover effects. This is the strategy that Dobbie and Yang pursue, even though it means discarding their earlier work. The value of this strategy rests on a judgment about how convincing this new research is, relative to the prior but incomplete evidence. Given the high quality of Dobbie, Goldin, and Yang (2018), it might have been valuable to build on rather than replace that earlier work, even though doing so would have yielded an entirely different paper. The strength of the counterargument to this rests on the present paper exploiting a credible design to uncover convincing and statistically precise causal estimates.

RESEARCH DESIGN STILL MATTERS To make progress in measuring the aggregate effects, Dobbie and Yang (implicitly) assume that spillovers are limited to within a county. As a result, the total effect of a policy can be measured by comparing county-wide outcomes in those areas that are (quasi-)randomly assigned to a new treatment (such as eliminating cash bail) and then comparing them to a control group with no such policy change. So far, so good.

This does not eliminate the need for a credible research design. Indeed, pushing the analysis to a higher level of aggregation simply shifts the research design challenge from the one set in Dobbie, Goldin, and Yang (2018) of finding individuals who were quasi-randomly assigned to pretrial detention to one of finding counties which were quasi-randomly assigned to a new treatment limiting the use of pretrial detention.

Unfortunately, the present paper does not deliver on this score. The central regressions correlate the change in various economic outcomes over the period 2000–2010 across twenty-four large counties, with the corresponding change in county-wide rates of pretrial detention among felony defendants. Some of the specifications include controls, but these control variables are all in levels, while the dependent variable is in changes. As such, there exists an enormous number of potentially confounding social, economic, political, legal, or crime-related variables whose changes are not controlled for, any of which may be responsible for the observed correlation between changes in economic conditions and changes in the probability of a felony arrest leading to pretrial detention.

Program evaluation research that purports to estimate causal effects tends to follow a pretty standard script in which the authors describe their design and why the particular variation that they have isolated might be considered exogenous. But this paper makes no such case.

The independent variable of interest is the change over a decade in the share of felony defendants in a county who are detained before their trial, and the paper offers no explanation about what might be driving this variation. Changes in rates of pretrial detention might be driven by exogenous changes in policy, but the authors provide no evidence of this. It's also possible that the variation in pretrial detention reflects judges responding to an array of broader social, cultural, economic, legal, or political forces. It's also possible that the detention rate might vary even if detention policies don't change, perhaps due to composition effects: the independent variable is the aggregate rate of pretrial detention, which might change if the mix of defendants accused of drug, property, and other crimes changed (and this compositional effect could change the aggregate even if judges didn't change how they treated defendants within any specific category). As such, anything that changes the mix of arrestees might be driving the variation in the independent variable, including changes in criminal opportunities, changes in the alternative labor market opportunities available to potential criminals, changes in the supply of potential victims, changes in policing policy, or changes in demographics.

The point is that there is a long list of factors that might drive variation in pretrial detention rates, and many of these factors likely also shape the outcome variables that Dobbie and Yang study, such as employment, poverty, or intergenerational mobility. To the extent that these factors are not controlled for in the analysis, they are omitted variables that potentially bias the empirical findings.

Moreover, while the paper reports estimates both with and without controls, it's important to note that while both the dependent variables (like the change in employment) and the independent variable of interest (the change in detention rates) are first difference or "change" variables, the control variables (like mean household income, the unemployment rate, the share of the population in various demographic and education groups—all measured at baseline) are included only as levels. Thus the regressions contain no controls for changes in economic, crime, or other factors over time.

The paper offers the usual warning that the "analysis is exploratory in nature," to be interpreted with "an abundance of caution," and that "these county-level specifications should not be interpreted as precise or causal estimates" because it analyzes "county-level changes in detention rates that are likely endogenous," and so "one should be cautious in interpreting β_1 as a causal effect." These are warnings well worth heeding, although the authors appear not to do so, as the analysis then turns to the sort of counterfactual policy analysis that only makes sense if these new estimates are interpreted as a causal effect.

All of this presents a difficult trade-off for policy analysts. Is it better to base policy on well-identified estimates that omit any measure of spillovers or on estimates which do incorporate spillover effects but are likely biased? One might caution that basing policy recommendations on these new estimates rather than on Dobbie, Goldin, and Yang (2018) involves an unpleasant trade-off between some insight into spillovers and a study marred by omitted variable bias. It's easy to imagine that the resulting estimates are even less reliable for policymakers.

Indeed, as I demonstrate below, a careful assessment of the magnitudes suggests that the new estimates of the total effect of pretrial detention are implausibly large, uncomfortably imprecise, and yield a pattern across racial groups that seems quite improbable. This pattern suggests these estimates are influenced more by omitted variable bias than by the spillover effects the paper purports to estimate.

THE MAGNITUDES ARE INCREDIBLE To assess the magnitudes of the new Dobbie and Yang estimates, I will focus on their analysis of the effects of

Figure 1. Employment Rate of 25- to 44-Year-Olds

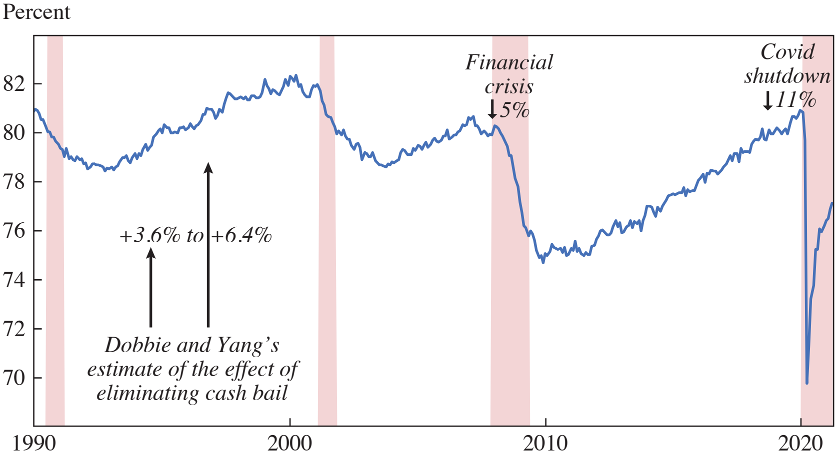

Source: Author's calculations using data from the Current Population Survey and the paper.

pretrial detention on the employment-to-population rate of 25- to 44-year-olds. I narrow my focus partly for brevity and partly because it is easier (at least for a labor economist) to interpret the magnitudes of employment changes. Many of the concerns that I raise apply in equal measure to Dobbie and Yang's analysis of changes in poverty rates and intergenerational mobility.

Table 2 of the paper shows the key result: regressing the change in the employment rate on the change in the detention rate yields a coefficient of –0.206 (without controls) or –0.115 (with controls). To interpret this magnitude, I'll focus on the key policy experiment that Dobbie and Yang analyze—the elimination of cash bail—which they argue would reduce pretrial detention rates from an average of 41.3 percent down to 10 percent. It follows that their regression predicts that this would raise the employment rate by $-0.206 \times (10\% - 41.3\%) = 6.4$ percentage points (and using the alternative coefficient from the regression with controls yields an effect of 3.7 percentage points). To be clear, these are effects on the employment rate in percentage points, and so relative to a typical prime-age employment-to-population ratio of around 80 percent, these estimates imply that employment levels would rise by 4.5 to 8 percent.

These numbers are implausibly large. To give some context, figure 1 shows the aggregate employment rate of 25- to 44-year-olds since 1990, highlighting periods of recession. The figure also superimposes arrows showing the estimated effect of eliminating cash bail, so as to facilitate comparison with the size of business cycle fluctuations. The central estimates suggest that cash bail has a larger effect on aggregate employment than the 2001 "tech wreck" recession, and it is roughly comparable to the 2008 financial crisis, which was (at the time) thought to be a once-in-a-century shock.

While it is difficult to make direct comparisons of the employment effects of an incarceration effect to those from a financial crisis, the more recent pandemic-related shutdown provides a more directly relevant yardstick. The pandemic and associated lockdowns—which effectively forced millions of people to stay at home and shuttered the service sector of the economy—caused a sharp decline in the employment rate in which the prime-age employment rate fell by 11 percentage points between its peak in January 2020 and the deepest part of the trough in April 2020. Dobbie and Yang's estimates suggest that the institution of cash bail has an effect on employment that is of a similar order of magnitude, albeit about half as large. It is hard to believe that pretrial detention, which involves briefer lockdowns of a much smaller fraction of the population, could have employment effects that are of the same order of magnitude as an economy-wide shutdown.

Moreover, these calculations reflect the average result of this policy shift across the whole country. In those areas where pretrial detention rates are higher, eliminating cash bail would yield an even larger decline in detention rates, and so the authors' estimates effectively suggest that some states might enjoy much larger employment gains. For instance, online appendix table A2 suggests that the cross-county standard deviation of detention rates in 1990 was 17 percentage points, which is roughly half the average decline in detention rates if cash bail were eliminated. This implies that a county whose initial detention rate was one standard deviation higher than the average would be forecast to experience roughly a 50 percent larger effect than the average effects outlined above, and a county that was two standard deviations higher would have an effect that is twice as large as the average effect shown in figure 1.

THE IMPRECISION IS UNHELPFUL One response might be to counter that perhaps these counterfactual analyses involve taking the point estimates too seriously. Thus, rather than focusing on the point estimates, it might be worth focusing on the confidence intervals that surround them. Here, the regression without controls (which yielded a coefficient of –0.206) came with an estimated standard error of 0.109. Applying the resulting 95 percent confidence interval to the earlier extrapolation of the effects of eliminating cash bail suggests that it would lead the employment rate to change somewhere between a –0.2 percentage point decline (which is only small relative to the numbers discussed earlier, but still a very large effect given the scale of pretrial incarceration) and an implausibly large 13.1 percentage point rise (in the specification without controls). The corresponding confidence interval from the specification including controls runs from a decline of –0.8 percentage points (which would be a large decline!) to an 8.0 percentage point rise.

Focusing on the confidence intervals yields a range that includes estimates that are no longer obviously implausible, but at the cost of suggesting that these estimates fail to falsify virtually any plausible effect.

The root cause of this statistical imprecision is not surprising: these regression results come from analyzing changes in only twenty-four counties over only a single time period. Moreover, the specification that includes controls adds eleven control variables, leaving very few degrees of freedom. A few power calculations at the beginning of the project might have led to the authors to look for an alternative empirical strategy.

ESTIMATED RACIAL DISPARITIES ARE TOO SMALL Dobbie and Yang probe beyond these aggregate effects and explore the differential effects of pretrial detention on employment (and poverty) rates by race and ethnicity. This analysis involves the same regressions as before, except the dependent variable is no longer the prime-age employment rate but rather the prime-age employment rate for a specific racial or ethnic group.

Importantly, the independent variable is the same in the regression analyzing Black employment as it is in the regression analyzing white employment: it is the county-wide detention rate averaged across all races, rather than a race-specific detention rate. This matters greatly for interpreting the coefficient estimates. If the social and economic processes that lead detention to affect employment are similar for Black people as for white people, then one might expect changes in the race-specific detention rate to have similar effects on Black and white employment. But changes in the aggregate detention rate—which is what Dobbie and Yang analyze—would then be expected to have quite disparate impacts on Black versus white populations, because Black people are dramatically overrepresented among detainees. After all, a policy that largely eliminates pretrial detention would lead a larger share of the Black population to avoid detention.

To get a sense of the relevant magnitudes, in the US population there are roughly five times more non-Hispanic white people than Black people, but among the population of detainees, there are 2.4 times more Black than non-Hispanic white people. Together, these numbers suggest that, on average, a Black person is twelve times more likely to be detained than a white person. Thus, reducing or eliminating pretrial detention might be expected to have a much larger—perhaps twelve times larger—effect on the Black employment rate than the white employment rate.

Yet the coefficients that Dobbie and Yang report in table 2 suggest the effect of changes in the aggregate pretrial detention rate on the Black employment rate is only about one and a half to three times as large as the effect on the white employment rate (depending on which regression specification you prefer). This differential is surprisingly small, perhaps indicating the influence of other omitted variables.

To be a bit more precise, the idea here is simplest when thinking about the direct effect of detention on a defendant. Aggregating up to the level of a whole community, this direct effect would be expected to have an effect on the Black community that is twelve times larger, because Black defendants are twelve times more likely to be affected by a policy change. (Here, I'm following the authors in assuming the effects are linear.) The spillover effects are more complicated, because they might spread from a defendant of one race to a broader community that may be racially mixed. The more racially homogeneous one's community, the more likely it is that the spillover effects would also have a disproportionate effect by race. Consider first the social spillovers that concern Dobbie and Yang. If the Black and white communities never interacted, the social spillovers of pretrial detention would also be twelve times larger within the Black community, because a typical Black member of the community would be twelve times more likely to have a friend affected by changes in detention policy. If there were complete integration, then both Black and white defendants would be equally likely to have friends who were affected by changes in detention policy, and so there would be identical effects in the two communities. For the economic spillovers described below, the total number of jobs in the economy doesn't change and so one group's employment gains are another's losses. Thus, if the direct effect of eliminating pretrial detention were to raise Black employment rates twelve times more than white employment rates, and total employment is unchanged, then (assuming Black and white workers compete for the same jobs) Black and white workers would be roughly equally likely to have their employment prospects shaped by these spillover or equilibrium effects. These economic spillovers would lead the sum of the direct and indirect effects of eliminating or reducing pretrial detention to boost Black employment rates (largely through the direct effect) but decrease white employment rates (as white workers lose their jobs to Black workers).

The point here might be stated more simply as follows: pretrial detention has a radically disparate impact on Black and white communities, yet the results in the paper show only mild disparities. This suggests either problems with the estimates (like omitted variable bias) or that the results reflect a more subtle social process than that outlined by the authors.

A THEORY-INFORMED PRIOR I have argued that the empirical strategy pursued by Dobbie and Yang does not yield much insight into spillover effects, both because of the myriad ways in which the estimates are likely confounded by omitted variables and because of the imprecision of their estimates. While it is easy to harp on the problems with identification, the more important problem is how to constructively provide policy advice based on our limited knowledge. A careful causal study based on plausibly exogenous variation which accounted for spillover effects would be incredibly helpful. But in its absence, economists must still provide useful advice.

I would advise starting from a theory-informed prior about the likely sign and magnitude of the spillover effects, and this is where standard models of the labor market might be helpful. But first, a key fact to bear in mind is that pretrial detention has only a small effect on the number of days a defendant spends behind bars. Indeed, in Dobbie, Goldin, and Yang (2018), based on the judge leniency natural experiment, the authors estimated that an exogenously assigned period of pretrial detention leads a defendant to spend only an extra one to two weeks behind bars (see appendix table A12). Basically detainees are more likely to spend time behind bars before their trial—on average, an extra week or two—but at trial they're often granted sentences equal to time served, and so there is no effect on their post-disposition period of incarceration. With this fact in mind, the next question is what our standard labor market models predict would follow from changes in pretrial detention policy.

Competitive labor markets. The simplest approach might be to consider a competitive labor market. The labor demand curve is given by the marginal revenue product of labor, and the labor supply curve is dictated by the marginal utility of leisure; in equilibrium the wage adjusts to bring these into balance at a point where the quantity of labor demanded is equal to the quantity supplied.

In this framework, the presence or absence of pretrial detention will have no first-order effects on total employment, because both labor demand and labor supply are largely unchanged. The marginal product of workers is basically unaffected by the presence or absence of pretrial detention, and likewise the marginal utility of leisure is unaffected.

That's the first-order prediction with labor demand "basically unaffected" because there are no impacts on the marginal product of the vast majority of the workforce who are not arrested and hence not subject to detention. There is a fraction of detainees who will be affected, but even for them, the period of pretrial detention is so brief that it is unlikely to have noticeable effects on productivity. All told, the effect on labor demand is second order because there is only a small impact on the productivity of only a small fraction of the workforce. Similar reasoning suggests there are only second-order effects on labor supply. While labor supply is affected to the extent that there is an incapacitation effect—it's impossible to work while behind bars—on any given day the total share of the population that is undergoing a period of pretrial detention is tiny (while a significant proportion of the population are arrested at some point, because the period of pretrial detention is so short, few workers are detained on any given day). Moreover, many of those people at risk of pretrial detention are already largely detached from the workforce (according to Dobbie and Yang, "only 32 percent [of detainees] are employed in the year prior to arrest"). As such, a robust majority of the people who might be released were cash bail eliminated would not count as part of the labor supply under any detention regime.

Adding frictions. Of course, evaluating nonemployment in a perfectly competitive framework is somewhat limited given that model has no meaningful role for unemployment. A somewhat richer framework might allow for the sort of labor market frictions that create unemployment. One simple reduced-form approach is to posit that those frictions lead the real wage to get "stuck" above the level that would equate labor supply and labor demand. This simple formulation is a stand-in for a range of frictions, from minimum wage laws or other frictions that directly push the wage up, to union wage pressure or wage bargaining that creates a quasi-labor supply curve above labor supply, to efficiency wage concerns that create a quasi-labor demand curve above labor demand.

This simple framework yields the same stark insight: the presence or absence of pretrial detention has no first-order effects on total employment. Labor demand and labor supply don't shift (or barely shift) for the reasons articulated above. And pretrial detention does not directly shape any of the frictions laid out above (it won't affect the minimum wage, union wage demands, the no-shirking condition, etc.), and so it won't affect the real wage. As such, there's no effect on aggregate employment.

This prediction is still consistent with the key finding of Dobbie, Goldin, and Yang (2018) that pretrial detention reduced the detainee's future employment prospects. The reconciliation of no aggregate employment effect, even with a substantial direct effect on detainees, is relatively straightforward, and it is all about the existence of an offsetting economic spillover. Once frictions create unemployment, they also create a metaphorical queue of potential workers at the factory gate looking for work. The factory owner now has more qualified applicants for each job than they need to hire, and so they can effectively discriminate against former detainees at no cost. After all, each detainee they don't hire can be replaced by hiring an equally talented non-detainee at the exact same wage. The simple point is that if pretrial detention policy doesn't change the number of jobs, then it won't have any effect on total employment, even as it affects who has those jobs.

This simple verbal model is akin to the ranking assumption of Blanchard and Diamond (1994), where employers who receive multiple acceptable applications might arbitrarily hire one set of candidates rather than another. In their setup, employers rank applicants by their unemployment duration; in the present case they might rank applicants by their criminal detention records instead. Blanchard and Diamond (1994) embed this assumption in a general equilibrium search and matching model that both incorporates matching frictions and dispenses with perfect competition in favor of Nash bargaining, and they still find that such "hiring rules do not affect how many are hired; but they determine who is hired, thus affecting the distribution of unemployment, as well as wages" (421).

CONCLUSION Dobbie and Yang have made a convincing case that pretrial detention is an important institution worthy of further study. Their earlier analysis pushed the issue to the forefront, showing quite large negative consequences at the individual level. The present paper seeks to go a step further, incorporating analysis of possible spillover effects.

Their new estimates suggest very large negative spillovers that yield extraordinarily large macroeconomic effects. But I'm not convinced, as I find the identification strategy weak and undefended, the racial pattern at odds with what one might expect, and the estimates imprecise.

The question a policy analyst is left with is whether to rely on these new estimates or use Dobbie, Goldin, and Yang (2018) as a more reliable starting point when analyzing the economic consequences of pretrial detentions. My sense is that the low statistical power in the present study combined with the risk that these new estimates are severely confounded by omitted variables means that point estimates could be quite some distance from the truth. Instead, I would suggest starting with the more credible estimates in Dobbie, Goldin, and Yang (2018), understanding that they need to be supplemented with some insight into likely spillovers. On this score, standard models of the labor market suggest that spillover effects likely operate to somewhat attenuate the direct effects. This perspective suggests that Dobbie, Goldin, and Yang (2018) had identified a plausible upper bound on the total effects of pretrial detention, with zero as the corresponding lower bound. This yields a set of bounds on the total effects of pretrial detention that is smaller, narrower, and more plausible than in the present paper.

REFERENCES FOR THE WOLFERS COMMENT

Baird, Sarah, Aislinn Bohren, Craig McIntosh, and Berk Ozler. 2014. "Designing Experiments to Measure Spillover Effects." Working Paper 6824. Washington: World Bank. https://documents1.worldbank.org/curated/en/266671468089953232/pdf/WPS6824.pdf.

Blanchard, Olivier Jean, and Peter Diamond. 1994. "Ranking, Unemployment Duration, and Wages." *Review of Economic Studies* 61, no. 3: 417–34.

Chodorow-Reich, Gabriel. 2019. "Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?" *American Economic Journal: Policy* 11, no. 2: 1–34.

Dobbie, Will, Jacob Goldin, and Crystal S. Yang. 2018. "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges." *American Economic Review* 108, no. 2: 201–40.

Frandsen, Brigham, Lars Lefgren, and Emily Leslie. 2019. "Judging Judge Fixed Effects." Working Paper 25528. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w25538.

GENERAL DISCUSSION Will Dobbie responded to comments from Justin Wolfers on the data, clarifying that it was case-weighted and incorporated base changes. Dobbie said that, with this considered, Wolfers and the authors had the same numbers for their calculations. Dobbie agreed completely with Wolfers that the changes-on-changes identification strategy should not be interpreted with the same confidence as the micro estimates from Dobbie, Goldin, and Yang.[1]

Dobbie contested Wolfers's comparison of the results to those of Agan and Starr as not perfect because in that study, the characteristics of the people remained constant, while in this paper, the people are able to potentially avoid a scarring activity that changes their characteristics.[2]

1. Will Dobbie, Jacob Goldin, and Crystal S. Yang, "The Effects of Pretrial Detention on Conviction, Future Crime, and Employment: Evidence from Randomly Assigned Judges," *American Economic Review* 108, no. 2 (2018): 201–40.

2. Amanda Agan and Sonja Starr, "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment," *Quarterly Journal of Economics* 133, no. 1(2018): 191–235, https://doi.org/10.1093/qje/qjx028.

Janice Eberly asked about the comparison between eliminating cash bail and banning the box. She further wondered if there might be other policies that would work toward alleviating the effects of pretrial detention.

Dobbie pointed out two ways to aid disadvantaged groups that are most affected by these policies. First, changing the cash bail system is an example of a policy that prevents scarring. Dobbie, Goldin, and Yang find that the impact on safety from lenient judges is minimal, while the economic impacts for the individual are large. Dobbie argued that the elimination of cash bail is an obvious policy to undertake because it makes these individuals more employable with minimal societal downside.

Dobbie said that the other set of policies that would be useful to alleviate the employment effects of a criminal record are policies that attempt to support reintegration into the economy despite scarring. Dobbie cites ban the box as one policy that fits into this category. Dobbie noted that the way that employers use signals like criminal record in the evaluation of employment applications is suboptimal, underscoring the importance of policies that improve the reintegration process. Dobbie also mentioned two other policies suggested by Conrad Miller: increased liability protection for employers and increasing the amount of objective information provided in the hiring process.

Dobbie has other ongoing work showing that workers with and without a criminal record are equally productive. Employers were receptive to being told about their mistake in avoiding workers who were equally productive.

Wolfers responded that he did not think ban the box was a perfect analogy for Dobbie's pretrial policy, but instead wanted to consider the general equilibrium with a straightforward model. Since the time that people spend in jail before going to trial is about two weeks, Wolfers thinks that the effect must be mostly from signaling. Wolfers gave the example of putting the letter *L* on the foreheads of 10 percent of the population and measuring the labor market effects. Clearly, the group that is discriminated against will face negative consequences, but Wolfers said that his strong prior, based on most models of the labor market, is that the overall employment effect must be close to zero if the effect comes from signaling. Wolfers acknowledged that there could be second-order effects that create an overall employment effect and emphasized the challenges in applying micro analysis that is very well identified to understanding what will happen in equilibrium.

Dobbie said that he would hesitate to accept that the employment effects are zero because there was also skepticism about the micro estimates in the work of Dobbie, Goldin, and Yang, showing that even though pretrial detention averages two weeks, labor market effects are still present after four years. Dobbie agreed that this paper does not completely bridge the gap between the micro and macro effects of pretrial detention but also expressed optimism that it will be possible to make credible macro estimates in the future.

Miller commented that comparing pretrial detention policy to ban the box is useful for interpreting how the labor market is responding to criminal history. Miller noted that there is something contradictory about the way that employers respond to a criminal record as shown in ban the box but do not dig into the arrest record, which would allow a deeper understanding of the context for a conviction or lack of conviction.

Steven Davis brought up two issues. First, there is the question of why employers respond so strongly to a conviction. Davis said that Wolfers made an important point on this issue, that employers often cannot discriminate against convicted employees in the form of lower wages, which might make hiring them more desirable. Davis noted minimum wage, collective bargaining, and the threat of lawsuits as impediments to this potential employer incentive for taking a chance on new hires with convictions. Second, Davis agreed with Wolfers's point that the micro effects of detention are likely to be diluted in equilibrium; however, Davis suggested that there are likely to be longer-term effects on future human capital accumulation for workers who are scarred by pretrial detention. Due to this scarring, they will be more likely to have future convictions rather than pursuing human capital gains through education and on-the-job training. This effect can cumulate over time for the individual and persist in equilibrium, because a segment of the population becomes permanently less productive and thus less employable.

Erica Groshen first raised the finding from her Cornell colleagues' work that the information that employers have on criminal records is often inaccurate.[3] Accounting for the deficiencies in these records could make the overall labor market impacts worse. Groshen also said that considering displaced workers is another way to think about the impact.[4] If workers faced with pretrial detention lose their jobs and face scarring, it could have an impact that mirrors that of displaced workers, particularly those

3. Martin Wells, Erin York Cornwell, Linda Barrington, Esta Bigler, Hassan Enayati, and Lars Vilhuber, *Criminal Record Inaccuracies and the Impact of a Record Education Intervention on Employment-Related Outcomes*, Cornell Criminal Records Panel Study 2020-01; https://ecommons.cornell.edu/handle/1813/103780.

4. See, for example, Henry S. Faber, "Employment, Hours, and Earnings Consequences of Job Loss: U.S. Evidence from the Displaced Workers Survey," *Journal of Labor Economics* 35, no. S1 (July 2017): S235–S272.

with low levels of education. Groshen suggested that this would imply larger declines in wages and potentially lead to estimates of employment effects that were more in line with those from the paper than what Wolfers would predict.

Caroline Hoxby recounted Wolfers's description of the simple process by which one worker moves to unemployment and is replaced by a worker who was previously unemployed, commenting that under ban the box legislation, Agan and Starr found that there were increases in discrimination against Black workers. Hoxby noted that considering the racial distribution of employed and unemployed workers could lead to exaggerated spillover effects relative to ignoring the racial composition of those groups.