

Male-biased Demand Shocks and Women's Labour Force Participation: Evidence from Large Oil Field Discoveries

By STEPHAN E. MAURER† and ANDREI V. POTLOGEA‡

†*University of Konstanz and Centre for Economic Performance*

‡*University of Edinburgh*

Final version received 25 February 2020.

Do male-biased labour demand shocks affect women's labour market outcomes? To study this question, we examine large oil field discoveries in the southern USA from 1900 to 1940. We find that oil wealth has an overall positive effect on female labour force participation that is driven by single women. While oil discoveries increase demand for male labour and raise male wages, they do not drive women out of the tradable goods sector or the labour force. Our findings suggest that the absence of any crowding out effects of oil wealth can be explained by compensating forces such as demand effects within the tradable sector, or by income effects that lead to growth in the non-tradable sector.

INTRODUCTION

Across the globe and over time, differences in female labour force participation are substantial (Olivetti 2013). Across sectors, gender ratios vary considerably, even within countries. For example, in 2008, women accounted for only 9% of all the hours worked in construction in the USA, but for 76% of all the hours worked in health services (Ngai and Petrongolo 2017). Because of these differences in economic specialization by gender, sector-level shocks can be gender-biased and affect men and women very differently.

One salient example of a gender biased shock to labour markets is the discovery of mineral wealth. Oil or natural gas discoveries can be considered male-biased labour demand shocks that typically lead to increased wages and improved labour market prospects for men. However, as argued by Ross (2008), oil wealth may be not only beneficial to men, but also detrimental to women's involvement in the labour market: Dutch disease effects can harm other tradable sectors that employ women and thus reduce demand for female labour. Moreover, by increasing male wages, oil discoveries may reduce married women's labour supply, as predicted by standard models of household labour supply (Ashenfelter and Heckman 1974; Becker 1981). Thus both supply of and demand for female labour may decrease, resulting in a decrease in female employment.

In this paper, we use the discoveries of large oil fields in the southern USA between 1900 and 1940 as a series of large and exogenous male-biased labour demand shocks to study the implications of such shocks for female labour and marriage market outcomes. We use data on the location and discovery of large oil fields, coupled with county-level census data, to compare the evolution of labour markets in oil-rich counties relative to baseline counties. Using a difference-in-differences research design, we confirm that oil discoveries indeed constitute substantial male-biased labour demand shocks. They lead to significant increases in male wages, substantial in-migration flows and a (relative) reallocation of male labour from agriculture to manufacturing, services and oil mining. However, the discovery of mineral wealth does not drive women out of the labour force. If anything, female labour force participation even increases in the aftermath of oil discoveries, particularly for single women. Oil appears to slightly weaken women's

propensity to work in manufacturing, but this is more than compensated for by relative growth of the service sector, especially personal services and trade.

We also study marriage markets and their interaction with the labour market. We find that oil wealth has no effect on women's general propensity to be married, and only a small effect on the propensity of young women to be married. Thus even though married women have a substantially lower propensity to participate in the labour market than single women, the effect of oil on labour markets mediated by the marriage market is negligible. Overall, our findings suggest that the effects of oil wealth on women are small and largely benign. We find little evidence for the crowding out mechanisms highlighted by some of the previous literature, and also find indirect evidence that supports the presence of other forces that tend to improve women's labour market prospects in the aftermath of oil discoveries. These forces include demand growth within the tradable goods sector, as well as income effects that lead to growth in the non-tradable services sector, in which women have a comparative advantage.

Our paper belongs to a sizeable literature that has analysed the implications of resource shocks and other demand shocks for local economies and local labour markets. Corden and Neary (1982) provide a theoretical discussion of the effects of a booming tradable goods sector on prices and employment in other tradable good sectors in a small open economy. Carrington (1996) uses the construction of the Trans-Alaska Pipeline as an exogenous labour demand shock to study the Alaskan labour market. Sachs and Warner (1999, 2001) provide cross-country evidence for a negative effect of resource abundance on economic growth, and Papyrakis and Gerlagh (2007) find similar effects at the US state level. On the other hand, Pratt (1980) describes how large oil discoveries in Texas led to sustained growth and development.

We confirm the findings of a series of a recent papers that have found positive effects of resource wealth on local economies (e.g. Black *et al.* 2005; Michaels 2011; Marchand 2012; Aragon and Rud 2013; Furchtgott-Roth and Gray 2013; Allcott and Keniston 2014; Fetzer 2014). Like them, we find that oil discoveries lead to sustained and substantial population growth and to a shift of the economy from agriculture into other sectors. In our case, these sectors are oil mining, manufacturing and services. However, the focus of this literature has typically been on the more general implications of resource wealth in terms of employment and welfare, and gender differences have not been analysed in much detail.

At the same time, the vast changes in female labour force participation over the last one hundred years have sparked considerable interest among labour economists. Research has highlighted the importance of both supply-side and demand-side factors, as well that of changing social norms. Reasons for supply shifts include medical advances (Goldin and Katz 2002; Albanesi and Olivetti 2016); labour-saving technological progress in the production of household goods (Greenwood *et al.* 2005; Jones *et al.* 2015), declining childcare costs (Attanasio *et al.* 2008), the persistent effects of male wartime mobilization (Acemoglu *et al.* 2004), and changes in the age structure (Huang 2018). On the demand side, the literature has documented the importance of structural transformation (Akbulut 2011; Voigtländer and Voth 2013; Ngai and Petrongolo 2017) and trade integration (Do *et al.* 2016; Gaddis and Pieters 2017) as shifters of gender-specific demands. Finally, Fernandez (2011, 2013) and Alesina *et al.* (2013) present evidence for the importance of social norms. Klasen and Pieters (2015) analyse the determinants of continued low female labour force participation in India, and find evidence of adverse factors operating through both demand-side and supply-side channels, as well as through conservative social norms.

To the best of our knowledge, Michael Ross (2008, 2012) was the first to connect resource wealth to female labour force participation. Several papers have followed since then (Kang 2009; Norris 2010; Simmons 2016; Aragon *et al.* 2016; Kotsadam and Tolonen 2016). However, the earlier part of this literature has focused on cross-country comparisons. Such comparisons suffer from the potential problem of not being able to control satisfactorily for differences in institutions, culture and other unobservables at the country level. Our setting is able to address this issue.

We rely on within-country variation in oil wealth that is plausibly exogenous. This allows us to avoid many of the shortcomings of cross-country studies such as the difficulty to adequately control for differences in institutions and culture. Using information about the timing of oil field discoveries, we can employ a clean difference-in-differences research design, comparing the evolution of oil-rich counties before and after discoveries to that of non-oil counties over the same time period. Oil discoveries as a source of variation in labour demand are also not subject to endogeneity concerns that arise in connection with other measures such as contemporaneous oil production. Closest to our paper in this respect are two contemporaneous studies by Kotsadam and Tolonen (2016) and Aragon *et al.* (2016). The former uses mine openings and closures in Africa, and finds a negative effect on female labour force participation. The key difference in our paper is that these mines are more short-lived than the large oil fields that we consider—with an average lifespan of 10 years, they lead to local boom and bust cycles. The oil wealth in our sample, on the other hand, had very broad and long-lasting growth implications (Michaels 2011). In this respect, our paper is more similar to that of Aragon *et al.* (2016), who, however, look at a persistent demand shock that is biased against males, namely the collapse of the UK coal industry. Our setting is therefore arguably closer to the theoretical mechanisms described by Ross (2008).

Finally, the case of the southern USA offers us the advantages of analysing consistent, high-quality data from a region that at that time still relied heavily on agriculture. This is important because much of the discussion surrounding the link between resource abundance and gender issues has concentrated on developing, often heavily agrarian economies.

The rest of this paper is organized as follows. Section I briefly reviews some of the channels through which oil discoveries could influence women's labour market outcomes. The data and empirical strategy of our paper are then discussed in Section II. Section III presents the results and several robustness checks, and Section IV concludes.

I. CONCEPTUAL FRAMEWORK

Before we proceed to our main analysis, in this section we provide a brief overview of the various mechanisms through which oil wealth may affect the labour market outcomes of women. We focus on describing channels operating directly through factor markets, and in particular through the labour market.¹ Building on existing work (see, for example, Ross 2008, 2012), we organize our discussion in three parts: (1) demand-side channels linking oil wealth and female labour market outcomes; (2) supply-side channels; and (3) the role of migration in mediating the effect of mineral wealth on women's labour market outcomes. A formal model that captures these various mechanisms is presented in Online Appendix A.

Demand-side channels

To analyse the effects of oil wealth on the demand side of local labour markets, it is useful to separate labour demand into two components, namely labour demand coming from the tradable sector and labour demand coming from the non-tradable sector. Focusing on the former, we note that a substantial literature (e.g. Ross 2008; Michaels 2011) has documented that large mineral endowments encourage economic specialization in natural resource extraction and associated sectors. These activities typically display heavily male labour forces, and can thus be expected to have little direct impact on the demand for female labour.

However, the growth of these activities can still have important indirect effects on the demand for female labour. In particular, the expansion of the extractive sector tends to push up the local prices of a wide array of production inputs, including male labour and land, that are in imperfectly elastic local supply. In turn, this may lead to two opposing effects on the demand for female labour originating from the tradable sector. First, increases in the prices of inputs shared by the oil and local tradable sectors will encourage substitution of these inputs with others whose prices have not been directly affected by the growth in the oil sector, for example, female labour. This leads to a positive ‘substitution effect’ on the local demand for female labour. On the other hand, if some of the inputs used by the oil sector are difficult to substitute in the production processes of other local tradable sectors, then increased prices of these inputs will raise the production costs of these sectors. This in turn renders these sectors uncompetitive on the national market and causes them to shrink, leading to a reduction in the local demand for female labour. We call this a ‘scale effect’ of oil discoveries on the tradable sector demand for female labour, which captures the standard crowding-out mechanism highlighted by the resource curse literature. All in all, the overall effect of oil wealth on the demand for female labour in the tradable sector is ambiguous and depends on the relative strength of the substitution and scale effects.

By contrast, the effect of oil discoveries on the demand for female labour in the non-tradable sector is more clear-cut. As the non-tradable sector is not subject to competition from outside the local economy, the negative scale effect identified for the tradable sector is not present. Indeed, the increase in purchasing power brought about by oil discoveries (as male wages increase) is likely to increase demand for local non-tradables, leading to a positive scale effect on the demand for female labour. Moreover, the substitution effect may also be operational, potentially further increasing demand for female labour. Thus oil wealth can be expected to increase the demand for female labour in the non-tradable sector.

Taken together, the overall effect of oil discoveries on the demand for female labour is therefore ambiguous. Should the negative scale effect in the tradable sector dominate both the substitution effect in that sector and the positive effect on the demand for female labour in the non-tradable sector, then the overall demand for female labour will decline in the aftermath of oil discoveries. This scenario, in which mineral wealth shocks crowd women out from the labour market, was first highlighted by Ross (2008, 2012) and can be described as a ‘gender-biased resource curse’. However, as our analysis indicates, this is just one of multiple scenarios that may prevail in the aftermath of oil discoveries. In general, mineral wealth may have negative, zero or positive effects on the demand for female labour, such that the actual effect of natural resource endowments remains an empirical question.

Supply-side channels

Large oilfield discoveries can also be expected to affect the supply of female labour. One potential channel is via women's reservation wages, which can be expected to increase in the aftermath of oil discoveries. A substantial literature in labour economics (e.g. Ashenfelter and Heckman 1974) has documented the existence of a negative elasticity of married women's labour supply with respect to their husbands' wages. In this context, if the discovery of oil leads to an increase in male wages, then one might expect to see a decline in women's labour supply and labour force participation. Moreover, while this mechanism should directly affect only married women, if women are forward-looking and their reservation wages are a function of expected future household income, then the effects could also extend to single women.

Another channel through which natural resource shocks may affect female labour supply is via their effects on the marriage market. If oil discoveries increase male wages and reduce unemployment, then they may increase the supply of 'marriageable men' and thus increase the thickness of local marriage markets. This in turn may increase marriage rates and reduce the age at first marriage for women. As the labour supply of married women is typically lower than that of unmarried women, this would be an additional force through which mineral wealth may depress female labour supply.

Taken together, the above mechanisms suggest that we can expect oil discoveries to reduce the female labour supply, which, absent countervailing effects on labour demand, should reduce women's involvement in the labour market.

The role of migration

A key limitation of our analysis so far is that it ignored the role of migration. This omission would be defensible in settings in which the units of interest can be reasonably considered as closed from the perspective of migration, such as in the cross-country analyses that are prevalent in the resource curse literature. However, we study the effects of oil discoveries in a within-country setting in which migration can be quantitatively important.

Adding migration to the analysis does not qualitatively change our conclusions regarding the potential effects of natural resource shocks on the demand for female labour. As long as the oil sector and other tradable sectors share at least one input that is an imperfectly elastic local supply, the 'scale' (or crowding out) and substitution effects remain operational. Similarly, with increased in-migration brought about by the discovery of mineral wealth, demand for local services (and other local non-tradables) can be expected to increase, with positive implications for the demand for female labour. Overall, the effect of oil discoveries on the demand for female labour remains ambiguous.

Our analysis of the effects of oil discoveries on the individual labour supply of women also remains qualitatively unchanged.² As long as migration does not render local male labour supply perfectly elastic, oil discoveries can be expected to bring about increases in male wages and a corresponding increase in female reservation wages. Moreover, even in the case in which male migratory responses are large, if in-migration into oil-rich countries is gender-biased, then this may raise female bargaining power in marriage markets, thus increasing their share of marital surplus and their reservation wages. Moreover, a rise in the share of marital surplus accruing to women could raise female marriage rates and lower the typical age at first marriage. In turn, this provides an additional mechanism leading towards a decrease in individual labour supply for women, given that married women in our setting typically display lower levels of attachment to

the labour market than single women. Taken together, our conclusion that large natural resource shocks are likely to depress the individual labour supply of women is robust to the addition of migration to the analysis.

The main novelty brought by the addition of migration to the analysis concerns the effect of oil discoveries on the *aggregate* female labour supply, over and above their effect on *individual* female labour supplies. To see this additional channel, consider a setting where the net effect of oil discoveries on the aggregate demand for female labour is zero (for instance because oil does not share any inputs with the tradable sector and the expenditure share of the subset of the non-tradable sector that employs women is negligible) and the effect of oil wealth on female individual labour supply is zero (say because the elasticity of female reservation wages with respect to male wages is zero) but marriage markets have to clear in equilibrium. In such a setting, any male migratory response to an oil shock is also accompanied by a female migratory response, such that marriage markets are kept in equilibrium. This migratory response increases the local female working-age population, but under conditions of constant aggregate labour demand, and an already binding (and constant) reservation wage does not commensurately increase the female labour force. Thus holding all else constant, female labour force participation mechanically declines in the aftermath of oil discoveries due to the operation of the migration channel.

II. DATA AND EMPIRICAL STRATEGY

In order to identify major oil field discoveries, we draw on the dataset compiled by Guy Michaels (2011), which lists all counties in the southern and the south-western USA that are situated above an oil field of 100 million barrels or more before any oil was extracted. We will refer to these counties simply as oil-rich counties. In addition, the dataset also contains information about the earliest major discovery in each county. We treat a county as having discovered its oil wealth after the first major discovery has taken place. Figures 1–5 give an overview of the geography of oil discoveries between 1900 and 1940 based on county shape files from NHGIS (Minnesota Population Center 2016). There are only relatively few oil discoveries in the early years, with most discoveries happening in the 1920s and 1930s. In terms of geographic scope, we follow Michaels (2011) and include all counties within 200 miles of the oil-rich counties in Texas, Louisiana and Oklahoma. This gives us a sample of 774 counties in total.³

For our outcome variables, we use the individual-level US census data available from the Integrated Public Use Microdata Series (IPUMS) for 1900 to 1940 (Ruggles *et al.* 2015). This includes four novel, recently made available preliminary full count (100%) samples for 1910, 1920, 1930 and 1940. For 1900, we use the largest samples available, which contains 5% of the population. We generally focus on the part of the population that is of working age, defined to be between 17 and 65.

Most of our variables of interest are available for our whole period of analysis. This includes labour force participation,⁴ the sector in which an individual works, gender, race and marriage status. Employment status conditional on labour force participation was not recorded in the censuses of 1900 and 1920, but is available for the remaining three census decades. For earnings, we observe annual wage income in 1940 only. As a proxy for wages in the years 1900–30, we use an ‘occupational earnings score’ variable that is available for 1900–40. This variable assigns to each individual the percentile rank of his occupation in terms of median earnings in 1950. While this measure does not capture wage changes within a given occupation, it will capture wage changes due to movements




FIGURE 1. Map of oil discoveries 1900.

Notes: Oil-abundant counties are shown in dark grey, neighbours of oil-abundant counties in mid-grey, and other counties in light grey. White areas indicate counties not yet created or dropped due to substantial area changes relative to 1940 (see note 2 for details).




FIGURE 2. Map of oil discoveries 1910.

Notes: See Figure 1.

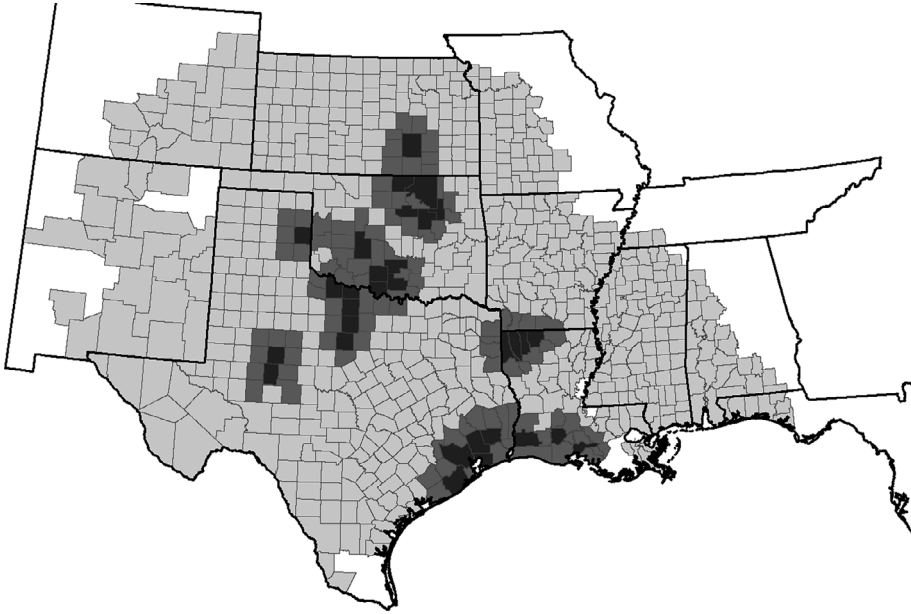


FIGURE 3. Map of oil discoveries 1920.
Notes: See Figure 1.

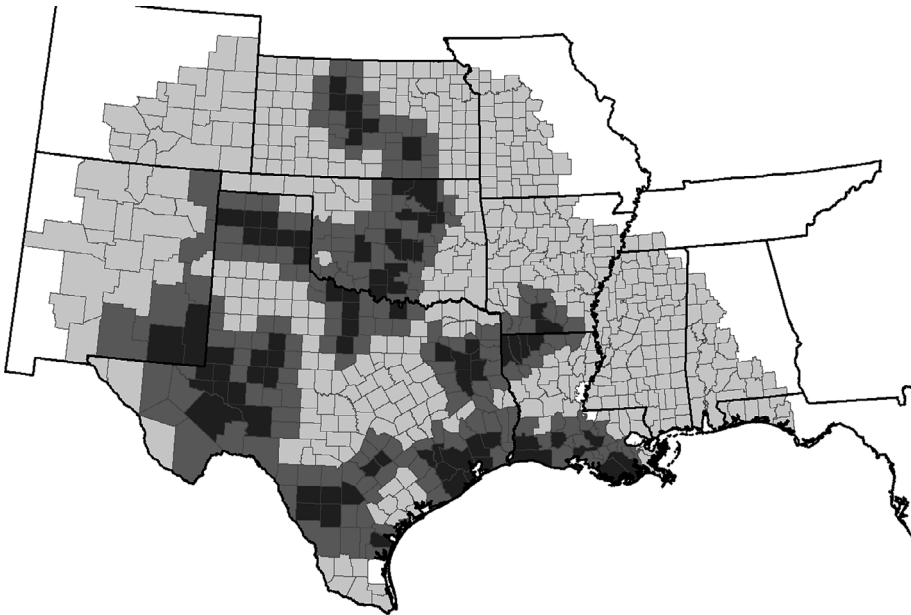


FIGURE 4. Map of oil discoveries 1930.
Notes: See Figure 1.

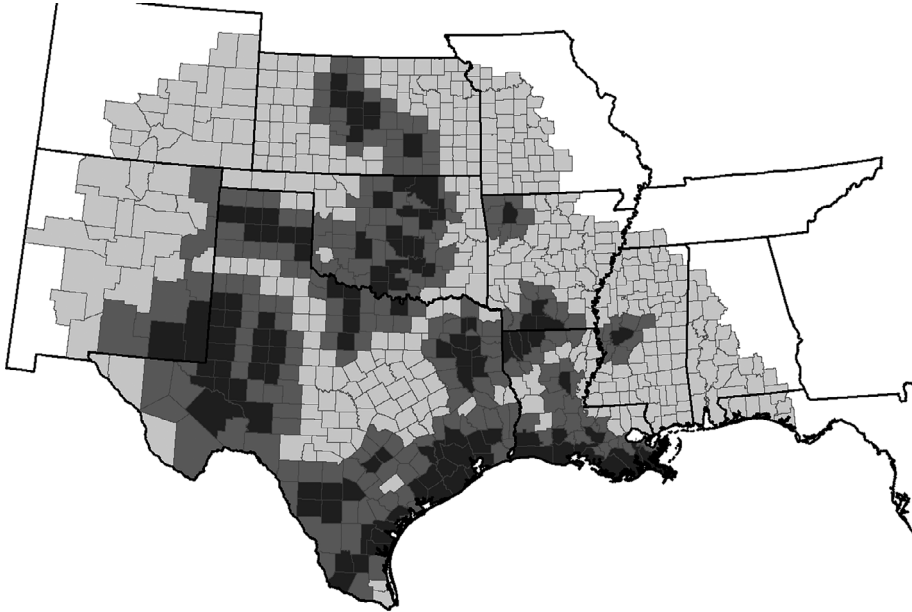


FIGURE 5. Map of oil discoveries 1940.

Notes: Oil-abundant counties are shown in dark grey, neighbours of oil-abundant counties in mid-grey, and other counties in light grey.

across occupations. Because of this, it also offers the additional advantage of being robust to local price changes that would drive a wedge between nominal and real wages.

The staggered nature of oil discoveries across space and over time lends itself quite naturally to a difference-in-differences research design. The basic regression that we run is of the form

$$(1) \quad y_{ct} = \alpha_c + \tau_t + \beta \text{DiscoveredOilField}_{ct} + X'_c \gamma_t + u_{ct},$$

where y_{ct} denotes outcome y (e.g. the labour force participation rate) in county c and year t . $\text{DiscoveredOilField}$ is a dummy that equals 1 if county c is oil-rich and at least one of its major oil fields has already been discovered, and 0 otherwise. τ and α are year and county fixed effects. X is a vector of control variables that vary at the county level. In line with Michaels (2011), we control for several geographic features that might be spuriously correlated with oil wealth: longitude, latitude, aridity, average annual rainfall, distance to the closest navigable river, and distance to the ocean. All of these variables vary at the county level only, but we allow for them to have time-varying effects.

Since our key variation lies at the county \times year level, we aggregate the individual census data to this level.⁵ We do not weight, and instead follow Black *et al.* (2005) in treating every county as a unit of observation. Given our setting, serial correlation in the error terms is a potential problem (Bertrand *et al.* 2004), as is spatial correlation. To address both issues, we cluster standard errors at the level of 1×1 degree cells (Bester *et al.* 2011). When calculating standard errors, we use the procedure of Young (2016) that calculates an effective degrees of freedom correction for clustered and robust standard errors. These effective degrees of freedom are also used when performing hypothesis tests.

The identifying assumption of our difference-in-differences approach is that in the absence of oil discoveries, counties with major oil deposits would have evolved in the same way over time as those without oil discoveries. In Online Appendix B, we report several checks to assess the plausibility of this assumption.

For wages, we have data only from 1940, requiring a cross-sectional specification with stronger identifying assumptions similar to those of Michaels (2011):

$$(2) \quad y_c = \beta \text{DiscoveredOilField}_c + X'_c \gamma + s_s + u_c.$$

In this case, $\text{DiscoveredOilField}_c$ codes whether county c has an oil field that was already discovered by 1940. X is the same vector of control variables as above, and the s are state fixed effects. The more demanding identification assumption behind this specification is that counties without oil wealth and counties that had not yet discovered their oil wealth by 1940 are a valid counterfactual for those counties with already discovered oil fields in 1940. In Online Appendix B, we show several robustness checks, the results of which support this assumption.

The setting of our analysis is the south-western USA five to nine decades after the American Civil War. After the war, the southern USA remained a predominantly agricultural region that lagged considerably behind the rest of the country and caught up only very slowly (Wright 1986; Caselli and Coleman 2001). This is also borne out by the county-level summary statistics displayed in Table 1 separately for counties with and without major oil fields. We show the labour force shares of four major sectors

TABLE 1
COUNTY-LEVEL SUMMARY STATISTICS

	Counties without major oil fields			Counties with major oil fields		
	Mean	S.D.	Obs.	Mean	S.D.	Obs.
Agriculture labour force share 1900	0.720	0.194	475	0.702	0.218	113
Manufacturing labour force share 1900	0.037	0.051	475	0.041	0.064	113
Service labour force share 1900	0.201	0.138	475	0.217	0.164	113
Oil mining labour force share 1900	0.000	0.001	475	0.000	0.000	113
Agriculture labour force share 1940	0.500	0.176	591	0.413	0.197	167
Manufacturing labour force share 1940	0.063	0.057	591	0.061	0.054	167
Service labour force share 1940	0.296	0.120	591	0.355	0.128	167
Oil mining labour force share 1940	0.005	0.018	591	0.056	0.078	167
Female labour force participation rate 1900	0.169	0.137	474	0.152	0.117	113
Female labour force participation rate 1940	0.197	0.072	603	0.201	0.065	171
Male labour force participation rate 1900	0.934	0.042	475	0.933	0.057	113
Male labour force participation rate 1940	0.864	0.141	603	0.876	0.141	171
Black population share 1900	0.204	0.262	475	0.213	0.234	113
Black population share 1940	0.157	0.210	603	0.132	0.152	171
Urban population share 1900	0.084	0.174	475	0.066	0.162	113
Urban population share 1940	0.197	0.228	603	0.276	0.236	171
Marriage share, all women, 1900	0.688	0.087	474	0.707	0.110	113
Marriage share, all women, 1940	0.725	0.037	603	0.734	0.047	171
Marriage share, women aged 16–25, 1900	0.491	0.144	470	0.513	0.170	111
Marriage share, women aged 16–25, 1940	0.510	0.069	603	0.537	0.077	171

(agriculture, manufacturing, services, oil mining) for the beginning and end of our sample. As can be seen, agriculture is by far the most important sector in 1900, but loses importance over the next 40 years. Manufacturing and services, on the other hand, are slowly gaining ground. Over the whole time period, our area of observation is a very rural region. Even by 1940, the vast majority of the population in our dataset live outside cities. In addition, the summary statistics also show a very low level of female involvement in the labour force. Female labour force participation was around 15% in 1900, and increased slightly between then and 1940. This is no southern peculiarity: Across the whole USA, female labour force participation rose from 20.6% in 1900 to 25.8% in 1940 (Goldin 1990, ch. 2). However, there was considerable heterogeneity along two dimensions: black women were much more likely to work than white women (Goldin 1977), and single women had much higher labour force participation than married ones (Goldin 1990).

All in all, our basic setting is thus a still heavily rural and agricultural economy that is slowly moving towards more ‘modern’ sectors, and which over the whole period of analysis displays relatively low female labour force participation rates. What is also visible from the summary statistics is that oil-rich and non-oil-rich counties are very similar in 1900. In Online Appendix B, we test this more thoroughly. Running cross-sectional regressions of our outcomes in 1900 on a dummy for being oil-rich, we find that there are no important differences between oil-rich and non-oil-rich counties before oil wealth was discovered.

III. RESULTS

In 1900, the economy of the southern USA was thus still predominantly agricultural and almost no major oil field had yet been discovered, as can be seen from Figure 1. To set the stage for our main analysis relating resource shocks to gender outcomes, we first check whether oil discoveries indeed represent important developments for the discovering counties. Table 2 shows difference-in-differences results for the log of the working-age population, female population share, and the share of the labour force employed in oil mining, manufacturing, services and agriculture. We do not find any evidence for a resource curse, but rather detect substantial growth in counties where major oil field discoveries take place: population increases by around 42%, and the labour force shifts from agriculture into oil mining and, to a lesser extent, services and manufacturing. These findings confirm the results of the analysis of Michaels (2011) for the same region that used across-county comparisons. By using the timings of the discoveries, our analysis is based on the weaker identifying assumption of common trends instead. The results show that oil discoveries constitute substantial shocks to a county’s economy that lead to growth and structural transformation. Interestingly, the population growth seems gender-balanced, with the female population share staying unchanged.

Since many of the mechanisms through which oil wealth may affect female labour market outcomes operate via the labour market outcomes of men, in Table 3 we present the implications of oil discoveries for men. We find no effect of oil discoveries on male labour force participation, which is perhaps unsurprising given that baseline male labour force participation was almost 90%. However, we find a large increase in male wages in the aftermath of oil discoveries. Oil wealth is associated with an increase in male wages of about 28 log points. Importantly, this wage increase for men is broad based and can be shown not to be solely driven by the emergence of high-wage jobs in oil mining.⁶ Coupled

TABLE 2
BROAD ECONOMIC EFFECTS OF OIL DISCOVERIES

Variables	ln(population) (1)	Female pop. share (2)	Share of labour force employed in			
			Oil (3)	Manufacturing (4)	Agriculture (5)	Services (6)
Discovered oil field	0.351*** (0.069)	0.000 (0.003)	0.065*** (0.010)	0.009* (0.005)	-0.105*** (0.023)	0.037*** (0.010)
Observations	3597	3597	3581	3581	3581	3581
Mean dep. var.		0.479	0.009	0.051	0.612	0.253

Notes

Number of clusters is 202. Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

*** ** * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

TABLE 3
OIL DISCOVERY EFFECTS FOR MENS

Variables	Labour force participation rate (1)	Employment rate (2)	Average ln (wage income) (3)	Average occupational earnings score (4)	Share of labour force employed in			
					Oil (5)	Manufacturing (6)	Agriculture (7)	Services (8)
Discovered oil field	0.000 (0.009)	-0.011** (0.005)	0.279*** (0.040)	5.840*** (1.159)	0.076*** (0.012)	0.012** (0.005)	-0.113*** (0.024)	0.031*** (0.010)
Observations	3597	2232	758	3581	3581	3581	3581	3581
Clusters	202	202	201	202	202	202	202	202
Mean dep. var.	0.895	0.951	5.992	26.623	0.011	0.057	0.653	0.195
Years	1900–40	1910, 1930, 1940	1940	1900–40	1900–40			

Notes

Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

*** ** * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

with the finding of a large migratory response for men⁷ in the aftermath of large oilfield discoveries, the evidence is consistent with oil wealth representing a substantial positive shock to the demand for male labour.

Moreover, we also find that oil discoveries have a sizeable impact on the sectoral allocation of male labour in affected counties. Perhaps unsurprisingly, oil wealth is associated with an increase of 7.6 percentage points in the share of the local male labour force working in the oil sector. Manufacturing and services experience more modest increases of 1.2 and 3.1 percentage points. The shift in the composition of male employment towards oil is sizeable: as there is virtually no oil mining prior to oil discoveries, our findings suggest that 7.5% of the whole male labour force of affected counties shifts towards the oil sector in the aftermath of discoveries. By contrast, the

propensity of men to work in agriculture declines by 11.3 percentage points in oil-rich counties.

Having established the implications of major oil discoveries for men, we now proceed to study our main outcomes of interest, namely the labour market outcomes of women. Our main results are presented in Table 4. Interestingly, we find that oil discoveries are associated with a 1.4 percentage point increase in female labour force participation in the affected counties. This result stands in contrast with some of the previous findings in the literature, including those of Ross (2008) and Kotsadam and Tolonen (2016). While the magnitude of the point estimate is modest, it is important to note that our setting features very low female labour force participation. Over the whole sample, the mean labour force participation rate is only 19.2%. Compared to this, our results indicate that the probability of a working-age woman being in the labour force increases by about 7% after oil discoveries.

We also find that oil discoveries are associated with an increase in female wages. However, at only 7.2 log points, the positive effect of oil wealth on female wages is much weaker than that for men. Nevertheless, coupled with the large migratory response to oil discoveries observed for women, our results again indicate that oil discoveries represented an overall positive shock for the local demand for female labour. Moreover, this increase in aggregate demand was sufficiently large not only to absorb the additional supply of female labour brought in by migration, but also to yield small increases in both female wages and the rate of female labour force participation.

Given that our wage findings above are based only on studying a cross-section for 1940, Tables 3 and 4 also include results for the ‘occupational earnings score’ for men and women, respectively. This variable is available for the entire period 1900–40, and captures changes in wages resulting from occupational transitions, but unfortunately does not capture changes in wages within occupations. Our results suggest that for men, the observed increase in wages was accompanied by occupational upgrading, with the

TABLE 4
OIL DISCOVERY EFFECTS FOR WOMEN

Variables	Labour force participation rate (1)	Employment rate (2)	Average ln (wage income) (3)	Average occupational earnings score (4)	ln(female labour force) (5)	ln(female workers in)		
						Agriculture (6)	Manufacturing (7)	Services (8)
Discovered oil field	0.014** (0.006)	−0.004 (0.004)	0.072*** (0.023)	−1.800** (0.803)	0.461*** (0.085)	0.148* (0.083)	0.077 (0.116)	0.492*** (0.092)
Observations	3596	2232	758	3542	3542	3454	2802	3517
Clusters	202	202	201	202	202	202	202	202
Mean dep. var.	0.192	0.943	5.541	23.868				
Years	1900–40	1910, 1930, 1940	1940	1900–40	1900–40			

Notes

Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

*** indicate $p < 0.01$, ** indicate $p < 0.05$, * indicate $p < 0.1$, respectively.

average man moving up in the 1950 median earning distribution by 5.8 percentile ranks. To illustrate the magnitude of this effect for men, going up 5.8 ranks roughly corresponds to half the distance between a porter (occupational earnings score 28) and a salesman/sales clerk (occupational earnings score 40.4). For women, the effect of oil wealth on the occupational earnings score is small and negative. The average woman moves down 1.8 percentile ranks in the ordering of occupations based on 1950 median earnings. These findings suggest that any positive effect of oil discoveries on female earnings took place within occupations, and that the new jobs created for women in the aftermath of oil discoveries were predominantly in lower-paid occupations.

To better understand the mechanisms behind our headline findings above, columns (5)–(8) of Table 4 show results for absolute growth in the female labour force and in the three broad sectors agriculture, manufacturing and services. The main takeaway of our findings from this exercise is that we find no evidence of women being crowded out of any part of the labour market. Female employment increases for all three broad sectors, albeit the increase in the case of manufacturing is not statistically significant. This suggests not only that the crowding-out force hypothesized in our conceptual framework is weak or absent, but also that there are forces operating in the direction of boosting female employment for at least two of the broad sectors (i.e. services and agriculture).

One plausible candidate mechanism for our findings is demand effects stemming from increases in local purchasing power brought about by oil wealth. Indeed, we find some support for this mechanism in Table 4, where we find that the service sector (which is the closest proxy for the non-tradable sector featured in our conceptual discussion) expands female employment the fastest. In fact, it is the only sector that increases female employment at a faster rate than the growth in the overall female labour force. Moreover, in Online Appendix B, we analyse the service sector response in more detail and find that the entire growth in female service sector employment is driven by the personal service and trade sectors. This suggests that income effects operating via final demand rather than demand linkages from other industries are the driver of this growth.⁸ Final demand linkages are also a plausible explanation for the growth in agricultural sector employment that we document. In the context of our conceptual framework, agriculture may be most easily regarded as a tradable sector. However, due to transportation costs and perishability concerns, some parts of the agricultural sector might behave more like a non-tradable sector.⁹

It is useful to complement the discussion above with an analysis of the evolution of the sectoral composition of female employment in the aftermath of oil discoveries. We present the results of such an analysis in columns (1)–(4) of Table 5. We confirm that local economies affected by oil discoveries experience a small shift of their female employment towards the services sector, with the share of services in the overall female labour force increasing by about 3 percentage points. However, this point estimate is marginally statistically insignificant. The most striking finding is the relative decline of the manufacturing sector, which experiences a decline in its share of the female labour force of 1.3 percentage points. While not large in absolute value, this is a sizeable reduction in the importance of manufacturing, given the low overall importance of manufacturing as an employer of women. Across our whole sample, only 3% of women work in manufacturing, meaning that the oil effect constitutes a decrease of more than a third. We also find a statistically significant increase in the share of the female labour force employed in the oil sector. However, at only 0.2 percentage points, this effect is economically negligible. Overall, the patterns of sectoral labour reallocation are much more muted for women than for men.

TABLE 5
SECTORAL EMPLOYMENT SHARES FOR WOMEN

Variables	Share of female labour force employed in			
	Oil (1)	Manufacturing (2)	Agriculture (3)	Services (4)
Discovered oil field	0.002*** (0.000)	-0.013*** (0.005)	-0.011 (0.020)	0.029 (0.018)
Mean dep. var.	0.000	0.029	0.326	0.623

Notes

3540 observations. Number of clusters is 202. Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

***, **, * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

As outlined in our theoretical discussion, another mechanism that may explain the rise in demand for female labour that we detect in the data is the substitution effect, where increases in the prices of production inputs that are substitutes for female labour drive growth in female labour demand. We can provide a partial test for this channel by checking for evidence of substitution of female labour for male labour (which we know is a production input that becomes more expensive due to oil wealth) in the aftermath of oil discoveries. We perform this test by tracking the evolution of gender ratios for the three broad sectors of agriculture, manufacturing and services. The results of this exercise are presented in Table 6 and suggest that the substitution effect is not an important driver of our headline findings. Female to male ratios do not increase in any of the the three broad sectors analysed, and in fact decline significantly in the manufacturing sector. This latter finding is likely to be driven by a change in the composition of manufacturing due to input–output linkages: male-intensive manufacturing activities that either supply intermediates to the oil sector or use oil intensively in their production are likely to grow significantly when oil wealth is discovered.¹⁰ It is important to note, however, that we cannot rule out the substitution channel being important with respect to other inputs (e.g. land) that may have gone up in price in the aftermath of oil discoveries.

We thus find an overall positive effect of oil wealth on the demand for female labour. However, given that oilfield discoveries were very large shocks, the magnitudes of the labour market effects for women appear relatively modest. Both the sectoral reallocation and wage effects are relatively small and much weaker than for men, while the positive effect on female labour force participation can still be regarded as modest.¹¹ One potential explanation for this is that oil discoveries trigger a series of quantitatively important opposing mechanisms (crowding-out effects in part of the tradable sector, positive demand effects for the non-tradable sector, substitution effects, and so on) that net out to produce only a small aggregate effect on women's labour market outcomes.

Another potential explanation is that women in the time and region of our analysis were not very responsive to labour market shocks, perhaps due to social norms and institutions. As Goldin (1990, 2006) shows, female labour force participation in the USA was still very low and only slowly rising during the first third of the 20th century. This is also reflected in our sample, where the female labour force participation rate over the whole time period is just below 20%. This is consistent with a majority of women being potentially far away from the participation margin, and thus not reacting much to shifts

TABLE 6
GENDER RATIOS WITHIN SECTORS

Variables	Female-to-male ratios in		
	Agriculture (1)	Manufacturing (2)	Services (3)
Discovered oil field	0.006 (0.007)	-0.067** (0.030)	-0.015 (0.025)
Observations	3540	3409	3524
Mean dep. var.	0.111	0.161	0.621

Notes

Number of clusters is 202. Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

***, **, * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

in the labour market. This argument, however, also suggests that there may be an important source of heterogeneity in women's responsiveness to labour market shocks: marriage status. Marriage was a crucial determinant of labour force participation in the USA in the early 20th century. Until the 1920s, it was typical for women to exit the labour force on marriage (Goldin 2006), and in many places, employers had explicit 'marriage bars' that prevented the hiring or led to the firing of married women. Not surprisingly then, in 1900, across the whole country, only 5.6% of all married women were working, as opposed to 43.5% of all single women, and the former share had increased to only 13.8% by 1940 (Goldin 1990). We would thus expect single women to be more responsive to labour market shocks than married women, and proceed to check for this.

Table 7 shows labour force participation results separately for single and married women. Consistent with the notion that married women were far from the participation margin, we do not find any effect of oil discoveries on the labour force participation of married women. The point estimate is positive, but below 1 percentage point and not statistically significant. For single women, however, we find an increase of 3.1 percentage points, which is both economically and statistically significant. Compared to an overall mean of 41%, a 3.1 percentage point increase is an increase of 7.5%. In other words, our headline finding of a positive labour force participation effect of oil wealth is entirely driven by single women, with married women largely unresponsive.

Finally, we study marriage markets and their interaction with the labour market. As we have shown, male wages increase substantially after oil discoveries. This could increase the local supply of 'marriageable' men and thereby increase marriage rates (Buckley 2003; Jelnov 2016). Increased marriage rates in turn may have a dampening effect on female labour force participation given that married women are less likely to work than single women.

One challenge in assessing the plausibility of this channel is the problem that there is a large 'stock' of already married women who are not affected by oil discoveries at all. We therefore need to focus on 'marginal' women whose decision about when and whether to marry can still be affected by oil wealth. These will typically be younger women, among whom the marriage rate and thus the stock of married women is lower. In Table 8, we show the effects of oil discoveries on a woman's probability of being

TABLE 7
EFFECTS OF OIL ON FEMALE LABOUR FORCE PARTICIPATION BY MARRIAGE STATUS

Variables	Female labour force participation rate	
	(1)	(2)
Discovered oil field	0.031*** (0.011)	0.009 (0.006)
Sample	Single women	Married women
Observations	3575	3596
Mean dep. var.	0.408	0.107

Notes

Number of clusters is 202. Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

***, **, * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

TABLE 8
OIL DISCOVERIES AND MARRIAGE RATES

Variables	Share of women married		
	(1)	(2)	(3)
Discovered oil field	0.001 (0.005)	0.014* (0.008)	0.022** (0.009)
Sample	All women	Women ≤ 30	Women ≤ 25
Observations	3596	3594	3590
Mean dep. var.	0.717	0.608	0.513

Notes

Number of clusters is 202. Standard errors, clustered at the 1×1 degree level, in parentheses. Standard errors and p -values are based on bias-adjusted standard errors and effective degrees of freedom corrections as in Young (2016).

***, **, * indicate $p < 0.01$, $p < 0.05$, $p < 0.1$, respectively.

married. Column (1) shows results for all women and illustrates our stock problem: among all women, the marriage rate is close to 70% (and many of those not married are widowed already), and we do not find an effect of oil wealth on marriage rates. For women aged 30 or below, or 25 or below, we find that oil increases the probability of being married by a statistically significant 1.4–2.0 percentage points. This effect, while providing some suggestive evidence in support of the ‘marriageable men’ view of the marriage market, is, however, too small to have any substantive impact on the labour market outcomes of women.

All in all, our results indicate that oil wealth did not negatively affect women’s labour market prospects in the southern USA during the first half of the 20th century. If anything, oil wealth caused a moderate increase in female labour force participation and female wages. The evidence is consistent with oil wealth being associated with an increase in demand for female labour, with demand effects brought about by an increase in local purchasing power in the aftermath of oil discoveries a likely channel. Moreover, most women are unresponsive to the oil discovery shocks: for the 70% of women in our sample who are married, the effects of oil discoveries on labour market outcomes are

negligible. Thus the effects of oil wealth that we identify are driven completely by single women. Given this difference between married and single women's responses to oil discoveries, a promising area for future research involves uncovering the role of social norms and institutions in shaping the responsiveness of female labour market outcomes to economic shocks.

It is important to note that our findings regarding the effects of natural resource wealth on female labour market outcomes are of course context-specific and may not generalize to other settings. Ross (2008) notes that the negative effect of oil on female labour force participation identified in his cross-country study seems to be absent in countries where women are commonly employed in the service sector. This is in line with our findings for the southern USA, which is clearly a setting where women had access to jobs in the service sector. Thus our results may be taken to be informative for settings featuring largely agricultural and relatively underdeveloped regional economies during a time period when oil production was still relatively labour-intensive and in a context in which prevailing social norms allow for women to be employed in service activities.¹²

In Online Appendix B, we perform a battery of robustness checks to determine whether our results can be attributed to oil discoveries in a causal way. The results of these checks generally support the validity of our empirical analysis. To address the question of spatial sorting driving our results, we use the birth states of individuals to identify some migrants. We find that for men, selective migration plays a role, whereas for women, this is much less of a concern. We also show that our results hold when controlling for changes in the age and race structure of the labour force. In addition, we use a leads and lags specification to analyse whether the effects of oil discoveries arise only after the respective oil fields are discovered. We also show that counties with and without major oil deposits were very similar in 1900, when nearly no oil was discovered. Moreover, we show that our qualitative conclusions are generally robust to including county-specific linear trends and to dropping all counties without large oil fields. Finally, using the methodology of Pei *et al.* (2019), we argue that oil discoveries are not systematically associated with the presence of the boll weevil, another important local economic shock during our period of analysis (Lange *et al.* 2009; Ager *et al.* 2017).

IV. CONCLUSION

In this paper, we have analysed the implications of a male-biased labour demand shock in the tradable sector for the labour market outcomes of women. Our particular focus has been to test the hypothesis that mineral wealth may have negative implications for women's involvement in the labour market.

Guided by a simple theoretical framework, we use oil discoveries in the south-western USA between 1900 and 1940 to empirically estimate the effect of oil wealth on women's labour and marriage market outcomes. Contrary to previous studies (e.g. Ross 2008), we do not find a negative relationship between oil wealth and female labour force participation. If anything, female labour force participation increases modestly, with the effect being driven by single women. Moreover, women in oil-rich counties display higher wages than those in baseline counties, though the effect is much smaller than that for men. The effects of oil wealth on marriage markets are small and concentrated on young women who display a higher propensity of being married. As a result, any effects of oil discoveries on labour markets mediated by the marriage market are negligible.

The pattern of findings outlined above is surprising given the large migratory response to oil discoveries that we identify for both men and women. It indicates that oil

wealth is associated with an increase in the (local) aggregate demand for female labour. This increase in demand is sufficient not only to absorb the additional female labour brought in by migration but also to yield a small increase in the local female labour force participation. To better understand the mechanisms driving this finding, we investigate the evolution of the sectoral composition of local economies in the aftermath of oil discoveries, as well as the evolution of the gender ratios within sectors. We find no evidence of crowding-out effects of oil wealth: female employment in all broad sectors, including tradable sectors such as agriculture and manufacturing, grows in the aftermath of oil discoveries. Moreover, the growth of female employment does not seem to be driven by the substitution of male labour with female labour, as gender ratios in all broad sectors change very little. Oil wealth is also associated with a small decline in the propensity of women to work in manufacturing and an increased propensity to work in the services sector, in particular in personal services and trade. Put together, these findings are consistent with increased demand driving growth in female employment in both the tradable and non-tradable sectors.

Overall, our results indicate that oil wealth did not act as a constraint on women's labour market opportunities in the southern USA during the first half of the 20th century. If anything, mineral wealth likely had a small positive effect on women's involvement in the labour market. In this context, further investigation of alternative explanations for the persistently low levels of female labour participation observed in the southern USA (and in the USA more broadly) during this period is warranted. In particular, a promising area for further research concerns the role played by social norms and institutions in shaping both labour market outcomes and the response of labour markets to economic shocks.

ACKNOWLEDGMENTS

We thank Guy Michaels, Steve Pischke, Giacomo Ponzetto, Esteban Aucejo, Ghazala Azmat, Michel Azulai, Nitika Bagaria, Sascha Becker, Dan Black, Florian Blum, Zelda Brutti, Francesco Caselli, Vasco Carvalho, Alejandro Cuñat, Marta De Philippis, Jeremiah Dittmar, Richard Hornbeck, Kilian Huber, Yu-Hsiang Lei, Alan Manning, Andrea Matranga, Rachel Ngai, Raphaël Parchet, Frank Pisch, Roland Rathelot, Ferdinand Rauch, Pedro Souza, Andrea Tesei, Jaume Ventura, Alwyn Young, and seminar participants at the London School of Economics, University of Oxford, Institut d'Economia de Barcelona, Halle Institute for Economic Research, Ludwig Maximilian University of Munich, University of Konstanz, the 2014 EDP Jamboree, the 15th Trans-Atlantic Doctoral Conference, the 2016 SAEe Simposio, the 2017 VfS Development Group Meeting, and the 2017 VfS Annual Meeting for their helpful comments and suggestions. We are particularly indebted to Guy Michaels and Paul W. Rhode for sharing their datasets with us. All remaining errors are our own.

Potlogea acknowledges financial support from the Spanish ministry of Economy and Competitiveness, grant number ECO2011-28965, scholarship reference BES-2012-060112.; Maurer thanks the LSE for its PhD Student Scholarship. The authors declare that they have no conflict of interest. A previous version of this paper was circulated under the title 'Fueling the gender gap? Oil and women's labour and marriage market outcomes'.

NOTES

1. As we want to focus our discussion on labour market linkages across sectors, we abstract from other intersectoral linkages, such as input–output linkages.
2. In the previous section, as we were ignoring migration, we did not make a distinction between individual female labour supply and aggregate female labour supply, as in a simplified setting with homogeneous female labour they move in lockstep.

3. One potential issue is that especially in the western part of the sample, populations were growing and counties therefore often changed, were split up, or were newly created. In order to address this, we compare the area of each county in each census decade to its area in 1940, and drop all observations from a given county-year cell if the absolute difference of the county's area in that year compared to 1940 exceeds 5%. In addition, we drop all counties from Oklahoma in 1900, as this state at this point was still largely unorganized and divided into the Oklahoma and Indian Territory. From the panel specifications, we also drop all counties that have only one observation, as they do not contribute anything to the panel results.
4. Starting from 1940, the US census changed its definition of labour force participation. Before 1940, being part of the labour force was defined based on whether individuals reported an occupation and were thus considered 'gainful workers'. In 1940, this was changed to a more modern definition that counted people as belonging to the labour force if they were employed or unemployed but actively looking for work. Goldin (1986) discusses some implications of these different definitions for the measurement of female labour force participation. She notes that there is no clear theoretical relationship between the two concepts. In our case, we will always control for census decade fixed effects, which should capture uniform changes in measurement.
5. When aggregating, we weight individual observations by the person weights provided by the IPUMS in order to improve the representativeness of our sample.
6. In Table B19 of Online Appendix B, we report results for specifications in which we drop all individuals working in the oil extraction sector. While the wage growth result weakens somewhat, it stays economically and statistically significant: outside the oil industry, men gain 22 log points in wages. In the same table, we also show results for the 1940 occupational earnings scores, and for the 1940 wage distribution percentiles.
7. Given the sizeable immigration in the wake of oil discoveries, another concern is to what extent our wage results are driven by spatial sorting. In Online Appendix B, we use the birth states of individuals to identify long-distance migrants. We find that for men, selective migration does indeed play a role, but cannot by itself fully account for our male wage findings.
8. This mechanism linking oil wealth to female employment growth in the non-tradable sector is similar to the link between high-skilled labour demand and the demand for low-skill-intensive non-tradables documented by Manning (2004) and Mazzolari and Ragusa (2013).
9. For example, Fogel (1962) states that dairy products were rarely traded inter-regionally.
10. Input-output linkages from the oil sector may also play a role in driving growth in the demand for female labour. However, given the pattern of our results, where female employment increases most slowly in manufacturing and most rapidly in services, this channel is unlikely to be first-order.
11. Note, however, that women's migratory response is large, which in turn may account for the small effects that we find for labour force participation and wages. This migratory response for women is, however, also difficult to interpret, as it is not clear whether it is a response to a labour market shock or a marriage market shock.
12. Generalizing our findings even to the entirety of the USA is not straightforward, given that the South has traditionally tended to display more conservative social norms on a series of issues, including gender (Rice and Coates 1995; Twenge 1997).

REFERENCES

- ACEMOGLU, D., AUTOR, D. H. and LYLE, D. (2004). Women, war, and wages: the effect of female labour supply on the wage structure at midcentury. *Journal of Political Economy*, **112**, 497–551.
- AGER, P., BRUECKNER, M. and HERZ, B. (2017). The boll weevil plague and its effect on the southern agricultural sector, 1889–1929. *Explorations in Economic History*, **65**, 94–105.
- AKBULUT, R. (2011). Sectoral changes and the increase in women's labour force participation. *Macroeconomic Dynamics*, **15**, 240–64.
- ALBANESI, S. and OLIVETTI, C. (2016). Gender roles and medical progress. *Journal of Political Economy*, **124**, 650–95.
- ALESINA, A. F., GIULIANO, P. and NUNN, N. (2013). On the origins of gender roles: women and the plough. *Quarterly Journal of Economics*, **128**, 469–530.
- ALLCOTT, H. and KENISTON, D. (2014). Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. NBER Working Paper no. 20508.
- ARAGON, F. M. and RUD, J. P. (2013). Natural resources and local communities: evidence from a Peruvian gold mine. *American Economic Journal: Economic Policy*, **5**, 1–25.
- , ——— and TOEWS, G. (2016). Resource shocks, employment, and gender: evidence from the collapse of the UK coal industry. Working Paper.
- ASHENFELTER, O. and HECKMAN, J. (1974). The estimation of income and substitution effects in a model of family labour supply. *Econometrica*, **42**, 73–85.
- ATTANASIO, O., LOW, H. and SANCHEZ-MARCOS, V. (2008). Explaining changes in female labour supply in a life-cycle model. *American Economic Review*, **98**, 1517–42.

- BECKER, G. S. (1981). *A Treatise on the Family*. Cambridge, MA: Harvard University Press.
- BERTRAND, M., DUFLO, E. and MULLAINATHAN, S. (2004). How much should we trust difference-in-differences estimates? *Quarterly Journal of Economics*, **119**, 249–75
- BESTER, C. A., CONLEY, T. G. and HANSEN, C. B. (2011). Inference with dependent data using cluster covariance estimators. *Journal of Econometrics*, **165**, 137–51.
- BLACK, D. A., MCKINNISH, T. and SANDERS, S. G. (2005). The economic impact of the coal boom and bust. *Economic Journal*, **115**, 449–76.
- BUCKLEY, P. D. (2003). Essays in applied microeconomics: signaling, marriage hazard rates, and ability sorting. PhD dissertation, Department of Economics, Massachusetts Institute of Technology.
- CARRINGTON, W. J. (1996). The Alaskan Labour market during the pipeline era. *Journal of Political Economy*, **104**, 186–218.
- CASELLI, F. and COLEMAN, J. W. (2001). The U.S. structural transformation and regional convergence: a reinterpretation. *Journal of Political Economy*, **109**, 584–616.
- CORDEN, M. W. and NEARY, P. J. (1982). Booming sector and de-industrialisation in a small open economy. *Economic Journal*, **92**, 825–48.
- DO, Q.-T., LEVCHENKO, A. A. and RADDATZ, C. (2016). Comparative advantage, international trade, and fertility. *Journal of Development Economics*, **119**, 48–66.
- FERNANDEZ, R. (2011). Does culture matter? In J. Benhabib, M. O. Jackson and A. Bisin (eds), *Handbook of Social Economics*, Vol. 1A. Amsterdam: North-Holland.
- (2013). Cultural change as learning: the evolution of female labour force participation over a century. *American Economic Review*, **103**, 472–500.
- FETZER, T. R. (2014). Fracking growth. CEP Discussion Paper no. 1278, London School of Economics.
- FOGEL, R. W. (1962). A quantitative approach to the study of railroads in American economic growth: a report of some preliminary findings. *Journal of Economic History*, **22**, 163–97.
- FURCHTGOTT-ROTH, D. and GRAY, A. (2013). The economic effects of hydrofracturing on local economies: a comparison of New York and Pennsylvania. Growth and Prosperity Report no. 1, Manhattan Institute for Policy Research, New York.
- GADDIS, I. and PIETERS, J. (2017). The gendered labour market impacts of trade liberalization: evidence from Brazil. *Journal of Human Resources*, **52**, 457–89.
- GOLDIN, C. (1977). Female labour force participation: the origin of black and white differences, 1870 and 1880. *Journal of Economic History*, **37**, 87–108.
- (1986). The female labor force and American economic growth, 1890–1980. In S. L. Engerman and R. E. Gallman (eds), *Long-term Factors in American Economic Growth*, Chicago, IL: University of Chicago Press.
- (1990). *Understanding the Gender Wage Gap: An Economic History of American Women*. New York: Oxford University Press.
- (2006). The quiet revolution that transformed women’s employment, education, and family. *American Economic Review: Papers and Proceedings*, **96**, 1–21.
- and KATZ, L. F. (2002). The power of the pill: oral contraceptives and women’s career and marriage decisions. *Journal of Political Economy*, **110**, 730–70.
- GREENWOOD, J., SESHADRI, A. and YORUKOGLU, M. (2005). Engines of liberation. *Review of Economic Studies*, **72**, 109–33.
- HUANG, C. (2018). Why are US women decreasing their labour force participation if their wages are rising? *Economic Inquiry*, **56**, 2010–26.
- JELNOV, P. (2016). The marriage age U-shape: what can we learn from energy crises? Working Paper.
- JONES, L. E., MANUELLI, R. E. and MCGRATTAN, E. R. (2015). Why are married women working so much? *Journal of Demographic Economics*, **81**, 75–114.
- KANG, A. (2009). Studying oil, Islam, and women as if political institutions mattered. *Politics & Gender*, **5**, 560–8.
- KLASSEN, S. and PIETERS, J. (2015). What explains the stagnation of female labour force participation in urban India? *World Bank Economic Review*, **29**, 449–78.
- KOTSADAM, A. and TOLONEN, A. (2016). African mining, gender, and local employment. *World Development*, **83**, 325–39.
- LANGE, F., OLMSTEAD, A. L. and RHODE, P. W. (2009). The impact of the boll weevil, 1892–1932. *Journal of Economic History*, **69**, 685–717.
- MANNING, A. (2004). We can work it out: the impact of technological change on the demand for low-skilled workers. *Scottish Journal of Political Economy*, **51**, 581–608
- MARCHAND, J. (2012). Local labour market impacts of energy boom-bust-boom in western Canada. *Journal of Urban Economics*, **71**, 165–74.

- MAZZOLARI, F. and RAGUSA, G. (2013). Spillovers from high-skill consumption to low-skill labour markets. *Review of Economics and Statistics*, **95**, 74–86.
- MICHAELS, G. (2011). The long term consequences of resource-based specialisation. *Economic Journal*, **121**, 31–57.
- MINNESOTA POPULATION CENTER (2016). National Historical Geographic Information System: version 11.0 [database]. Minneapolis, MN: University of Minnesota; available online at <http://doi.org/10.18128/D050.V11.0> (accessed 28 February 2020).
- NGAI, R. L. and PETRONGOLO, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, **9**, 1–44.
- NORRIS, P. (2010). Perhaps petroleum perpetuates patriarchy? A response and critique to Ross. Unpublished manuscript, Harvard Kennedy School of Government.
- OLIVETTI, C. (2013). The female labour force and long-run development: the American experience in comparative perspective. NBER Working Paper no. 19131.
- PAPYRAKIS, E. and GERLAGH, R. (2007). Resource abundance and economic growth in the United States. *European Economic Review*, **51**, 1011–39.
- PEI, Z., PISCHKE, J.-S. and SCHWANDT, H. (2019). Poorly measured confounders are more useful on the left than on the right. *Journal of Business and Economic Statistics*, **37**, 205–16.
- PRATT, J. A. (1980). *The Growth of a Refining Region*. Greenwich, CT: Jai Press.
- RICE, T. W. and COATES, D. L. (1995). Gender role attitudes in the southern United States. *Gender and Society*, **9**, 738–50.
- ROSS, M. L. (2008). Oil, Islam and women. *American Political Science Review*, **102**, 107–23.
- (2012). *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations*. Princeton, NJ: Princeton University Press.
- RUGGLES, S., GENADEK, K., GOEKEN, R., GROVER, J. and SOBEK, M. (2015). Integrated Public Use Microdata Series: version 6.0 [dataset]. Minneapolis, MN: University of Minnesota; available online at <http://doi.org/10.18128/D010.V6.0> (accessed 28 February 2020).
- SACHS, J. D. and WARNER, A. M. (1999). The big rush, natural resource booms and growth. *Journal of Development Economics*, **59**, 43–76.
- and ——— (2001). The curse of natural resources. *European Economic Review*, **45**, 827–38.
- SIMMONS, J. W. (2016). Natural resource wealth and women's economic and political power in the U.S. states. *Comparative Political Studies*, **49**, 115–52.
- TWENGE, J. M. (1997). Attitudes toward women, 1970–1995: a meta analysis. *Psychology of Women Quarterly*, **21**, 35–51.
- VOIGTLÄNDER, N. and VOTH, J. (2013). How the West 'invented' fertility restriction. *American Economic Review*, **103**, 2227–64.
- WRIGHT, G. (1986). *Old South, New South*. New York: Basic Books.
- YOUNG, A. (2016). Improved, nearly exact, statistical inference with robust and clustered covariance matrices using effective degrees of freedom corrections. Working Paper.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

- A. Theoretical framework
- B. Further results and robustness checks
- C. Variable definitions and data sources