American Economic Journal: Applied Economics 2022, 14(2): 228–255
https://doi.org/10.1257/app.20190131

Recessions, Mortality, and Migration Bias: Evidence from the Lancashire Cotton Famine†

By Vellore Arthi, Brian Beach, and W. Walker Hanlon*

We examine the health effects of the Lancashire Cotton Famine, a sharp downturn in Britain's cotton textile manufacturing regions that was induced by the US Civil War. Migration was an important response to this downturn, but as we document, migration also introduces a number of empirical challenges, which we overcome by introducing a new methodological approach. Our results indicate that the recession increased mortality among households employed in the cotton textile industry. We also document localized spillover effects on households providing nontradable services in the areas affected by the recession. (*JEL* E32, I12, J63, N13, N33, N63, N93)

We examine the health consequences of the Lancashire "Cotton Famine," a large, temporary, and negative economic shock to the cotton textile manufacturing regions of England and Wales caused by the US Civil War.[1] On the eve of the war, cotton textile production was Britain's most important industrial sector, employing 2.3 percent of the total population and accounting for 9.5 percent of the manufacturing workforce. This sector, however, was entirely reliant on raw cotton imports, and 70 percent of those imports came from the US South. The Civil War disrupted this flow of cotton, generating a sharp and geographically concentrated economic contraction that displaced hundreds of thousands of mill workers.

* Arthi: University of California Irvine and NBER (email: varthi@uci.edu); Beach: Vanderbilt University and NBER (email: brian.beach@vanderbilt.edu); Hanlon: Northwestern University and NBER (email: whanlon@northwestern.edu). Ilyana Kuziemko was coeditor for this article. We thank James Feigenbaum, James Fenske, Joe Ferrie, Marco Gonzalez-Navarro, Tim Hatton, Taylor Jaworski, Amir Jina, Shawn Kantor, Carl Kitchens, Adriana Lleras-Muney, Doug Miller, Grant Miller, Christopher Ruhm, William Strange; audiences at the 2017 ASSA Annual Meeting, 2017 NBER Cohort Studies Meeting, 2017 PAA Annual Meeting, 2017 SDU Workshop on Applied Microeconomics, 2018 All-California Labor Economics Conference, and 2018 NBER DAE Spring Meeting; and seminar participants at Columbia University, Cornell University, University of Essex, Florida State University, University of Michigan, Princeton University, Queen's University, Queen's University Belfast, RAND, University of Toronto, University of California Davis, and University of Warwick for helpful comments. For funding, we thank the UCLA Rosalinde and Arthur Gilbert Program in Real Estate, Finance and Urban Economics, the California Center for Population Research, the UCLA Academic Senate Faculty Research Grant Fund, and the National Science Foundation (CAREER Grant No. 1552692). We are grateful to the UK Data Archive and Campop for providing data used in this project. This study builds on a previous NBER Working Paper (No. 23507), "Estimating the Recession-Mortality Relationship when Migration Matters."

† Go to https://doi.org/10.1257/app.20190131 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] Historians often refer to this event as the "Cotton Famine," where the term "famine" is used metaphorically to describe the dearth of cotton inputs. In this paper, we largely avoid this term since it can be misleading in a study focused on health.

The magnitude of this economic shock, and its importance in British history, has attracted attention from researchers for many years (Arnold 1864; Watts 1866; Ellison 1886; Henderson 1934; Farnie 1979). Recent studies have examined the impact on marriage rates (Southall and Gilbert 1996), poor relief (Boyer 1997; Kiesling 1996), innovation (Hanlon 2015), and long-run city growth (Hanlon 2017). Despite this attention, the human costs of the cotton shock remain debated. Even contemporary reports are contradictory: some observers remarked on the "wan and haggard look about the people," while at the same time local health officers reported a "lessened death-rate throughout nearly the whole of the [cotton] districts."[2]

Qualitative accounts from the time suggest that many displaced workers chose to migrate in search of work elsewhere—and as we show, accounting for this migration is a key challenge to estimating the health effects of this downturn. In the first part of our paper, we corroborate these contemporary reports by providing evidence of substantial and systematic out-migration from the cotton textile districts in response to the cotton shortage.[3]

While migration is a natural response to changes in local economic conditions, the existing literature on recessions and health offers little guidance for how to overcome the empirical challenges introduced by migration.[4] The fundamental issue is as follows. A typical mortality rate calculation normalizes death counts by the area's underlying population. Population counts, however, are generally only well measured in census years (i.e., decennially), whereas death counts are reported more frequently (e.g., annually). Thus, if recessions induce migration, and if these movements are not perfectly captured in intercensal population estimates, unobserved migration can change the size and composition of a location's true at-risk population relative to what is observed, generating a spurious change in mortality rates that we will misinterpret as reflecting the true impact of local shocks on health. A second issue introduced by migration is spillovers: to the extent that individuals migrate toward areas offering better economic opportunities, we are likely to observe migration between treatment and putative control locations, which has the potential to bias coefficient estimates obtained in panel-data regressions.

We adopt an empirical strategy that leverages two features of this setting in order to overcome the identification challenges introduced by migration. The first feature is the plausibly exogenous timing and spatial incidence of the shock, which allows us both to cleanly identify the cohorts exposed to the downturn and to better isolate and correct for spatial spillovers due to migration. The temporal component of the economic shock was short, sharp, and generated by outside forces that were largely

[2] The first quote comes from Dr. Buchanan, *Report on the Sanitary Conditions of the Cotton Towns*, Reports from Commissioners, British Parliamentary Papers, Feb.–July 1863, p. 301. The second quote is from Arnold (1864, 156).

[3] Our most conservative estimates suggest the population of cotton textile–producing regions fell by 2.2 percent during the downturn. As a point of comparison, Fishback, Horrace, and Kantor (2006) report that 11 percent of the US population moved during the Great Depression, with 60 percent of moves occurring within state.

[4] The existing literature tends to assume that migration is not a meaningful threat to inference. Lindo (2015), however, shows that estimates of the recession-mortality relationship differ depending on the level of aggregation in the analysis (e.g., whether we examine county- versus state-level data). Lindo posits that this may be due to migration, but he is not able to rule out other possibilities. The features of our setting allow us to construct estimates of the recession-mortality relationship that differ only based on whether they account for migration. Thus, we are able to explicitly test the extent to which migration can undermine inference.

unexpected. Meanwhile, because the shock was transmitted through the cotton textile industry, its direct effects were concentrated in locations where the firms in that industry clustered, a spatial pattern due to underlying natural endowments. A second and equally important feature of our setting is that it allows us to draw on comprehensive, individually identified, and publicly available census and death records for all of England and Wales. We link these sources to construct a large sample of longitudinal microdata that allows us to follow individuals across time and space. We leverage these features to answer two main questions. First, what impact did this recession have on health, and through what channels? Second, would our estimates of these effects fundamentally differ if we were unable to overcome the bias introduced by migration?

To answer these questions, we begin by defining the cohorts directly at risk of exposure to the downturn: those residing in major cotton textile–producing areas of Britain on the eve of the US Civil War, as enumerated in the 1861 British census.[5] We then link those individuals to deaths occurring during the downturn (1861–1865) regardless of where those deaths occurred.[6] This process produces an individual-level longitudinal dataset that allows us to hold the size and composition of cohorts fixed and thus to accurately identify mortality patterns for the group initially resident in cotton locations, relative to residents of other locations, irrespective of where they may have subsequently migrated and died. Conducting a similar linking exercise for the 1851 census (linked to 1851–1855 deaths) allows us to adopt a difference-in-difference framework to recover a causal effect. Next, we deal with the potential spillovers between migrant-sending and migrant-receiving areas in the following way. First, we provide evidence that during the downturn, large numbers migrated out of the cotton districts and into nearby noncotton districts, mostly within a 25 km radius. Given this spatial concentration, we then separately estimate the mortality effects of the cotton shortage on each of these sets of districts, relative to a third set of more distant control districts that offer a cleaner counterfactual.

Our analysis generates three main sets of findings. First, we show that the cotton shortage had an adverse impact on mortality for the population initially residing in cotton districts at the time of the shock, especially for the elderly. We estimate that the shock generated around 24,000 excess deaths within the cotton textile districts, equal to 9.5 percent of total deaths. Around 10,000 of these occurred among those aged 55 or over, an increase in deaths of 18.8 percent for that age group. This substantial increase in mortality stands in contrast to existing research on modern developed economies. That literature, which we discuss below, consistently finds that health improves during recessions. Our findings indicate that this relationship may be very different in settings with weaker social safety nets and higher baseline mortality.

[5] The 1861 British census was taken just before the onset of the US Civil War. Historical evidence makes it clear that people in both the United States and abroad failed to anticipate the severity of the conflict (one contemporary observer, Arnold (1864), wrote that the bombardment of Fort Sumter "took the world by surprise" (p. 40)), and there is little evidence that the British economy was substantially affected until late 1861 or early 1862.

[6] Our linking approach, which we discuss further in Section IIF, follows seminal papers in this literature (e.g., Ferrie 1996; Abramitzky et al. 2012, 2014; Feigenbaum 2015, 2016; and Bailey et al. 2020).

Second, we provide new evidence examining the impact of the cotton shortage on those households reliant on the industry for employment and those households that did not work in the cotton textile sector but resided in locations where it was the main employer. This is possible given the richness of our longitudinal microdata, which contain detailed information on occupations and family structure. Our direct visibility into the household is novel in this literature, and our results show both that cotton workers, and the family members of cotton workers, experienced substantial mortality increases as a result of the shock. However, we also show substantial effects among noncotton households residing in the cotton textile areas. Thus, in addition to treatment through employment, we observe substantial treatment through location. Digging deeper, we find evidence that the effect of the shock on noncotton households in cotton regions was particularly severe for those providing nontraded local services as well as those working in sectors sharing input-output linkages to the cotton textile sector. This evidence provides a richer view of how a shock to one important industry can ripple through a local economy.

Finally, we document the importance of our empirical approach for overcoming the bias introduced by migration. Our methods allow us to isolate the impact of migration from other factors, enabling us to provide the first direct evidence of the impact of unobserved migration on estimates of the recession-mortality relationship. We find that this impact is substantial in our setting: while our main linked microdata results show that the downturn raised mortality rates, when we intentionally ignore migration, by inferring treatment status based on the location of death—and thus adopting a data structure similar to what is commonly used in the literature—we fail to recover this effect. Indeed, in some cases, we find the opposite result. Thus, addressing migration bias substantially alters the conclusions that we draw, as the recession would have looked much healthier had we not adequately dealt with these issues.[7]

The methodological approach that we apply to deal with the impact of migration has the potential to be useful for studying the relationship between recessions and mortality in other settings where migratory responses are prevalent. Work on modern developed countries suggests that recessions improve health through channels such as increasing exercise, reducing smoking and alcohol use (Ruhm 2000; Ruhm and Black 2002; Ruhm 2005), and freeing up time to care for children and the elderly (Dehejia and Lleras-Muney 2004; Ruhm 2000; Aguiar, Hurst, and Karabarbounis 2013; Stevens et al. 2015). A number of these studies use aggregate-data methods following Ruhm (2000). Our results suggest that, in cases where migration is a meaningful margin of adjustment, it is important to deal with this source of bias in order to accurately measure the recession-mortality relationship. To that end, the techniques we introduce offer a simple and intuitive solution for researchers faced with similar challenges.

Our results also extend our understanding of the relationship between recessions and mortality into a historical setting characterized by high baseline mortality rates, a poor infectious disease environment, limited medical care, and weak social safety

[7] This offers an explanation for the disparate assessments of local versus national contemporaries, the former of whom described a reduction in deaths in the cotton districts and the latter of whom attested to considerable suffering among out-of-work cotton operatives and their families.

nets. While there is a large literature on the relationship between business cycles and health, most of the evidence on how temporary income fluctuations affect health across the age distribution comes from analysis of developed countries. Much less evidence is available from low-income settings (Miller and Urdinola (2010) being a notable example), and only a few studies (Fishback, Haines, and Kantor 2007; Stuckler et al. 2012) examine the impact of recessions on mortality in historical contexts such as the one we consider.[8] Augmenting this existing evidence is useful because it can help us begin to map out how and why the recession-mortality relationship varies within and across settings. In addition, our ability to harness extremely rich data, and to deal with potential migration bias concerns, enables us to push our results beyond what has been possible in these previous studies—by, for example, separating out "occupation" and "location" effects.

I. Empirical Setting

A. *The Timing and Incidence of the Cotton Shortage*

The cotton textile industry was the largest and most important industrial sector of the British economy during the nineteenth century. For historical reasons, British cotton textile production was geographically concentrated in the northwest counties of Lancashire and Cheshire, which held over 80 percent of the cotton textile workers in England and Wales in 1861.[9] This concentration, which dates back to at least 1830, is thought to be driven by the location of rivers, which were used for power; access to the port of Liverpool; and a history of textile innovation in the eighteenth century (Crafts and Wolf 2014). Figure 1 depicts this spatial distribution by plotting the share of employment accounted for by the cotton textile industry in each district using data from the Census of Population of 1851.

Because Britain did not produce cotton, the success of its cotton textile industry was dependent on reliable access to imported raw cotton—and in the run-up to the US Civil War, 70 percent of these inputs came from the US South (Mitchell 1988). The war prompted a sudden and dramatic rise in world cotton prices, sharply reducing British imports of US cotton and causing a sharp drop in British cotton textile production. These effects are depicted in Figure 2. During the US Civil War period, other cotton-producing countries such as India, Egypt, and Brazil rapidly increased their output, and British inventors produced new technologies to make use of these new sources of supply (Hanlon 2015). Nevertheless, these increases were not large enough to offset the lost US supplies, although they did contribute to the rapid rebound in imports after 1865.[10]

[8] There is, of course, a related historical literature on longer-term income fluctuations and mortality (i.e., the Malthusian Trap). Our paper differs substantially from this literature in that we are focused on economic fluctuations occuring over short time scales, while that literature focuses on changes over long periods.

[9] Calculation based on data collected by the authors from the 1861 Census of Population reports.

[10] Consistent with this, alternative proxies for industry output (firms' raw cotton consumption and variable operating costs (excluding cotton)) exhibit a similar pattern. See Hanlon (2015) and Mitchell and Deane (1962) on cotton consumption and Forwood (1870) for wage and cost data.

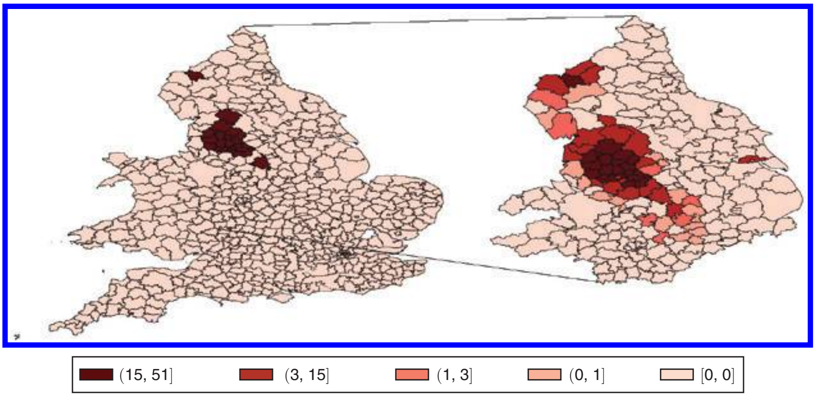

Figure 1. Spatial Distribution of Cotton Textile Industry

Notes: Data on the geography of the cotton textile industry are calculated from the 1851 Census of Population. Shaded in the map of England and Wales are districts with over 10 percent of employment in cotton, while the inset shows the percent of employment in cotton textiles in the core cotton region, with darker colors indicating a greater share of employment in cotton.

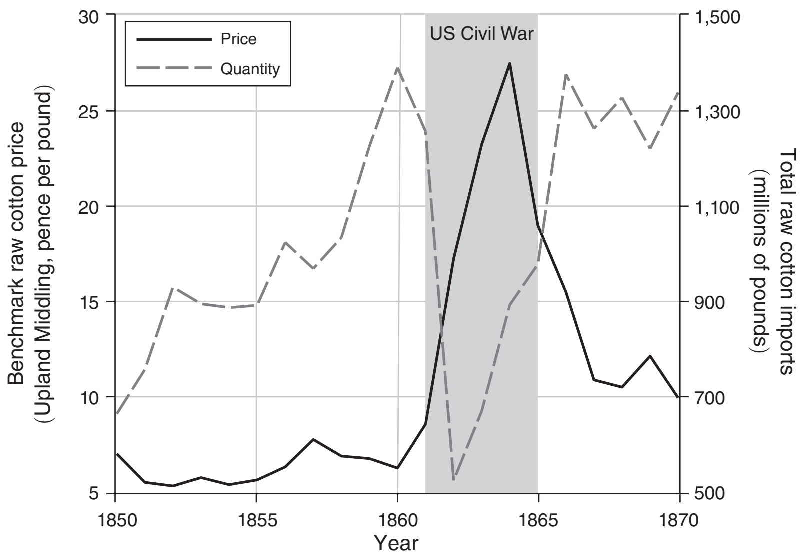

Figure 2. Cotton Prices and Imports

Notes: Import data from Mitchell (1988). Price data, from Mitchell and Deane (1962), are for the benchmark Upland Middling variety.

The direct effects of the US Civil War were largely confined to the cotton textile sector and the districts where it was located, and there is little evidence of a broader reallocation of economic activity. One indicator of this is that there was little effect on imports or exports other than those associated with textiles (see online Appendix A.2). Another factor was that the cotton textile industry had very weak input-output connections (Thomas 1984; Horrell, Humphries, and Weale 1994). Almost all inputs were imported, with the exception of machinery (which

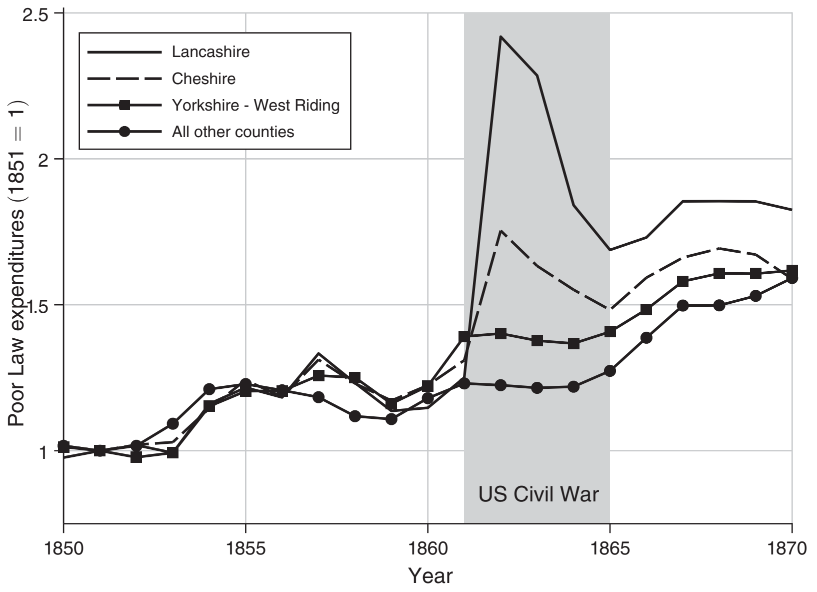

Figure 3. The Spatial Incidence of the Cotton Shock

Note: Data are collected by the authors from the annual reports of the Poor Law Board.

was produced in the cotton textile districts) and coal. Downstream, some output was sold to clothing-producing firms, though much was exported or sold directly to households. As a consequence, the cotton shortage did not lead to a larger nationwide recession (Henderson 1934, 20).

Figure 3 offers additional support for this conclusion. This graph describes the expenditures by local Poor Law boards in the main cotton textile counties (Lancashire and Cheshire) across the study period. For comparison, we also present data for nearby Yorkshire County, which was not heavily dependent on cotton textile production, as well as for the remainder of the country. During the downturn, we see an increase in Poor Law expenditures in the cotton textile areas, while the remainder of the country was largely unaffected. Online Appendix A.1 shows that similar patterns are observed if we focus on the number of able-bodied relief seekers rather than Poor Law expenditures.

B. *Responses to the Cotton Shortage*

During the downturn, workers in the affected areas adopted a variety of coping mechanisms. Reports indicate that at the height of the recession (winter 1862), roughly 500,000 persons in cotton-producing regions depended on public relief funds, with over 270,000 of these supported by the local Poor Law boards and an additional 230,000 reliant on private charities (Arnold 1864, 296).[11] This relief, however, differs

[11] Additional relief programs included public works employment for unemployed cotton workers, though most public works employment began in 1863, after the worst of the crisis had passed. See Arnold (1864) for a discussion of public works.

sharply from the social safety nets of today. Poor Law funds were associated with pauperism and only provided for the barest level of subsistence. They also required "labour tests" such as rock-breaking, which workers found demeaning. Indeed, there is evidence that workers tried to avoid drawing on this stigmatized source of support (Kiesling 1996; Boyer 1997). Instead, displaced workers tended to respond by reducing consumption and dipping into any available savings. Once their savings were depleted, workers pawned or sold items of value, including furniture, household goods, clothing, and bedding (Watts 1866, 214; Arnold 1864). Many eventually turned to poor relief, but others migrated in search of work elsewhere.[12]

One way to examine migration patterns is to study the evolution of population, population growth rates, and net migration rates across decades using census data. These patterns are given in Figure 4. The top-left panel describes the evolution of log population in cotton districts, nearby districts, and all other districts from 1851–1881. The top-right panel describes the growth in district population across each decade, normalized by the 1851–1861 change (the decade preceding the downturn).[13] The bottom panel describes implied net in-migration rates over the same period.[14] This figure reveals three important patterns. First, it shows a substantial slowdown in population growth in the cotton textile districts in the decade spanning the cotton shortage.[15] This change appears to be driven by both increased out-migration and decreased in-migration (a conclusion supported by the bottom panel as well as additional evidence in online Appendix A.3). Second, we observe an acceleration in population growth in nearby districts, which we define here as noncotton districts within 25 km of a cotton district. Meanwhile, there is little change in the population growth trend in districts beyond 25 km. These patterns are consistent with short-distance migration from cotton textile districts during the downturn. Third, these changes essentially disappear after 1871, highlighting the temporary nature of the shock.

These implied migration flows were meaningfully large. In terms of magnitude, had the population of the cotton districts grown from 1861–1871 at the same rate that it grew in 1851–1861, these districts would have had 54,000 additional residents in 1871, a figure equal to 2.2 percent of the districts' 1861 population. Similarly, if nearby districts had grown in 1861–1871 at the rate they grew during 1851–1861, they would have had 61,000 fewer residents, which is equal to 4 percent of the

[12] Watts (1866), for example, describes how, "The trade of Yorkshire has received such an impetus during the famine … many thousands of operatives have only crossed Blackstone Edge [which divides Yorkshire from Lancashire] (p. 226–27)". Arnold (1864) described how "thousands had passed to east and south" (p. 470).

[13] As noted in footnote 5, the 1861 census was taken in April, the same month that the US Civil War began, and so it should be thought of as a clean prewar population observation.

[14] Implied net migration is calculated as the difference between the observed population count in a district in a given census year and the population that we would have expected in that district-year given the population in the previous census plus all births and less all deaths in the intervening years. We then divide by initial population to create rates. This conceptual approach has been used in studies of migration such as Fishback, Horrace, and Kantor (2006) and Bandiera, Rasul, and Viarengo (2013).

[15] Note that overall population growth in the cotton areas remains positive from 1861–1871, as it does in all other locations. This growth reflects the very high rate of fertility in all locations, which continued until the fertility transition that began at the end of the 1870s (Beach and Hanlon 2020). Given this strong underlying forcing factor, the impact of migration is visible mainly in the changes in the growth rates shown in the top right panel.

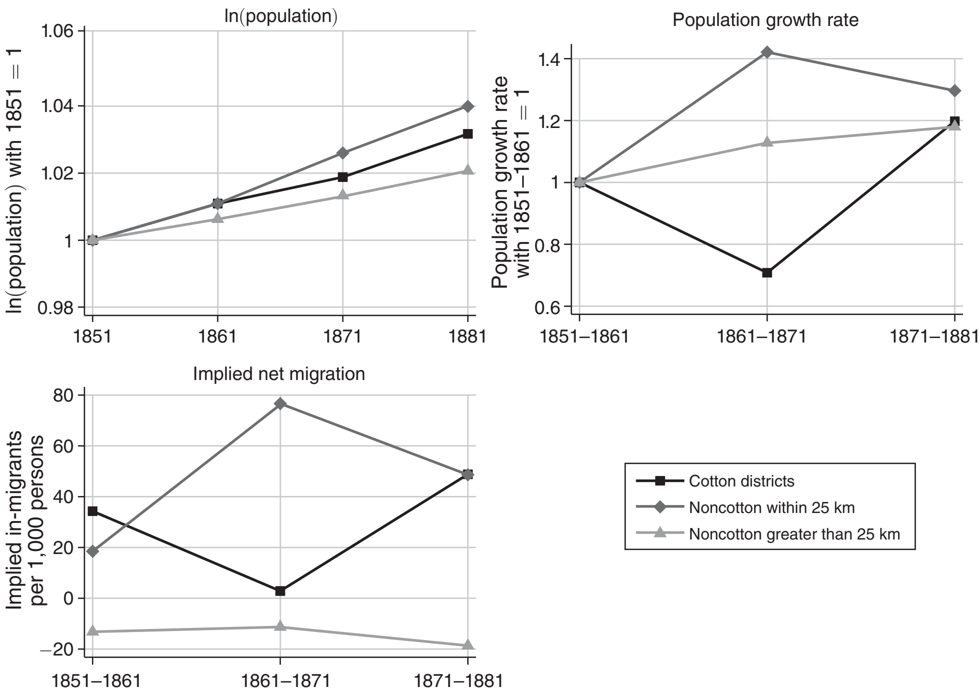

Figure 4. Migration Response to the Cotton Shortage

Notes: This graph describes population dynamics for all cotton districts, all noncotton districts within 25 km of a cotton district, and all remaining noncotton districts. Cotton districts are defined as those districts with more than 10 percent of employment in cotton textile production in 1851. The population growth rate for each group of districts is normalized to one in 1851–1861. Implied net migration is given as the difference between the terminal census estimate of population and the postcensal (i.e., initial census population less intervening deaths, plus intervening births) estimate of population, all divided by initial population (such that positive values represent in-migrants).

Sources: Data for the top two panels are from the Census of Population. Data for the bottom panel are from the Census of Population and the Registrar General's reports of annual vital statistics.

districts' 1861 population. Note that these figures will understate the migration response if some migrants returned between 1865 and 1871.[16]

There is also some evidence that migration away from the cotton textile districts during the US Civil War was selective. Online Appendix Figure 8 shows that young adults were somewhat more likely to migrate. However, the change in population in the 20–39 age group accounts for only about three-fifths of the overall change in population of the cotton districts between 1861 and 1871. Thus, a substantial amount of migration likely occurred among other segments of the population as well.

Migratory responses of the sort documented here have two important implications for our analysis. First, there are good reasons to expect that this migration impacted health in very real ways. For instance, the cotton textile districts were the least intrinsically healthy locations in Britain at this time because they were

[16] These patterns are consistent with the city-level experiences documented in Hanlon (2017).

highly industrialized, densely populated, and heavily polluted.[17] Thus, those leaving the cotton districts likely enjoyed some protective effects of migration that will work against the results that we find here, causing the recession we study to appear healthier in our results than it actually was.[18] Second, migration poses a number of empirical challenges, largely related to the mismeasurement of population size and composition. In the mortality analysis that follows, we discuss both these substantive and methodological concerns related to migration and develop an approach to estimation wherein the spurious health effects related to unobserved migration can be stripped from the real health effects of the downturn.

II. Mortality Analysis

How did health respond to this temporary local shock? Contemporary reports suggest a number of channels through which the cotton shortage affected health.[19] Some local registrars—the officials responsible for compiling death records—described a reduction in deaths in the cotton districts. One such official attributed this to "more freedom to breathe the fresh air, inability to indulge in spirituous liquors, and better nursing of children."[20] Notably, these are some of the same channels modern studies cite as an explanation for the procyclical mortality relationship they find.[21] However, other reports indicate that the inability to afford food, clothing, and shelter negatively affected health, particularly for the elderly. The effect of reduced income is illustrated by the reappearance of typhus—a disease spread by lice and strongly associated with poverty—in Manchester in 1862, after many years of absence. These conflicting reports highlight the fact that the net effect of the cotton shortage on mortality is ambiguous ex ante.

A. *Methodological Issues Introduced by Migration*

One thing contemporary reporters cannot tell us, however, is whether mortality among those initially resident in cotton districts increased during the US Civil War. This is because local registrars had visibility only into the health of individuals currently living (and dying) in their district. Given the substantial migration response we have documented, the fact that these officials were unable to track individuals over time and space poses a problem to us as well. To see why, consider the following estimating equation:

$$MR_{dt} = \beta SHOCK_{dt} + X_{dt}\Gamma + \phi_d + \eta_t + \epsilon_{dt}, \tag{1}$$

[17] The crude mortality rate in cotton districts was over 26 deaths per thousand, compared to 25.7 in nearby districts and 23.2 nationwide.

[18] This differs from the US Great Migration, where Black migrants moved toward rather than away from more urban, industrialized, and polluted locations (Black et al. 2015).

[19] See online Appendix A.4 for details.

[20] Quoted from the *Report of the Registrar General*, 1862.

[21] See Dehejia and Lleras-Muney (2004) and Ruhm (2000); see Aguiar, Hurst, and Karabarbounis (2013) on freeing up time for breastfeeding, childcare, exercise, and other salutary activities; see Stevens et al. (2015) on raising the quality of eldercare; and see Ruhm and Black (2002) and Ruhm (2005) on limiting the capacity for unhealthy behaviors such as smoking and alcohol use.

where MR_{dt} is the mortality rate in a given location (i.e., district) d, η_t and ϕ_d are a full set of time period and location fixed effects, $SHOCK_{dt}$ is an indicator equal to one if district d is a cotton district and time t is the shock period (1861–1865), and X_{dt} is a set of district-level controls. This equation closely follows the existing literature examining the impact of business cycles on health within a panel framework.[22]

While this equation is a natural starting point, migration may affect estimates obtained from equation (1) in two key ways. First, migration may cause the dependent variable, MR_{dt}, to be systematically mismeasured. Second, migration-induced spillovers may affect results through the comparison, implicit in equation (1), between treated and control locations. Below, we discuss each of these potential channels for bias and how they are addressed in our analysis.

On the first point, migration may affect estimates obtained from equation (1) through mismeasurement of the true at-risk population. Migration changes both the size of the population, which appears in the denominator used to calculate the mortality rate, as well as the composition of the population, which determines the population's average mortality risk, in ways that are unobservable to the researcher. If some migration is unobserved, the population denominator used to calculate the mortality rate will be incorrect.[23] Further, even if overall population flows are perfectly observed, migration may still be selective, which will cause the underlying mortality risk faced by the population in a given location to be different from what is observed.

Linked individual-level longitudinal data offer a solution to these issues. By fixing individuals to their location at the onset of the shock, their deaths can be correctly attributed to their experience of the shock whose effects we are trying to estimate, irrespective of where these deaths ultimately occur. Thus, this approach ensures that the population represented in the denominator of the mortality rate (i.e., the population at risk) corresponds to the group of people whose deaths appear in the numerator.[24] Accordingly, we modify our specification of interest to

$$\left(\frac{MORT_{dt}}{POP_{dt}}\right) = \beta SHOCK_{dt} + X_{dt}\Gamma + \phi_d + \eta_t + \epsilon_{dt}, \tag{2}$$

[22] Within that literature, this estimating equation is most similar to Miller and Urdinola (2010), who use coffee price shocks and spatial variation in coffee cultivation as an exogenous shock to local economic conditions in Colombia. As in that paper, we do not use $SHOCK_{dt}$ as an instrument for unemployment because suitable unemployment data do not exist. In our setting, the best proxy available to us is the number of Poor Law relief seekers, but it is not consistently available for the entire study period. Another reason we prefer this explanatory variable to annual unemployment rate fluctuations is that it presents a more plausibly exogenous shock to local economic conditions (particularly in the presence of migration), one that enables us to cleanly identify and track the specific group of individuals exposed to the downturn whose effects we wish to estimate.

[23] If people migrate to locations offering better economic conditions, and if migration is not fully captured by intercensal population estimates, then unobserved out-migration will lead to an artificially high population denominator and fewer observed deaths in the numerator because of a smaller at-risk population. Conversely, unobserved in-migration will lead to an artificially low population denominator but more observed deaths because the true at-risk population has increased. Thus, the unobserved relocation of individuals from one region to another can mechanically generate the false appearance of health change where there has been none.

[24] In other words, this approach holds fixed the size and composition of the population at risk.

where POP_{dt} is the population in a district d at the beginning of period t (in our empirical setting, at the 1851 or 1861 census) and $MORT_{dt}$ is the number of deaths among that population during the period (i.e., from 1851–1855 or 1861–1865).

It is worth noting that migration may have very real effects on mortality. For example, migration may affect mortality in both migrant-sending and migrant-receiving areas through congestion effects (e.g., disease contagion, strain on fixed local resources, or labor market competition). Alternatively, the act of migration can change underlying population health by, say, depleting the migrant's health stocks or by relocating people across locations with different intrinsic conditions. If, for example, people move to healthier locations, then migration will have a real and beneficial impact on health. While estimates obtained from a linked-data approach will purge the spurious impact of migration on observed mortality patterns, they will capture—alongside the direct effects of the recession on mortality—any *real* effects of recession-induced migration on mortality.

On the second point, migration can affect results obtained from equations (1) and (2) by generating spillovers from treated to control locations, thus violating the assumptions necessary for causal inference in a difference-in-difference approach. This issue can be addressed if migrant-sending and migrant-receiving locations can be identified and compared to a third set of locations that were not contaminated by spillovers.[25] To operationalize this intuition, we modify our specification to separately estimate the impact of the shock on migrant-receiving districts,

$$(3) \quad \left(\frac{MORT_{dt}}{POP_{dt}}\right) = \beta SHOCK_{dt} + \gamma RECEIVING_{dt} + X_{dt}\Gamma + \phi_d + \eta_t + \epsilon_{dt},$$

where $RECEIVING_{dt}$ is an indicator equal to one for districts receiving migrants from the treated districts during the treatment period. In the case of the cotton shortage, most migration occurred to nearby locations. Thus, in our setting, $RECEIVING_{dt}$ will simply be an indicator variable (or variables) identifying districts within a specified radius of cotton districts. With these modifications in hand, we now turn our attention to constructing our linked dataset.

B. *Constructing Our Linked Sample*

To estimate the relationship between recessions and mortality in the presence of unobserved migration, we require individual-level longitudinal data that identify both an individual's place of residence at the beginning of the recession and whether that individual died within the specified recession period thereafter. Our linked sample relies on two main data sources that allow us to recover this information. The

[25] An alternative approach is to aggregate to higher geographic levels. For instance, one could combine migrant-sending and migrant-receiving areas and, in essence, treat the two areas as a single unit. This type of aggregation ignores the fact that the various local labor markets within the aggregated study area are likely experiencing dramatically different economic conditions, which may undermine the researcher's ability to recover the causal effect of economic conditions on mortality.

first is individually identified death records for the entire population of England and Wales over the years 1851–1855 (our control period) and 1861–1865 (the recession period). The second is the full-count British census for the years 1851 and 1861. Because census enumeration took place in April of 1861, just as the US Civil War began and before it had any meaningful effect on the British economy, this means that we can identify deaths in the cohort of individuals actually exposed to the cotton shortage.

We obtain census microdata from the UK Data Archive.[26] In addition to preserving the structure of the household, these data include individual names, location at the time of enumeration, age, and some additional information.[27] Our deaths data come from the records of the General Registrar's Office (GRO),[28] which we have collected for the years 1851–1855 and 1861–1865. These data include information on the decedent's first and last name, age, and location of death. Further details on the deaths data and how they were obtained can be found in online Appendix B.1.

We construct our longitudinal dataset by linking the census and death records. A valid link is defined as one where the first name and last name are an exact match between the GRO data and the census and where the inferred birth year is no more than five years apart.[29] We allow for a five-year threshold, which is standard in the linking literature (see, e.g., Abramitzky et al. 2021), because neither data source explicitly asks about birth year. The census asks for the individual's age at the time of enumeration, while the death index reports age at time of death. Because the assigned birth year depends on when these events occur relative to the individual's birth month, it is natural to expect some disagreement. We allow the threshold to span five years to account for any other misreporting of age (e.g., age heaping). This strategy yields a final sample of 150,792 deaths (or about 7.1 percent of all deaths) for the 1851–1855 period and 126,509 deaths (or 5.8 percent) for the 1861–1865 period.[30]

While our linking strategy attempts to follow the best practices in the literature (e.g., Ferrie 1996; Abramitzky, Boustan, and Eriksson 2012, 2014), there are some important differences with respect to linking in our setting. First, we are linking people over relatively short periods of time, never more than five years. This means that name changes, such as those due to marriage, are less common. As a result, women are well-represented in our linked sample. A second advantage is that the name information provided in the British census is likely more accurate than contemporaneous US census records. One reason for this is that there were few recent foreign migrants in Britain, who may have changed their names as they assimilated. A second reason is that the British procedure for collecting the census differed in that households filled out their own census forms rather than verbally providing

[26] Schurer and Higgs (2020a, b).

[27] Individual-level microdata are not available from the next closest censuses, in 1871 or 1841.

[28] General Registrar's Office (2019).

[29] A consequence of this five-year threshold is that the first name, last name, and inferred birth year must be unique within a five-year window.

[30] There are a number of potential reasons why our linking rate differs across these two periods, including differences in the care that went into collecting the census, differences in the way the data were transcribed, changes in name uniqueness over time, emigration, etc.

their information to an enumerator.[31] Because of this, we refrain from using name cleaning algorithms like Soundex. The main disadvantage relative to the existing literature is that we are not able to leverage birthplace information, as that is not reported in the death index. For this reason, the next section summarizes results from a battery of empirical tests to illustrate the reliability of our linked sample.

As a check on the impact of the linking procedure on our results, in online Appendix D.3 we present a second set of findings. These results are obtained from a different set of underlying death records, linked to census data using a different procedure. The results obtained from this alternative sample are similar to those obtained from our preferred data, which we view as strong evidence that the specific nature of our linking procedure is not driving our results.

C. *Assessing Our Linked Sample*

One way to check whether our linked sample is reasonable is to see how the probability of finding a link declines as the distance between the death location and enumeration location rises. This analysis, presented in online Appendix B.2.1, shows that deaths are much more likely to be matched to individuals previously enumerated in the same district and that the chance of observing a link falls off rapidly and fairly smoothly as the distance between the death district and the census enumeration district increases. These results are consistent with the intuition that migration tends to decline with distance and suggest that our linking procedure is performing well. As a point of comparison, we can also link between the full 1851 and 1861 censuses. The distribution of distances between district of enumeration in 1851 and that in 1861 looks nearly identical to what we see when in our sample of linked deaths. Finally, we can plot the distribution of distances when we randomly link census records within the 1851 census. There we see a very different "hump-shaped" pattern, suggesting that the previous results are not simply a mechanical artifact of the linking procedure.

A second way to assess the quality of our links is to run a falsification test. We classify every individual in our linked dataset as having died in the five years following enumeration. If we are correct, then if we were to look for these individuals in the subsequent census (i.e., five to ten years after we say they died), we should not find any of them. Unfortunately, the 1871 microdata are not yet digitized, so we can only run this test for those we classify as having died between 1851 and 1855. Of the 150,792 individuals we classify as dead, 17.27 percent (or 26,041) link to a record in 1861 (same first name, same last name, and age within a 5-year threshold). However, the advantage of this exercise is that we can also leverage birthplace information, and it turns out that if we require the birthplace to also match, then only 8.67 percent of our "dead" sample can be linked to the 1861 census. These will be upper bound estimates if families recycled names following the deaths of relatives within a five-year period. This suggests that an upper bound on our false positive

[31] Enumerators still visited every household to check and collect the forms and assisted households in the completion of the form when necessary.

rate is between 8.67 and 17.27 percent.[32] Overall, our false positive rate is on the lower end of what is obtained in other linked papers (see Bailey et al. 2020 and Abramitzky et al. 2021). Note that in our difference-in-difference framework, these false positives likely work against us by pushing our mortality coefficients toward zero.

In addition to linking accuracy, it is also important to know whether the mortality patterns in our linked sample are representative. One way to test this is to generate results assigning our linked deaths to the district in which they occur and then to compare these to results obtained from comprehensive data on all deaths in England and Wales—data, taken from the Registrar General's reports,[33] in which deaths are reported in aggregate form by district of occurrence (henceforth, "aggregate data"). We present these results later, in Table 5. These results show that we are able to recover estimates that are both practically and statistically equivalent to those from aggregate data. The fact that we can reproduce the results obtained with comprehensive aggregate data when we structure our linked sample to mimic the structure of these aggregate data suggests that our linked sample is reasonably representative overall.

The main dimension on which our linked sample of deaths differs from aggregate mortality is in the age distribution. In our linked sample, young children, and particularly infants, are underrepresented. This is a mechanical consequence of our procedure since an infant death in, say, 1865 can never be linked to someone alive in the 1861 census. We take two approaches to dealing with this issue. One approach is simply to analyze different age groups separately. Alternatively, when estimating effects across all age groups, we reweight our linked sample so that its age distribution is representative of that in the corresponding aggregate deaths data.

In online Appendix B, we check our linked sample against aggregate deaths data on the dimensions of socioeconomic status and sex. Drawing on the occupation data listed in the census, we see the share of deaths among white- versus blue-collar workers in the linked sample are very similar to those generated from aggregate mortality data. Thus, for the working population, our linked sample appears to be quite representative in terms of socioeconomic status. In terms of gender, women are slightly overrepresented in our linked sample, where they account for 51.4 percent of deaths in the 1851 period versus 49.2 percent in the aggregate data and 50.5 percent of deaths in the 1861 period versus 48.8 percent in the aggregate data. This is most likely due to women's names being more unique than those of men (Rossi 1965).

[32] Imposing various other criteria provides us with information on how to lower our false positive rate. The first new criterion that we impose is that a "linkable" record should have a unique first name, a unique-*sounding* last name (as determined by NYSIIS codes), and a unique age (within five years). This criterion lowers our false positive rate range to between 7.72 and 13.16 percent (where the lower bound estimate requires that the potential links have the same birthplace). If we further require that a "linkable" record be one with a unique-*sounding* first and last name, the range falls to 7.50 to 12.38 percent. As a second set of criteria, we could require that the distance between place of enumeration in 1851 and place of death be within a certain threshold. If we set that threshold to 200 km, then our false positive rate range is between 8.19 and 16.35 percent. When the threshold is 100 km, the range becomes 7.57 to 15.23 percent. With a 50 km threshold, the false positive rate is between 6.88 and 14.15 percent. Finally, if we impose that there be no migration, the false positive rate is between 5.31 to 11.80 percent. These criteria will form the basis of various robustness checks.

[33] Registrar General of England and Wales (1851–1871)

D. *Estimation Strategy*

Building upon the empirical framework introduced in Section IIA, our estimating equation of interest is the following difference-in-difference specification:

$$(4) \qquad \left(\frac{MORT_{dt}}{POP_{dt}}\right) = \beta COTDIST_d \times POST_t + \sum_{i \in \{25,50,75\}} \gamma_i NEAR_d^i \times POST_t + X_{dt}\Gamma + \phi_d + \eta_t + \epsilon_{dt}.$$

The variable POP_{dt} is the population in a district d at the time of enumeration (i.e., 1851 or 1861), and $MORT_{dt}$ is the number of deaths among that group of people during the period of interest (i.e., from 1851–1855 or 1861–1865).[34] The variable $COTDIST_d$ is an indicator for whether district d is a cotton district, and $POST_t$ is an indicator for the 1861–1865 period.[35] The variables $NEAR_d^{25}$, $NEAR_d^{50}$, and $NEAR_d^{75}$ are indicator variables equal to 1 if district d is within 0–25 km, 25–50 km, or 50–75 km from a cotton district. The inclusion of these variables is informed by the spatial concentration in migration that we documented in Section IB. The vector X_{dt} is a set of additional district-level controls.

This equation deals with both migration-induced mismeasurement of mortality rates and spillovers between migrant-sending and -receiving areas, such that β reflects the impact of the cotton shortage on the mortality rate of the treated population, regardless of where they died. However, one challenge with estimating equation (4) is that because our data do not include unique individual identifiers (e.g., a social security number), we are not able to link every death back to a census record. To see how this affects our analysis, let $MORT_{dt}$ be the number of deaths of individuals initially resident in district d, and let λ be the share of these deaths that we are able to match back to census records. What we can observe in our linked data is $\widetilde{MORT}_{dt} = MORT_{dt}\lambda$. Substituting out $MORT_{dt}$ in equation (4) and reorganizing, we have

$$(5) \qquad \left(\frac{\widetilde{MORT}_{dt}}{POP_{dt}}\right) = \tilde{\beta} COTDIST_d \times POST_t + \sum_{i \in \{25,50,75\}} \tilde{\gamma}_i NEAR_d^i \times POST_t + X_{dt}\tilde{\Gamma} + \phi_d + \eta_t + \epsilon_{dt},$$

where $\tilde{\beta} = \beta\lambda$ and $\tilde{\gamma}_i = \gamma_i\lambda$. This shows that we can obtain β estimates by multiplying the $\tilde{\beta}$ coefficients (and standard errors) obtained from our linked data by the linking rate λ. To ease the interpretation of our results, we make this adjustment in all of our main analysis tables.

[34] While our approach collapses microdata to the district-of-origin level, thus creating district-of-origin cohort mortality rates, an alternative approach is to run logit or probit regressions at the individual level.

[35] Cotton textile districts are defined as those with greater than 10 percent of employment in cotton textiles in 1851, a decade before the US Civil War, although in robustness exercises we also consider continuous measures of cotton employment. The location of industry was relatively persistent, and so results are similar when using the spatial distribution of industry in 1861.

One identification assumption in our analysis is that the probability of linking should not be correlated with the treatment. Our analysis approach will not be biased due to variation in linking rates across locations that are fixed over time nor by changes in linking rates over time that are common across locations. However, we may worry that there were time-varying changes in the probability of linking. The most plausible violation of this assumption is that migration generated by the shock may have made it more difficult to link cotton district residents observed in the 1861 census to deaths over the period 1861–1865, say, because they moved abroad. However, if individuals who emigrated are less likely to be linked (i.e., because they left Britain altogether) and emigration increased from the cotton districts during the shock, then this will bias the estimated effect of the shock downward since it will cause the number of linked deaths among those initially resident in the cotton districts to *understate* the true number of deaths. Thus, if anything, this form of bias will work against the countercyclical results that we find.

There are a few other points worth mentioning about our empirical specification. First, in the main text, we report standard errors clustered by district. To address the possibility of spatial correlation, our main results also report *p*-values from a permutation test that provides an alternative assessment of statistical significance while respecting the spatial structure of our data. For a detailed description of our permutation test as well as a discussion of why we prefer this to alternative methods such as clustering or spatial standard errors, see online Appendix C.1. Second, when looking at all-age mortality results, we control for the share of different age groups in each district, which naturally influence total mortality. We also include initial district population as a control because the period we study saw substantial improvements in sanitary technology, which were most important in larger cities with high population density. In addition, we control for what we call the "linkability rate," which is given by the number of individuals enumerated in the census in a district that can be uniquely identified by their first name, last name, and age (within a five-year window), divided by the total number of individuals enumerated in the census in that district. We calculate this rate for each district-by-period cell. Third, we follow the conventions of existing literature and weight all regressions by population, although as we show in our robustness checks, weighting does not affect the results.

As a final point, it is worth noting that, while our linked sample allows us to identify deaths among both migrants and stayers, it is not possible to separately assess the mortality rates for these two groups, and so, to comment on the causal impact of migration on health. This is because we are not able to observe the population of migrants; we only observe migrants in the linked sample conditional on their death.

E. *Summary Statistics*

Table 1 presents summary statistics for key analysis variables. The first two rows highlight Britain's high mortality during this period. As a point of comparison, the population-weighted mean death rate across the two periods, as calculated from the aggregate Registrar data, is 23.2 deaths per 1,000 persons, and even in the aggregate data we see a few districts with mortality rates of 40 deaths per 1,000 persons or higher. Relative to the rates observed in aggregate data, our data exhibit more

Table 1—Summary Statistics: Linked Data Assigned to District of Enumeration

	Mean	Standard deviation	Min	Max	Observations
MR in 1851	27.801	8.513	0.513	46.517	538
MR in 1861	21.379	6.394	0.150	42.619	538
Cotton dist. ind.	0.045	0.207	0	1	538
Nearby (0–25 km) ind.	0.048	0.215	0	1	538
Cotton employment share	0.017	0.071	0	0.513	538
Under 15 pop. share 1861	0.36	0.021	0.201	0.416	538
Age 15–54 pop. share 1861	0.52	0.028	0.468	0.705	538
Over 54 pop. share 1861	0.12	0.021	0.069	0.183	538

variation across districts due to the fact that they are based on a (linked) sample of deaths rather than all deaths. In addition, Table 1 shows that cotton districts make up 4.5 percent of the sample, while an additional 4.8 percent of the sample are nearby districts within 25 km of the cotton districts. In terms of age, the population was fairly young, with 36 percent under age 15 compared to 12 percent over 54.

F. *Main Results*

Before turning to our main regression results, we examine patterns in the raw data to help fix ideas about the results that follow. First, we see evidence that residents of cotton textile districts faced a greater mortality risk during the downturn. Over the 1851–1855 period, 6.2 percent of our linked deaths originated from cotton districts. During the 1861–1865 period, however, 7 percent of our linked sample originated from cotton districts. Second, we see evidence consistent with an increase in migration during the downturn. Among the linked deaths from the 1851–1855 period, 73.2 percent of individuals who were enumerated in a cotton district also died in a cotton district. In the 1861–1865 period, however, this figure falls to 67.4 percent.

Next we examine these patterns within our formal regression framework. Table 2 presents our main findings. Note that these coefficients have been adjusted for the linking rate, such that they can be interpreted as the change in the mortality rate in the population per 1,000 persons per year.[36]

Column 1 presents our simplest specification, while the results in column 2 include controls for district population density, the population shares of individuals in different age groups, and a control for whether the district had more people with "linkable" (more unique) names. While these controls do predict the change in mortality rates, particularly the age controls, they do not have a substantial impact on the cotton district coefficient. In column 3, we address the possibility that the effects of the shock may have spilled over into nearby districts by separately estimating the impact on noncotton districts in various distance bands around the cotton districts. These show some marginally statistically significant evidence of adverse spillover effects in the population initially residing in the nearest set of districts.

[36] When adjusting by the linking rate, we use the average linking rate across the full data sample.

Table 2—Baseline Effects of the Shortage Using Linked Data

	DV: Deaths per 1,000 individuals (per year)			
	(1)	(2)	(3)	(4)
Cotton District × Cotton Shortage	2.194	2.024	2.534	
	(0.463)	(0.519)	(0.605)	
Cotton Emp. Share × Cotton Shortage				6.766
				(1.654)
Nearby (0–25 km) × Cotton Shortage			1.054	0.785
			(0.597)	(0.576)
Nearby (25–50 km) × Cotton Shortage			0.191	0.081
			(0.623)	(0.628)
Nearby (50–75 km) × Cotton Shortage			0.586	0.509
			(0.656)	(0.645)
District controls		Yes	Yes	Yes
Observations	1,076	1,076	1,076	1,076
R^2	0.022	0.395	0.398	0.397
Permutation test p-values for effect on cotton districts				
p-values	0.119	0.089	0.050	0.002

Notes: Underlying sample includes 277,057 linked deaths. Standard errors in parentheses are clustered at the district level. Deaths are assigned to the district of initial residence (i.e., district of census enumeration). Regressions are weighted by district population. All regressions include period fixed effects and district fixed effects. District controls include ln(population density), share of individuals enumerated in the census with a "linkable" name, the share of the population in each of the following age categories (under 15, 15–54, and over 54, with 15–54 as the omitted category), and region-by-period fixed effects.

Consistent with this, including these controls results in an increase in the cotton district coefficient. Finally, column 4 replaces the cotton district indicator variable with a continuous measure of treatment: the share of district employment in cotton textiles. Here we observe similar results, which are stronger in terms of statistical significance.[37] Note that, in this specification, the spillover effect in nearby districts appears weaker, a result that may reflect the presence of a few cotton textile workers in those areas.

At the bottom of the table, we report p-values from a permutation test of the key explanatory variable. Our permutation test, described in detail in online Appendix C.1, provides an alternative approach to constructing confidence intervals that involves iterating across a large set of potential placebo cotton districts, estimating results, and then comparing estimates based on the true cotton districts to the distribution of placebo coefficients. Because our placebo sets of cotton districts incorporate the bunched spatial pattern observed in the true cotton districts, these placebo results can help address potential concerns about spatial correlation.

Table 3 breaks these results down by age group. The clearest pattern here is the mortality increase experienced by older residents of the cotton districts. We observe substantial effects for adults over 25, which become statistically significant starting at age 45. In contrast, for the two younger groups, the estimated effect is very

[37] In terms of magnitude, a one SD increase in the explanatory variable in column 3 leads to an increase in deaths of 0.525 per thousand, while a one SD increase in the explanatory variable in column 4 leads to an increase in deaths of 0.477 per thousand.

TABLE 3—DECOMPOSING THE CHANGE IN MORTALITY BY AGE

	DV: Deaths per 1,000 individuals (per year)						
Age group:	Under 15 (1)	15–24 (2)	25–34 (3)	35–44 (4)	45–54 (5)	55–64 (6)	Over 64 (7)
Cotton District × Shortage	0.224	0.171	0.894	1.512	3.066	6.740	13.477
	(1.078)	(0.551)	(0.678)	(0.939)	(1.086)	(1.861)	(3.899)
Nearby (0–25 km) × Shortage	0.777	0.355	0.897	0.988	1.231	4.028	4.418
	(1.149)	(0.471)	(0.693)	(1.077)	(1.129)	(2.109)	(3.333)
Nearby (25–50 km) × Shortage	−0.811	0.387	1.005	0.649	0.377	2.045	2.805
	(1.342)	(0.443)	(0.655)	(0.910)	(1.038)	(1.619)	(2.813)
Nearby (50–75 km) × Shortage	0.241	0.317	1.214	1.848	1.481	2.199	−0.371
	(1.211)	(0.561)	(0.678)	(1.069)	(1.207)	(1.616)	(3.259)
Observations	1,076	1,076	1,076	1,076	1,076	1,076	1,076
R^2	0.130	0.084	0.087	0.089	0.122	0.180	0.317
Linked deaths	75,795	19,390	20,784	22,180	25,095	32,489	80,545
Permutation test p-values for effect on cotton districts							
p-values	0.476	0.422	0.206	0.227	0.089	0.003	0.028

Notes: Standard errors, clustered at the district level, are in parentheses. Deaths are assigned to the district of initial residence (i.e., district of census enumeration). Regressions are weighted by district population. All regressions include district fixed effects, period fixed effects, and controls for ln(population density), the share of individuals (within each age category × district of enumeration × census cell) that have a "linkable" name, and region-by-period fixed effects.

close to zero. Recall that we should be cautious in interpreting the coefficient for the under-15 age group since our linking procedure will mechanically miss many deaths among infants and young children, a group that contributed a large share of the deaths in this category.

The by-age pattern of effects is consistent with contemporary reports describing the health effects of the shock (see online Appendix A.4). For example, assessments from local health officials at the time suggest that the health of young children during this period improved when working mothers in cotton textiles—a heavily female industry—lost their jobs and were able to spend more time on breastfeeding, household hygiene, and childcare. For young children, this likely offset the adverse effects of material deprivation, an explanation in line with similar recent findings in modern Colombia (Miller and Urdinola 2010). For further discussion of infant health, see online Appendix Table 14, where we examine results on births and infant mortality using aggregate data.

To put these magnitudes in context, our preferred all-age results in column 3 of Table 2 imply that the cotton shock generated 24,418 excess deaths in the cotton textile districts from 1861–1865, equal to 9.5 percent of total deaths in the cotton districts over this period.[38] Using the age group regressions, we estimate roughly 10,191 deaths among those aged 55 and over (an 18.8 percent increase in deaths in that age group), 7,402 among adults aged 15–54 (an 11.2 percent increase), and

[38] To provide an alternative benchmark, we can think about the excess deaths over the period 1861–1865 among the cohort initially residing in cotton textile districts as being equivalent to roughly twice the number of deaths from diarrhea in these districts over the preceding 5-year period—or, to compare to some of the other leading causes of disease of the time, 86 percent of the deaths from tuberculosis, 66 percent of the deaths from other respiratory causes, or 209 percent of the deaths from scarlet fever in these districts over 1856–1860.

2,595 among those under 15 (1.9 percent increase). Thus, the shock appears to have substantially elevated mortality over the period 1861–1865 among the population initially residing in cotton districts. That the adverse mortality effects were strongest among older adults is consistent both with contemporary reports of a rise in respiratory ailments, to which the elderly are especially vulnerable, and with the temporary accentuation of seasonal patterns in mortality found during the 1861–1865 period.[39]

Online Appendix D.1 presents results from several robustness exercises. These show that results are similar regardless of whether we weight our regressions by district population; drop outlier locations such as Manchester, Liverpool, or Leeds; or omit the foreign-born from our linked deaths sample (see online Appendix Table 8). We also consider more restrictive linking approaches and samples with fewer false positive links (online Appendix Table 9). That analysis, too, produces similar results.

While the evidence laid out in Section IIC establishes that our linked sample is largely representative of the universe of deaths over the period in question, we conduct an additional validation exercise to show that our results are not influenced by features of our linked data. Specifically, in online Appendix D.3, we present results using an entirely different linked database based on both alternative death data (taken from the freeBMD website)[40] and an alternative linking method (where links are based entirely on unique first and last name combinations). Results obtained using this linked dataset show patterns similar to those documented in our main results. We view the fact that this alternative linked sample yields similar results as strong evidence that our linking procedure is not driving the results that we document.

Because this alternative linked dataset spans additional years not available to us via the GRO, it also allows us to generate a number of additional results—among them, results on the phenomenon of harvesting, which speaks to the overall mortality costs of the recession (see online Appendix D.3). Here, however, we focus on one particular advantage of this alternative linked dataset: namely, the fact that it also covers the years from 1855–1860.[41] This allows us to generate placebo results wherein we treat 1856–1860 as a placebo shock period and then estimate mortality effects in the cotton districts over that period as compared to the 1851–1855 period. These findings, reported in online Appendix Table 12, show no evidence of increased mortality in the cotton districts during the placebo period. This indicates that our results are not due to differences in underlying mortality trends in the cotton textile districts prior to the US Civil War.

G. *Results by Sector of Employment*

Next, we examine the extent to which our results are driven by the experience of those directly reliant on income from the cotton textile industry (effects through “employment”), versus the extent to which they are driven by broader local distress, which might have affected families who resided in cotton textile districts, even

[39] These results also fit with some existing studies, such as Stevens et al. (2015), that show that recession-induced changes in the mortality risk of older adults are responsible for much of the effect of business cycles on total mortality.

[40] FreeBMD (2018)

[41] We were unable to collect data on additional years via the GRO due to restrictions on the use of their site. See online Appendix B.1 for further details.

Table 4—Decomposing Effects by Sector of Employment

	DV: Deaths per 1,000 individuals (per year)			
	(1)	(2)	(3)	(4)
Cotton District × Cotton Shortage	1.443		1.292	
	(0.714)		(0.732)	
Head Employed in Cotton × Cotton Shortage		1.416	0.860	
		(0.510)	(0.542)	
Head Employed in Cotton × Cotton Dist. × Shortage				2.476
				(0.704)
Head Employed in Non-Tradeables × Cotton Dist. × Shortage				3.032
				(0.866)
Head Employed in Linked IO × Cotton Dist. × Shortage				1.713
				(0.981)
Head Employed in Other × Cotton Dist. × Shortage				0.595
				(0.733)
Head Employed in Tradeable Manuf. × Cotton Dist. × Shortage				0.764
				(0.945)
Head Employed in Transport × Cotton Dist. × Shortage				0.320
				(0.671)
Head Outside Labor Force × Cotton Dist. × Shortage				–0.936
				(1.444)
Observations	32,677	32,677	32,677	32,677
R^2	0.023	0.023	0.023	0.024

Notes: Two-way clustered standard errors (by district and occupation) are reported in parentheses. All regressions include period and district fixed effects; a series of indicators for whether the district is 0–25, 25–50, or 50–75 km from a cotton textile district in the post period; region-by-period fixed effects; ln(population density); the share of the population that is under 15; and the share of the population that is over 54. Deaths are assigned to the initial district of residence (i.e., census enumeration district). Regressions are weighted by district × industry population.

though they themselves were dependent on other sectors (effects through "location"). To examine these issues, we take advantage of the fact that our census data include rich occupation information, and we classify families based on the sector that the household head worked in. These occupations, which are closer to industry identifiers than what we would call occupations today, allow us to study how mortality varied by sector of employment.

As a starting point, consider the share of total deaths accounted for by cotton textile workers and their households in our linked data. Deaths of cotton textile workers (i.e., those reporting cotton textile work as their occupation) accounted for 1.20 percent of all linked deaths among employed workers in our sample from 1851–1855, but this rose to 1.58 percent in 1861–1865, a 32 percent increase. Similarly, if we focus on entire households, members of cotton households accounted for 1.14 percent of deaths in 1851–1855 but 1.42 percent from 1861–1865, a 25 percent increase. These figures provide preliminary evidence that mortality among cotton workers may have been increasing during the cotton shock, though we acknowledge that in these raw data the increase could be due in part to an increase in the share of cotton workers in the population.

To provide more rigorous evidence on the incidence of the shock across sectors and locations, we organize our data into occupation-by-district bins and run regressions looking at how the shock affected mortality for households with heads linked to specific sectors (e.g., cotton textiles) as well as those located in the cotton districts, while including a full set of occupation and district fixed effects. While we focus on the household head occupation in these results, in online Appendix D.2, we present additional evidence looking only at the employed population and using a worker's own occupation.

The results are presented in Table 4. We begin, in column 1, with a specification that is comparable to the results in Table 2, except that our data are now organized into occupation-by-location bins and we now control for occupation fixed effects. The coefficient estimate shows that we still find evidence of elevated mortality in the cotton districts when using this specification, though the results are somewhat attenuated. In column 2, we instead look at how mortality changed among households with a head employed in the cotton textile sector (regardless of location) during the shock period. Here we see strong evidence that mortality increased among cotton households. In column 3, we study both the impact of being in a cotton textile district and the impact of being a cotton textile household. These estimates suggest that most of the increase in mortality is explained by being in a location experiencing the shock, though there is some (not statistically significant) evidence that even in those locations, cotton textile households were worse off than others.

In the last column, we then consider effects across all types of households in the cotton textile districts, separated into the type of occupation held by the household head. Since these occupation groups span all of the households in the cotton textile districts with observed occupations, we do not include a separate cotton district effect in the regression. The results show that within the cotton textile districts, we see strong evidence of an increase in mortality among cotton textile households as well as equally strong effects among households that produced nontraded services. There is also somewhat weaker evidence of increased mortality among households in industries linked to cotton textile production through input-output connections (e.g., textile finishing and clothing). We do not see statistically significant effects among households working in industries that mainly produced other traded goods, transportation, or those households with a head outside of the labor force.

We break these results down further in online Appendix D.2, where we present two additional sets of results. First, we show that the mortality increase in households dependent on the cotton textile sector during the shock was larger than the change observed among almost every other occupation group. Second, we break estimates of effects within the cotton textile districts down into more disaggregated occupation groups.

These results shed new light on how a shock to one important local industry can ripple across other sectors of a local economy. To our knowledge, this is the first study within this literature that is able to separate these "employment" and "location" effects.[42]

[42] Note, however, that some papers in this literature do touch on similar themes. For instance, Sullivan and von Wachter (2009) focuses on an individual's experience of job loss (a pure "employment" channel), while the impact on traffic fatalities found in Ruhm (2000) implies more general "location" effects. Though perhaps not perfectly analogous, Stevens et al. (2015), too, identifies spillovers through the labor market—a "location"-channel result—showing that because recessions raise the quality of eldercare staff, the health of elderly people rises.

TABLE 5—AGGREGATE-TO-LINKED COMPARISON: DOES MIGRATION HAVE A MEANINGFUL IMPACT ON RESULTS?

	DV: Annual deaths per 1,000 individuals (per year)						
	Under 15 (1)	Age 15–24 (2)	Age 25–34 (3)	Age 35–44 (4)	Age 45–54 (5)	Age 55–64 (6)	Over 64 (7)
Panel A. Actual aggregate data (Drawn from Registrar's Reports)							
Cotton District	−4.660	−1.325	−1.511	−0.499	−0.624	0.267	2.094
× Cotton Shortage	(1.019)	(0.284)	(0.346)	(0.496)	(0.632)	(0.793)	(1.807)
Panel B. Aggregate-like linked data (Links assigned to district of death)							
Cotton District	−4.248	−1.361	−0.952	0.052	−2.035	2.104	4.989
× Cotton Shortage	(3.035)	(0.654)	(0.975)	(1.399)	(1.918)	(3.285)	(6.991)
Different from panel A? (*p*-value)	0.889	0.961	0.596	0.702	0.485	0.587	0.674
Panel C. Preferred, migration-corrected linked data (Links assigned to district of enumeration)							
Cotton District	0.224	0.171	0.894	1.512	3.066	6.740	13.477
× Cotton Shortage	(1.078)	(0.551)	(0.678)	(0.939)	(1.086)	(1.861)	(3.899)
Different from panel A? (*p*-value)	0.000	0.018	0.002	0.049	0.003	0.002	0.009

Notes: Standard errors, clustered at the district level, are in parentheses. Regressions are weighted by district population. All regressions include controls for ln(population density), proximity to cotton (0–25 km, 25–50 km, and 50–75 km), and the share of the population in each period-by-age-by-place-of-enumeration cell with a linkable name. Regressions also include district fixed effects, period fixed effects, and region-by-period fixed effects.

III. Does Migration Matter?

Finally, we ask, How important is it that we were able to adjust for migration in this setting? That is, we examine whether intentionally failing to account for recession-induced migratory responses fundamentally alters our conclusions as to the health impact of this historical downturn. By doing so, we provide the first direct evidence of the impact of unobserved migration on estimates of the recession-mortality relationship. Our main approach follows the methodology applied thus far to the linked microdata, comparing total deaths in 1861–1865 to deaths in 1851–1855, but instead uses data taken directly from aggregate district death counts transcribed from the Registrar General's reports.[43] This allows an apples-to-apples comparison between results obtained using the linked data and those generated from the more commonly available aggregate data.

As a first step, we compare results obtained from these aggregated reports to the results obtained from analogously organized linked data—that is, linked data wherein deaths are assigned to the location of death rather than to the residence at the time of enumeration. Because the aggregate data are assigned to the location of death, we should expect these results to be similar. Panel A of Table 5 reports results obtained using aggregate data. Panel B reports results using "aggregate-like" data, i.e., our linked data in which deaths are assigned based on the location of death. As in our main results, to allow comparability, the linked coefficient estimates have

[43] These data cover the same districts used in the linked data analysis and in fact are the same data used to test the representativeness of the linked microdata (for more on these data, see Section IIC as well as online Appendix B.3). As in the linked analysis, population data come from the census and are available every ten years starting in 1851.

been inflated by the linking rate, so that these estimates reflect β rather than $\tilde{\beta}$. The similarity between the estimates in panel A and panel B is striking. Moreover, these estimates are statistically indistinguishable, as given by the *p*-values at the bottom of panel B. This tells us that our linked sample can recover the results obtained from aggregate data; i.e., the linked deaths appear to be representative of aggregate deaths. This provides a powerful check confirming the quality of our linked data.

Next, consider the results obtained using our preferred linked data, where deaths are assigned to each person's district of residence at the time of enumeration rather than their district of death. These estimates are in panel C. Note that the only difference between the results in panel B and those in panel C is whether deaths are assigned to location of death or initial location of residence. Thus, a comparison between the results in panels B and C provides a direct test of the impact of migration on our estimates. Clearly, migration has a substantial impact; the results that account for migration, in panel C, indicate a much more severe mortality effect than the results suggested by panel B.

Finally, having illustrated the ability of the linked data to mimic aggregate data, we compare the results obtained from aggregate data (panel A) to our preferred linked approach (panel C), which accounts for migration bias by assigning deaths to each person's district of residence at the onset of the shock. This comparison provides a direct test of the extent to which an analysis based on aggregate data can recover the true impact of the cotton shortage on the relevant population at risk of exposure. Our results show that the equality of these coefficient estimates can be strongly rejected, suggesting that, in the presence of a substantial migration response, an analysis based on aggregate data would fail to accurately recover the impact of the cotton shock on mortality—and in fact, would lead to meaningfully inaccurate results. That we cannot reject equality when our linked deaths intentionally ignore migration but can reject equality when we do account for migration is particularly telling. This means that the difference is due to migration rather than to differences between the linked sample and the aggregate data. The direction of this bias is also noteworthy: the analysis based on aggregate data makes the cotton shock appear much healthier than we know it to be from our migration-corrected linked data—in many cases, even flipping the sign of the estimated relationship.

IV. Conclusion

We examine the mortality consequences of the Lancashire Cotton Famine, a recession in Britain's cotton textile–producing regions that was precipitated by the US Civil War. In addition to its intrinsic historical interest as one of the defining crises of industrializing Britain, two features of this setting are of particular significance to the study of the recession-mortality relationship. First, ours is a setting with limited social safety nets and high baseline mortality. Accordingly, evidence on the mortality impact of this recession helps deepen our understanding of the interplay between economic conditions and health in low-income settings, particularly across the age distribution. Second, and perhaps related to the limited safety nets of the time, the cotton shortage was a recession that generated a systematic migratory response. While migration is a natural means of coping with an income shock, it also poses threats to inference, which have largely been ignored by the existing literature on recessions and health. In this paper, we offer an empirical strategy that overcomes these issues and allows us to recover clean causal estimates of the mortality impact of the cotton downturn even in the presence of migration.

Our results are twofold. First, we find robust evidence that the cotton shortage increased mortality among those initially resident in cotton districts, both across the age distribution and particularly for the elderly. Thus, one conclusion to draw from our results is that recessions can cause a substantial increase in mortality in a poor setting. In addition, we show that economic shocks hitting one industry can have substantial local impacts beyond just the families reliant on that industry for employment. In the setting we consider, these local spillover effects appear to be concentrated among providers of nontraded local services.

Second, we show that an analysis that does not explicitly deal with migration would have led us to conclude that this recession improved health, when in reality, health deteriorated in response to the downturn. This is true not only when we analyze aggregate data using the empirical methodology standard in this literature, but it is also true when we reorganize our linked microdata to intentionally ignore migration by defining exposure based on the district of death. These results illustrate that large migratory responses can pose a meaningful threat to inference and suggest that it is important to account for such responses in order to accurately recover the impact of a localized shock on mortality. We present an approach that can help deal with the bias arising from a recession-induced migratory response. Given that migration remains a key margin of adjustment to local shocks in many settings, these methods may be useful for the broader literature studying the relationship between recessions and mortality.

REFERENCES

Abramitzky, Ran, Leah Platt Boustan, and Katherine Eriksson. 2012. "Europe's Tired, Poor, Huddled Masses: Self-Selection and Economic Outcomes in the Age of Mass Migration." *American Economic Review* 102 (5): 1832–56.

Abramitzky, Ran, Leah Platt Boustan, and Katherine Eriksson. 2014. "A Nation of Immigrants: Assimilation and Economic Outcomes in the Age of Mass Migration." *Journal of Political Economy* 122 (3): 467–506.

Abramitzky, Ran, Leah Platt Boustan, Katherine Eriksson, James Feigenbaum, and Santiago Pérez. 2021. "Automated Linking of Historical Data." *Journal of Economic Literature* 59 (3): 865–918.

Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis. 2013. "Time Use during the Great Recession." *American Economic Review* 103 (5): 1664–96.

Arnold, R. Arthur. 1864. *The History of the Cotton Famine: From the Fall of Sumter to the Passing of the Public Works Act*. London, UK: Saunders, Otley, and Co.

Arthi, Vellore, Brian Beach, and W. Walker Hanlon. 2022. "Replication data for: Recessions, Mortality, and Migration Bias: Evidence from the Lancashire Cotton Famine." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E128521V1.

Bailey, Martha J., Connor Cole, Morgan Henderson, and Catherine Massey. 2020. "How Well Do Automated Linking Methods Perform? Lessons from U.S. Historical Data." *Journal of Economic Literature* 58 (4): 997–1044.

Bandiera, Oriana, Imran Rasul, and Martina Viarengo. 2013. "The Making of Modern America: Migratory Flows in the Age of Mass Migration." *Journal of Development Economics* 102: 23–47.

Beach, Brian, and W. Walker Hanlon. 2020. "Culture and the Historical Fertility Transition." Unpublished.

Black, Dan A., Seth G. Sanders, Evan J. Taylor, and Lowell J. Taylor. 2015. "The Impact of the Great Migration on Mortality of African Americans: Evidence from the Deep South." *American Economic Review* 105 (2): 477–503.

Boyer, George R. 1997. "Poor Relief, Informal Assistance, and Short Time during the Lancashire Cotton Famine." *Explorations in Economic History* 34 (1): 56–76.

Conley, T.G. 1999. "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92 (1): 1–45.

Crafts, Nicholas, and Nikolaus Wolf. 2014. "The Location of the UK Cotton Textiles Industry in 1838: A Quantitative Analysis." *Journal of Economic History* 74 (4): 1103–39.

Dehejia, Rajeev, and Adriana Lleras-Muney. 2004. "Booms, Busts, and Babies' Health." *Quarterly Journal of Economics* 119 (3): 1091–1130.

Donald, Stephen G., and Kevin Lang. 2007. "Inference with Difference-in-Differences and Other Panel Data." *Review of Economics and Statistics* 89 (2): 221–33.

Ellison, Thomas. 1886. *The Cotton Trade of Great Britain: Including a History of the Liverpool Cotton Market and of the Liverpool Cotton Brokers' Association*. London, UK: Effingham Wilson, Royal Exchange.

Farnie, D.A. 1979. *The English Cotton Industry and the World Market 1815–1896*. Oxford, UK: Clarendon Press.

Feigenbaum, J. 2015. "Intergenerational Mobility during the Great Depression." Unpublished.

Feigenbaum, J. 2016. "A Machine Learning Approach to Census Record Linking." Unpublished.

Ferrie, Joseph P. 1996. "A New Sample of American Males Linked From the 1850 Public Use Micro Sample to the Manuscript Schedules of the 1860 Federal Census of Population." *Historical Methods: A Journal of Quantitative and Interdisciplinary History* 29 (4): 141–56.

Fishback, Price V., William C. Horrace, and Shawn Kantor. 2006. "The Impact of New Deal Expenditures on Mobility during the Great Depression." *Explorations in Economic History* 43 (2): 179–222.

Fishback, Price V., Michael R. Haines, and Shawn Kantor. 2007. "Births, Deaths, and New Deal Relief during the Great Depression." *Review of Economics and Statistics* 89 (1): 1–14.

Forwood, William B. 1870. "The Influence of Price upon the Cultivation and Consumption of Cotton During the Ten Years 1860–70." *Journal of the Statistical Society of London* 33 (3): 366–83.

Free UK Genealogy CIO. 2018. "FreeBMD." FreeBMD. https://www.freebmd.org.uk/ (accessed April 1, 2018).

General Registrar's Office. 2019. "General Registrar's Office Death Index." https://www.gro.gov.uk/ (accessed January 7, 2019).

Hanlon, W. Walker. 2015. "Necessity is the Mother of Invention: Input Supplies and Directed Technical Change." *Econometrica* 83 (1): 67–100.

Hanlon, W. Walker. 2017. "Temporary Shocks and Persistent Effects in the Urban System: Evidence from British Cities after the U.S. Civil War." *Review of Economics and Statistics* 99 (1): 67–79.

Henderson, W.O. 1934. *The Lancashire Cotton Famine 1861–1865*. Manchester, UK: Manchester University Press.

Horrell, Sara, Jane Humphries, and Martin Weale. 1994. "An Input-Output Table for 1841." *Economic History Review* 47 (3): 545–66.

Kiesling, L. Lynne. 1996. "Institutional Choice Matters: The Poor Law and Implicit Labor Contracts in Victorian Lancashire." *Explorations in Economic History* 33 (1): 65–85.

Lindo, Jason M. 2015. "Aggregation and the Estimated Effects of Economic Conditions on Health." *Journal of Health Economics* 40: 83–96.

Miller, Grant, and B. Piedad Urdinola. 2010. "Cyclicality, Mortality, and the Value of Time: The Case of Coffee Price Fluctuations and Child Survival in Colombia." *Journal of Political Economy* 118 (1): 113–55.

Mitchell, B.R. 1988. *British Historical Statistics.* Cambridge, UK: Cambridge University Press.

Mitchell, B.R., and Phyllis Deane. 1962. *Abstract of British Historical Statistics*. Cambridge, UK: Cambridge University Press.

Olivetti, Claudia. 2014. "The Female Labor Force and Long-Run Development." In *Human Capital in History: The American Record*, edited by Leat Platt Boustan, Carola Frydman, and Robert A. Margo, 161–97. Chicago, IL: University of Chicago Press.

Registrar General of England and Wales. 1851–1871. "Annual Report of the Registrar-General of Births, Deaths, and Marriages in England." http://www.histpop.org/ (accessed January 1, 2016).

Rossi, Alice S. 1965. "Naming Children in Middle-Class Families." *American Sociological Review* 30 (4): 499–513.

Ruhm, Christopher J. 2000. "Are Recessions Good for Your Health?" *Quarterly Journal of Economics* 115 (2): 617–50.

Ruhm, Christopher J. 2005. "Healthy Living in Hard Times." *Journal of Health Economics* 24 (2): 341–63.

Ruhm, Christopher J., and William E. Black. 2002. "Does Drinking Really Decrease in Bad Times?" Journal of Health Economics 21 (4): 659–78.

Schurer, K., and E. Higgs. 2020a. "Integrated Census Microdata (I-CeM), 1851–1911." UK Data Service. http://doi.org/10.5255/UKDA-SN-7481-2.

Schurer, K., and E. Higgs. 2020b. "Integrated Census Microdata (I-CeM) Names and Addresses, 1851–1911: Special Licence Access." UK Data Service. http://doi.org/10.5255/UKDA-SN-7856-2.

Southall, Humphrey, and David Gilbert. 1996. "A Good Time to Wed?: Marriage and Economic Distress in England and Wales, 1839–1914." Economic History Review 49 (1): 35–57.

Southall, H.R., D.R. Gilbert, and I. Gregory. 1998. "Great Britain Historical Database: Labour Markets Database, Poor Law Statistics, 1859–1939." UK Data Service. http://doi.org/10.5255/UKDA-SN-3713-1.

Stevens, Ann H., Douglas L. Miller, Marianne E. Page, and Mateusz Filipski. 2015. "The Best of Times, the Worst of Times: Understanding Procyclical Mortality." American Economic Journal: Economic Policy 7 (4): 279–311.

Stuckler, David, Christopher Meissner, Price Fishback, Sanjay Basu, and Martin McKee. 2012. "Banking Crises and Mortality during the Great Depression: Evidence from U.S. Urban Populations, 1929–1937." Journal of Epidemiology and Community Health 66 (5): 410–19.

Sullivan, Daniel, and Till von Wachter. 2009. "Job Displacement and Mortality: An Analysis Using Administrative Data." Quarterly Journal of Economics 124 (3): 1265–1306.

Thomas, Mark. 1984. "An Input-Output Approach to the British Economy, 1890–1914." PhD diss. Oxford University.

UK Parliament. 1868. "Parliamentary Papers Report 0455153." UK Parliamentary Archives. https://archives.parliament.uk/ (accessed January 1, 2016).

Watts, John. 1866. *The Facts of the Cotton Famine*. London, UK: Simpkin, Marshall, and Co.