Quarterly Journal of Political Science, 2012, 7: 69–87

Comment

Oil, Islam, Women, and Geography: A Comment on Ross (2008)*

Matthew Groh[1] and Casey Rothschild[2]

[1] *Innovations for Poverty Action; matthew.groh@gmail.com*
[2] *Wellesley College; crothsch@wellesley.edu*

ABSTRACT

In ''Oil, Islam, and Women,'' Michael Ross (2008a) develops a gendered Dutch Disease theory, which points to oil wealth as a potential explanation for the slow progress towards gender equality in the Middle East. He then presents empirical analysis in support of this theory and concludes that ''women in the Middle East are underrepresented in the workforce and in government because of oil — not Islam'' (p. 107). This brief comment re-examines Ross's data and finds that they do not justify his conclusion: upon closer examination, his data do not provide evidence that oil rents causally affect female labor force participation rates via the gendered Dutch Disease. We argue that, in fact, his data are as or more consistent with Islam playing an important role in explaining the lagging female labor force participation rates than they are with oil playing an important role.

* The authors thank Alexandre Debs and participants in the Middlebury College Senior Honors Workshop for helpful comments, and Middlebury College for partial funding. The views expressed herein are those of the authors and do not necessarily reflect the views of Innovations for Poverty Action.

Supplementary Material available from:
http://dx.doi.org/10.1561/100.00011036_supp
MS submitted 19 April 2011; *final version received 1 November 2011*
ISSN 1554-0626; DOI 10.1561/100.00011036

In ''Oil, Islam, and Women,'' Michael Ross suggests that ''oil, not Islam, is at fault'' for women's slow progress towards gender equality in the Middle East[1]. According to the *gendered Dutch Disease* (GDD) model he proposes as the underlying causal mechanism, natural resource booms lead to an appreciating currency. This induces a contraction of textile and other industries that provide the typical entrée for women into the labor force, and this, in turn, impedes progress towards greater female political influence. Ross buttresses this theoretically novel and compelling case with a regression-based empirical analysis of data from the World Bank to conclude that ''petroleum perpetuates patriarchy,'' and Islam does not.[2]. This comment re-examines Ross's empirical analysis and finds that it does not, in fact, support this conclusion.

In light of the Arab Spring and the associated ''unparalleled opportunit[ies] to incorporate a broader interpretation of women's rights'' in newly drafted constitutions (Economist, 2011), it is particularly important to understand the deep reasons for slow progress towards gender equality in the region. Ross's article has been an influential contribution to a broader debate about these reasons (Sharabi, 1988; Landes and Landes, 2001; World Bank, 2004; Inglehart and Norris, 2003a). It received the 2009 Heinz Eulau award for the best article in the American Political Science Review and spawned a substantial follow-up literature, much of it critical. Some of the criticism has questioned Ross's underlying theory (Norris, 2009; Charrad, 2009); some of it has broadly accepted Ross's conclusion that oil matters, but critiqued his conclusion that oil is *all* that matters or that oil is what matters *the most* (Adida *et al.*, 2011; Alexander and Welzel, 2011; Gorman, 2009; Ingvaldsen, 2010; Price, 2011, World Bank, 2011). By and large, the literature has left the impression that Ross's basic empirical findings are robust. This comment formally critiques Ross's empirical methods and results to dispel that impression.

Ross's empirical case for ''the main implication of [his] model... [that]... *A rise in the value of oil production will reduce female participation in the labor force*''[3] (Ross's italics) is built on two sets of regressions: a between estimator based on cross-sectional regressions using time averaged data, and a set of fixed-effect, first-differenced regressions based on panel data. Ross finds statistically significant effects of Oil Rents Per Capita (henceforth

---

[1] Ross (2008a, p. 107).
[2] Ross (2008a, p. 120).
[3] Ross (2008a, p. 110).

''Oil Rents'') on Female Labor Force Participation in both econometric models. Our critique of his empirical evidence is two-fold. First, we show that the significant coefficient on Oil Rents in Ross's between regressions appear to be driven entirely by inter-regional differences omitted from his empirical analysis. A closer examination of Ross's data appears to instead suggest something about the Arabian Peninsula *other* than oil — possibly historically-driven proclivities toward religious or cultural conservatism — are driving his results.

Second, we argue that Ross's time-series (first-differenced fixed effects) regressions are simply not well suited to test his hypothesis that oil perpetuates patriarchy: they exploit short-run intra-country variation to identify what is, at heart, a longer run inter-country mechanism. Unsurprisingly, this means that these estimates are not robust to plausible modifications. Moreover, we show using supplemental regressions on real exchange rates that whatever is driving the apparent significance of oil in these regressions, it does not appear to be Ross's hypothesized causal mechanism — the GDD.

Cross-National Regressions

Ross's central evidence that oil, not Islam, is the driver of persistent gender inequalities in the Middle East is a set of coefficients from a series of cross-country between regressions of Female Labor Force Participation rates on Oil Rents using country-level variables time-averaged over the 1993–2002 period. Ross also includes, as explanatory variables, log income (and its square), proportion of the population of working age, a Middle East and North Africa (MENA) region dummy,[4] a Communist dummy (for countries with a communist legal system at any point post-1960), and an Islam variable which measures the normalized fraction of a given country's population which is Muslim.

Ross's baseline results[5] appear in Column (1) of Table 1. The Oil Rents coefficient is statistically significant, from which he concludes that ''[H]igher oil rents are linked to lower rates of female labor force participation'' (p. 115). We believe that Ross's data do not support this conclusion. Figures 1a and 1b illustrate the crux of our argument. The remaining columns of Table 1 formalize it.

---

[4] Algeria, Bahrain, Egypt, Iran, Iraq, Jordan, Kuwait, Lebanon, Libya, Morocco, Oman, Qatar, Saudi Arabia, Syria, Tunisia, United Arab Emirates, and Yemen.

[5] Ross (2008a, p. 114, Table 2, Column 4).

Table 1. Cross-national regressions on female labor force with regional effects.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Income	−1.864	−1.979	−2.140	−2.137	−2.454	−2.449	−2.483
(log)	(0.878)	(0.892)	(0.886)	(0.897)	(0.831)	(0.838)	(0.845)
Income	2.122	2.197	2.410	2.419	2.700	2.698	2.723
squared (log)	(0.824)	(0.832)	(0.831)	(0.837)	(0.775)	(0.778)	(0.781)
Working	−0.350	−0.317	−0.366	−0.376	−0.364	−0.367	−0.362
age	(0.142)	(0.144)	(0.132)	(0.137)	(0.133)	(0.138)	(0.146)
Oil Rents	−0.210	0.009	−0.014	−0.009	−0.014	−0.028	−0.025
per capita	(0.055)	(0.064)	(0.036)	(0.061)	(0.035)	(0.057)	(0.057)
MENA	−0.326	−0.297					
	(0.117)	(0.119)					
MENA		−0.272					
Interaction		(0.100)					
Rest			−0.525	−0.524			
of MENA			(0.336)	(0.335)			
Rest of MENA				−0.550			
interaction				(0.291)			
Peninsula			−2.433	−2.479	−2.232	−2.255	
			(0.321)	(0.379)	(0.273)	(0.326)	
Peninsula				0.006		0.020	
interaction				(0.080)		(0.078)	
Islam	−0.139	−0.159	−0.154	−0.150	−0.232	−0.231	−0.232
	(0.116)	(0.117)	(0.114)	(0.117)	(0.081)	(0.082)	(0.082)
Communist	0.286	0.276	0.304	0.309	0.319	0.320	0.318
bloc	(0.104)	(0.104)	(0.102)	(0.104)	(0.100)	(0.101)	(0.103)
Constant	−0.012	0.026	0.125	0.127	0.083	0.081	0.082
	(0.060)	(0.061)	(0.068)	(0.066)	(0.060)	(0.058)	(0.058)
Number of	167	167	167	167	167	167	160
observations							

Note: Dependent variable is female nonagricultural labor force participation, 1993–2002. Standard errors in parentheses. Column 1 replicates Column 4 from Table 2 of Ross (2008a). All variables are standardized. Column (7) drops countries in the Arabian Peninsula.

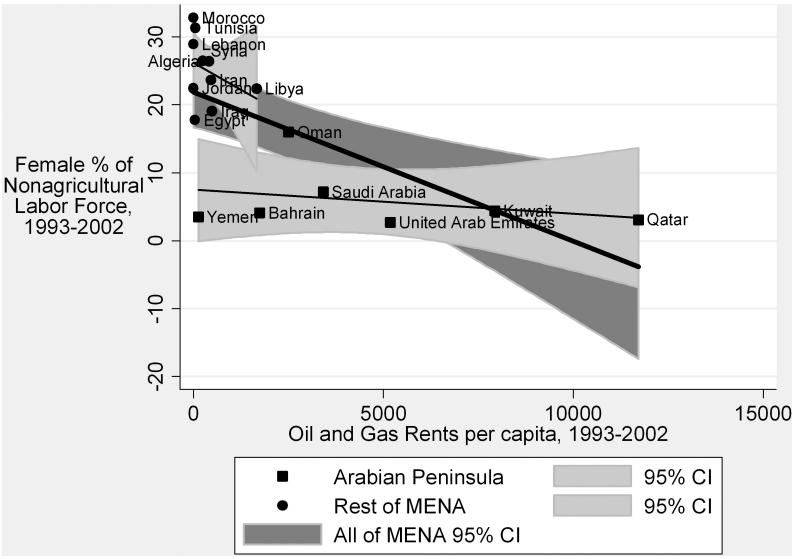

Figure 1a. Oil Rents and Female Labor Force Participation in the middle east.
Notes: Data from Ross (2008a).

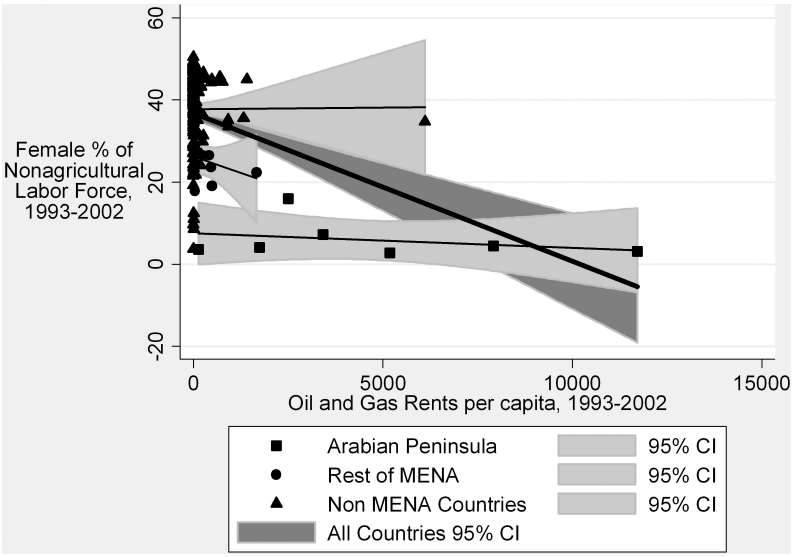

Figure 1b. Oil Rents and Female Labor Force Participation Worldwide.
Notes: Data from Ross (2008a).

Figure 1a re-creates Ross's Figure 3, a scatter plot of average per-capita oil and gas rents versus average female labor force participation over the 1993–2002 period in the MENA region. This scatter plot reveals a clear negative correlation between the two variables, providing a clean qualitative depiction of Ross's ''Oil Rents cause low Female Labor Force Participation'' results. Figure 1a superimposes on this plot *three* best-fit lines and their associated confidence intervals: one for all the data points, one for countries on the Arabian Peninsula, and one for the remaining MENA countries. The latter two confidence intervals are consistent with horizontal lines: there is no evidence of a robust correlation between Oil Rents and Female Labor Force Participation on the Arabian Peninsula. There is only weak evidence of an effect in the rest of MENA. Comparing the confidence intervals for the two sub-regions with the line of best fit for the entire dataset indicates that the apparent overall downward slope is driven primarily by differences *across* regions, not within regions. In other words, the apparent downward slope is a result of the fact that Oil Rents and Female Labor Force Participation are correlated (in opposite directions) with geographical region. Figure 1b contains an analogous plot for the entire dataset, breaking it up into three regions: the Arabian Peninsula, the rest of MENA, and the rest of the world. Again, the qualitative downward slope appears to be driven entirely by *inter*-regional effects; there is no robust evidence of any *intra*-regional effects of oil.

Table 1 formally establishes that this graphical intuition extends to Ross's between regressions. Ross's key takeaways from the regression we replicated in column (1) are (i) that the oil coefficient is negative and significant and (ii) that the coefficient on Islam is statistically indistinguishable from zero. We test the robustness of these conclusions in the other six columns. Column (2) adds an oil–MENA interaction. This allows the oil rent effect (on Female Labor Force Participation) to vary by region. The direct Oil Rent effect is insignificant, while the MENA–oil interaction is negative and statistically significant. This suggests that the significance of oil rents in Ross's regression (Column (1)) is being driven entirely by something *within* MENA.

Columns (3) and (4) probe further by breaking MENA down into the Arabian Peninsula[6] and the rest of MENA. Column (3) includes regional effects but not oil–region interactions, and Column (4) includes both. They

[6] Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, United Arab Emirates, and Yemen.

both tell the same story: the significance of Oil Rents disappears (and the coefficient drops in size dramatically relative to Column (1)), and the coefficient on the Peninsula dummy is negative and statistically significant. In Column (4), the coefficient on the Rest of MENA–Oil interaction is negative and marginally significant, but a joint test for the significance of the three oil terms shows them to be jointly insignificant ($p = 0.16$). In short, the apparent oil effects in Column (1) appear to really be an Arabian Peninsula effect, just as Figure 1b suggests.

Since the Rest of MENA specific terms in Columns (3) and (4) are statistically indistinguishable from zero, we also considered specifications which contain only terms for the Peninsula; these are reported in Columns (5) and (6). These yield the same basic conclusions: the Oil Rents coefficient is small and statistically indistinguishable from zero. More interestingly, in these specifications the coefficient on Islam is statistically significantly negative. Column (7) runs the same regression as Column (6), but excludes the Arabian Peninsula countries. Unsurprisingly — given that most of the variation in Ross's Islam variable occurs outside of the Arabian Peninsula (the countries of which have uniformly high percentage Muslim populations) — this exclusion does not materially change the coefficient estimates.

A closer analysis of Ross's between data and regressions therefore lead us to reach a conclusion nearly opposite to his: we find little evidence that oil matters, *per se*, in driving female labor force participation and some mild evidence that Islam does.

Fixed-effect First-differenced Regressions

Ross's other significant piece of empirical evidence that oil rents retard progress towards gender equality is based on the following fixed-effects, first-differenced regression model.

$$\Delta Y_{i,t} = \alpha_i + \beta \Delta x_{i,t-1} + \eta_{i,t}. \tag{1}$$

In Equation (1), i and t index country and year, respectively, and Δ denotes first (time) differences, so that for any variable $z_{i,t}$, $\Delta z_{i,t} \equiv z_{i,t} - z_{i,t-1}$. Y is the Female Labor Force Participation rate, x is the time-varying subset of explanatory variables used in the between regressions discussed above, and α_i are country-specific fixed effects. The error terms $\eta_{i,t}$ follow an AR(1)

process.[7] Ross's evidence is based on the significance of the Oil Rents coefficient in various specifications of Equation (1).

We offer two critiques of this evidence. First, we argue that the empirical specification of Equation (1) is poorly adapted to testing Ross's basic hypothesis, so that even robustly significant Oil Rent coefficients should not be interpreted as providing support for Ross's hypothesis. Second, we argue that the coefficients are not robustly significant: as in the preceding analysis of Ross's between regressions, the coefficient is sensitive to the inclusion of regional effects.

A Mismatch between Theory and Empirical Specification

Ross's GDD theory suggests that countries which experience sustained oil booms can expect to have relatively stagnant female labor force participation rates. As discussed in Frankel (2010), we would expect this relationship to be relatively slow-moving, and we would expect it to be driven mainly by differences in Oil Rent levels or in long-term Oil Rent trends across (otherwise similar) countries. But the coefficient on Oil Rents in estimates of Equation (1) cannot be driven by these differences: Oil Rent *levels* are differenced out, and across-country differences in Oil Rent *trends* are absorbed by the fixed effects α_i. (Because of first differencing, the α_i are country specific linear time-trends in Equation (1).)

Instead, the coefficients on Oil Rents in Equation (1) are identified off of short term differences *from* country-specific trends in Oil Rent growth rates. That is, a significant negative coefficient on Oil Rents in estimates of Equation (1) indicates that years in which Oil Rents grew faster than usual for a given country tended to be immediately followed by years in which Female Labor Force Participation rates grew more slowly than usual for that country. This sort of short-run-differences-from-trend variation — although perhaps interesting in its own right — does not provide a good test of the underlying GDD theory.

Table 2 replicates Ross's table of estimates of Equation (1).[8] Column (2) is the baseline regression on his full 169 country data set. Columns (3) and (4) are robustness checks: column (3) drops the two most influential countries

---

[7] In Ross's notation $\eta_{i,t} = \varepsilon_{i,t} - \varepsilon_{i,t-1}$. We use this alternative notation to clarify precisely *what* follows the AR(1) process in his specification: it is $\eta_{i,t}$, not $\varepsilon_{i,t}$. We determined this by exactly replicating his regression results using his data.

[8] Ross (2008a, p. 113 Table 1).

Table 2. Pooled time-series cross-national regressions, with first differences and fixed effects (from Ross 2008a, Table 1).

	(1)	(2)	(3)	(4)
ΔIncome (log)	−0.011	−0.039	−0.014	−0.051
	(0.032)	(0.033)	(0.027)	(0.047)
ΔIncome squared (log)	0.017	0.049	0.021	0.021
	(0.033)	(0.033)	(0.028)	(0.048)
ΔWorking Age	0.115	0.115	0.066	0.177
	(0.025)	(0.025)	(0.024)	(0.013)
ΔOil Rents		−0.026	−0.017	−0.049
		(0.006)	(0.007)	(0.011)
Constant	0.034	0.033	0.010	0.154
	(0.007)	(0.007)	(0.005)	(0.069)
Number of observations	5234	5234	5168	5395

Note: Dependent Variable is Female Labor Force Participation, 1960–2002. Standard errors in parentheses. Country fixed effects are used in each estimation. In column 3, the two most influential countries have been dropped from the sample. In column 4, year dummies were included in place of the AR(1) process.

(Saudi Arabia and Kuwait), and column (4) uses time fixed effects instead of the AR(1) error process. Column (1) is the same as Column (2), *sans* the Oil Rents regressor. The coefficients on Oil Rents are negative and statistically significant in each specification.

Ross interprets this significance as evidence supporting his GDD theory. The preceding discussion suggests that the significance of the Oil Rent coefficient is actually driven by something else. This is testable. According to Ross's GDD theory, the connection between Oil Rents and Female Labor Force Participation is intermediated by real exchange rates: movements in Oil Rents drive (real) currency appreciation, and this appreciation then crowds out industries such as textiles through which females would otherwise have entered the labor force. Insofar as the significant coefficient in Ross's within regressions is driven by the GDD, we should see robust evidence of the intermediating influence of real exchange rates. Moreover, this evidence should be present in regressions which are based on the same basic identifying variation that Ross's regressions employ. We therefore ran

Table 3a. Intermediate steps — the effects of oil rents on real exchange rates.

	(1)	(2)	(3)	(4)	(5)
ΔIncome (log)	−2.618	−2.846	−3.451	−3.917	
	(6.293)	(6.391)	(6.719)	(6.382)	
ΔIncome squared (log)	3.555	3.822	4.524	6.155	
	(6.346)	(6.479)	(6.900)	(6.474)	
ΔWorking age	−0.561	−0.556	−0.560	0.820	
	(1.910)	(1.910)	(1.968)	(1.981)	
ΔOil rents		−0.279	−0.534	−0.229	−0.426
		(1.364)	(1.910)	(1.377)	(1.463)
Constant	8.708	8.706	8.737	0.791	7.813
	(1.323)	(1.323)	(1.341)	(8.449)	(1.365)
Number of observations	4116	4116	4056	4116	5566

Note: Dependent Variable is Δ Real Exchange Rate, 1970–2006. Real exchange rate data from USDA (http://www.ers.usda.gov/Data/Macroeconomics/). Standard errors in parentheses. Country fixed effects are used in each estimation. In column 3, Kuwait and Saudi Arabia have been dropped from the sample. In column 4, year dummies were included in place of the AR(1) process.

first-differenced fixed-effect regressions of Real Exchange Rates on Oil Rents and of Female Labor Force Participation Rates on Real Exchange Rates.[9] Tables 3a and b present the results of these two regressions. Columns (1)–(4) in each table mirror the four specifications presented in Table 1. Column (5) of each table drops the covariates to focus on the key coefficients.

The key coefficients in Tables 3a and 3b — on Real Exchange Rates and on Female Labor Force Participation, respectively — are small and statistically indistinguishable from zero in all five specifications. There is no evidence that Oil Rents are driving Real Exchange Rates, nor is there evidence that Real Exchange Rates are driving Female Labor Force Participation rates. More precisely, there is no evidence that short-run within-country differences from country-specific trends in Oil Rents are driving

[9] We used real exchange rates with the U.S. dollar provided by the U.S. Department of Agriculture (http://www.ers.usda.gov/Data/Macroeconomics/). It includes real exchange rates for 190 countries between 1970 (post Bretton-Woods) and 2010.

Table 3b. Intermediate steps — the effects of exchange rates on female labor force participation.

	(1)	(2)	(3)	(4)	(5)
ΔIncome (log)	−0.061	−0.077	−0.102	−0.086	
	(0.050)	(0.052)	(0.051)	(0.052)	
ΔIncome squared (log)	0.042	0.056	0.085	0.061	
	(0.051)	(0.052)	(0.052)	(0.052)	
ΔWorking age	0.184	0.182	0.180	0.221	
	(0.015)	(0.016)	(0.015)	(0.016)	
ΔReal exchange		0.000	0.000	0.000	0.000
		(0.000)	(0.000)	(0.000)	(0.000)
Constant	0.019	0.030	0.010	0.303	0.020
	(0.010)	(0.011)	(0.010)	(0.070)	(0.010)
Number of observations	4487	4101	4041	4101	5023

Note: Dependent Variable is Δ Female Labor Force Participation, 1970–2006. Real exchange rate data from USDA (http://www.ers.usda.gov/Data/Macroeconomics/). Standard errors in parentheses. Country fixed effects are used in each estimation. In column 3, Kuwait and Saudi Arabia have been dropped from the sample. In column 4, year dummies were included in place of the AR(1) process.

short-run within-country differences from country-specific trends in Female Labor Force Participation Rates. It therefore appears unlikely that the GDD is the causal mechanism for the significant Oil Rents coefficients in this set of Ross's regressions.[10]

Sensitivity

Given the highly-specialized variation that identifies the Oil Rents coefficient in Equation (1), it is worth briefly exploring the robustness of those coefficient estimates.

The systematic differences between countries on the Arabian Peninsula, countries in the rest of MENA, and countries in the rest of the world are

[10] This critique of Ross's results is consistent with the broader empirical literature on the Dutch Disease. Magud and Sosa's (2010) meta-analysis of this literature indicates inconsistent empirical evidence of robust correlations between natural resource shocks and real currency appreciation.

apparent in the cross-sectional data reported Figure 1b and Table 1. We first test whether the inclusion of regional differences impacts the coefficients on Oil Rents in the within regressions reported in Table 2. We use several specifications and report results in Table 4. Column (1) contains Ross's baseline specification with an AR(1) error process. Column (2) replicates this specification with White (1980) standard errors, which allow for arbitrary serial correlation and across-time heteroskedasticity in the errors (Arellano,

Table 4. Pooled time-series cross-national regressions with region–oil interactions (fixed effects).

	(1)	(2)	(3)	(4)	(5)
ΔIncome (log)	−0.039	−0.042	−0.042	−0.040	−0.040
	(0.033)	(0.044)	(0.044)	(0.045)	(0.045)
ΔIncome	0.049	0.020	0.019	0.016	0.016
squared (log)	(0.033)	(0.050)	(0.050)	(0.051)	(0.051)
ΔWorking	0.115	0.141	0.141	0.141	0.141
age	(0.025)	(0.037)	(0.037)	(0.037)	(0.037)
ΔOil rents	−0.026	−0.046			
	(0.006)	(0.011)			
Outside MENA			0.020	0.021	
interaction			(0.049)	(0.049)	
MENA			−0.047		
interaction			(0.011)		
Peninsula				−0.053	−0.053
interaction				(0.007)	(0.007)
Rest of MENA				0.121	
interaction				(0.021)	
Non-peninsula					0.089
interaction					(0.025)
Constant	0.033	0.048	0.048	0.048	0.049
	(0.007)	(0.001)	(0.001)	(0.001)	(0.001)
Number of observations	5234	5395	5395	5395	5395

Note: Dependent variable is Female Labor Force Participation, 1960–2000. Standard errors in parentheses. Robust standard errors except column 1 (AR(1) errors).

1987); it confirms that Ross's results are not sensitive to the rather specific error structure he employs. Using the same robust error structure, columns (3)–(5) modify the regressions to include the regional effects. Column (3) allows the effect of oil within MENA to differ from the effect outside of MENA. Column (4) allows for distinct effects of oil in each of the three regions discussed above (the Arabian Peninsula, the rest of MENA, and the non-MENA countries). Column (5) pools the latter two regions but allows for differences between the Peninsula and the rest of the world.

Statistically significant Oil Rents effects are present in all specifications, but the regional patterns suggest a more complicated effect than the baseline regressions in columns (1) and (2) indicate. Columns (4) and (5) show a statistically significantly negative effect of Oil Rents on Female Labor Force Participation on the Arabian Peninsula, but not elsewhere. Indeed, in Column (5) the effects outside of the Peninsula are statistically significantly *positive*. An F-test easily rejects the equality of the effects on the Peninsula and off the Peninsula ($p = 0.0001$) in Column (5).

Columns (1)–(5) of Table 5 report the results of the same set of regressions with random instead of fixed country effects. The random effects versions of the columns (3)–(5) regression also include regional dummies (the results are insensitive to dropping them). These regressions support all of the conclusions of the fixed effects regressions.

Random effects estimates gain efficiency relative to Ross's estimates by incorporating between-country as well as within-country variation in growth rate trends. That is, if countries with higher Oil Rents growth rate trends also tend to be countries with slower Female Labor Force Participation growth rate trends, these random effects estimates will pick this up in the coefficient on Oil Rents, while in fixed effects estimates they would be absorbed in the country-specific fixed trends α_i. Ross discusses reasons to be concerned by identifying using this type of variation.[11] Hausman (1978) provides a formal way to test these concerns, and in all five cases, the random effects regressions "pass" this test: this indicates that *if* the fixed effects estimates are valid, then there is no reason to reject the validity of the random effects regressions. (See the bottom row of Table 5 for the results of the tests.) Indeed, a comparison of Ross's baseline fixed effects regression and the random effects analog in Column (1) of Tables 3 and 4, respectively, reveals virtually identical coefficient estimates.

---

[11] Ross (2008a, p. 112).

Table 5. Pooled time-series cross-national regressions with region–oil interactions (random effects).

	(1)	(2)	(3)	(4)	(5)	(6)
ΔIncome	−0.049	−0.058	−0.061	−0.061	−0.060	−0.058
(log)	(0.032)	(0.040)	(0.040)	(0.039)	(0.039)	(0.039)
ΔIncome	0.059	0.037	0.041	0.039	0.038	0.036
squared (log)	(0.033)	(0.045)	(0.045)	(0.044)	(0.044)	(0.044)
ΔWorking	0.114	0.146	0.145	0.145	0.145	0.145
age	(0.021)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
ΔOil	−0.028	−0.047				
rents	(0.006)	(0.013)				
MENA			−0.048			
Interaction			(0.013)			
Outside MENA			0.020	0.021		
interaction			(0.088)	(0.088)		
Outside			−0.531	−0.325		
MENA			(0.219)	(0.209)		
Peninsula				−0.055	−0.055	−0.054
interaction				(0.013)	(0.013)	(0.013)
Peninsula				0.521	0.826	1.052
				(0.492)	(0.451)	(0.460)
Rest of MENA				0.118		
interaction				(0.027)		
Outside					0.087	0.087
peninsula					(0.034)	(0.034)
interaction						
Islam						−0.121
						(0.048)
Constant	−0.030	−0.032	0.451	0.245	−0.061	−0.075
	(0.055)	(0.046)	(0.215)	(0.204)	(0.043)	(0.043)
Number of	5395	5395	5395	5395	5395	5395
observations						
Hausman test	$p = 0.37$	$p = 0.52$	$p = 0.48$	$p = 0.42$	$p = 0.35$	
p-value						
χ^2(DF)	$\chi^2(4) = 4.29$	$\chi^2(4) = 3.26$	$\chi^2(5) = 4.47$	$\chi^2(6) = 6.02$	$\chi^2(4) = 5.58$	
statistic						

Note: Dependent variable is Female Labor Force Participation, 1960–2000. Standard errors in parentheses. Robust standard errors except column (1) (AR(1) errors). Hausman test in columns (3)–(5) is based on the (non-reported) random effects regressions containing the exact same variables as in columns (3)–(5) of Table 3.

Beyond providing additional identifying variation, random effects estimates are useful because they allow us to incorporate time-invariant regressors. We have, accordingly, included regional dummies in Columns (3)–(5). We also included the time-invariant measure of Islam that Ross uses in his between regressions. Column (6) in Table 5 reports the results of

including this variable in the specification of Column (5). The coefficient on Islam is negative, significant, and larger in magnitude than the Oil Rents coefficient: Islam *does* seem to matter.

In summary, a more careful look at Ross's fixed effects, first differenced regressions reveals three things, each of which cuts against his conclusions. First, the evidence of an effect of Oil Rents on Female Labor Force Participation is significantly weaker than his baseline results suggest: evidence of an effect is specific to countries on the Arabian Peninsula. Second, the effect does not appear to be driven by the GDD, since there is no evidence supporting the intermediate causal steps underlying it. Third, there is evidence that ''Islam'' — even measured bluntly as percentage of Muslims in each country — *does* negatively affect female labor force participation.

Discussion

In a special section of *Politics & Gender* devoted to a discussion of Ross's 2008 article, Ross writes: ''My [2008] article suggests that oil wealth does a better job [than Islam] of explaining a) why the Middle East is different from other regions and b) why the status of women varies so dramatically among Middle Eastern countries.''[12] Our analysis of his data indicates nearly the opposite: oil does not appear to explain why the labor force participation rates of women varies so dramatically among countries, while Islam does appear to have some predictive power.

Ross's (2008a) statistical results appear to be driven largely by omitted regional differences, specifically differences between countries on the Arabian Peninsula and in the rest of the world. The importance of accounting for regional differences in explaining the relationship between oil and female empowerment has been noted before. For example, the World Bank (2011) — drawing on Do *et al.* (2011) analysis of oil wealth and female labor force participation — observes that: ''while oil has a dampening effect on female labor force participation on average across the world, rates of female labor force participation in MENA countries are well below what their oil endowments alone would imply.''[13] Following Rauch and Kostyshak (2009), who argue that ''from a socioeconomic point of view, the Arab world is too

[12] Ross (2009, p. 576).
[13] World Bank (2011, p. 9).

diverse to be a useful aggregate,''[14] we have used Ross's data to show that regional variation *within* the Middle East and North Africa is important as well. The literature points to a number of possible underlying explanations for these regional differences.

First, Charrad (2009) argues that historical patriarchal kinship networks are a causal factor behind lagging female empowerment in the Middle East. These networks, and the cultural norms associated with them (i) existed well before the discovery of oil, (ii) were particularly strong on the Arabian Peninsula. (Saudi Arabia, Kuwait, and the United Arab Emirates are Charrad's first three examples of countries with ''a long history of strongly patriarchal structures''[15]), and (iii) are likely, insofar as culture is strongly persistent, to be associated with lagging present-day female labor force participation.

By an accident of geography, oil rents (per capita) turned out to be particularly high on the Peninsula. This ''accidental'' historico-geographical correlation between recently discovered oil and the deep cultural history would explain the large scale pattern of Figure 1b and Ross's closely related regression results: countries with particularly high per-capita oil rents also tend to be countries with lagging female labor force participation rates. Under this telling, however, oil has no *causal* effect on these participation rates. Indeed, as the region-specific trend-lines in Figure 1b and our region-adjusted versions of Ross's regressions indicate, *within* regions with more homogenous cultural histories, contemporary oil rents appear to be minimally- or uncorrelated with female labor force participation.

Oversimplifying this argument somewhat: if oil really were the primary driver of lagging female labor force participation in oil-rich Qatar, Kuwait, the United Arab Emirates, and Saudi Arabia, then we would expect oil-poor but historically and culturally similar Yemen to have significantly higher rates of female labor force participation. Instead, consistent with Charrad's hypothesis, it has comparable rates.

Second, and consistent with Charrad's argument, Alesina *et al.* (2011) argue that patriarchal norms are literally rooted in a region's soil. They argue that indigenous plough use entrenched gendered work norms, and they show that a significant amount of present-day cross-country variation in female labor force participation rates can be explained by regional variation in agricultural heritage.

---

[14] Rauch and Kostyshak (2009, p. 166).
[15] Charrad (2009, p. 548).

Landes and Landes (2001) and Inglehart and Norris (2003a, b) posit a third, related, explanation for lagging female labor force participation that is also consistent with our analysis of Ross's data: they argue that Islamic beliefs play a central role in lagging female empowerment in the Middle East. Ross dismisses the importance of Islam largely on the basis of the statistically insignificant coefficient on his percentage of Muslim residents variable (*viz*, Table 5, Column 1). As Norris (2009) notes, however, ''this measure ... does not take into account important variations among Muslim societies.'' This variation is associated with geography. The five divergent schools of Islam within the Middle East are naturally sorted by region: Hanifi in the Arab Middle East, Maliki in North Africa, Shafi in the southern peninsula, Hanbali in Saudi Arabia, and Ja'fari in Iran (World Bank, 2011). This suggests that our geography dummies may be picking up on some aspect of religious beliefs that Ross's blunt measurement is missing.

They might, for example, be picking up on fundamentalism, which Blaydes and Linzer (2008) have argued inhibits progress towards gender equality. We do not have a good measure of ''fundamentalism'' to include in our regressions and formally test this hypothesis, but it is likely that such a variable would indeed be strongly correlated with our Arabian Peninsula dummy. Informally, the Arabian Peninsula is home to Wahhabism, the most fundamental form of Islam. More formally, the Comparative Values Survey of Islamic Countries found that 88% of the 1413 respondents from Saudi Arabia (the only country in the survey on the Arabian Peninsula) agreed or strongly agreed with the statement ''a good government should implement only the laws of Shari'a'', compared with 66% (of 10,764 respondents) and 54% (of 4721 respondents) in the non-Peninsular MENA and non-MENA countries surveyed (respectively: Algeria, Egypt, Iraq, Jordan; and Bangladesh, Indonesia, Nigeria, Pakistan). Similarly, in the Arab Barometer Survey (ABS, 2005), 49% of 908 respondents from Yemen (again, the only country in the survey on the Arabian Peninsula) strongly agreed with the statement that ''If a Muslim converts to another religion, he must be punished by execution'' compared with only 33% (of 4719 respondents) in the non-Peninsular MENA countries (Algeria, Jordan, Lebanon, Morocco, Palestine).

Fourth and finally, Morrison (2009) makes the complementary argument that oil rents are neither pro-democratic nor anti-democratic *per se*, but rather that oil rents act as a regime stabilizing force — i.e., a force which facilitates maintaining the status quo power structures. In the present

context, this suggests that the presence of high oil rents would have helped ''lock in'' pre-existing low levels of female empowerment in the countries on the Arabian Peninsula — consistent with oil mattering, but not through a gendered Dutch Disease. Our read of Ross's data is consistent with this. As we document above, his data indicate little evidence that oil affects female labor force participation once regional effects are included. When we replicated Ross's regressions of Female Parliamentary seats on Oil Rents (*viz* Ross, 2008a, Table 4), however, Oil Rents remained significant even after including regional effects — although the regional effects are quantitatively important and do reduce the magnitude of the effect.

Conclusions

We have argued that the empirical support for Ross's claim that ''oil, not Islam'' is at fault for the lagging progress towards gender equality in the Middle East is quite weak. We find Ross's theoretical argument that natural resource wealth could play a role in gender equality dynamics eminently reasonable, and it seems plausible that oil plays *some* role in these dynamics. Our point is simply that Ross's empirical work does not support his underlying theory: his data do not provide robust evidence that low rates of female labor force participation in the Middle East are driven by a gendered Dutch Disease; they do not provide much evidence that oil is an important driver of female labor force participation rates at all; and they provide some mild evidence that Islam *is*. The relative importance of kinship ties, agricultural history, oil, Islam, and other factors, and the mechanisms through which they affect gender equality is still an open question. Concluding that''[t]he persistence of patriarchy in the Middle East has relatively little to do with Islam, but much to do with the region's oil-based economy''[16] is premature.

References

Adida, C. L., D. Laitlin, and M.-A. Valfort. 2011. "Gender, Generosity and Islam: A Perspective from France." *APSA 2011 Annual Meeting Paper*.

Alesina, A., P. Giuliano, and N. Nunn. 2011. "Fertility and the Plough." *American Economic Review* 101(3): 499–503.

Alexander, A. and C. Welzel. 2011. "Islam and Patriarchy: How Robust Is Muslim Support for Patriarchal Values?" *World Values Research* 4(2): 40–70.

[16] Ross (2008a, p. 120).

Arab Barometer Survey. 2005. <http://www.arabbarometer.org>.
Arellano, M. 1987. "Computing Robust Standard Errors for Within-Group Estimators." *Oxford Bulletin of Economic Statistics* 49: 431–434.
Blaydes, L. and D. Linzer. 2008. "The Political Economy of Women's Support for Fundamentalist Islam." *World Politics* 60(4): 576–600.
Charrad, M. 2009. "Kinship, Islam, or Oil: Culprits of Gender Inequality." *Politics & Gender* 5(4): 546–553.
Comparative Values Survey of Islamic Countries. 1999–2006. http://www.thearda.com/Archive/Files/Descriptions/ISLAMVAL.asp.
Do, Q-T., A. Levchenko, and C. Raddatz. 2011. "Engendering Trade." *World Bank Policy Research* Working Paper 5777.
Frankel, J. 2010. "The Natural Resource Curse: A Survey." *NBER* Working Paper 15836.
Gorman, B. 2009. "The Green Glass Ceiling: Gender Inequality and Wahhabi Political Influence." *Manuscript*.
Hausman, J. 1978. "Specification Tests in Econometrics." *Econometrica* 46(6): 1251–1271.
Inglehart, R. and P. Norris. 2003a. "The True Clash of Civilizations." *Foreign Policy* 135: 62–70.
Inglehart, R. and P. Norris. 2003b. *Rising Tide*. New York: Cambridge University Press.
Ingvaldsen, M. 2010. "Low Female Labor Force Participation in the Gulf: Cultural Preferences or Necessary Consequences of Large Oil Rents." *Manuscript*.
Landes, D. and R. Landes. 2001. "Girl Power: Do Fundamentalists Fear Our Women?" *New Republic* October 8: 20–23.
Magud, N. and S. Sosa. 2010. "When and Why Worry About Real Exchange Rate Appreciation? The Missing Link Between Dutch Disease and Growth." *IMF Working Papers* 1–32.
Morrison, K. 2009. "Oil, Nontax Revenue, and the Redistributional Foundations of Regime Stability." *International Organization* 63: 107–138.
Norris, P. 2009. "Petroleum and Patriarchy: A Response to Ross." *Politics & Gender* 5(4): 553–560.
Price, A. 2011. "Constraints and Opportunities: the Shaping of Attitudes Toward Women's Employment in the Middle East." The Ohio State University. *PhD Dissertation*.
Rauch, J. E. and S. Kostyshak. 2009. "The Three Arab Worlds." *Journal of Economic Perspectives* 23(3): 165–188.
Ross, M. 2008. "Oil, Islam, and Women." *American Political Science Review* 102(1): 107–123.
Ross, M. 2008. "Replication data for: Oil, Islam, and Women," http://hdl.handle.net/1902.1/14307 UNF:5:fsZ56s2dvxP26at+iCdOhg== V1.
Ross, M. 2009. "Does Oil Wealth Hurt Women? A reply to Caraway, Charrad, Kang, and Norris." *Politics & Gender* 5(4): 575–582.
Shane, M. "Real Historical Annual Exchange Rates for Baseline Countries and Regions, 1970–2009." ERS International Macroeconomic Database. http://www.ers.usda.gov/Data/ExchageRates/.
Sharabi, H. 1988. *Neopatriarchy: A Theory of Distorted Change in Arab Society*. New York: Oxford University Press.
White, H. 1980. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* 48: 721–746.
Women and the Arab Awakening. October 15, 2011. *The Economist*.
World Bank. 2004. *Gender and Development in the Middle East and North Africa*. Washington, DC: World Bank.
World Bank. 2011. *Gender Equality and Development in the Middle East and North Africa Region*. Washington, DC: World Bank.