Taxing Identity
Theory and Evidence from Early Islam

Mohamed Saleh* Jean Tirole*

July 3, 2020

Abstract

A ruler who does not identify with a social group, whether on religious, ethnic, cultural or socioeconomic grounds, is confronted with a trade-off between taking advantage of the out-group population's eagerness to maintain its identity and inducing it to "comply" (conversion, quit, exodus or any other way of accommodating the ruler's own identity). This paper first nests economists' extraction model, in which rulers are revenue-maximizers, within a more general identity-based model, in which rulers care also about inducing people to lose their identity, both in a static and an evolving environment. The paper then constructs novel data sources to test the implications of both models in the context of Egypt's conversion to Islam between 641 and 1170. The evidence comes in support of the identity-based model.

Keywords: Islam, poll tax, identity taxation, Laffer curve, legitimacy.

JEL numbers: D82, H2, N45, Z12.

*Toulouse School of Economics and Institute for Advanced Study in Toulouse, 1 Esplanade de l'Université, Toulouse 31080, France. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 669217 - ERC MARKLIM). Financial support from the ANR-Labex IAST is also gratefully acknowledged. We are grateful to participants at numerous conferences and seminars and to Ran Abramitzky, Robert Barro, Roland Bénabou, Dora Costa, Barry Eichengreen, James Fenske, Roberto Galbiati, Oded Galor, Avner Greif, Walker Hanlon, Richard Hornbeck, Emir Kamenika, Timur Kuran, Alessandro Lizzeri, Joel Mokyr, Thomas Piketty, Giacomo Ponzetto, Emmanuel Saez, Noam Yuchtman, Fabrizio Zilibotti, and four referees for their very helpful comments. We thank Amirreza Ahmadzadeh, Paul-Henri Moisson, Norhan Muhab, Roxana Pozo, and Cyril Thomson, for their excellent research assistantship.

"Muhammad was sent as a prophet and not as a tax collector."

Umar II, the Umayyad Caliph from 717 to 720 CE

1 Introduction

1.1 Motivation and main insights

Hostility toward populations on the ground of their religious, ethnic, linguistic, cultural, economic, political, or sexual-orientation identity, is commonplace. While a voluminous literature covers rulers' violent (non-price) policies against these "unwanted" populations,[1] the non-violent (price) approach of taxing identity has received much less attention. Yet, the taxation of identity may be fully assumed: A poll (head) tax was levied by the early Arab Caliphate, and by subsequent Muslim-ruled polities, on their non-Muslim subjects up to the mid-19th century. Or it may be more subdued, as when local governments discriminate among neighborhoods when locating amenities, or when countries restrict access to public goods to permanent residents or citizens, or (in dictatorships) members of the ruling party.[2]

Taxing identity exposes rulers to a conflict between extracting agents' willingness to pay for keeping their identity and inducing them to lose it (convert, assimilate, quit the organization or the country...). There are two views on how rulers solve this dilemma. The economists' typical view is that rulers, especially in pre-modern polities, are revenue-maximizers.[3] Although less explored by economists, identity taxation is accommodated by this *extraction model*, since the maximization of tax revenue involves tolerating some level of identity maintenance.

While the extraction model has much merit, a second view of identity taxation is that rulers care not only about money, but also about inducing people to lose their identity, even at the expense of lower tax revenue. This may be due to an ideological mission

[1]For example, Voigtländer and Voth (2012) and Anderson et al. (2017) study anti-Semitic persecutions.

[2]Examples of identity taxation abound. Various European polities imposed a tax on Jews up to the 1800s. Romans levied a poll tax from which citizens were exempted, until Roman citizenship became universal under Emperor Caracalla. The Reformation was characterized by a shift from identity taxation, the tithe imposed by the Catholic church on its adherents, to secular taxation (Dittmar and Meisenzahl 2020). In constitutional countries, taxes can be targeted less explicitly toward unwanted populations. For instance, the 1942 one-off Varlik Vergisi (wealth) tax in Turkey was imposed on all citizens' fixed assets (Artunç and Agir 2017). While on paper a non-discriminatory tax, it affected most severely non-Muslims, who controlled a large portion of the economy, and led to their exodus. Communist countries used Communist Party membership to allocate positions. Local and national governments' policies with respect to the provision of local public goods for migrants (training, housing, bureaucratic hassle, intolerance toward harassment...) is yet another example.

[3]For example, see De Long and Shleifer (1993) and Besley and Persson (2011).

to win converts or to a political goal to expand a ruler's support base. Accordingly, we study the paradigm in which a ruler has both revenue and identity objectives. In this *identity-based model*, which nests the extraction model, the ruler levies two taxes: a uniform tax, which mechanically has no impact on conversions and therefore is purely extractive, and a discriminatory one levied on those who maintain their identity, which does affect conversions. A straightforward implication of this more general model is that, absent delegation problems, the discriminatory tax lies on the downward-sloping side of the corresponding Laffer curve.

Our historical context is taxation in the aftermath of the Arab conquest of the then-Coptic Christian Egypt in 641 CE, until the fall of the Fatimid Caliphate in 1170. The Arab Caliphate levied both a discriminatory (poll) tax on religion, imposed on non-Muslims (initially all Egyptians) and removed upon conversion to Islam, and a non-discriminatory (uniform) one on land that was paid regardless of the taxpayer's religion.[4] While this system is consistent with the extraction model, the identity-based model was (implicitly) endorsed on empirical grounds by pioneering historians such as Wellhausen (1902), Becker (1902), Bell (1910), and Grohmann (1932); they postulated that tax-induced conversions led to a loss in poll tax revenue, which is only possible under the identity-based model. Indeed, faced with a deteriorating poll tax revenue, Umar II, who was renowned for his piety, called for more conversions at the cost of a lower tax revenue, suggesting an identity-based motive. While the Arab Caliphate enforced its tax system throughout all its conquered territories, including Iran and India, we limit ourselves in this paper to Egypt, because it is where local-level papyrological records on poll tax payments, and medieval narratives on churches, poll tax hikes, and conversion waves, survived.

The historical context offers a number of advantages to study identity taxation. First, authorities automatically validated conversions to Islam, in contrast with situations in which identity compliance (e.g., permanent residency or naturalization) can be rejected. In non-automatic-acceptance situations, one would need a theory as to how applications are accepted or rejected as a function of the applicant's and the authority's identity strength. They are also more challenging to study empirically, because of the complexity of the two-step sorting process. Second, both identity taxation and conversions are observed (by the Caliphate and by us), unlike situations of implicit identity taxation (e.g., allocation of public goods) where it is more difficult to observe their identity basis. Third, there were two forms of taxation: discriminatory and uniform, which generates interesting dynamics of how the uniform tax evolves over time in re-

[4] Between 641 and circa 750, non-Muslim landholders also paid a higher land tax rate. By 750, the difference in land tax rate was abolished (see the end of the introduction and Sections 2 and 4.3).

sponse to the changes in the discriminatory tax base. Finally, conversion to Islam was irreversible because of the death penalty on apostates. While definitive exit from the tax base is a common feature in identity taxation contexts, the Caliphate institutionalized the irreversibility of the conversion decision.

Cross-sectional analysis. Section 3 develops the framework. Taxation is delegated by the ruler/central authority (CA) to local authorities (LAs). Districts differ in the identity strengths of the local collector (the LA) and of the population. LAs levy a uniform and a discriminatory tax. The CA has no local presence and can only request a transfer from the LA.

An extractive ruler does tax agents' identity. Maximal extraction requires maximizing separately revenues on the uniform and the discriminatory taxes.[5] In the identity-based paradigm, the CA's fiscal motivation is two-fold: extract as much revenue as possible and induce conversions. The latter motivation alters the discriminatory tax, which induces conversions, but not the uniform tax, which does not and remains purely extractive. This introduces a divergence relative to the extraction model, the consequences of which we investigate theoretically and empirically.

Section 3 thus explores the cross-sectional differences between the extraction and the identity-based models. Here, the interesting focus is on agency. The delegation of the collection of taxes to local tax collectors is of no consequence in the extraction model, at least if the ruler has enough information on local conditions: the CA and LAs both aiming at maximizing revenue creates congruence between them. Not so in the identity-based model, as the local authorities may not share the ruler's identity preferences. A case in point is early Islamic Egypt, in which Copts rather than Arabs administered tax collection in many districts. The main theoretical result here is that the discriminatory tax still lies on the downward-sloping side of the Laffer curve, but ceteris paribus an LA with a stronger identity strength levies a higher discriminatory tax, induces more conversions and raises less revenue.[6] We also demonstrate that the population's own identity strength increases the discriminatory tax rate and revenue, but mitigates conversions, an implication that is common across both models.

To test apart the cross-sectional implications of the two models, we exploit the local variation in early Islamic Egypt across *kuras* (Egypt's administrative units) in the

[5] When there is no threat of rebellion: under a threat of tax-induced rebellion, the CA must tone down its demands and the designs of the uniform and discriminatory taxes become intertwined.

[6] Under a threat of rebellion, discriminatory tax revenue is an inverted-U in LA's identity strength. We emphasize that converts can rebel, but nonetheless we focus on the no-rebellion-threat case in the cross-sectional analysis, because our (indirect) evidence on poll tax revenue comes from 1200 (and 1375), when the rebellion threat had largely subsided (converts were already 84% by 1200, which is a strong predictor of a low threat from Proposition 5).

identity strength of LAs, controlling for the identity strength of Copts.[7] We think of the CA as either Egypt's governor or the Caliph (see the time-series part for a discussion). We think of the LAs, not only as *kura* headmen, but as the entire local bureaucracies (village headmen, tax collectors, accountants, land surveyors). We measure LAs' identity strength by a *kura*-level dummy variable that takes value 1 if an Arab tribe settled permanently in 700–969. This variable arguably captures the level of Arabization of the LAs: In *kuras* where Arabs settled, they were more likely to replace Coptic LAs, whereas in non-Arab-settled *kuras* Coptic LAs remained in power. Arabization of LAs was constrained by (1) the number and spatial distribution of Arab tribes, and (2) the availability of Arabs with enough highly specialized human capital that was necessary to replace the entire local bureaucracies. We measure Copts' identity strength by a village-level[8] dummy variable =1 if the village is believed, according to pre-641 local Coptic legends (recorded in Coptic narratives), to have been visited by the Holy Family in its legendary biblical flight to Egypt. We also provide, as a robustness check, an alternative measure: a dummy variable =1 if a pre-641 Coptic saint or martyr is documented to have lived in the village according to another Coptic narrative.

We construct novel data, based on medieval Coptic narratives and papyrological tax records, in order to measure our three outcomes: (1) The proportion of converts is measured by a village-level dummy variable =1 if the village did not have any Coptic churches or monasteries in 1200. We conduct a number of robustness checks to test the validity of this measure. (2) The poll tax rate is measured by the individual-level annual poll tax payment in 641–1100, localized at the *kura*-level, based on papyrological poll tax registers and receipts, that survived for only 4 out of 42 *kuras* (11% of villages in 1315 and 14% of the population in 1897). However, sample selection appears to be quasi-random: *kuras* (and villages) in the poll-tax sub-sample do not differ with respect to most observables from out-of-the-sample areas. The main exception is (exogenous) geography: papyri were more likely to survive in hotter and drier areas. (3) While we do not have local-level data on poll tax revenue per capita, we impute it at the *kura* level for the poll-tax sub-sample, by multiplying the proportion of villages with at least one church or monastery in 1200 with the average poll tax payment. We also provide a second piece of (suggestive) evidence based on the village-level total tax transfer per unit of land in 1375, post the Arab Caliphate period.

Our evidence on the impact of LAs' identity strength comes in support of the identity-based model. The findings on outcomes (2) and (3) are only suggestive, though: we do not conduct econometric analysis for these outcomes, due to the small number

[7] We do not have a panel dataset that traces *kuras* over time, though.

[8] There are 1,782 villages located within 42 *kuras*.

of *kuras* in the poll-tax sub-sample, and thus our analysis relies on a (statistical) comparison of means.[9] We find that villages located in Arab-settled *kuras* are less likely to have Coptic churches and monasteries in 1200 (14% versus 22% in non-Arab-settled *kuras*). This finding is robust to the inclusion of a host of pre-641 control variables, and to an instrumental variable strategy for Arab settlement. We also document that the average individual poll tax payment is 27% higher, and that the imputed poll tax revenue per capita is halved, in Arab-settled *kuras*. No-church villages in 1200 faced a lower total tax transfer per unit of land and were less likely to be granted to Mamluks (the ruling military elite) in 1375. Furthermore, as implied by both models, we find that the Holy-Family-visit villages are more likely to have churches and monasteries in 1200, and that Holy-Family-visit *kuras* have higher individual poll tax payments, on average (though not statistically significant), and higher imputed per-capita poll tax revenues. We obtain similar results when we use the saint-martyr measure. Finally, we discuss a number of alternative interpretations of Arab settlement, where we argue that they are unlikely to hold.

Time-series analysis. Section 4 explores the time-series implications of the extraction and identity-based models. Rulers and agents are forward-looking, and conversions are permanent (apostasy assumption). The first key theoretical result concerns conversions, discriminatory tax rate, and discriminatory tax revenue. Under the extraction model, all conversions occur at date 1, and we do not expect any conversions or discriminatory tax hikes to occur thereafter.[10] The identity-based model allows for later poll tax hikes and conversion waves, though, where the equilibrium exhibits a sufficient-statistic property: In particular, date-t outcomes are determined by the highest ruler identity (and the lowest budget need) so far, a form of ratcheting. A statistical implication of this result is that earlier rulers are more likely to order discriminatory tax hikes and conversion waves than later rulers and that the extraction model's relevance increases over time. Yet, both the possibility that later rulers have stronger identity and the time-decreasing agency cost between the CA and LAs (due to, say, the increasing Arabization of LAs) imply that later rulers can still impact policy.

The second key result of the time-series analysis pertains to the effect of the threat of rebellion on the uniform tax under both models, offering a much richer picture than a mere capping of a tax below the revenue-maximizing one to prevent a tax-induced

[9]Imputed poll tax revenue, being at the *kura* level for only 4 *kuras*, does not allow a statistical test. Our second (econometric) evidence on outcome (3), based on village-level total tax transfer in 1375, is only suggestive, because it comes from the post-Caliphate period.

[10]The extraction model allows for later discriminatory tax hikes (but not conversion waves) if there is an increase in external threat, an implication shared with the identity-based model.

rebellion. Even in an otherwise fully stationary environment, the threat of rebellion subsides over time, enabling the ruler to raise the uniform tax. The intuition is that converts have less to gain from rebelling as they have already given up on their identity. This holds even though agents realize that by converting they lose their option value of costlessly keeping their identity in the future in case of a successful rebellion.

To test the first result, we exploit the time-series variation in the CA's identity strength (c_t) and budget needs (B_t) across Egypt's rulers, controlling for external and internal threats. Our analysis is based on two datasets, one at the governor level (N = 122), and the other at the Caliph level ($N = 65$), spanning 530 years from 641 to 1170.[11] We measure CA's identity strength by (1) governor's hostility (=1 if governor is hostile toward non-converts), based on medieval Coptic narratives, and (2) Caliph's piety (=1 if the Caliph is not known for drinking alcohol), based on medieval Muslim narratives. Budget needs are measured by a dummy variable =1 if there is no military expedition by the Caliphate under ruler t. We control for external threats by a dummy variable =1 if there was an upcoming foreign attack under ruler t, and for the internal threat of rebellion by ruler t's start year (time trend).

We employ medieval Coptic narratives to measure the incidence of a poll tax hike and of a conversion wave under ruler t. We first note that we observe poll tax hikes and conversion waves after "date 1" (defined as 641–661), unlike what is predicted by the extraction model. They became less frequent over time, though, which is also consistent with the identity-based model. We then study econometrically if (later) poll tax hikes and conversion waves can be indeed explained by historical booms in c_t, or busts in B_t, as implied by the identity-based model. In accordance with the theory, and given the binary coding of both c_t and B_t, we distinguish between the "contemporaneous" effects of each of c_t and B_t, and the "historical legacy" effect: the number of previous high-c_t rulers (n^c_{t-1}), and the number of previous low-B_t rulers (n^B_{t-1}). The identity-based model then implies that the probability of poll tax hikes and of conversion waves is higher among high-c_t (and low-B_t) rulers, and is decreasing in each of n^c_{t-1} and n^B_{t-1}, whereas the extraction model implies no effect of any of these variables.

Our findings come in support of the identity-based model. First, hostile governors (and pious Caliphs) were indeed 54% (37%) more likely to order poll tax hikes, and 45% (36%) more likely to induce conversion waves. We trace the smaller effects of Caliphs to the governor-Caliph agency relationship that may have implied that governors had more influence on tax policy and conversions. We explain the contemporane-

[11]The two levels were identical when Egypt was independent, in 868–905 and 935–1170. In the empirical analysis, t denotes the numbering of governors/Caliphs, and not the calendar year. We weight rulers by the length of their tenure: long-tenure rulers are expected to have more influence on tax policy.

ous effect by two factors that arguably left a leeway for poll tax hikes and conversions in later periods: (1) the high B_t faced by earlier high-c_t rulers, that probably suppressed conversions, and (2) the Arabization of LAs over time. Second, we mostly fail to find effects of budget need busts, suggesting that CA's identity strength was the prevailing determinant of taxation under the Arab Caliphate. Third, while we detect a negative effect of n^B_{t-1}, we fail to find a robust (negative) effect of n^c_{t-1}, probably due to the high collinearity of the two variables. This suggests that the probability of poll tax hikes and conversions declined over time, which is consistent with the time-increasing relevance of the extraction model. While we do not have high-frequency data on poll tax revenue, the fact that it decreased over time, yet at a decreasing rate, is consistent with this interpretation.

We then use the second theoretical result with respect to the uniform tax, to study why the Caliphate removed the cap on the uniform tax only circa 750, and not before. We cannot study the timing of this reform econometrically, because it was a Caliphate-wide one-time policy change, and we thus rely on a theory-based interpretation of history. This policy change is important on both theoretical and historical grounds. Theoretically, the removal of the cap on the uniform tax results from a dynamic "divide-and-rule" policy. Historically, this policy change created the canonical form of Islamic taxation, where both converts and non-converts are subject to the same land tax rate, but only differ on the poll tax. Both the extraction and identity-based models explain the delayed tax reform by the high threat of rebellion early on, which then declined over time as conversions increased. While, as predicted by theory, the composition of rebels changed over time to include both converts and non-converts, the suppression of the tax revolts enabled the Caliphate to preserve the new tax system.

Extensions. Finally, Section 5 discusses two main extensions: persecutions and emigration. First, the agency approach, being based on a potential conflict between the LAs and the CA, also may rationalize in well-defined circumstances the use of inefficient, non-price instruments such as persecutions. Persecutions can be also rationalized as a signalling device. These insights shed light on the relative role of persecutions (Mikhail 2004, El-Leithy 2005) versus taxation (Frantz-Murphy 2004, Rapoport 2018) in inducing conversions, which has generated a debate among historians. Persecutions were rare under the Caliphate, but became more common under the Bahri Mamluks in 1250–1354. While both agency and signalling can potentially explain persecutions under the Caliphate (when the poll tax was decentralized), only signalling can explain the Mamluks' violent approach (when the poll tax was centralized). Second, we study emigration as an alternative to conversion. We show that emigration is accommodated by our model. While irrelevant to Copts (who never emigrated from Egypt), emigration has been relevant in many other contexts, and we discuss some of these historical examples.

1.2 Related literature

The paper is related to a few strands of literature. It differs from the optimal taxation literature in at least two ways: the agency problem and the hysteresis effects associated with exit from the tax base. The paper shares with Becker (1957)'s theory of discrimination the feature that decision-makers have a possible distaste for minority membership. Acemoglu (2006)'s ruler taxes a constituency beyond the revenue-maximizing level so as to weaken the latter. The focus of the two papers is markedly different, as are the conclusions; for example, revolts are triggered by a soft tax treatment in Acemoglu and a tough one in our paper.[12] Taxes may also lie on the downward-sloping side of the Laffer curve for the taxation of externalities and internalities (e.g. tobacco or pollution). The interaction between Pigouvian taxation and revenue-generating distortive taxation is well known in public finance (Sandmo 1975, Bovenberg and de Mooij 1994, Bovenberg and Goulder 1996). Furthermore, our modeling, properly reinterpreted, also covers the design of "sin taxes" (O'Donoghue and Rabin 2006). For a hyperbolic consumer with present bias, consumption today brings immediate benefit and a delayed cost. The optimal tax may lie on the downward-sloping side of the Laffer curve. This literature however ignores agency in tax collection as well as issues related to the tax structure and to the specific dynamics of taxation and rebellion under ratcheting of compliance (apostasy, costly return. . .); it thereby cannot guide the empirical strategy employed in this paper.

A large literature studies optimal taxation with non-utilitarian welfare functions (e.g. Fleurbaey and Maniquet 2011). Saez and Stantcheva (2016) derive optimal taxation in an environment that is not necessarily welfarist (in particular, social welfare weights can depend on individual or aggregate characteristics which do not enter individuals' utilities). Their focus is on allowing various considerations, such as counterfactuals (what

[12]In his model, the ruling elite not only aims at extracting rents from the output of an enterprising middle-class via a tax on its output, but also may try to achieve other goals with the tax, thus exceeding the maximal extraction tax rate. First, the elite may itself own firms and taxing the middle-class output discourages middle-class production and reduces the market wage. As Acemoglu emphasizes, this result hinges on limited tax instruments (for example, a tax on labor hired by the middle-class firms could take care of limiting competition for labor). By contrast, we study optimal taxation. Second, the middle class may rely on its financial power to rebel. That reason is complementary to our section on rebellion, which is based on manpower rather than money; as a consequence, the agents rebel when ill-treated by the ruler in this paper, while they rebel when well-treated and therefore empowered in Acemoglu's contribution. Overall, both the rationales for hurting and the focus differ between the two papers.

would have happened in the absence of taxes?), horizontal equity, libertarianism, equality of opportunity concerns, and poverty alleviation, to matter per se, independently of their consequences on the taxpayers' utility. Much work has also been devoted to investigate the impact of altruism on optimal taxation (e.g. Diamond 2006, Farhi and Werning 2010, Kaplow 1995). These two literatures investigate neither the taxation of unwanted populations, nor its dynamic evolution as unwanted population members convert or leave the polity or organization.

The paper contributes to the economics of religion. One primary focus of this literature has been to identify the causal impact of religious beliefs on economic outcomes (Barro and McCleary 2003, Botticini and Eckstein 2005, Becker and Woessmann 2009). Instead of emphasizing the causality from religion to economics, our paper documents how economic incentives can alter the religious affiliation. In this respect, our paper contributes to a recent empirical literature that attempts to elicit the willingness to pay to maintain one's identity (or beliefs) (Augenblick et al. 2016, Delavande and Zafar 2018). Another line of this literature explores how political authorities co-opt religious ones to preserve legitimacy under the threat of rebellion (Greif and Tadelis 2010, Chaney 2013, Belloc et al. 2016, Rubin 2017, Cantoni et al. 2018). While less central to our explanation of discriminatory taxation and conversions, maintaining ruler's legitimacy via altering the religious composition of taxpayers is our preferred explanation of the removal of the cap on the uniform tax rate.

Our paper also speaks to the economic history of the Middle East. Michalopoulos et al. (2018) show that the spread of Islam across countries today is correlated with proximity to pre-600 trade routes, and with lower land productivity and higher land inequality, where they argue that Islam's redistributive institutions mitigated the incentives for predation in highly unequal areas. Like these authors, we investigate the geographical spread of Islam but focus on other determinants such as the local identity strengths of authorities and population. Indeed, while trade routes may explain the spread of Islam in territories that Islam reached by trade (e.g., southeast Asia and sub-Saharan Africa), they are less likely to explain the spread of Islam in territories that Islam reached by conquest (which comprises the whole of the Arab Caliphate). They do not explain either the local variation in conversions within Egypt, which is all proximate to trade routes. Furthermore, we account for the local variation in land productivity and land inequality in the empirical analysis of conversions (see Section 3.2.3). Finally, we also investigate the timing of the conversion to Islam.

Second, Kuran (2012) argues that the *waqf* system, a tax-exempt religious endowment that enabled owners to protect their property rights against state confiscation, was one reason behind the economic stagnation of the region, as it induced owners to lock in capital in unproductive investments. Although this theory does not (directly) address the tax policy and conversion question, we may be under-estimating the discriminatory tax if state confiscation was an additional tax that targeted Copts, besides the poll tax. Indeed, Copts were not allowed to form *waqfs*, which suggests that they were less protected from confiscation than converts. Nevertheless, we do not think that state confiscation was a major component of the discriminatory tax in our context for two reasons: (1) it targeted Coptic elites, unlike the poll tax which was levied on everyone, and (2) confiscations are less relevant for the Arab Caliphate than for the Mamluk period.[13]

2 Historical background

Islamization of Egypt. Arabs conquered Egypt in 641. On the eve of the Conquest, the vast majority of Egyptians were Coptic Christians, while non-Coptic Christians and Jews formed two small (urban) minorities. During the centuries that followed, non-Muslims shrank from 100% of the population in 641 to 16% in 1200. Two pieces of evidence further suggest that the largest decline occurred by 969 (pre-Fatimid). First, poll tax revenue fell sharply by 680, despite the constant *de jure* poll tax rate, suggesting a rapid decline in the proportion of non-Muslims from 100% in 641 to 42% in 680, and further to 33% in 786 and 23% in 813 (Courbage and Fargues 1997). Second, Bulliet (1979) uses paternal lineages of prominent individuals in Iran to identify the approximate date at which an ancestor, a son of a non-convert, adopted an Arabic name and converted to Islam. He documents an "S-curve" of conversions which he extrapolates to other regions including Egypt, finding that conversions peaked in the 9th century.

There is a consensus among historians that the Islamization of Egypt was driven by Copts' conversions to Islam, rather than by population replacement, via either of (1) Arab immigration and Copt emigration, (2) Muslims' higher fertility and/or lower mortality, and (3) inter-marriages between Muslim males and Coptic females (which result by Islamic law in Muslim offspring). Appendix Section B.1 suggests that (1) the number of Arab settlers was small relative to Egypt's population, (2) Muslims (both Arabs and converts) did not have higher fertility or lower mortality than non-converts, and (3) inter-marriages were rare. Conversion to Islam was observed by the state.[14] Being Muslim was an "absorbing state" owing to three Islamic laws on (1) the death

[13] Arabs did not confiscate land in Egypt upon the Conquest (unlike in the Levant and Iraq), and the vast majority of land remained in the hands of the local (Coptic) population. Only public domain land and royal Byzantine land was confiscated. Under the Arab Caliphate, state confiscation of Copt property took place under al-Mutawakkil (847–861), Ahmad ibn Tulun (868–905), and al-Hakim (996–1021).

[14] A papyrological list of converts in 700–900 reveals that a convert had to declare his new faith in front of the authorities, adopt an Arabic name, and become a client of an Arab patron.

penalty on apostates, (2) the offspring of a Muslim male being automatically Muslim, and (3) Muslim females' obligation to marry only Muslim males.[15]

Islamic tax system. Arabs taxed religion and land. Upon the Conquest, they imposed on Egyptians a discriminatory tax (τ) that was removed upon conversion to Islam, and a uniform tax (λ) on land that was levied regardless of religious affiliation. Up to circa 750, τ was made up of two components: (1) a poll tax, an annual per head cash tax on free adult males, and (2) the positive difference in land tax rate between the rate on Copts (*kharaj*) and the rate on Muslims (variously called *ushr*, *zakat*, *sadaqa*). λ was thus equal to the *ushr*. Importantly, the *ushr* was capped at 10% of land's yield according to *hadith* (prophet's sayings).[16]

Circa 750 (date uncertain), Caliphs, backed by jurists, raised the land tax on Muslims from the *ushr* to the *kharaj* rate. They further removed any preexisting country-specific treaty-based upper bound on *kharaj*, by denying the historical existence of peace treaties in most of the conquered territories, including Egypt. Consequently, from that date on, τ equated the poll tax, and λ equated the *kharaj* land tax. As a result, the land tax increased sharply circa 750. It then fluctuated over time in response to economic shocks, but never went back to its pre-750 level. This change marks the establishment of the *canonical* Islamic tax system that remained in place, until the abolition of the poll tax in 1856.[17]

Tax administration. Caliphs delegated poll and land tax collection to Egypt's (fiscal) governors, who decided on the total budget that was used to pay the tribute to the Caliphate,[18] and to finance the salaries of Egypt's top officials, the army, the police, the judiciary, and the bureaucracy. Importantly, poll and land taxes were *not* raised to finance local public goods, which were financed instead by ad hoc levies and corvée labor.

[15] Because Egyptians were mostly Copts in 641, and Egypt's Muslims are mostly descendants of Copts who converted to Islam, we use the terms "Copt" and "non-Muslim," and "convert" and "Muslim," interchangeably.

[16] Due to the lack of papyrological evidence on the *ushr* tax before 750, it has been argued that Muslim landholders actually paid no land tax before 750.

[17] We abstract from other discriminatory taxes in our definition of τ: (1) special taxes on non-Muslims (up to 750) levied for specific uses (e.g., military expenses, lodging for officials, governor's expenses, local public projects), because they were irregular ad hoc levies, (2) military conscription on Muslims (up to 833), because (a) it was in return for a state stipend, and (b) it was *not* widespread in Egypt, (3) (non-state) taxes/subsidies administered by religious organizations (churches, monasteries, mosques), because (a) we do not have evidence on their magnitudes, and (b) they were not enforced by the state. Similarly, we abstract from the expansion in λ tax base after 858 to include pasture, weir, and various crops and products.

[18] When Egypt was a province of the Caliphate in 641–868 and 905–935. Caliphs generally appointed two governors: military and fiscal; the two positions could have been held by the same person. In 868–905 and 935–1170, governors and Caliphs were identical, because Egypt was independent then.

Governors/Caliphs delegated tax administration to the local authorities of each *kura*. The total budget was allocated across *kuras* based on population size (among other characteristics), and transfer demands were issued to each *kura*. Local tax revenues were then sent to Egypt's capital city. Poll and land tax rates varied locally, because local authorities had discretionary power on tax rates, or at least, on the level of tax enforcement.

3 Cross-sectional analysis

3.1 Theory

(a) Model description

There is a continuum of districts, indexed by $i \in [0,1]$, each with a mass 1 of *agents*. Agents are initially endowed with the same identity (say, the Coptic religion). There is a second identity (say, Islam) that the agents can embrace, abandoning the initial one.

Each district i is run by a *local authority* (LA), which optimally collects discriminatory and uniform taxes $\{\tau_i, \lambda_i\}$. The discriminatory tax rate $\tau_i \geq 0$ is levied solely on agents who choose to maintain their identity. The associated tax revenue is thus equal to τ_i times the fraction of such agents. In contrast, the uniform tax λ_i will for notational simplicity denote the tax revenue or effective tax rate paid by all agents. Our preferred interpretation of λ_i is the extractive one: λ_i stands for the maximum revenue that can be obtained through a uniform tax (say, a tax on land); but we are agnostic about the determination of the district-i uniform tax λ_i (indeed, in the case of Egypt, Islamic law capped the uniform/land tax until 750) and therefore keep its formulation as general as possible.

An individual agent is characterized by a parameter $\theta \in \mathbb{R}$, measuring his willingness to pay for keeping his initial identity (his "identity strength"). Hence, a type-θ agent in district i keeps his identity if and only if $\theta \geq \tau_i$. Willingnesses to pay in district i are distributed according to cdf $F(\theta - r_i)$. Thus, identity is more pregnant in a district with a higher r_i . We assume that the distribution F is smooth, has density f, and satisfies the monotone hazard rate property: $f(\theta)/[1-F(\theta)]$ is strictly increasing. This assumption will in particular guarantee that objective functions are strictly quasi-concave.[19]

[19]The parameter θ should be thought of as the net WTP for keeping identity. In our application, embracing Islam may create an option value for the convert (or his lineage) from possibly adhering to the new, not-yet-experienced religion. This option value is to be subtracted from the gross benefit from remaining Copt. In particular, the support of F may include negative values of θ, i.e. agents who would convert even in the absence of discriminatory tax.

The departure from the extraction model is that taxes embody an identity-related motive: The district-i LA incurs (dis)utility $c_i \in \mathbb{R}$ times the fraction of agents preserving their identity; depending on the district, the cost c_i can be positive or negative. The parameters r_i and c_i thus measure the agents' and the LA's *identity strengths* in district i. There is a smooth joint distribution $G(c_i, r_i)$ on district characteristics.

A *central authority (CA)* has an indirect extraction motive: It cares about transfers from local authorities. And, like the local authorities, it also has a nonfinancial objective: It incurs disutility $c > 0$ per agent preserving his identity. The CA relies on local authorities to collect taxes. It can only demand a district-specific transfer T_i; that is, it does not have any local presence that would allow it to interfere with local tax collection, and the LA is residual claimant for its revenue. The LA therefore has a direct extraction motive. For simplicity, there is no asymmetry of information between the CA and the LA regarding $\{r_i, c_i\}$.

Objective functions

An *agent* of type θ living in district i has objective function:

$$U_i(\theta) = -\lambda_i - \min\{\tau_i, \theta\}.$$

Because *LA i* has a (dis)taste c_i (or taste if c_i <0) per agent preserving his identity and is residual claimant for its revenue once it has transferred T_i, its utility is:

$$V_i = [\lambda_i + R_i(\tau_i) - T_i] - c_i[1 - F(\tau_i - r_i)] = \lambda_i + R_i^a(\tau_i) - T_i$$

where $R_i(\tau_i) = \tau_i[1 - F(\tau_i - r_i)]$ denote the discriminatory tax revenue, and $R_i^a(\tau_i) = (\tau_i - c_i)[1 - F(\tau_i - r_i)]$ is an "adjusted tax revenue" that accounts for the LA's (dis)taste for the agents' keeping their identity.

The *CA* receives taxes in amount $\int_0^1 T_i di$, and has welfare in the absence of rebellion threat $W = \int_0^1 [T_i - c[1 - F(\tau_i - r_i)]]di$. We similarly define the CA's adjusted tax revenue as $R_i^c(\tau_i) = (\tau_i - c)[1 - F(\tau_i - r_i)]$. If each district's collected tax is equal to the transferred revenue (which will be a non-trivial implication of symmetric information[20]), this expression can alternatively be rewritten as:

$$W = \int_0^1 [\lambda_i + R_i^c(\tau_i)]di.$$

(b) Equilibrium tax and revenue

LA's program. Faced with transfer demand T_i, the LA in district i solves:

$$\max_{\{\tau_i\}} R_i^a(\tau_i) - T_i$$

subject to the revenue-collection constraint:

$$\lambda_i + R_i(\tau_i) \geq T_i.$$

We define three tax rates that will play a central role in what follows. Let

$$\tau_i^m \equiv \arg\max\{R_i(\tau_i)\}$$

[20]Under moral hazard, symmetric information may not preclude the agent from enjoying rents, as in the efficiency wage model.

denote the extraction-maximizing (or monopoly) discriminatory tax in district i, yielding the maximally extractive revenue R_i^m,

$$\tau_i^a(c_i) \equiv \arg\max\{R_i^a(\tau_i)\}$$

denote the preferred discriminatory tax of a LA with identity strength c_i, and

$$\tau_i^c \equiv \tau_i^a(c)$$

denote the CA's preferred discriminatory tax. Trivially, $\tau_i^a(c_i) > \tau_i^m$ if and only if $c_i > 0$. In particular, the CA's preferred tax τ_i^c is on the decreasing-revenue side of the Laffer curve.

Proposition 1 ***(equilibrium discriminatory tax and revenue).***

For any given district i:

(i) Regardless of the sign of c_i, the district-specific equilibrium discriminatory tax satisfies $\tau_i \geq \tau_i^m$ (that is, it is weakly on the decreasing-revenue side of the district's Laffer curve), is almost everywhere differentiable, and where so, satisfies $\frac{\partial \tau_i}{\partial r_i} \in (0,1)$ and $\frac{\partial \tau_i}{\partial c_i} \in [0,1)$. The conversion rate, $F(\tau_i - r_i)$, is weakly increasing in the LA's strength of identity, c_i, and decreasing in the agents' identity strength r_i.

(ii) District i's transfer T_i is equal to its revenue, $\lambda_i + R_i(\tau_i)$. It is invariant to c_i and is equal to $\lambda_i + R_i^m$ for $c_i < 0$, decreases with c_i in $(0,c)$, and is invariant to c_i for $c_i > c$ (see Figure 1(b)).[21] *It is strictly increasing in the agents' identity strength r_i.*

(iii) There is no delegation cost if and only if $c_i \geq c$.

The intuition behind Proposition 1 can be grasped from Figure 1(a). Let us separate local authorities' types into three categories: *zealous* when $c_i > c$, *soft* when $c_i \in (0,c)$, and *counterattitudinal* when $c_i < 0$. From CA's viewpoint, a zealous authority puts too much weight on inducing agents to surrender their identity and too little on revenue; it is easy to control them, as a revenue requirement at the level of the revenue the CA would raise itself forces the local authority both to lower its discriminatory tax and to raise more revenue; there is no agency/delegation cost.

In contrast, the other local authorities are not fervent enough. The CA would like them to raise their discriminatory tax, but faces an agency problem. In the case of a soft LA, an increase in its discriminatory tax leads to less revenue; however, a lower transfer requirement does not induce it to increase its tax; rather, it chooses to pocket the difference between revenue and transfer. Technically, the set of implementable discriminatory taxes is $[\tau_i^m, \tau_i^a(c_i)]$. In sum, the outcome for soft local authorities is dictated by the latter's preferences. In the case of counterattitudinal local authorities, the implementable set is symmetrically $[\tau_i^a(c_i), \tau_i^m]$. The CA prefers the highest tax that

[21] A corollary is that an increase in the CA's identity strength, c, leads to more conversions and a lower total revenue.

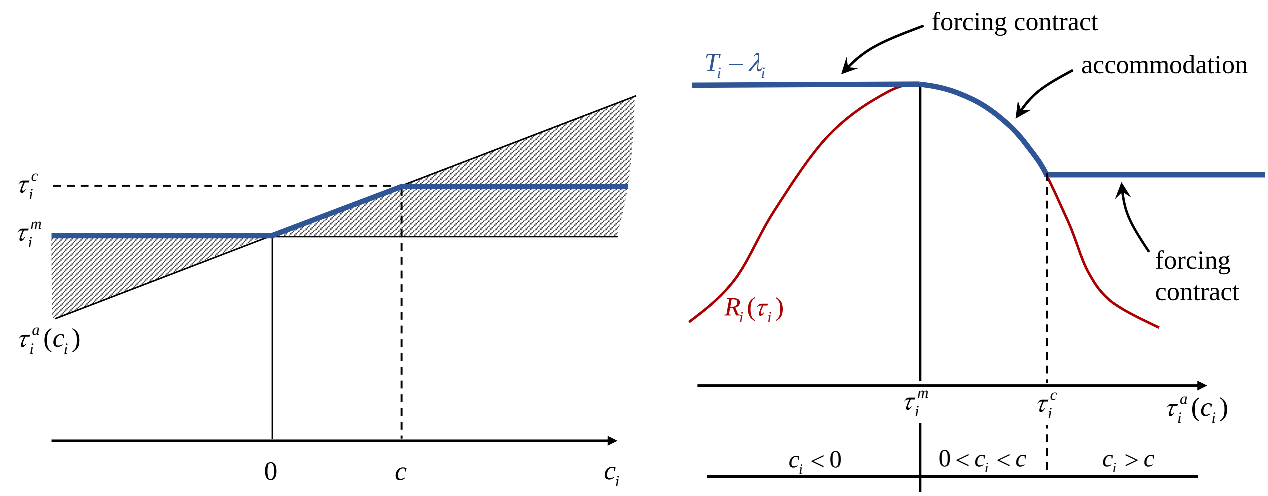

(a) Implementable (hatched area) and optimal delegated tax (thick line)

(b) Revenue and local authority identity strength

Figure 1: **Delegated tax collection and Laffer curve (no rebellion)**

can be implemented, the extractive tax, to any below it, because it brings both more revenue and more agents to surrender their identity.

The comparative statics with respect to agent religiosity are straightforward. District i's demand for maintaining identity is $1 - F(\tau_i - r_i)$. Under the monotone-hazard-rate condition, the demand elasticity, $\tau_i f(\tau_i - r_i)/[1 - F(\tau_i - r_i)]$, is decreasing and so the discriminatory tax is increasing in identity strength r_i, while overall conversions, $F(\tau_i - r_i)$, decrease. Those properties are not affected by the agency problem.

We compare outcomes in this model, which combines extractive and identity motives, with the classical extraction-only model. The latter corresponds to $c_i = 0$ for all i.[22]

Corollary 1 ***(horserace with extraction model).***
(i) Both the extraction model and the identity-based one predict that the discriminatory tax, the revenue and the transfer all grow with the district's agent identity r_i, and that the conversion rate decreases with r_i.
(ii) In contrast with the identity-based model, the discriminatory tax and revenue and the overall transfer in the extraction model do not vary with either local or central authority's strength of identity.

The threat of rebellion will be studied in much detail in the time-series analysis, where it is most interesting. Appendix A however builds the possibility of rebellion

[22] When $c_i = 0$ for all i, the CA's identity is irrelevant, as districts maximally extract regardless of the revenue request.

into the cross-sectional analysis. A rebellion, if successful, kicks out the rulers and results in (say) the absence of taxes. It requires cooperation among the various districts, as each district is too small to challenge the CA's rule. The CA can change the agents' incentive to rebel by reducing the revenue demands it imposes on local authorities. For example, and referring to Figure 1(b), it can moderate its revenue demands in districts run by counterattitudinal tax collectors. Appendix A shows that under the threat of rebellion, the discriminatory tax τ_i is still increasing in the LA's identity strength c_i. The discriminatory tax revenue (R_i) is now inverted-U shaped in the LA's identity strength (c_i), with a peak for a secular LA $(c_i = 0)$.

3.2 Empirics

3.2.1 Cross-sectional implications of the extraction and identity-based models

Table 1 summarizes these implications. There are two determinants: identity strength of the LA (c_i) and identity strength of agents (r_i). While the implications of the two models are the same with respect to r_i, they differ on c_i. By construction, the extraction model implies that the discriminatory tax rate (τ_i), the proportion of converts (F_i), and the discriminatory tax revenue (R_i), are all insensitive to c_i, whether there is a rebellion threat or not. On the contrary, the identity-based model implies that τ_i and F_i are both increasing in c_i. Without rebellion threat, R_i is decreasing in c_i, because LAs with $c_i > 0$ operate on the downward-sloping side of the Laffer curve of τ_i. Under the rebellion threat, though, R_i is an inverted-U in c_i.

Table 1: **Cross-sectional implications of the extraction and identity-based models**

Outcome	Extraction model $\frac{\partial}{\partial c_i}$	Extraction model $\frac{\partial}{\partial r_i}$	Identity-based model $\frac{\partial}{\partial c_i}$	Identity-based model $\frac{\partial}{\partial r_i}$
Discriminatory tax (τ_i)	0	$\in (0,1)$	$\in (0,1)$	$\in (0,1)$
Prop. converts (F_i)	0	$-$	$+$	$-$
Discriminatory tax revenue (R_i)	0	$+$	$-^{\dagger}$	$+$

† Under the identity-based model, R_i is decreasing in c_i if there is no threat of rebellion, and is an inverted-U under the threat of rebellion.

3.2.2 Measuring c_i and r_i

To test the two models apart, we exploit the local variation in early Islamic Egypt in c_i, controlling for r_i. We define the CA as Egypt's governors or Caliphs.[23] We define

[23] We abstract in the cross-sectional empirical analysis from the agency problem that may arise between Caliphs and governors, because our focus here is on the implications of the agency problem between the

the LAs as *kura* headmen and the local bureaucracies, including village headmen, land surveyors, scribes, and tax collectors. While we measure c_i at the *kura* level i, we measure r_{ji} at the village level j located within *kura* i. With respect to outcomes, we measure (1) F_{ji} at the village level j, (2) τ_{hi} at the individual level h located within *kura* i (i.e. not localized at the village level), and (3) we impute R_i at the *kura* level i in 1200, and we observe T_{ji} (tax transfer) at the village level in 1375. We first explain how we measure c_i and r_{ji}.

Identity strength of LA (c_i). We measure c_i by a dummy variable that takes value 1 if at least one Arab tribe settled permanently in *kura* i between 700 and 969. Before the Arab Conquest in 641, LAs were all Coptic Christians. While Arabs initially kept the status quo, they soon attempted to Arabize the LAs. Arabization was constrained, though, by the spatial distribution of Arab settlers, and more importantly, by the number of Arabs with enough highly specialized human capital to replace, not only *kura* headmen, but the entire local bureaucracies.[24] This is supported by administrative evidence. Individual-level (non-localized) data on occupational titles and religious affiliation, constructed from the Egyptian papyri dating between 641 and 969, reveal that LAs were partially Arabized during this period: Muslims (Arabs and converts) came to occupy jobs at all bureaucracy levels, and to be over-represented among high bureaucrats.[25] Non-convert Copts kept being over-represented among mid-low bureaucrats, though (tax collectors, scribes, land surveyors) (Saleh 2018).

We argue that Arab settlement measures the Arabization of LAs. Between 641 and 969, Arab tribes settled in certain *kuras* but not others, first temporarily during the spring season (in 641–700), and then permanently (in 700–969).[26] In *kuras* where Arabs settled, they were more likely to replace Copts in the LA. Consequently, these *kuras* faced a larger share of Arabs in the LA, compared to *kuras* where Arabs did not settle. We think of Arab-settled *kuras* as characterized by $c_i > 0$, with Arab administrators either *zealous* ($c_i > c$) or *soft* ($0 < c_i < c$), and of non-Arab-settled *kuras* as characterized by $c_i < 0$, with *counterattitudinal* Coptic administrators. While c_i is continuous in theory, our empirical measure is dichotomous, as we do not have a measure

CA (Caliph/governor) and the LAs. The Caliph-governor agency problem will be relevant in interpreting the time-series empirical evidence, though.

[24] Coptic mid-low bureaucrats were difficult to replace, because they had highly specialized human capital in land surveying, measuring Nile level, and agriculture in rural Egypt. Many Copts remained as *kura* headmen, too. Basilios, the Coptic head of *Aphrodito* circa 710, is a well-known example.

[25] While Muslim high-level bureaucrats in this dataset are almost certainly Arabs, we cannot separate Arabs from converts at mid-low levels of the bureaucracy, because converts had to adopt an Arabic name and became clients of Arab tribes. This pooling will over-estimate the share of Arabs in mid-low LAs.

[26] Arabs were 58% more likely to settle permanently after 700 in *kuras* where they settled temporarily before 700.

of c_i among Coptic and Arab administrators. This is not a concern for most implications in Table 1, which are monotonic in both c_i and r_i, except for the inverted-U R_i in c_i that is implied by the identity-based model under the threat of rebellion, which we will not be able to disentangle. We will come back to this point in Section 3.2.5.

Panel (A) of Appendix Figure B.1 shows the location of *kuras* where Arab tribes settled in 700–969. Arabs were more likely to settle in the eastern and western Delta than in the central Delta, and in the northern Nile Valley than in the southern Valley.

Identity strength of agents (r_{ji}). We measure r_{ji} at the village level by a dummy variable that takes value 1 if it is believed, according to pre-641 local Coptic legends, that a village j, located within *kura* i, had been visited by the Holy Family (henceforth, HF) during its legendary biblical flight to Egypt. The HF visit legend has been an important element of popular Coptic Christianity until today. Villages on this list are mentioned in an apocryphal book that is attributed to Theophilus, the patriarch of Alexandria in 384–412. Although the book's authorship and date are both doubtful, with some scholars attributing the book to an unknown author in the post-641 period, most of the local legends that the book's author compiled likely date from the pre-641 period, as the HF visit was mentioned in Coptic sources during the Roman and Byzantine eras. We think of the HF status as the upper tail of r_{ji}. Specifically, the HF-visit villages contained "miraculous" sites that Jesus and/or Mary were believed (among locals) to have created or touched, such as hand-prints, footprints, trees, and wells, which might have instilled a particularly strong sense of Coptic identity among the local population before 641. Panel (B) of Appendix Figure B.1 shows the location of villages for which the HF visit status =1, which is the case for 24 villages (1.3%) of 1,782 villages in 1315.

To address the limitations of the HF visit measure, we employ, as a robustness check, an alternative measure of r_{ji}: a dummy variable =1 if a Coptic saint or martyr spent (part of) their lives in village j between 49 CE, the customary date of establishing the Coptic Church of Alexandria by Saint Mark, and 641 CE. Like the HF measure, the saint-martyr measure captures the upper tail of r_{ji}: it is equal to 1 for 30 villages (1.7% of villages). We constructed this measure from the Coptic Synaxarium, the major medieval liturgical Coptic book that compiles biographies of saints and martyrs arranged according to days of the Coptic calendar year. According to Coptic beliefs, these local saints and martyrs performed miracles, and were (mostly) tortured to death by either Roman (pagan) governors, or Byzantine (non-Coptic Christian) governors, in defense of their Coptic Christian faith. Importantly, the two measures are weakly positively correlated ($\rho = 0.14$) at the village level,[27]

[27] Among 1,782 villages, there are 1,730 villages for which both measures are equal to 0, and 4 villages

3.2.3 Impact on conversions (F_{ji})

Measuring F_{ji}. We employ a dummy variable F_{ji}=1 if village j in *kura* i did not have any Coptic church or monastery circa 1200, based on Abul-Makarim's Coptic chronicle. We argue that F_{ji} captures the proportion of converts between 641, when all villages were 100% Copt, and 1200, the end of the Arab Caliphate period. The presence of an operating Coptic church or monastery is an indicator of the local presence of a sizable non-convert population.

Our measure is valid under the following assumptions: (1) the universe of villages is observed in 641 (no post-641 villages), (2) every village had at least one Coptic church or monastery in 641, (3) the conversion of the vast majority of a village's population resulted in the disappearance of its churches and monasteries, via desertion or transformation into mosques, (4) the list of churches and monasteries in 1200 is complete, and (5) there is no differential movement of converts and non-converts across villages.

These assumptions are supported by a number of observations. In support of (1), we define the universe of villages based on the 1315 cadastre's village list.[28] Most of these villages existed before 641 (Ramzi 1954). As a robustness check, we further restrict our analysis to a subset of villages mentioned in Byzantine-period sources, that was compiled by the French archaeologist Amélineau (1850–1915).[29] The results are qualitatively similar to the main findings (Appendix Table B.4). In support of (2), Amélineau's villages are quite large (mean population in 1897 is 5,900, compared to 2,700 in non-Amélineau villages). Hence, they are most likely to have had at least one church or monastery in 641. In support of (3), our measure is negatively correlated ($\rho = -0.29$) with the actual number of non-convert Coptic households in 1245 among villages in *Fayum kura*, based on al-Nabulsi's *Fayum* cadastre. We also estimate the regression using the individual-level population census samples in 1848 and 1868, and we obtain similar results (Appendix Table B.5). In support of (4), Abul-Makarim's list is the most complete enumeration of churches and monasteries in medieval Egypt. It has more entries, and geographic coverage, than any other list. We obtain similar results if we use the list of churches and monasteries in 1500 (Appendix Table B.5). In support of (5), (a) rural-rural migration was outlawed: papyrological administrative records reveal that "fugitives" who fled their villages were forced to go back, (b) (tax-induced) rural-urban migration is unlikely because cities were controlled by Arab LAs.[30]

---

for which they are both equal to 1.

[28]The earliest extant comprehensive list of Egyptian villages dates to the 1298 cadastre, but it is not digitized.

[29]This is not an exhaustive list of pre-641 villages, though; it only includes villages that were large enough to be mentioned in the Byzantine sources.

[30]In 1848, when mobility restrictions and the poll tax were both still enforced, the proportion of rural-

Panel (C) of Appendix Figure B.1 shows the spatial distribution of villages that did not have any Coptic church or monastery in 1200. According to this measure, converts were already in the majority by 1200: 84% of Egyptian villages did not have any Coptic church or monastery by then. But there was considerable spatial heterogeneity; for example, Coptic churches and monasteries were more likely to survive in the central Delta and the southern Valley.

Basic specification. Our objective is to examine whether the effect of c_i on F_{ji} is consistent with the extraction model or with the identity-based model. Recall that the implications of both models are the same with respect to r_{ji}. We first estimate the following model:

$$F_{ji,1200} = \beta_0 + \beta_1 c_{i,700-969} + \beta_2 r_{ji,641} + X_{1,i}\beta_3 + X_{2,ji}\beta_4 + \varepsilon_{ji} \quad (1)$$

where $F_{ji,1200}$ is a dummy variable =1 if there is not any Coptic church or monastery in village j, located within *kura* i, circa 1200. The main regressor is $c_{i,700-969}$ =1 if an Arab tribe settled in *kura* i in 700–969. The second regressor is $r_{ji,641}$ =1 if village j is believed, according to pre-641 local Coptic legends, to have been visited by the Holy Family.

We include two sets of control variables. First, the vector $X_{1,i}$ includes Byzantine-period *kura*-level controls: (1) the logarithm of urban population of *kura* i circa 300; using urbanization as a proxy for income is standard in history, as urban populations were richer on average,[31] and (2) a dummy variable =1 if there was a Byzantine garrison in *kura* i circa 600, which captures military resistance to the Arab Conquest. Second, the vector $X_{2,ji}$ includes geographic village-level controls: (3) FAO-GAEZ cereals suitability index, which is the maximum value of the suitability indices of barley, wheat, beans, and maize, under irrigation and intermediate input level,[32] (4) mean temperature, (5) temperature range, (6) slope, and (7) rainfall. Standard errors are clustered at the *kura* level, the level of aggregation of Arab settlement. We also report spatial-autoregressive standard errors as a robustness check in Appendix Table B.6.

Two remarks are in order. First, one concern that arises is the potential correlation between c_i and r_{ji}: Arab tribes may have chosen where to settle based on the baseline identity strength of Copts. We thus estimate an alternative specification, in which we interact Arab settlement with the HF visit status, and the results are qualitatively similar to the original results, while the interaction term itself is statistically insignificant (Appendix Table B.8). Second, the FAO-GAEZ cereals suitability index accounts for the

---

rural cross-*kura* immigrants is not statistically different between Muslims and Copts (5.7% versus 6.1%).

[31] Urban population is defined as the sum of the population of Greek cities (metropolis) and of *nome* capitals (Egypt's administrative units during the Roman period).

[32] FAO-GAEZ does not provide crop suitability measures under irrigation and *low* input level.

local variation in land productivity, which is suggested by Michalopoulos et al. (2018). To further account for the impact of land inequality on conversions, we include, as a robustness check, a *kura*-level dummy variable =1 if there is at least one *autopract* estate circa 600; the *autopragia* was a privilege granted to large landholders in late Byzantine Egypt allowing them to pay taxes directly to the capital city and to collect taxes in their constituencies. It can be thus used to measure land concentration in each *kura*. The results are qualitatively similar to the main results (Appendix Table B.8).[33]

The null hypothesis (H_0) on β_1 is the extraction model, which implies that $\frac{\partial F_i}{\partial c_i} = 0$. The alternative hypothesis (H_1) is the identity-based model, which implies that $\frac{\partial F_i}{\partial c_i} > 0$. H_0 on β_2 is that it is equal to 0. H_1 is that $\frac{\partial F_i}{\partial r_i} > 0$, which is consistent with both models.

Instrumental variable for Arab settlement. The identification assumption in equation (1) is that the cross-*kura* variation in Arab settlement is exogenous to baseline characteristics of *kuras*, which may be driving conversions. This assumption may be violated due to (1) reverse causality: Arab settlers may have settled in *kuras* with larger convert populations, and (2) omitted variables: Arab settlement choice may have hinged on other unobservable pre-641 characteristics of *kuras* that also account for the variation in conversions. To deal with the potential endogeneity of Arab settlement, we employ an instrumental variable (IV) strategy, where we estimate the following first-stage regression:

$$\begin{aligned} c_{i,700-969} = \; & \alpha_0 + \alpha_1 DistancetoArish_i + \alpha_2 BorderDesert_i \\ & + \alpha_3 (DistancetoArish_i \times BorderDesert_i) + X_{1,i}\alpha_4 + X_{2,i}\alpha_5 + v_i \end{aligned} \tag{2}$$

where $DistancetoArish_i$ is the distance between the capital of *kura i* and *Arish*, a small town close to Egypt's northeastern border, that was the first to be captured by Arabs in 639 due to its proximity to the Arab peninsula (Conquest was by land from the northeast), $BorderDesert_i$ =1 if *kura i* borders desert land. All *kuras* border hinterland, except these in central Delta.

We argue that the IVs are relevant. Column (1) of Appendix Table B.7 suggests that Arabs were more likely to settle in *kuras* that are both closer to *Arish* and bordering desert. This is confirmed by historical evidence. For one, proximity to *Arish* determined the extent to which Arabs were willing to travel, although there were exceptions.[34] For another, Arabs preferred *kuras* that bordered desert, where they first settled temporarily during the spring season in 641–700, to practice hunting and horse riding in a similar environment to that of the Arab peninsula. Starting from circa 700, they settled in these

[33] We do not include this variable in the basic set of controls because it is missing for half of the *kuras*.

[34] Regardless of the distance to *'Arish*, Arabs were more likely to settle closer to frontier cities such as *Aswan* in the south and Alexandria in the north. Also, Arabs were more likely to settle in western Delta than in central Delta (that is actually closer to *'Arish*), due to the former's proximity to desert.

kuras permanently.

Furthermore, we argue that the IVs are exogenous, as they are determined by geography. They arguably satisfy the exclusion restriction, conditional on controls. Columns (2)-(9) of Appendix Table B.7 reveal that the IVs are not correlated with most Byzantine-period and geographic characteristics, with the exception of urban population circa 300 and temperature.

Findings. The findings are in Table 2. Column (1) reveals that the probability of conversion to Islam in 641–1200 is higher in Arab-settled *kuras*: whereas 22% of villages located in *kuras* where Coptic administrators remained in power ($c_i < 0$) had Coptic churches or monasteries in 1200, the proportion is only 14% in *kuras* where Arab tribes settled in 700–969 ($c_i > 0$). Because all *kuras* were 100% Copt before 641, this finding suggests that *kuras* where Arabs settled witnessed relatively more conversions to Islam between 641 and 1200.

Column (2) shows that HF-visit villages (at the upper tail of r_{ji}) were more likely to have Coptic churches or monasteries in 1200 at 75%, compared to only 15% in non-HF-visit villages. We obtain similar results when we use the saint-martyr measure (Appendix Table B.3).

Including both regressors and control variables in columns (3)-(5) yields similar results to those in columns (1) and (2). The IV results in column (6) indicate that the coefficients of Arab settlement and of the HF visit retain their magnitudes and statistical significance.

The positive effect of Arab settlement on conversions is consistent with the identity-based model, but not with the extraction model, which implies that conversions are insensitive to c_i (Table 1). The negative effect of the HF visit is consistent with both models.

3.2.4 Impact on discriminatory tax rate (τ_{hi})

Measuring τ_{hi}. We employ Egypt's papyrological tax registers and receipts in 641–1100 to measure the annual poll tax payment (τ_{hi}) made by taxpayer h, located in *kura* i, in 641–1100 (N= 408 individual taxpayers).[35] Poll tax records survived in only 4 out of 42 *kuras*, all located in the Nile Valley: *Hermopolis* (N= 77), *Aphrodito* (N= 314), *Fayum* (N= 7), and *Ihnas* (N= 10) (Appendix Figure B.1, panel (D)).[36] Furthermore,

[35] Appendix Figures B.2 shows example pictures of the secondary sources that we used to construct our dataset. Initially, the Arabic term (*jizya*) meant "tax in cash" that included both the poll tax and the cash land tax. The term was later confined to the poll tax during the 8th century. Poll tax payments are clearly identified in our tax papyri sample, though.

[36] We exclude 143 individual poll tax records with missing location.

Table 2: **Local determinants of conversions to Islam in 641–1200**
Dependent variable: =1 if no Coptic church or monastery in village j circa 1200

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) 2SLS
=1 if Arab settlement (c_i)	0.083		0.082	0.078	0.076	0.115
	(0.033)**		(0.034)**	(0.031)**	(0.032)**	(0.052)**
=1 if Holy Family visit (r_{ji})		-0.585	-0.584	-0.594	-0.624	-0.623
		(0.084)***	(0.082)***	(0.083)***	(0.089)***	(0.088)***
Byzantine *kura*-level controls?	No	No	No	Yes	Yes	Yes
Geographic village-level controls?	No	No	No	No	Yes	Yes
Obs (villages)	1782	1782	1782	1782	1751	1751
Clusters (*kuras*)	42	42	42	42	42	42
R^2	0.01	0.03	0.04	0.04	0.05	0.05
Mean dep. var. in control	0.78	0.85	0.78	0.78	0.78	0.78
KP Wald F-stat						16.65

Notes: Standard errors clustered at the *kura* level are in parentheses. Byzantine-period controls are: (1) the logarithm of urban population in *kura* i circa 300, and (2) a dummy variable =1 if there was a Byzantine garrison in *kura* i circa 600. Geographic controls are: (3) FAO-GAEZ suitability index to the cultivation of barley, wheat, beans, and maize, under irrigation and intermediate input level, (4) mean temperature, (5) temperature range, (6) slope, and (7) rainfall. * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Sources: See Appendix Section B.2.

95% of the sample comes from *Hermopolis* and *Aphrodito*. These four *kuras* include 11% of the total number of villages in the 1315 cadastre, and 14% of the population in the 1897 census. Appendix Figure B.3 shows the histogram of τ_{hi} in each *kura*.

The poll tax sub-sample is subject to two caveats. First, both the small number of clusters (*kuras*), and their concentration in the Nile Valley, raise a natural concern about the national representativeness of the sample. Two remarks bolster our confidence in the sample, though: (1) Sample selection appears to be quasi-random. Appendix Table B.9 reveals that villages/*kuras* in the poll-tax sub-sample are not statistically different with respect to most observables, in comparison to out-of-sample villages/*kuras*. The main exception is (exogenous) geographic characteristics: villages in the poll-tax sub-sample have higher temperature, greater temperature range, less rainfall, higher slope, and higher likelihood of bordering desert, than out-of-sample villages. This confirms a long-known fact in Greco-Roman, Coptic, and Arabic papyrology: papyri are more likely to survive in the Nile Valley due to its dry and hot climate.[37] (2) We re-estimate

[37]Random events further uncovered papyri in specific locations within the Valley. The tax papyri of *Aphrodito* were discovered in 1901 by local farmers while digging a well. The papyri were then distributed among farmers, and the remaining documents ended up in museums, including the British Museum.

equation (1) for villages in the poll-tax sub-sample, and the results are qualitatively similar to those for the full sample (Appendix Table B.10).

The second caveat about the poll-tax sub-sample is that papyri are typically dated within a range (century or longer).[38] We thus pool papyri from all four locations and date them between 641 and 1100. This raises a concern that we may be confounding the cross-sectional effect of c_i on τ_{hi} with its time-series effect. The latter effect may arise due to *kura*-specific changes in c_i, or Egypt-level changes in CA's identity strength, over time. In the absence of panel data on c_i and τ_i over time, we cannot rule out this concern. However, we note that our finding that τ_{hi} is higher, on average, in Arab-settled *kuras* holds if we limit our sample to the pre-Fatimid period (the difference is not statistically significant, though). This mitigates the concern that LAs (and c_i) may have changed under the Fatimids.

Evidence. We do not estimate a regression model for τ_{hi}, because of the small number of clusters (*kuras*) in the poll-tax sub-sample. However, we provide suggestive evidence by examining the difference in mean τ_{hi} (1) between Arab-settled and non-Arab-settled *kuras* in 700–969 ($\Delta_{c_i}\bar{\tau}$), and (2) between HF-visit and non-HF-visit *kuras* ($\Delta_{r_i}\bar{\tau}$).[39] The downside is that we are not able to control for potentially confounding variables. The standard error of the difference in means is clustered at the *kura* level. As the small number of clusters may bias the standard errors downwards (Cameron et al. 2008), we estimate the p-value using the Wild Cluster Restricted (WCR) bootstrap (Roodman et al. 2018).

The null hypothesis on $\Delta_{c_i}\bar{\tau}$ is the extraction model, which implies that it is equal to 0. The alternative hypothesis is the identity-based model, which implies that it is greater than 0. The null hypothesis on $\Delta_{r_i}\bar{\tau}$ is that it is equal to 0, while the alternative hypothesis is that it is greater than 0, which is consistent with both models.

The findings are shown in Table 3. Taxpayers in *kuras* of *Hermopolis*, *Fayum*, and *Ihnas*, where Arabs settled in 700–969 ($c_i > 0$), paid on average a higher poll tax rate in 641–1100 by 0.29 dinar, 27% more than the average poll tax in *Aphrodito*, where Arabs did not settle and the LA thus remained Coptic ($c_i < 0$). The difference is statistically significant, and the magnitude is economically sizable: It amounts to 3% of the annual wage of manual low-skilled workers in 661–969, who constituted the low-income poll tax bracket, and to 29% of the *de jure* annual poll tax on this bracket (=1 dinar). This result is consistent with the identity-based model, and not the extraction model (Table

[38] *Hermopolis*'s sample is from 731–1100, *Aphrodito* 703–733, *Fayum* 641–1005, and *Ihnas* 701-900.

[39] Because τ_{hi} is only localized at the *kura* level, we aggregate r_{ji} to the *kura* level, where we define r_i as a dummy variable =1 if $\bar{r}_i = \sum r_{ji}/n_i > \text{median}(\bar{r}_i)$. The *kura*-level HF-visit =1 for *Hermopolis* and *Ihnas*, whereas the saint-martyr measure =1 for *Hermopolis* and *Fayum*.

1). Furthermore, taxpayers in *Hermopolis* and *Ihnas* that are believed to have been visited by the Holy Family, and thus had higher r_i, paid a higher poll tax (as implied by both models), yet the difference is not statistically significant.[40] We obtain similar results when we use the saint-martyr measure (diff = 0.28 [p-value = 0.15]).

Table 3: **Local determinants of the poll tax rate in 641–1100**

	=0			=1			
	N	Mean	SD	N	Mean	SD	Diff
Arab settlement in 700–969	314	1.07	1.27	94	1.36	1.09	0.29 (0.009)
Holy Family visit	321	1.08	1.26	87	1.36	1.12	0.29 (0.111)

Notes: p-value of the difference in means in parentheses. It is estimated using Wild Cluster Restricted (WCR) bootstrap, with clustering at the *kura* level, Webb weights, and 999,999 replications.
Source: See Appendix Section B.2.

3.2.5 Impact on discriminatory tax revenue (R_i)

The null hypothesis (H_0) on the effect of c_i on R_i is that it is equal to 0, as implied by the extraction model. H_0 is also consistent with the identity-based model under the threat of rebellion, since failing to reject H_0 with our binary measure of c_i (comparing LAs with $c_i < 0$ and $c_i > 0$) is consistent with an inverted-U R_i. The alternative hypothesis (H_1) is that $\frac{\partial R_i}{\partial c_i} < 0$, which is implied by the identity-based model without a rebellion threat. H_0 on the effect of r_i is that it is equal to 0, whereas H_1 is that $\frac{\partial R_i}{\partial r_i} > 0$, which is consistent with both models.

There are no local-level data on R_i or T_i under the Arab Caliphate. We thus provide two (indirect) pieces of evidence, based on (1) imputing R_i in 1200, and (2) observing T_{ji} in 1375. However, because both R_i and T_{ji} are imputed/observed when the internal threat of rebellion had subsided (recall that converts were already 84% in 1200), we focus on the implications of the extraction and identity-based models under no rebellion threat.

Evidence 1. We impute R_i for the poll tax sub-sample *kuras*: $R^{imp}_{i,1200} = \bar{\tau}_{i,641-1100} \times (1 - \bar{F}_{i,1200})$, where $R^{imp}_{i,1200}$ is the imputed poll tax revenue in dinar per capita for *kura* i in 1200, $\bar{\tau}_{i,641-1100}$ is the average τ_{hi} in *kura* i in 641–1100, $\bar{F}_{i,1200}$ is the proportion of

[40] We also estimate Lee's bounds of the effects of c_i and r_i under non-random selection of the poll-tax sub-sample. To do this, we first aggregate τ_{hi} to the *kura* level as $\bar{\tau}_i = \sum \tau_{hi}/n_i$, where we define the sample selection variable =1 if *kura* i is in the poll-tax sub-sample, and =0 if not. We then weight each *kura* by a frequency weight that is equal to its population size in 1897. The estimated Lee's bounds of the effects of Arab settlement and of the HF-visit are [0.27,0.29] and [0.08,0.29], respectively.

villages in *kura* i that have no church or monastery in 1200.[41] Panel (E) of Appendix Figure B.1 shows the map of $R^{imp}_{i,1200}$.

We first compare $R^{imp}_{i,1200}$ between *Fayum*, which received Arab settlers in 700–969, and *Aphrodito*, which did not. Both *kuras* have the same HF-visit status (=0). We find that by 1200, *Fayum* had half of *Aphrodito*'s poll tax revenue per capita (0.13 dinar versus 0.27 dinar), implying that $\Delta_{c_i} R^{imp}_i < 0$.[42] This is consistent with the identity-based model without rebellion threat. Second, we compare $R^{imp}_{i,1200}$ of *Ihnas* and *Hermopolis*, where the HF-visit status =1, to that of *Fayum*, where the HF-visit status =0. All three *kuras* have the same Arab settlement status (=1). We find that by 1200, *Fayum* had a slightly lower per-capita poll tax revenue (0.13 versus 0.15). This suggests that $\Delta_{r_i} R^{imp}_i > 0$, which is consistent with both models. When we use the saint-martyr measure, we obtain stronger results: *Ihnas*, where the saint-martyr status =0, had lower $R^{imp}_{i,1200}$ than *Fayum* and *Hermopolis* (0.08 versus 0.18).

Evidence 2. Next, we examine if c_i and r_{ji} under the Arab Caliphate had an impact on total tax transfer, post the Caliphate period. For this purpose, we construct data on $\tilde{T}_{ji}$, total tax transfer (*'ibra*) per unit of land in village j within *kura* i, from the Mamluk-period cadastres of 1315 (land area) and 1375 (total tax transfer), the earliest extant cadastres with local-level data on tax transfer.[43] In 1375, the state estimated each village's average yearly total tax revenue, when granting its tax collection right to a beneficiary (LA). Assignment of villages to LAs was a function of $\tilde{T}_{ji}$: Mamluks (ruling elite) were granted high-$\tilde{T}_{ji}$ villages, as a compensation for their services to the Sultan, whereas non-Mamluk LAs paid $\tilde{T}_{ji}$ in advance.

Mamluk and non-Mamluk LAs were presumably extractive ($c_i = 0$); they mainly cared about the village's tax worth. The identity-based model thus predicts that Mamluks would keep the Caliphate's tax policy unaltered (ratcheting). The findings in Appendix Table B.11 reveal three things: (1) High-F_{ji} villages in 1200 were *less* lucrative

[41] If we weight villages by their population size in 1897, we obtain similar results for Arab settlement and the saint-martyr measure, but the (positive) difference between HF and non-HF *kuras* disappears.

[42] (1) Because $N = 4$, we cannot conduct a statistical test of the difference. (2) The result holds if we control for the saint-martyr measure instead, i.e. comparing *Ihnas* and *Aphrodito* where the saint-martyr =0. (3) Village-level data from *Fayum* reveal that it had a small number of non-converts, and a low poll tax revenue by 1245 (Rapoport 2018).

[43] We do not observe R_{ji}, though. Furthermore, our measure ($\tilde{T}_{ji}$) captures T_{ji} only if population size per unit of land, and yield per unit of land, are both held constant for all j. Let q_{ji} denote the amount of land, z_{ji} the average yield per unit of land, n_{ji} the number of inhabitants, total tax transfer is thus $T^{Tot}_{ji} = q_{ji} z_{ji} \lambda_{ji} + n_{ji} \tau_{ji} (1 - F_{ji})$. In theory, we normalize $q_{ji} = z_{ji} = n_{ji} = 1$. We observe $\tilde{T}_{ji} = \frac{T^{Tot}_{ji}}{q_{ji}} = \lambda_{ji} z_{ji} + \frac{n_{ji}}{q_{ji}} R_{ji}$. Hence, $\tilde{T}_{ji} = T_{ji}$ only if z_{ji} and $\frac{n_{ji}}{q_{ji}}$ are the same for all j. Empirically, we control for z_{ji} by the FAO-GAEZ cereals suitability index, and for $\frac{n_{ji}}{q_{ji}}$ by the population size in 1897 (earliest census with digitized full-count village-level tabulations) divided by land area in 1315.

in 1375: they had a lower $\tilde{T}_{ji}$ and were less likely to be granted to Mamluks. We interpret this finding as consistent with the identity-based model. F_{ji} in 1200 is a "sufficient statistic" for the Caliphate's tax legacy: it is the outcome of both c_i and r_{ji} under the Caliphate and is measured at the end of the Caliphate period. The finding suggests that $\tilde{T}_{ji}$ in 1375 is decreasing in F_{ji} in 1200, which is itself increasing in c_i under the Arab Caliphate, and decreasing in r_{ji}. (2) The HF visit status (r_{ji}) has a positive effect on $\tilde{T}_{ji}$ and on the probability of assignment to Mamluk LAs, as expected. However, its impact disappears once we control for F_{ji}. (3) Arab settlement in 700–969 (c_i) does not have a statistically significant impact on $\tilde{T}_{ji}$, or on the probability of assignment to a Mamluk LA. This is possibly because being observed at the *kura* level, Arab settlement has less power in explaining $\tilde{T}_{ji}$ and Mamluk LAs' assignment, especially given that the latter two variables (observed at the village level) exhibit bigger within-*kura* variation than cross-*kura* variation.

3.2.6 Discussion of the cross-sectional evidence

Our findings indicate that Arab LAs in 700–969 imposed a higher poll tax in 641–1100 than Coptic LAs, induced more conversions to Islam by 1200, and thus faced lower poll tax revenue in 1200. No-church villages continued to face a lower tax transfer in 1375. These findings are consistent with the identity-based model. There are alternative interpretations, though, of Arab settlement, which we discuss below.

First, Arab LAs may have been more efficient tax collectors than Coptic LAs (higher state capacity in the extractive sense), either because of lower corruption (less leniency with taxpayers), or better information about taxpayers. Three remarks are in order: (1) Arab LAs' higher capacity in levying τ_i is not consistent with the finding that they collected lower R_i, (2) Arab LAs' lower corruption is actually included in our interpretation of c_i, so long as it is identity-based (i.e. targeting non-converts), (3) It is unlikely that Arab LAs had more information than Coptic LAs. If anything, Coptic LAs were initially more experienced than Arabs, although the information gap may have declined over time.

Second, Arab settlement may have a mechanical effect on the proportion of Muslims, via immigration. This is an unlikely interpretation for three reasons. (1) Historians agree that Egypt's Islamization was driven by conversions, and not by demographic factors (see Section 2). (2) Our measure of F_{ji}, Coptic churches and monasteries in 1200, actually depends on the absolute number of non-converts, rather than their population share. Hence, even if Arab settlement reflects a large Arab immigration wave, this will impact our measure only if Arabs turned churches into mosques, in spite of the existence of a large non-convert population (a scenario on which there is no historical evidence). (3) The other two outcomes (τ_{hi} and R_i) cannot be explained by the pure mechanical effect.

Third, Arabs may have coerced Copts to convert via persecutions. This is unlikely, though, because persecutions were relatively rare under the Arab Caliphate (see Section 5).

Fourth, Arabs may have persuaded Copts of the attractiveness of Islam, reducing r_i and inducing more conversions. But then these areas should have faced a lower, not a higher, τ_i, which is contrary to what we find.

4 Time-series analysis

4.1 Theory

(a) Modeling

Time is indexed by $t \in \{1, 2, \ldots, +\infty\}$. Because we now are interested in time-series rather than cross-section analysis, we presume for expositional simplicity that there is no heterogeneity among districts and that the LAs have the same identity preferences as the CA ($c_{it} = c_t \geq 0$ for all i). As we later show, results carry over if there is an identity wedge between CA and LAs. The implementable discriminatory taxes are therefore described by the interval $[\tau^m, \tau^a(c_t)]$. We allow parameters to vary across periods. Let $\{\lambda_t, \tau_t\}$ denote the uniform and discretionary taxes. We initially take λ_t as exogenous, an assumption that is consistent with revenue-maximization that occurs in the absence of threat of rebellion (λ_t is then the extractive uniform tax at date t). We assume that conversions are permanent. Reswitching may be costly either because of apostasy rules (as in the case of conversions to Islam) or because of the existence of human investments (Jewish intellectuals who left Germany for the US did not come back once politics in Germany returned to normal). Incentive compatibility and apostasy imply that each date t is characterized by a cutoff θ_t^* such that types $\theta \geq \theta_t^*$, and only them, have kept their identity up to date t (included).

The date-t CA's instantaneous objective function is (normalizing $r_i = 0$ for all i)

$$w_t = \lambda_t + \tau_t[1 - F(\theta_t^*)] - c_t[1 - F(\theta_t^*)],$$

where $c_t \geq 0$ is the date-t CA's identity strength (we allow this strength to vary over time). We assume for expositional simplicity that all parameters are deterministic. The results fully generalize if they are stochastic. We will first abstract away from the possibility of rebellion. The date-t CA (ruler)'s intertemporal welfare under discount factor $\beta \in [0, 1)$ is:

$$W_t = \sum_{k=0}^{+\infty} \beta^k [\lambda_{t+k} + (\tau_{t+k} - c_t)[1 - F(\theta_{t+k}^*)]]. \tag{3}$$

Because we allow the ruler to change over time, what will be chosen by ruler $t+k$ is evaluated from the point of view of ruler-t preferences. Agent θ's intertemporal welfare is

$$U = \sum_{t=1}^{+\infty} \beta^{t-1}[-\lambda_t - z_t(\tau_t - \theta)]$$

where $z_t = 1$ if the agent has maintained his identity up to date t (included) and $z_t = 0$ otherwise.

We will also assume that the CA must meet a per-period budget constraint at level B_t (for example, budgetary needs may be high because of a war),

$$\lambda_t + R_t \geq B_t, \tag{4}$$

where $R_t \equiv \tau_t[1 - F(\theta_t^*)]$ is the discriminatory tax revenue. This constraint may or may not be binding, but we assume that, in equilibrium, it can always be met through some choice of taxes.

Because identity switches are permanent, for all t

$$\theta_t^* \geq \theta_{t-1}^*. \tag{5}$$

(b) No internal or external challenge

Imagine first a world in which *both* rulers and agents are myopic ($\beta = 0$). Consider the tax that yields the CA's static optimum under the budget constraint:

$$\tau^*(c_t, B_t - \lambda_t) \equiv \arg_{\{\tau_t[1-F(\tau_t)] \geq B_t - \lambda_t\}} (\tau_t - c_t)[1 - F(\tau_t)].$$

τ^* is increasing in c_t and decreasing in $(B_t - \lambda_t)$. Being myopic, agent θ converts whenever he has not yet converted yet and $\tau_t > \theta$. Ruler t chooses

$$\tau_t = \max\{\tau^*(c_t, B_t - \lambda_t), \theta_{t-1}^*\} \tag{6}$$

To understand (6), suppose first that $\tau^*(c_t, B_t - \lambda_t) \geq \theta_{t-1}^*$ (as is the case for instance if there have been few or no conversions yet). By definition, $\tau^*(c_t, B_t - \lambda_t)$ yields the static optimum and cutoff $\theta_t^* = \tau^*(c_t, B_t - \lambda_t)$. Next, suppose that $\tau^*(c_t, B_t - \lambda_t) < \theta_{t-1}^*$. In the range $\tau_t \in [0, \theta_{t-1}^*]$, the demand for conversion is inelastic and so the objective function, $\lambda_t + [\tau_t - c_t][1 - F(\theta_{t-1}^*)]$, is strictly increasing in τ_t. In either case, $\theta_t^* = \tau_t$. It turns out that these strategies are still optimal when the players value the future:

Proposition 2 ***(dynamics of conversion)***. *For any $\beta \in [0,1)$, there exists a Markov perfect equilibrium in which both the ruler and the agents behave as if they were myopic. The date-t tax and cutoff are $\tau_t = \theta_t^* = \max_{1 \leq k \leq t} \tau^*(c_k, B_k - \lambda_k)$.*

The equilibrium can further be shown to be unique if the horizon is finite; and, under additional assumptions, under infinite horizon (the environment considered here) as well.

To grasp the intuition behind Proposition 2, note first that condition (6) implies that, like the cutoff, the discriminatory tax is weakly increasing over time: $\tau_{t+1} \geq \tau_t$ for all t.

And thus, if $\tau_t > \theta$, then $\tau_{t+k} > \theta$ for all $k > 0$. This means that all types below τ_t enjoy no option value of keeping their identity and so should convert. Conversely, types θ above τ_t get immediate net benefit $\theta - \tau_t > 0$ from not converting and keep their option value for the future. So the agents' optimal behavioral rule is myopic.

To understand why the date-t ruler optimally behaves myopically, assume away budget constraints and suppose that the date-$(t+1)$ ruler will have a stronger identity $(c_{t+1} \geq c_t)$. The only way for the date-t ruler to affect his successor's behavior is to induce even more conversions than the latter; but this strategy lowers the date-t ruler's payoff relative to playing myopically, both at date t and at date $(t+1)$ as well as the future dates. Conversely, suppose that the date-$(t+1)$ ruler is less eager to convert agents. Then picking the myopic optimum has a double benefit for the date-t ruler as this policy also forces the date-$(t+1)$ to select the date-t ruler's optimum. The formal proof of Proposition 2 follows the lines in Tirole (2016).

It can further be checked that, even if the CA does not set taxes itself, it can still, through a transfer demand T_t, induce aligned LAs to implement the policy described in Proposition 2.[44]

Corollary 2 ***(ratcheting and sufficient statistic).*** *Consider an economy with parameter sequence* $\{c_t, B_t, \lambda_t\}_{t\geq 1}$*. Then the tax base shrinks and the discriminatory tax increases over time. In particular:*
(i) If only c_t *varies, then* $\tau_t = \tau^*(\max_{1\leq k\leq t}\{c_k\}, B-\lambda)$.
(ii) If only $B_t - \lambda_t$ *varies, then* $\tau_t = \tau^*(c, \min_{1\leq k\leq t}\{B_k - \lambda_k\})$.

This corollary sheds light on the argument according to which the tax base shrinks and vanishes as agents convert. The apostasy assumption and its ratcheting corollary validate this "shrinking tax base" argument, but also show that it is not a foregone conclusion. Indeed, the discriminatory tax and tax revenue are constant in a stationary economy for the identity-based model. They are also constant for a non-stationary economy under the extraction model: In the extraction model, the ruler maximizes $\lambda_t + R(\tau_t, \theta_t^*)$ subject to $R(\tau_t, \theta_t^*) \geq B_t - \lambda_t$, where $R(\tau_t) = \tau_t[1 - F(\theta_t^*)]$. From our assumption that the budgetary need can always be met, then $\tau_t = \tau^m$, the monopoly level that maximizes $\tau[1 - F(\tau)]$, for all t.

Corollary 3 ***(time-series comparison with the extraction model).*** *In the extraction model, the tax base and the discriminatory tax are constant over time.*

[44] Either $\theta_{t-1}^* > \tau^a(c_t)$, and then LA i's objective function is $\lambda_t + \tau_t[1 - F(\theta_{t-1}^*)] - T_t$ for $\tau_t \leq \theta_{t-1}^*$, or the smaller $\lambda_t + \tau_t[1 - F(\tau_t)] - T_t$ for $\tau_t > \theta_{t-1}^*$. Strict quasi-concavity then implies that the LA's optimum is at $\tau_t = \theta_{t-1}^*$. For $\tau^a(c_t) \geq \theta_{t-1}^*$, the equilibrium policy can be decentralized by similarly setting a transfer demand $T_t = \max\{B_t, \lambda_t + \tau^a(c_t)[1 - F(\tau^a(c_t))]\}$.

Time-increasing relevance of extraction model. We can now formalize the related claims that (a) later rulers have less influence on tax policy and outcomes than earlier ones, and that (b) the extraction model becomes more and more relevant over time. In our model, the tax policy has two components: a uniform tax that is driven by extractive motives, and a discriminatory tax that reflects both extractive and identity considerations. The claims are based on the idea that over time the ruler (statistically) induces fewer conversions and a smaller decrease in discriminatory tax revenue than his predecessors. This might seem obvious as the group of agents having kept their identity shrinks over time. However, this is not a prediction of the extraction model; relatedly, the remaining non-converts have a very strong identity and are willing to pay much for keeping it.[45]

Delegation in time-series model. We assumed for expositional simplicity that LAs were congruent with the CA. However, even in the presence of (possibly district-specific) agency problems, the optimality of myopic behaviors and the ratchet property still hold.[46] This implies for example, a strong-identity earlier ruler may not have had the opportunity to convert as many Copts as he desired because of the strong presence of Copts among the LAs, leaving scope for further conversions by subsequent rulers who were not necessarily more religious.

The rest of this section assumes no-divergence and focuses on the (exogenous or endogenous) direct drivers of external and internal challenges, and not on variations in the exogenous parameters studied in Proposition 2 and its corollaries:

Assumption 1 *In the rest of the section, $c_t = c$, $\lambda_t = \lambda$, and $B_t = B$ for all t.*

(c) External threats

[45]Our claim is based on Proposition 2, which states that $\theta_t^* = \max_{1 \le k \le t} \tau^*(c_k, B_k - \lambda_k)$. Suppose that the joint distribution of the ruler's type c_t and of the net budgetary needs $B_t - \lambda_t$ is the same over time. This generates a distribution $H(\tau_t)$ on some interval $[\underline{\tau}, \overline{\tau}]$ for the date-t ruler's *desired* discriminatory tax $\tau^*(c_t, B_t - \lambda_t)$(the actual one, as we showed, may be constrained by previous choices), and a cumulative distribution function $H^{t-1}(\max_{1 \le k \le t-1} \tau^*(c_k, B_k - \lambda_k))$ for the highest discriminatory tax prior to date t. The expected number of conversions at date t is equal to $\int_{\underline{\tau}}^{\overline{\tau}} [\int_{\chi}^{\overline{\tau}} [F(\tau) - F(\chi)] dH(\tau)] dH^{t-1}(\chi)$, which after an integration by parts can be shown to be decreasing in t. Similarly, using the fact that the discriminatory tax is on the downward-sloping side of the Laffer curve, the expected reduction in discriminatory tax revenue from date $t-1$ to date t is $\int_{\underline{\tau}}^{\overline{\tau}} [\int_{\chi}^{\overline{\tau}} [R^c(\chi) - R^c(\tau)] dH(\tau)] dH^{t-1}(\chi)$ and is decreasing in t.

[46]Consider for example the case in which, at date t, the CA has identity c_t and the LAs identity c_t^{LA} (again, this identity could be district specific, at the cost of heavier notation). The date-t CA must then account for the implementability constraint $\tau_t \in [\tau^a(c_t^{LA}), \tau^m]$ if $c_t^{LA} \le 0$ and $\tau_t \in [\tau^m, \tau^a(c_t^{LA})]$ if $c_t^{LA} \ge 0$. Let (I) denote this implementability condition. Let $\tau^*(c_t, c_t^{LA}, B_t - \lambda_t) \equiv \arg\max_{\{(I), \tau_t[1-F(\tau_t)] \ge B_t - \lambda_t\}} (\tau_t - c_t)[1 - F(\tau_t)]$ denote the CA's desired static discriminatory tax, characterized in Proposition 1 and Figure 1. The equilibrium discriminatory tax in the time-series model is then $\tau_t = \max\{\tau^*(c_t, c_t^{LA}, B_t - \lambda_t), \theta_{t-1}^*\}$.

Suppose that there is conditional probability $x_t \geq 0$ that the ruler is evicted at date t for some external reason. When the ruler is evicted, taxes –or at least the discriminatory tax– are no longer collected.[47] We assume that the ruler cares not only about taxes and current conversions, but also about his "legacy", that is the number of converts in the future even if he is kicked out (so, we keep the payoff function described in (3). In particular, the date-t ruler still internalizes $-c[1-F(\theta^*_{t+k})]$ even if an external challenge has annihilated taxes by date $t+k$). The uncertainty about the ruler's perennity makes agents more reluctant to convert as doing so eliminates the option value of having kept one's identity. This option value was shown to be equal to 0 in Proposition 2, but is strictly positive here. Except for the sequence $\{x_t\}_{t\geq 1}$, all parameters are invariant as stated in Assumption 1, and we suppose that the budget constraint is never binding (the analysis can be generalized if that is not the case).

Proposition 3 ***(option value under external threats).*** *Let* $\tau^c \equiv \arg\max_{\{\theta\}}(\theta - c)[1 - F(\theta)]$ *and* $K_t \equiv \left(1 + \frac{\beta}{1-\beta}x_{t+1}\right)$. *In equilibrium,*[48] *the date-t discriminatory tax is* $\tau_t = K_t\tau^c$ *and the discriminatory tax revenue is* $R_t = K_t\tau^c[1-F(\tau^c)]$. *In particular, if* x_t *is weakly decreasing (increasing) over time, so are* τ_t *and* R_t. *All conversions occur at date 1.*

Corollary 4 ***(external threats: comparison with the extractive model).*** *The external-threats dynamics for the extractive model are identical with those of the identity model, except that the stable fraction of converts is* $F(\theta^m)$ *where* θ^m *solves* $\max\{\theta[1-F(\theta)]\}$.

(d) Internal threats and time-decreasing resistance

To facilitate the understanding of endogenously evolving internal challenges, we first gain intuition about the threat of rebellion by analyzing the static case (Proposition 4) and then state our main proposition (Proposition 5).

(d1) Static analysis of the rebellion threat. Let us first consider the static case and show that in the presence of a rebellion threat, the CA in general benefits from having the ability to cap the uniform tax at a level lower than the extractive level λ. Assume that it takes $[1-F(\hat{\theta})]$ rebels to topple the CA, and the individual cost of doing so is ρ. In the following, we will say that the threat of rebellion is low (resp. high) if $\hat{\theta}$ is low (high), that is if the number of required rebels is high (low); we could alternatively index the threat of rebellion by (minus) the cost ρ of rebelling. To avoid unnecessary notation, assume $\hat{\theta} \geq 0$. The no-rebellion constraint for taxes $\{\hat{\lambda} \leq \lambda, \hat{\tau}\}$ is that the rebellion cost

[47] The analysis would hold as long as the discriminatory tax is reduced in the future.

[48] As in Proposition 2, equilibrium uniqueness requires further assumptions in the case of an infinite horizon.

exceeds the marginal rebel's gain $G(\hat{\theta})$ from a successful rebellion:[49]

$$\rho \geq \hat{\lambda} + \min\{\hat{\tau}, \hat{\theta}\} \equiv G(\hat{\theta}).$$

Assumption 2 ***(relevant rebellion threat).*** $\lambda + \min\{\hat{\theta}, \tau^c\} > \rho$.

Recall that in the absence of rebellion threat, the CA's first best is $\hat{\lambda} = \lambda$ and $\hat{\tau} = \tau^c$. Were Assumption 2 violated, the threat of rebellion would be irrelevant and the first-best level of taxes $\{\lambda, \tau^c\}$ would prevail. We look at the optimal pair $\{\hat{\lambda} \leq \lambda, \hat{\tau}\}$ of taxes that the CA would like to implement. Let $\tilde{\tau} < \tau^c$ be uniquely defined as $\arg\max\{R^c(\hat{\tau}) - \hat{\tau}\}$ or $\tilde{\tau} + \frac{F(\tilde{\tau})}{f(\tilde{\tau})} = c$ (this is the optimal discriminatory tax when an increase in that tax must be offset 1-for-1 by a decrease in the uniform tax). The CA picks the discriminatory tax rate that maximizes $\hat{\lambda} + (\hat{\tau} - c)[1 - F(\hat{\tau})]$ subject to $\hat{\lambda} \leq \lambda$ (feasibility), $\hat{\lambda} + \min\{\hat{\tau}, \hat{\theta}\} \leq \rho$ (no-rebellion constraint) and $\hat{\tau} \in [\tau^m, \tau^c]$ (implementability). For the sake of simplicity, we do not put any lower bound at 0 for $\hat{\lambda}$ (uniform subsidies are feasible).

Finally, let $\theta^* \in [\tau^m, \tau^c]$ be defined by $\theta^* \equiv R^c(\tau^c) - R^c(\tau^*) + \tau^*$, where $\tau^* \equiv \max\{\tau^m, \tilde{\tau}, \rho - \lambda\} \in [\tau^m, \tau^c]$.

Proposition 4 ***(capping the uniform tax to thwart rebellion: the static case).*** *Under Assumptions 1 and 2,*

(i) For a low threat of rebellion ($\hat{\theta} < \theta^*$), the marginal rebel is a convert; the optimal policy for the CA is to reduce the uniform tax to $\hat{\lambda} = \rho - \hat{\theta} < \lambda$, and to keep the discriminatory tax at $\hat{\tau} = \tau^c$.

(ii) For a high threat of rebellion ($\hat{\theta} > \theta^*$), the optimal policy for the CA is to reduce both the uniform tax from λ to $\rho - \tau^*$ and the discriminatory tax from τ^c to τ^*. The marginal rebel is a non-convert.

(d2) Dynamic analysis of the rebellion threat. Suppose next that $t = 1, 2, ..., +\infty$ and that agents and the CA apply the same discount factor β to future utilities. The assumption that $T = +\infty$ is important here; for, with a finite horizon, the gain from a successful rebellion would decrease over time, generating an artificial increase over time in the relative cost of rebellion (expressed relative to future benefits). We assume that the cost of rebellion is $\rho/(1-\beta)$: while rebellion is a one-shot activity, we normalize its per-period cost to be ρ to facilitate the comparison with the static legitimacy model. The willingness to pay to keep one's identity is still θ per period.

Proposition 5 ***(far-sighted players and decreasing resistance).*** *Let $\tau^* \equiv \max(\tau^m, \tilde{\tau}, \rho - \lambda)$ and $\theta^* \equiv R^c(\tau^c) - R^c(\tau^*) + \tau^* \in (\tau^*, \tau^c)$. Under Assumptions 1 and 2:*

[49] Allowing for negative values of $\hat{\theta}$, this condition would be $\rho \geq \lambda + \min\{\max\{\hat{\theta}, 0\}, \tau\}$.

(i) If $\hat{\theta} < \theta^*$, the marginal rebel $\hat{\theta}$ converts at date 1. In the CA's optimal MPE, the CA backloads the uniform tax, charging a low uniform tax at date 1 and raising the uniform tax to $\min\{\lambda,\rho\}$ *once the threat of rebellion has subsided. The discriminatory tax is equal to τ^c in all periods.*

If $\hat{\theta} > \theta^*$ and $\rho - \tau^* \leq \lambda$, the marginal rebel $\hat{\theta}$ never converts. The discriminatory tax and the uniform tax are equal to $\tau_t = \tau^*$ and $\lambda_t = \rho - \tau^*$ for all t.

Despite the lack of commitment, the CA's per-period welfare is in both cases the same as in the static model, namely $\rho - \tau^* + R^c(\tau^*)$ for $\hat{\theta} \leq \theta^*$ and $\rho - \hat{\theta} + R^c(\tau^c)$ for $\hat{\theta} \geq \theta^*$.

(ii) The MPE maximizing the CA's payoff (characterized in part (i)) is furthermore coalition-proof à la Bernheim-Peleg-Whinston (1987) if $\tau^c \geq \beta \frac{\hat{\theta}+\lambda-\rho}{1-\beta}$ when $\hat{\theta} < \theta^*$, or if $\hat{\theta} > \theta^*$.

Intuition. Assume in a first step that all parties are myopic ($\beta = 0$); in particular, each generation cares about its own welfare, but apostasy creates a linkage between periods as conversions apply to future generations. A key insight is that, when the marginal rebel is a convert, the marginal rebel's incentive to rebel decreases over time, as depicted in Figure 2(a) in the two-period case. Earlier converts' gain from a successful rebellion is limited to the uniform tax and no longer includes the preservation of their foregone identity. Thus, suppose that the threat of rebellion is not too high: $\widehat{\theta} < \theta^*$. A myopic CA then selects $\{\widehat{\lambda} = \rho - \widehat{\theta}, \widehat{\tau} = \tau^c\}$ at date 1. All types $\theta \leq \tau^c$ including the marginal rebel convert at date 1. Because the marginal rebel cares only about the uniform tax from date 2 on, the no-rebellion constraint at date 2 and at any subsequent date t yields:

$$\lambda_t = \min\{\rho, \lambda\}$$

and the date $t \geq 2$ welfare becomes $\min\{\rho,\lambda\} + R^c(\tau^c)$. This is also the maximal welfare that can be obtained in any given period: The uniform tax cannot exceed $\min\{\rho,\lambda\}$ without triggering a rebellion, and $R^c(\tau^c)$ is the maximum adjusted revenue from the discriminatory tax. All conversions occur at date 1, as the discriminatory tax is constant at τ^c from date 1 on. But the uniform tax increases from $\lambda_1 = \rho - \widehat{\theta} < \min\{\rho,\lambda\}$ to $\lambda_2 = \lambda_3 = \cdots = \min\{\rho,\lambda\}$ once the threat of rebellion has decreased. In particular, it increases to equal the extractive tax if $\lambda \leq \rho$.

By contrast, when the marginal rebel is a non-convert in the static analysis ($\widehat{\theta} > \theta^*$, see Figure 2(b)), the threat of rebellion remains the same over time. The CA in each period must still satisfy $\widehat{\lambda}_t + \widehat{\tau}_t \leq \rho$ for each t. And so, $\widehat{\tau}_t = \tau^*$ and $\widehat{\lambda}_t = \rho - \tau^*$ for all $t \geq 1$. The equilibrium is stationary and replicates the static analysis in each period.

When agents are far-sighted ($\beta > 0$), one might guess that the agents' resistance in this case would no longer subside over time, as they internalize the fact that not

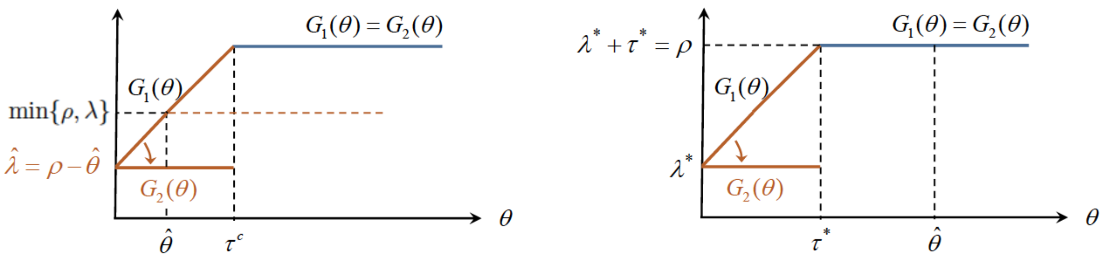

(a) Low rebellion threat (b) High rebellion threat

Figure 2: **Time-decreasing resistance**

Note: $G_t(\theta)$ = type θ's gain from a successful rebellion at date t.

rebelling will lead to an increase in future taxes. Interestingly, this is not the case. The reason has to do with the difference in objectives between marginal and inframarginal agents when the marginal rebel is a convert; the marginal rebel is then concerned solely with the discounted flow of uniform taxes; by contrast, agents who do not convert are affected by both the uniform and the discriminatory discounted taxes, as is the CA. The CA can soft-pedal uniform taxes and backload their flow so as to dissuade the marginal convert from rebelling. Put differently, he can divide and conquer the agent community. Once the resistance of the converts has been reduced, the CA can then increase the tax burden.

Finally, we note that the extraction model, which is a special case of the identity-based model[50], exhibits the same pattern regarding decreasing resistance:

Corollary 5 ***(comparison with the extraction model).*** *Under Assumptions 1 and 2,*
(i) If $\widehat{\theta} < \tau^m$, the marginal rebel is a convert; the discriminatory tax is equal to its extractive level τ^m in all periods, while the uniform tax is raised over time from $\rho - \widehat{\theta}$ in the first period to max$\{\rho, \lambda\}$ thereafter.
(ii) If $\widehat{\theta} \geq \tau^m$, the marginal rebel is a non-convert; the discriminatory tax is equal to its extractive level τ^m in all periods, while the uniform tax is constant at level $\rho - \tau^m < \lambda$ over time.

4.2 Empirics

4.2.1 Time-series implications of the extraction and identity-based models

We summarize these implications in Table 4. The outcomes are the same as in the cross-sectional analysis: F_t, τ_t, and R_t. We add a fourth one: the removal of the cap on

[50]It satisfies in particular $\tau^* = \tau^m$ when the threat of rebellion is binding.

λ_t circa 750. There are three exogenous determinants: (1) identity strength of CA (c_t), (2) budgetary need (B_t), and (3) probability of CA's eviction due to an external threat (x_t), in addition to an "endogenous" determinant (in the sense of being driven by other exogenous parameters in the model): (4) the (necessarily) declining internal threat of rebellion ($G_t(.)$).

Both models imply that in a stationary environment, all conversions must occur at date 1, and F_t, τ_t, and R_t remain constant thereafter. The two models differ, though, in their prediction (and explanation) of conversion waves and tax rises in non-stationary environments. On the one hand, the extraction model does not predict any new conversions (ΔF_t) after date 1. However, it implies that τ_t and R_t may change in response to external threats: as x_t increases (decreases), τ_t and R_t both go up (down), while F_t remains unchanged. On the other hand, while the identity-based model makes the same implications for x_t, it offers two additional explanations of (later) poll tax hikes and conversion waves, that are not implied by the extraction model: (1) an increase in c_t above the maximum historical level, or (2) a decrease in B_t below the minimum historical level; both factors result in a rise in τ_t and F_t, and a decline in R_t (however, a decrease in c_t, or an increase in B_t, leaves the tax policy unaltered). To sum up, the identity-based model implies a more gradual decline in F_t, in response to spikes in c_t or busts in B_t, than the extraction model. It also offers more reasons (other than x_t) for why τ_t may rise over time. Finally, according to both models, the decrease in $G_t(.)$ cannot explain changes in τ_t, F_t, and R_t, but can explain lifting the cap on λ_t. Notice that λ_t, being extractive, is insensitive to changes in c_t, B_t, and x_t.

Table 4: **Time-series implications of the extraction and identity-based models**

Δ Outcome	Extraction model				Identity-based model			
	B_t	c_t	x_t	$G_t(.)$	B_t	c_t	x_t	$G_t(.)$
	↑/↓	↑/↓	↑/↓	↓	↑/↓	↑/↓	↑/↓	↓
$\Delta\tau_t$	0	0	+/−	0	0/+	+/0	+/−	0
ΔF_t	0	0	0	0	0/+	+/0	0	0
ΔR_t	0	0	+/−	0	0/−	−/0	+/−	0
$\Delta\lambda_t$	0	0	0	+	0	0	0	+

Notes: (1) ↑: variable increases over time, ↓: variable decreases over time. (2) The effects of the exogenous parameters, c_t, B_t, and x_t are analyzed without the threat of rebellion. The decline in $G_t(.)$ is endogenously driven by past conversions.

4.2.2 Measuring c_t, B_t, x_t, and $G_t(.)$

To test the two models apart, we exploit the time-series variation in c_t and B_t under the Arab Caliphate, controlling for x_t and $G_t(.)$. We constructed two datasets, the first is at Egypt's governor level ($N = 122$), and the second at the (higher) Caliph level ($N = 65$), covering a period of 530 years from 641 to 1170. Two remarks are in order. First, the two levels (governors and Caliphs) are identical in 868–905 and 935–1170, when Egypt was independent.[51] Second, as we explain below, the two datasets differ in their measure of c_t. They use the same measures of B_t, x_t, $G_t(.)$, $\Delta\tau_t$, and ΔF_t, though, aggregated to either the governor level (dataset 1) or the Caliph level (dataset 2).

Identity strength of CA (c_t). We measure c_t by two binary measures ($\hat{c}_t$), a governor-level measure based on Coptic narratives, and a Caliph-level measure based on Muslim narratives. The first measure captures Egypt's (fiscal) governors' hostility toward non-converts, according to the portrayals of governors in Coptic contemporary chronicles. Specifically, we constructed a dummy variable that takes value 1 if a governor t is portrayed as hostile to non-converts, and 0 if neutral, unmentioned, or friendly.[52] We employ *The Chronicle of John of Nikiu* for the Rashidun period (641–661), and the *History of the Patriarchs of the Coptic Church of Alexandria*, the *Liber Pontificalis* (Book of the Popes) of the Coptic Church, for 661–1170.

A few remarks are in order. (1) Coptic chronicles offer an important advantage: They provide rare fine-grained information on the attitudes of Egypt's (fiscal) governors, even some of the least known ones, in 641–868 and 905–935, when Egypt was a province of the Caliphate. Unlike Caliphs (who were absentee rulers of Egypt), governors typically receive only a brief mention in Muslim narratives. (2) We construct our measure from explicit statements on attitudes (i.e. preferences), and not policies, of governors.[53] (3) *History of the Patriarchs* was compiled by a single author for the period 661–1000: Severus ibn al-Muqaffa, the bishop of *Hermopolis* in the late 9th century, which mitigates the concern about comparability across multiple authors. (4) Scholars agree that Severus "compiled" existing Coptic narratives, and did not "create" them, which mitigates the concern that governors' portrayals may merely reflect Severus's perceptions. (5) Although Coptic chronicles are often ideologically biased against Muslims, this bias is unlikely to vary systematically across governors.

[51] During the late Fatimid period (1074–1170), the Fatimid Caliph was a figurehead, and Egypt was de facto ruled by viziers. Thus, for this period, we treat the vizier as the effective governor/Caliph.

[52] Out of 122 governors, there are 15 who are coded as hostile, 12 coded as friendly, 28 coded as neutral, and 67 who are not mentioned. However, unmentioned governors have a much shorter tenure, on average: Mentioned governors ruled for a total of 374 years, or 71% of the period of 530 years (hostile: 119 years; friendly: 117 years; neutral: 138 years). We pool neutral, unmentioned, and friendly governors, because they are theoretically unlikely to generate poll tax rises and conversions. And indeed they do not differ statistically with respect to the likelihood of poll tax rises and conversions.

[53] Examples of hostile governors include Amr ibn al-As (641–646; 659–663), who "had no mercy on the Egyptians," al-Asbagh (685–705), who "was a hater of the Christians." Examples of friendly governors include al-Layth ibn al-Fadl (798–803) who "was a good man and favored the Christians."

Notwithstanding its advantages, this measure may raise an (arguably, inevitable) concern: Even if it is based on Coptic perceptions of governors' attitudes, these perceptions can themselves be shaped by governors' behavior, most importantly their tax policy. This can generate a spurious correlation between governors' hostility and changes in τ_t and F_t.

To address this concern, our second measure captures Caliphs' piety in the Muslim sense (i.e. independent of hostility toward non-converts), based on Muslim narratives. We focus on one aspect of piety that is arguably comparable across Sunni Caliphs in 641–969 and Ismaili Shiite (Fatimid) Caliphs in 969–1170: a dummy variable =1 if a Caliph is not known for drinking alcohol. Unlike the first measure, though, we are only able to measure Muslim piety at the Caliph level. We employ Sirhan (1978) (a secondary source based on Muslim medieval narratives) for 641–868, and two medieval sources: al-Dhahabi's *The Lives of Noble Figures* for 868–969 and al-Maqrizi's *History of the Fatimid Caliphs* for 969–1170.[54]

We plot the two measures of c_t in Appendix Figure B.4. Both measures suggest that c_t was high in 641–750, and that it declined in 750–868. The two measures deviate, though, after 868: While measure 1 suggests that c_t remained low, measure 2 indicates that c_t increased.

Budget needs (B_t). We measure (low) B_t by a dummy variable ($\hat{B}_t$) that takes value 1 if there is no major military campaign (whether offensive or defensive) by the Caliphate against a foreign empire, during the reign of governor t (dataset 1) or Caliph t (dataset 2). Funding military campaigns presumably necessitates a bigger budget. It is possible under the identity-based model, but not the extraction model, that c_t and B_t are positively correlated: more hostile governors, or more pious Caliphs, may want to conquer more. This will mitigate the impact of each of c_t and B_t on tax policy, though; in other words, we will be underestimating the true impact. Appendix Figure B.5 shows that military campaigns followed a U curve over time: They were more frequent in 641–750 due to the Caliphate's expansionary wars, became relatively less frequent between 750 and 868 when the Abbasid Caliphate consolidated power over the conquered territories, disappeared in 868–1100, before they increased again in 1100–1170 due to the (defensive) campaigns against the Crusaders.

External threats (x_t). We measure x_t by a dummy variable that takes value 1 if there is a threat of foreign attacks during the reign of ruler t. To create this variable, we

[54]al-Dhahabi is considered one of the main trusted sources among Sunni Muslims on the personal biographies of Muslim politicians and clergy. al-Maqrizi (despite being Sunni) is considered one of the main objective medieval sources on Fatimid Caliphs.

first constructed yearly data on foreign attacks on the Caliphate in 641–1170. We then defined an external threat as a dummy variable =1 during the 5 years preceding (but not including) the year when the actual attack takes place. Appendix Figure B.5 reveals that external threats also followed a U curve: They were more frequent in 641–750, due to the Byzantine (and Nubian) attacks. They subsided after 750, before they spiked in the late Fatimid period, due to the First and Second Crusades (1096–1099, 1147–1152), and the Crusaders' (failed) invasion of Egypt (1154–1169).

Threat of rebellion ($G_t(.)$). We control for the decline in $G_t(.)$ by a trend in time. We prefer to not employ the cumulative sum of past conversions (which theoretically drives the decline in $G_t(.)$) as a proxy, because it is endogenous from an empirical viewpoint.

4.2.3 Impact on discriminatory tax rate (τ_t) and conversions (F_t)

Measuring $\Delta\tau_t$ and ΔF_t. We employ Coptic chronicles to measure (1) $\Delta\tau_t$ by a dummy variable =1 if a poll tax rise is mentioned during the reign of governor t (dataset 1), or Caliph t (dataset 2), and =0 if no poll tax rise is mentioned, and (2) ΔF_t by a dummy variable =1 if there is a conversion wave to Islam among Copts during the reign of governor t (dataset 1), or Caliph t (dataset 2), and =0 if no conversion wave is mentioned. A no mention of a poll tax rise or a conversion wave implies that τ_t and F_t remained constant: Coptic chronicles do not mention decreases in the poll tax rate, or reverse conversions to Coptic Christianity.

Using Coptic chronicles to measure $\Delta\tau_t$ and ΔF_t offers two advantages: (1) They provide details on the actual poll tax policies of (fiscal) governors and Caliphs (the enforcement of which was delegated to the LAs), instead of the *de jure* poll tax rate which shows little variation between 641 and 1170. (2) They describe conversion waves that are omitted by Muslim sources, and at a higher frequency than Bulliet (1979)'s, Courbage and Fargues (1997)'s, and Saleh (2018)'s estimates of F_t (see Section 2).

Long-term trends. We depict the long-term trends of the incidence of poll tax rises and conversion waves in Appendix Figure B.6. First, we observe poll tax hikes and conversion waves in later episodes, and not only at "date 1," which can be roughly defined as the Rashidun period (641–661). While the trajectory of conversions is not consistent with the extraction-only model (recall that this model does not predict any new conversions after date 1), poll tax hikes may (or may not) be explained by external threats, according to both models. Second, both time series reveal a declining trend: tax rises and conversion waves became less frequent over time. As we have shown in Section 4.1, this may be a statistical consequence of the shrinking tax base that is implied by the identity-based model. However, it remains to be examined if poll tax hikes and conversions can be indeed explained by spikes in c_t, or busts in B_t, as implied by the identity-based model. We thus turn now to the formal evidence.

Empirical specification. We focus on the effects of c_t and B_t on the probability of poll tax hikes and conversion waves under ruler t, because they enable us to disentangle the two models (recall that the two models make the same implications for x_t and $G_t(.)$). We think of our binary measures, $\hat{c}_t$ and $\hat{B}_t$, as truncations at some level c^* or B^*: $\hat{c}_t = 1$ if $c_t \geq c^*$ and $\hat{c}_t = 0$ otherwise, and $\hat{B}_t = 1$ if $B_t \leq B^*$ and $\hat{B}_t = 0$ otherwise. Letting $c^{t-1} \equiv max_{1 \leq k \leq t-1} c_k$ and $B^{t-1} \equiv min_{1 \leq k \leq t-1} B_k$, and $\hat{c}^{t-1} \equiv max_{1 \leq k \leq t-1} \hat{c}_k$ and $\hat{B}^{t-1} \equiv max_{1 \leq k \leq t-1} \hat{B}_k$, the associated binary variables, Proposition 2 then implies that, ceteris paribus, a ruler with binary type $\hat{c}_t = 0 < \hat{c}^{t-1} = 1$, or $\hat{B}_t = 0 < \hat{B}^{t-1} = 1$, keeps the discriminatory tax and conversions unchanged; a ruler with $\hat{c}_t = 1 > \hat{c}^{t-1} = 0$, or $\hat{B}_t = 1 > \hat{B}^{t-1} = 0$, increases the discriminatory tax and induces conversions. When $\hat{c}_t = \hat{c}^{t-1}$, or $\hat{B}_t = \hat{B}^{t-1}$, ruler t may or may not increase the discriminatory tax and induce conversions. Suppose that c_t and B_t are each independent random draws from distributions $G(c_t)$ and $H(B_t)$, respectively.[55] Then, letting $n^c_{t-1} \equiv \sum_{1 \leq k \leq t-1} \hat{c}_k$ and $n^B_{t-1} \equiv \sum_{1 \leq k \leq t-1} \hat{B}_k$ denote the number of realizations $\hat{c}_k = 1$ and $\hat{B}_k = 1$ up to $t-1$, respectively, we have for instance for $n^c_{t-1} \geq 1$, $E[F_t - F_{t-1}] = \hat{c}_t \int_{c^*}^{+\infty} [F(\tau^a(c_t)) - F(\tau^a(c^{t-1}))] \frac{dG(c_t)}{1-G(c^*)} \frac{dG^{n^c_{t-1}}(c^{t-1})}{[1-G(c^*)]^{n^c_{t-1}}}$. So, in reduced form $E[F_t - F_{t-1}] = \hat{c}_t W(n^c_{t-1})$, where W is a decreasing function converging to 0 as n^c_{t-1} goes to infinity. To sum up, the probability of poll tax rises and conversion waves is increasing in each of $\hat{c}_t$ and $\hat{B}_t$, and is decreasing in each of n^c_{t-1} and n^B_{t-1}.

In both datasets, we observe that $\hat{c}^{t-1} = 1$ for every ruler $t \geq 2$ and $\hat{B}^{t-1} = 1$ for every ruler $t \geq 4$. Hence, we distinguish between two cases for $\hat{c}_t$ (for $t \geq 2$): (1) $\hat{c}_t = \hat{c}^{t-1} = 1$, and (2) $\hat{c}_t = 0 < \hat{c}^{t-1} = 1$. Similarly, there are two cases for $\hat{B}_t$ (for $t \geq 4$): (1) $\hat{B}_t = \hat{B}^{t-1} = 1$, and (2) $\hat{B}_t = 0 < \hat{B}^{t-1} = 1$. Hence, Proposition 2 implies the two following testable hypotheses:

Hypothesis 1. For each of $\hat{c}_t$ (for $t \geq 2$) and $\hat{B}_t$ (for $t \geq 4$), the identity-based model implies that the probability of poll tax hikes and conversions is higher, on average, under case (1) than under case (2). The extraction model implies no difference in average probability.

Hypothesis 2. The identity-based model implies that the probability of poll tax rises and conversions after date 1 is decreasing in each of n^c_{t-1} and n^B_{t-1}: earlier high-c_t (and low-B_t) rulers have more influence than later rulers. The extraction model implies that poll tax rises and conversions after date 1 are insensitive to both n^c_{t-1} and n^B_{t-1}.

[55] As noted in the discussion of budget needs in Section 4.2.2, B_t and c_t may be correlated. But this does not affect the reasoning below.

To test Hypotheses 1 and 2, we estimate the following model separately for poll tax hikes and conversion waves in 641–1170 starting from ruler $t \geq 4$. For each outcome, we estimate the model using both the governor-level and Caliph-level datasets:

$$outcome_t = \beta_0 + \beta_1 \hat{c}_t + \beta_2 n^c_{t-1} + \beta_3 \hat{B}_t + \beta_4 n^B_{t-1} + \beta_5 Z_t + \varepsilon_t, \qquad t \geq 4 \qquad (7)$$

where we examine two outcomes: (1) a dummy variable =1 if a poll tax rise is mentioned under ruler t, and (2) a dummy variable =1 if a conversion wave is mentioned under ruler t. The main regressors are: (1) $\hat{c}_t$ =1 if governor t is hostile to non-converts (dataset 1), or Caliph t is pious (dataset 2), and =0 otherwise: for $t \geq 2$, value 1 implies that $\hat{c}_t = \hat{c}^{t-1} = 1$, while value 0 implies that $\hat{c}_t = 0 < \hat{c}^{t-1} = 1$, (2) n^c_{t-1}, which captures the historical legacy of past high-c_t rulers, (3) $\hat{B}_t$ =1 if the Caliphate did not embark on a military expedition under ruler t, and =0 otherwise: for $t \geq 4$, value 1 implies that $\hat{B}_t = \hat{B}^{t-1} = 1$, while value 0 implies that $\hat{B}_t = 0 < \hat{B}^{t-1} = 1$, and (4) n^B_{t-1}, which captures the historical legacy of past low-B_t rulers.

We include in the vector Z_t the following control variables: (5) x_t: a dummy variable =1 if there is an expected foreign attack during the reign of ruler t, (6) a dummy variable =1 if there is at least one adverse Nile shock under ruler t (Chaney 2013), as a measure of negative economic shocks that may reduce taxpayers' ability to pay τ_t, λ_t, or both, and (7) $G_t(.)$: year when ruler t ascended to power, which captures the trend in time.

Four remarks are in order: First, the binary coding of each of c_t and B_t implies a more general specification that includes the following interaction terms: $\hat{c}_t \times n^c_{t-1}$ and $\hat{B}_t \times n^B_{t-1}$. When we do so, however, the results become noisy and statistically insignificant, due to the high multicollinearity between n^c_{t-1} and n^B_{t-1}.[56]

Second, equation (7) is a special case of the Almon finite distributed-lag model: $outcome_t = \gamma_0 + \sum_{s=0}^{t-1} \gamma_{1s} \hat{c}_{t-s} + \sum_{s=0}^{t-1} \gamma_{2s} \hat{B}_{t-s} + \gamma_3 Z_t + u_t$: (1) n^c_{t-1} and n^B_{t-1} assign equal lag weights (=β_2 and β_4, respectively) for all lags $s \geq 1$. This implies a "strong memory:" lag weight does not decay over time (as suggested by theory), (2) $\hat{c}_t$ and $\hat{B}_t$ imply a different lag weight (=β_1 and β_3, respectively) for the current ruler t ($s = 0$).

Third, we weight each dataset used in each regression by a frequency weight that is equal to the length of ruler t's tenure. This presumes that rulers who stayed longer in power had more influence on tax policy and conversions, which is realistic from a historical viewpoint: it generally took time for governors (dataset 1) to consolidate power, align the interests of the LAs with their own, and hence implement policies. The case is even stronger for Caliphs (dataset 2), who had to align, not only LAs' interests, but also governors' interests.

[56] An even more general specification is to allow c_t and B_t to be correlated by including a full set of interaction terms of $\hat{c}_t$, $\hat{B}_t$, n^c_{t-1}, and n^B_{t-1}. Again, this specification produces noisy and statistically insignificant estimates.

Fourth, we estimate Newey-West standard errors assuming that the error structure is heteroskedastic and serially correlated up to 15 lags (dataset 1) and 11 lags (dataset 2). We determined the number of lags (m) using Lazarus et al. (2018)'s rule of thumb: $m = (1.3)T^{1/2}$, with rounding up, where T is the number of governors/Caliphs.

According to Hypothesis 1, H_0 on β_1 and β_3 is the extraction model: both coefficients are equal to 0, whereas H_1 is the identity-based model: both coefficients are positive. According to Hypothesis 2, H_0 on β_2 and β_4 is the extraction model: both coefficients are equal to 0, while H_1 is the identity-based model: both coefficients are negative.[57]

Findings. The results are shown in Table 5. First, we find that the probability of a poll tax rise and of a conversion wave are both increasing in $\hat{c}_t$, which is consistent with the identity-based model. The contemporaneous effect of $\hat{c}_t$ retains its magnitude when controlling for n^c_{t-1} and n^B_{t-1}. In terms of magnitude, a hostile governor ($\hat{c}_t = \hat{c}^{t-1} = 1$)'s probability of ordering a poll tax rise is 54% higher than non-hostile governors ($\hat{c}_t = 0 < \hat{c}^{t-1} = 1$) (column 2). In a similar vein, a pious Caliph's probability of ordering a poll tax rise is 37% higher than non-pious Caliphs (column 4). The probability of inducing a conversion wave is 45% higher among hostile governors (column 6), and 36% higher among pious Caliphs (column 8). We trace the smaller effects of Caliph's piety to two reasons: (1) agency: Proposition 1, applied to the Caliph-governor agency relationship, implies that (fiscal) governors whose $c^{governor} < c^{Caliph}$ had more influence on tax policy than Caliphs, and (2) it is measured at a higher level (Caliphs) in 641–868 and 905–935, and thus exhibits less variation. In line with the theory, we trace the contemporaneous effect of $\hat{c}_t$ to two factors that arguably left a leeway for later rulers to influence tax policy and conversions: (1) the high budget needs of earlier high-c_t rulers, which may have suppressed earlier conversions, and (2) the Arabization of LAs, which increased the congruence of CA and LAs over time.

Second, we fail to detect an effect of $\hat{B}_t$, except for conversions in the Caliph-level dataset (column 8). This suggests that under the Arab Caliphate, CA's identity strength had more influence on tax policy and conversions than budget needs.

Finally, while we detect a negative impact of n^B_{t-1}, we fail to detect a robust (negative) impact of n^c_{t-1}, probably due to the high collinearity between the two variables. This suggests that the probability of poll tax hikes and of conversion waves declined

[57]The two models make the same implications for the control variables: (1) x_t has a positive impact on $\Delta\tau_t$ but not on ΔF_t, (2) $G_t(.)$ has no effect on either outcome, and (3) the effect of an adverse Nile shock is ambiguous, because it can cause a leftward shift in θ_t (ability to pay τ_t), λ_t, or both. We do not show the results on the control variables, though, in order to save space. They are mostly statistically insignificant.

over time, which is consistent with the time-increasing relevance of the extraction model.

Table 5: **Time-series determinants of $\Delta\tau_t$ and ΔF_t in 641–1170**

	=1 if discriminatory tax rise is mentioned under ruler t				=1 if conversion wave is mentioned under ruler t			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
=1 if gov. hostile ($\widehat{c}_t = \widehat{c}^{t-1} = 1$)	0.53	0.54			0.47	0.45		
	(0.15)***	(0.15)***			(0.19)**	(0.20)**		
=1 if Cal. pious ($\widehat{c}_t = \widehat{c}^{t-1} = 1$)			0.26	0.37			0.24	0.36
			(0.10)**	(0.15)**			(0.13)*	(0.12)***
$n^c_{t-1} \equiv \sum_{1\le k\le t-1} \widehat{c}_k$	-0.00	0.02	-0.02	0.03	0.00	-0.00	-0.02	0.07
	(0.01)	(0.03)	(0.01)***	(0.02)	(0.01)	(0.02)	(0.01)***	(0.02)***
=1 if no war ($\widehat{B}_t = \widehat{B}^{t-1} = 1$)		0.11		0.08		0.04		0.29
		(0.12)		(0.14)		(0.09)		(0.09)***
$n^B_{t-1} \equiv \sum_{1\le k\le t-1} \widehat{B}_k$		-0.02		-0.07		-0.00		-0.07
		(0.01)***		(0.03)**		(0.01)		(0.03)**
Controls?	No	Yes	No	Yes	No	Yes	No	Yes
Obs (governors/Caliphs)	121	119	64	62	121	119	64	62
Years	524	513	526	509	524	513	526	509
R^2	0.29	0.41	0.16	0.26	0.33	0.38	0.19	0.42
p-value (Breusch–Godfrey test)	0.03	0.06	0.08	0.05	0.00	0.00	0.02	0.02
Mean dep. var.	0.13	0.13	0.25	0.25	0.07	0.07	0.18	0.18

Notes: Data are at the governor level in columns 1, 2, 5, and 6, and at the Caliph level in columns 3, 4, 7, and 8. We omit ruler 1 in columns 1, 3, 5, and 7, and rulers 1–3 in columns 2, 4, 6, and 8. Governors and Caliphs are identical when Egypt is independent in 868–905 and 935–1170. Newey-West standard errors, assuming that the error structure is both heteroskedastic and autocorrelated up to 15 lags in the governor-level regressions, and up to 11 lags in the Caliph-level regressions, are in parentheses. Controls are (1) =1 if foreign attack expected, (2) =1 if adverse Nile shock, and (3) ruler's start year. Regressions are weighted by the length of ruler's tenure. H_0 for the Breusch-Godfrey test is no serial correlation up to 15 lags in columns 1–4, and 11 lags in columns 5–8. * $p<0.10$, ** $p<0.05$, *** $p<0.01$. A constant is included in all regressions.
Sources: See Appendix Section B.2.

Impact on discriminatory tax revenue (R_t). We only observe R_t at a few scattered points in time based on Muslim medieval narratives. We plot these numbers in Appendix Figure B.7. The take-away message is that the poll tax revenue declined rapidly through 813, and continued to decline albeit at a slower rate through 1090. This is consistent with the statistical implication of the identity-based model (shrinking tax base) and the time-increasing relevance of the extraction model: as F_t increased over time, poll tax revenue fell at a decreasing rate.

4.3 Impact on the uniform tax (λ_t)

Historical puzzle. As we demonstrated in Section 2, historians have long documented the tax reform circa 750, by which the Caliphate decided to impose the *kharaj* land tax rate on converts, which implies, in the language of our model, a removal of the cap on the uniform tax. This raises a puzzle for early Islam: the long delay between the Arab Conquest and the tax reform that lifted the cap on the uniform tax. While we cannot study the determinants of the timing of the tax reform econometrically, because it was a one-time Caliphate-wide policy change, we use theory to shed light on its potential cause(s). To be certain, there was a religious cost of lifting this constraint, but doing so a century earlier would have given the CA more leeway in raising finances, which was particularly valuable at a time of high budget needs caused by the expansionary Arab conquests in 641–750. One classic rationalization for the delay is that the poll tax revenue dwindled due to conversions. This argument however is inconsistent with the extraction model, or for that matter any theory that would not put discriminatory taxation on the downward sloping side of the Laffer curve. The hypothesis we propose is that the CA took advantage of the decreasing-resistance property. Consistent with this hypothesis, we note that F_t actually decreased by 750, implying a decline in the internal threat of rebellion. The Abbasid Caliphate thus became more daring to raise the uniform tax. Although the composition of rebels in tax revolts now included both converts and non-converts, the Abbasids eventually managed to suppress the tax revolts by violence, and thus kept the new tax system intact.[58]

Theory. We obtain two corollaries in the simple context of myopic agents and rulers. These corollaries also hold when $\beta > 0$. Consider Proposition 5: The uniform tax is initially low to avoid a rebellion, and so a tax reform is not necessary or at least yields low benefits. Once the threat of rebellion has decreased, though, the uniform tax is optimally raised, which may require a tax reform if the initial cap was low.

Corollary 6 ***(timing of tax reform).*** *Suppose that the uniform tax is initially capped by some level $\overline{\lambda}$, and that removing this cap, allowing any level of uniform tax up to the extraction level $\lambda > \overline{\lambda}$, generates some instantaneous cost $C > 0$ for the CA. Under Assumptions 1 and 2,*
(i) If the threat of rebellion is low ($\widehat{\theta} < \widetilde{\theta}$, for some $\widetilde{\theta}$)[59] and the cap on the uniform tax is binding ($\lambda < \rho - \widehat{\theta}$), then the tax reform occurs at date 1 if $\overline{\lambda} < \rho - \tau^c - C$ and at

[58]We observe 15 tax revolts under the Arab Caliphate, all erupted between 726 and 866. The first 5 revolts (726–783), included only non-converts. Starting from 783, though, rebels included both converts and non-converts.

[59]$\widetilde{\theta}$ in general differs from θ^*, as the cap affects the welfare in the two regions.

date 2 if $\overline{\lambda} \in (\rho - \tau^c - C$, min $\{\rho, \lambda\} - C)$ (it never occurs if c if higher).
(ii) If the threat of rebellion is high ($\widehat{\theta} > \widetilde{\theta}$), then the tax reform, if it ever occurs, always occurs at date 1.

Second, we have assumed for simplicity that the CA is well-informed about the threat of rebellion. As a consequence, rebellions constrain the tax system but do not occur on the equilibrium path. With imperfect information about the threat of rebellion, rebellions in general will occasionally occur in equilibrium. When there is little uncertainty, rebellions will be rare. To obtain results about the composition of the rebel group after date 1 (at date 1 all start non-converts, so only non-converts can rebel), we consider the limit of distributions of the rebellion parameters ρ and $\widehat{\theta}$ converging to the certainty case[60]. The intuition behind the following proposition can be grasped from Figure A.1 (i) and (ii). Suppose for instance that $min\{\rho, \lambda\} = \rho$ and that the marginal rebel is a convert; a small overestimation of the cost of rebellion will lead converts with types in roughly $[\widehat{\theta}, \tau^c]$ and non-converts with types $\theta \geq \tau^c$ to join the rebellion. Compare this with the case in which the marginal rebel is a non-convert. Then a small overestimation of the level of ρ will lead (almost) only non-converts to rebel.

Corollary 7 ***(composition of rebel group).*** *When the uncertainty about the cost of rebellion is small, at date 1, only non-converts rebel when a rebellion occurs. Later on: (i) If the threat of rebellion is low ($\widehat{\theta} < \theta^*$), actual rebellions involve both converts and non-converts. (ii) If the threat of rebellion is high ($\widehat{\theta} > \theta^*$), actual rebellions involve almost only non-converts (the fraction of rebels who are converts tends to 0 as the uncertainty vanishes).*

5 Extensions

This section discusses two extensions, persecutions and emigration. Appendix A develops other extensions: discrimination through discriminatory access to public goods and services, discriminatory empathy, social norms, and Malthusian ruler.
(a) *Persecutions*. Can the CA benefit from replacing a discriminatory tax with an alternative proselytic strategy such as coerced conversions or persecutions[61]? Given their

[60]Two comments are in order here. First, we keep the analysis informal. The notion of vanishing uncertainty is the same as in Nash's celebrated non-cooperative Nash demand game when the uncertainty about the size of the endowment vanishes. Second, the uncertainty could affect other parameters than ρ and $\widehat{\theta}$ without changing the analysis.

[61]This does not mean that forced conversions cannot result from our model. Consider the European-African slave trade (suggested to us by Itzchak Tzachi Raz); Europeans force-converted Africans to Christianity, arguing that they were saving their souls from eternal hell (the Africans' actual utility obviously differed from the Europeans' perception of it). Forced conversions can be understood in the

ignorance of individual preferences, their ability to reach their goals is constrained by incentive compatibility, the fact that agents with the strongest identity are necessarily less likely to convert. A straightforward generalization of the analysis in Stokey (1979) and Riley and Zeckhauser (1983) for our model shows that a CA obtains its highest welfare through a discriminatory tax, and so there is no restriction involved in assuming this particular approach to inducing conversions.[62] Because a discriminatory tax brings the most revenue to the ruler for any desired level of conversion, a question arises as to why rulers may use (or tolerate) a priori inferior non-price instruments such as persecutions. We offer two theories for why persecutions may arise.

(a1) *Agency*. The first theory flows directly from our agency model. The CA may find the LA too soft toward non-converts and too preoccupied with revenue. Allowing mob persecution may then be a second-best way of inducing more conversions.[63] Consider a CA with identity c and an LA with identity c_i. Express the cost of persecution borne by a non-convert in district i, $p_i \geq 0$, in terms of money, so that the agents' total cost of keeping their identity is $\tau_i + p_i$. Persecution does not bring any cash; it only serves to deter the agents from keeping their identity. The CA chooses the level of acceptable persecutions and the LA then collects taxes.[64]

Corollary 8 ***(agency and persecutions)****.* *Consider an economy with parameter sequence* $\{c_t, c_{it}, B_t, \lambda_t\}_{t\geq 1}$. *Then the tax base shrinks and the discriminatory tax increases over time:*

(i) Persecutions do not occur as long as the CA's identity is not much stronger than the LAs' identity: There exists a function c^* *satisfying* $c^*(\tilde{c}) > \tilde{c}$ *for all* $\tilde{c}$ *such that there are no persecutions* ($p_{it} = 0$) *if and only if* $c_t \leq c^*(c_{it})$.

(ii) The ruler is more likely to allow persecutions in districts with the weakest identity.

(a2) *Signaling*. The ruler may use persecutions as a signal. Because high discriminatory taxes by themselves are an efficient signal, one must explain why they are not the

following way in our model: due to their "benevolent" intent, Europeans had a very high utility of conversion (a high c), and so the solution may have been a corner solution with all converting to Christianity (an outcome equivalent to forced conversion). Of course for this to hold, either there must be an upper bound on the support of θ, or the Africans' wealth was limited so that they could not pay a large τ, or both.

[62] The easy observability of individual religious choices matters for the choice of instrument. The deterrence literature stresses that random monitoring calls for much higher penalties to affect behavior, because they are enforced with small probability. This raises the issue of risk aversion or limited liability. Therefore, with infrequent monitoring, non-price instruments, such as jail, the pillory or the death penalty, are more frequent.

[63] The level of persecution could be district specific (as here) or else uniform across districts (in which case only part (i) of the following corollary is relevant).

[64] Persecutions under the Arab Caliphate were ordered by the CA (Caliph or governor).

primary method of signaling. The CA may signal his congruence with the mob at a time he needs strong support from it; implicitly, a stronger-identity ruler must find it less distasteful to persecute the agents (sorting condition). The CA may alternatively use the enforcement of the law- or rather the lack of enforcement!- to signal that persecutions are approved by a majority of citizens as in Bénabou and Tirole (2020). This may designate a scapegoat and strengthen the in-group against a common enemy in a wartime situation for example.

Historians of early Islam are divided on whether conversions to Islam were tax-induced or persecution-induced. Historians who endorse the persecution-based narrative trace conversions to the violent suppression of tax revolts in the 9^{th} century (Mikhail 2004) which crashed non-converts' ability to rebel ever since, or to persecutions (and state-sponsored mob violence) under the Bahri Mamluks in 1250–1354 (El-Leithy 2005). To study persecutions in both periods, we collected information on persecution waves from Coptic and Muslim narratives, leading to the following observations: (1) Persecutions were rare under the Arab Caliphate taking place only under al-Mutawakkil (847–861) and al-Hakim (996–1021), but intensified under the Mamluks. (2) When they occurred, persecutions were always complementary to raising the poll tax (even under the Bahri Mamluks), and hence conversions that occurred during persecution episodes may have also resulted from increasing the poll tax. (3) While we do not observe conversion waves in episodes where the poll tax did not increase, we do observe conversion waves in episodes where persecutions did not occur. (4) Persecutions cannot explain why non-converts were richer, on average, than converts (Saleh 2018), unless we assume that violence was disproportionately directed towards the poor (in reality, persecutions targeted the rich). (5) We do not have localized data on persecutions under the Arab Caliphate, when persecutions occurred under a decentralized tax administration, and thus persecutions then may be consistent with signalling, agency, or both. (6) Persecutions under the Bahri Mamluks are consistent with signalling and not agency, as the poll tax collection was centralized between 1250 and 1315, hence the congruence in identity between the LA and the CA was likely to be high.

(b) *Emigration*. The model allows for emigration as a way for the unwanted population to comply with the ruler's identity. Suppose that polities do not allow agents to change identity, or that identity is inalterable (race, ethnicity), so an agent's choice is between paying the taxes and emigrating. Think of θ as the agent's willingness to pay to stay in the country. The remaining population corresponds to $\theta \geq \theta^* = \lambda + \tau \equiv \hat{\tau}$. Assuming the uniform tax λ (related to the productivity of land, say) remains constant, $V = \lambda + (\hat{\tau} - c)[1 - F(\hat{\tau} - r)] - T$. Thus a simple relabeling shows that our model captures emigration as well.

Although emigration is irrelevant to early Islamic Egypt,[65] it is prominent in many other historical cases. The first example is the *emigration of Zoroastrians from Iran*. The first wave of emigration occurred under the Sunni Samanids (819–999), an independent state of the Abbasid Caliphate. A second wave occurred under the Qajars (1789–1925). Both emigration waves were in response to the imposition of the poll tax and persecutions.

A second example is the *emigration of Christian minorities in Europe* in response to persecutions. Many Anabaptists (Dutch Mennonites) fled from the south of the Netherlands to Germany, England, and the north of the Netherlands, in response to the discriminatory measures put in place by Phillip II of Spain in 1566. Huguenots (French Calvinists) fled from France to England and other parts of the world, in response to Louis XIV's Edict of Fontainebleau, and the consequent persecution of Protestants.

In both examples, rulers were presumably characterized by $c > 0$. In other situations, though, a ruler is not a unitary unit, but an elite with potentially divided preferences. An important example is the *expulsion of the Moriscos from Spain in 1609*. Following the Christian reconquest of Spain that was completed in 1492, Spanish Muslims were first forced to convert to Christianity via a series of edicts between 1500 and 1525. Forced converts then were perceived to be "crypto-Muslims," until they were eventually expelled from Spain in 1609. Spain's Christian elites were divided on the value they attached to (religious) identity (c), though. On one side, the nobles preferred to exploit their Muslim vassals through forced labor services and a share of their harvest; in the language of our model, they were extractive rulers ($c = 0$). On the other side, the Church and the King attached higher value to religion by achieving religious (Christian) demographic homogeneity, even at the cost of economic loss ($c > 0$). From 1238, date of the conquest of Valencia by King Jaume I of Aragon, through 1525, when Muslims were forced to convert to Christianity, the nobility's extractive motives were prevalent as they succeeded in exploiting Muslims. They kept being so after 1525, but lost the battle in 1609 when the Moriscos were expelled from Spain.[66]

6 Conclusion

The paper made two contributions. It first developed a simple model of optimal one-shot and repeated taxation/extraction by a polity that trades off its hostility towards a group's identity and its reluctance to let exile, conversions or quits erode the contribution base. It provided a set of testable predictions on how discriminatory taxation and

[65] At Egypt's level, (non-convert) Copts rarely emigrated from the country (see Appendix Section B.1). At the local level, the state restricted rural-rural migration (see Section 3.2.6).

[66] See Chaney and Hornbeck (2015) for a detailed study of the economic impact of this episode.

the erosion of the contribution base are impacted by the ruler's and the governed's identity preferences and by agency in tax collection. Changes in these explanatory variables as well as uncertainty about the ruler's tenure generate interesting fiscal and identity dynamics. Finally, it showed that the permanent loss of identity dampens one's incentive to rebel, and that the threat of rebellion against fiscal extraction peters out over time, even when those who have altered their identity stay in the constituency (as is the case for religious conversions).

The second contribution is empirical/historical. The paper considered one particular historical event, the incentivized conversion of Egyptian Copts following the Arab conquest in the 7th century. Building on novel data sources, including tax papyri in 641-1100, list of churches and monasteries in 1200 (as a proxy for conversions), and proxies for LAs' and Copts' identity strengths, we first provided local-level support for the identity-based model, showing that a stronger enforcer identity increased conversions and the discriminatory tax, and reduced discriminatory tax revenue, suggesting taxation on the downward-sloping side of the Laffer curve and thus providing evidence for the identity-based model. Then, exploiting Coptic and Muslim medieval narratives, we constructed Egypt-level time-series proxies for governor's hostility toward non-converts, Caliph's piety, and budget needs, and we were able to trace discriminatory tax hikes and conversion waves in 641–1170. The Egypt-level time-series evidence comes in support of the identity-based model. Finally, using theory and history we were able to shed light on how the decline in the threat of rebellion, due to conversions, may have triggered the Caliphate-wide circa 750 tax reform lifting the cap on the non-discriminatory tax. Understanding the determinants of this reform matters because it is an attempt to endogenize a major "Islamic" institution: the canonical post-750 tax system, instead of treating it as "Islamic," exogenous, and ahistorical.

The theory can in principle be tested in a variety of historical environments where a discriminatory policy was used to induce an unwanted group to change its identity by adopting that of the ruling group, and where the optimal mix of discriminatory and uniform policies evolved in response to changes in taxpayers' identity composition. We mentioned some of these examples in the introduction.

Even though persecutions and emigration played a minor role in our historical context, we discussed how the identity-based model accommodates them. Persecutions are an interesting area of future study, with regard to both the agency problem and signalling, and to the substitutability/complementarity of price and non-price tools of discrimination. Emigration was prominent in many historical episodes, during which oppressed groups dwindled in size. Extending our exploratory theoretical treatment of persecutions and emigration, and performing empirical work along these lines, would be fascinating. For that, one will need to delve in greater depth into the foundations of the ruler's preference function. For example, does the ruler care primarily about population homogeneity? Or does he take a more religious stance of caring about conversions, and if so, how does he conceive his legacy (narrowly as the fraction of minority members in the polity, or broadly as his impact on worldwide conversions)? Particularly interesting would be the study of the strategic interaction, static and dynamic, among multiple rulers to offload or to the contrary attract the minority.

We view this paper as a first step toward further empirical and theoretical studies of optimal identity taxation with time-persistent status changes and their implications for the tax structure and the dynamics of ruler's legitimacy. We hope that it will stimulate empirical work building on other data sets, which will allow more structural estimations. We leave these promising alleys for research to future work.

References

Acemoglu, D. (2006). Modeling Inefficient Institutions. In Blundell, R., Newey, W., and Persson, T., editors, *Advances in Economic Theory, Proceeding of World Congress 2005*. Cambridge University Press, Cambridge. Chapter 9.

Anderson, R. W., Johnson, N. D., and Koyama, M. (2017). Jewish Persecutions and Weather Shocks: 1100–1800. *Economic Journal*, 127(602):924–958.

Artunç, C. and Agir, S. (2017). The Wealth Tax of 1942 and the Disappearance of Non-Muslim Enterprises in Turkey. mimeo.

Augenblick, N., Cunha, J. M., Dal Bó, E., and Rao, J. M. (2016). The Economics of Faith: Using an Apocalyptic Prophecy to Elicit Religious Beliefs in the Field. *Journal of Public Economics*, 141:38–49.

Barro, R. J. and McCleary, R. M. (2003). Religion and Economic Growth across Countries. *American Sociological Review*, 68(5):760–781.

Becker, C. H. (1902). *Beiträge zur Geschichte Ägyptens unter dem Islam*. Verlag Von Karl J. Trübner, Strasbourg.

Becker, G. S. (1957). *The Economics of Discrimination*. University of Chicago Press, Chicago.

Becker, S. O. and Woessmann, L. (2009). Was Weber Wrong? A Human Capital Theory of Protestant Economic History. *Quarterly Journal of Economics*, 124(2):531–596.

Bell, H. I. (1910). *Greek Papyri in the British Museum: Catalogue with Texts*, volume IV: The Aphrodito Papyri. London.

Belloc, M., Drago, F., and Galbiati, R. (2016). Earthquakes, Religion and Transition to Self-Government in Italian Cities. *Quarterly Journal of Economics*, 131:1875–1926.

Besley, T. and Persson, T. (2011). The Logic of Political Violence. *The Quarterly Journal of Economics*, 126(3):1411–1445.

Botticini, M. and Eckstein, Z. (2005). Jewish Occupational Selection: Education, Restrictions, or Minorities? *Journal of Economic History*, 65(4):922–948.

Bovenberg, A. and de Mooij, R. A. (1994). Environmental Levies and Distortionary Taxation. *American Economic Review*, 84:1085–1089.

Bovenberg, A. and Goulder, L. H. (1996). Optimal Environmental Taxation in the Presence of Other Taxes : General-Equilibrium Analyses. *American Economic Review*, 86:985–1000.

Bulliet, R. W. (1979). *Conversion to Islam in the Medieval Period: An Essay in Quantitative History*. Harvard University Press, Cambridge, MA.

Bénabou, R. and Tirole, J. (2020). Laws and Norms. mimeo.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Cantoni, D., Dittmar, J., and Yuchtman, N. (2018). Religious Competition and Reallocation: The Political Economy of Secularization in the Protestant Reformation. *The Quarterly Journal of Economics*, 133(4):2037–2096.

Chaney, E. (2013). Revolt on the Nile: Economic shocks, religion, and political power. *Econometrica*, 81:2033–2053.

Chaney, E. and Hornbeck, R. (2015). Economic Dynamics in the Malthusian Era: Evidence from the 1609 Spanish Expulsion of the Moriscos. *Economic Journal*, 90:1404–1440.

Courbage, Y. and Fargues, P. (1997). *Christians and Jews under Islam*. I. B. Tauris Publishers, London-New York. J. Mabro (Tr.).

De Long, J. B. and Shleifer, A. (1993). Princes and Merchants: European City Growth before the Industrial Revolution. *The Journal of Law and Economics*, 36(2):671–702.

Delavande, A. and Zafar, B. (2018). Information and anti-american attitudes. *Journal of Economic Behavior & Organization*, 149:1–31.

Diamond, P. (2006). Optimal Tax Treatment of Private Contributions for Public Goods with and without Warm Glow Preferences. *Journal of Public Economics*, 90:897–919.

Dittmar, J. E. and Meisenzahl, R. R. (2020). Public Goods Institutions, Human Capital, and Growth: Evidence from German History. *Review of Economic Studies*, 87(2):959–996.

El-Leithy, T. (2005). *Coptic Culture and Conversion in Medieval Cairo 1293-1524 AD*. PhD thesis, Princeton University, Princeton.

Farhi, E. and Werning, I. (2010). Progressive Estate Taxation. *Quarterly Journal of Economics*, 125(2):635–673.

Fleurbaey, M. and Maniquet, F. (2011). *A Theory of Fairness and Social Welfare*. Cambridge University Press, Cambridge.

Frantz-Murphy, G. (2004). Conversion in Early Islamic Egypt: The Economic Factor. In Conrad, L., editor, *Muslims and Others in Early Islamic Society*. Ashgate, Burlington, VA. pp. 323-330.

Greif, A. and Tadelis, S. (2010). A Theory of Moral Persistence: Crypto-Morality and Political Legitimacy. *Journal of Comparative Economics*, 38(3):229–244.

Grohmann, A. (1932). Aperçu de Papyrologie Arabe. *Etudes de papyrologie*, 1:23–95.

Kaplow, L. (1995). A Note on Subsidizing Gifts. *Journal of Public Economics*, 58:469–477.

Kuran, T. (2012). *The Long Divergence: How Islamic Law Held Back the Middle East*. Princeton University Press, Princeton.

Lazarus, E., Lewis, D. J., Stock, J. H., and Watson, M. W. (2018). HAR Inference: Recommendations for Practice. *Journal of Business & Economic Statistics*, 36(4):541–559.

Michalopoulos, S., Naghavi, A., and Prarolo, G. (2018). Trade and Geography in the Spread of Islam. *Economic Journal*, 128(616):3210–3241.

Mikhail, M. S. A. (2004). *Egypt from Late Antiquity to Early Islam: Copts, Melkites, and Muslims Shaping a New Society*. PhD thesis, University of California Los Angeles, Los Angeles.

O'Donoghue, T. and Rabin, M. (2006). Optimal Sin Taxes. *Journal of Public Economics*, 90:1825–1849.

Ramzi, M. (1994 [1954]). *al-Qamus al-Gughrafi lil Bilad al-Misriya min 'Ahd Qudama' al-Misriyin ila Sanat 1945 (Geographic Dictionary of Egyptian Localities from the Time of Ancient Egyptians to 1945)*. Egyptian General Book Authority, Cairo.

Rapoport, Y. (2018). *Rural Economy and Tribal Society in Islamic Egypt: A Study of al-Nabulusi's Villages of the Fayyum*. Brepols Publishers, Turnhout.

Riley, J. and Zeckhauser, R. (1983). Optimal Selling Strategies: When to Haggle, When to Hold Firm. *Quarterly Journal of Economics*, 98:267–287.

Roodman, D., MacKinnon, J. G., Nielsen, M. Ø., and Webb, M. D. (2018). Fast and Wild: Bootstrap Inference in Stata Using Boottest. mimeo.

Rubin, J. (2017). *Rulers, Religion, and Riches: Why the West Got Rich and the Middle East Did Not*. Cambridge University Press, Cambridge.

Saez, E. and Stantcheva, S. (2016). Generalized Social Welfare Weights for Optimal Tax Theory. *American Economic Review*, 106(1):24–45.

Saleh, M. (2018). On the Road to Heaven: Taxation, Conversions, and the Coptic-Muslim Socioeconomic Gap in Medieval Egypt. *Journal of Economic History*, 78(2):394–434.

Sandmo, A. (1975). Optimal Taxation in the Presence of Externalities. *The Swedish Journal of Economics*, 77:86–98.

Sirhan, J. (1978). *al-Musamara wal-Munadama 'ind al-'Arab hatta al-Qarn al-Rabi' al-Hijri (The Boon-Companionship in the Arab World until the 4th Century A.H.)*. PhD thesis, American University of Beirut, Beirut.

Stokey, N. (1979). Rational Expectations and Durable Goods Pricing. *Bell Journal of Economics*, 12:112–128.

Tirole, J. (2016). From Bottom of the Barrel to Cream of the Crop: Sequential Screening with Positive Selection. *Econometrica*, 84(4):1291–1343.

Voigtländer, N. and Voth, H.-J. (2012). Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany. *Quarterly Journal of Economics*, 127(3):1339–1392.

Wellhausen, J. (1927 [1902]). *The Arab Kingdom and its Fall*. University of Calcutta, Calcutta. M. G. Weir (Tr.).

Online Appendix

A Theory

A. *Proof of Proposition 1*

Let us first imagine that the LA's revenue-collection constraint is not binding. The monotone-hazard-rate condition implies that the LA's optimal discriminatory tax in district i, τ_i^a, is a weakly increasing function of r_i and c_i, with pass-through rates between 0 and 1.

Next we identify the discriminatory tax rates that are implementable by the CA through a transfer requirement. For $c_i < 0$, these are exactly described by the interval $[\tau_i^a(c_i), \tau_i^m]$: For any $\tau_i < \tau_i^a(c_i)$, the LA would raise the discriminatory tax to $\tau_i^a(c_i)$, raise more revenue and reach its optimal discriminatory taxation. For any $\tau_i > \tau_i^m$, the LA would lower the discriminatory tax to τ_i^m or below. Symmetrically, the implementable discriminatory tax rates for $c_i > 0$ are exactly described by the interval $[\tau_i^m, \tau_i^a(c_i)]$.

The upper bound on T_i for $c_i < 0$ is therefore equal to $\lambda_i + R_i^m$. Furthermore, setting $T_i = \lambda_i + R_i^m$ forces the LA to set discriminatory tax τ_i^m, which is as close to τ_i^c as the CA can get. Strict quasi-concavity of the latter's objective function then implies that this transfer requirement is optimal. For $c_i > 0$, the analysis is similar. For $c_i \geq c$, the CA can get its first best by setting $T_i = \lambda_i + R_i(\tau_i^c)$. The LA is then forced to moderate its discriminatory taxation so as to be able to raise enough revenue. Finally, for $c_i \in (0, c)$, the closest implementable tax rate (which is therefore optimal from strict quasi-concavity) is $\tau_i^a(c_i)$; the requested transfer is then $T_i = \lambda_i + R_i(\tau_i^a(c_i))$. ‖

B. *Threat of rebellion in the static model with district heterogeneity.*

Suppose that in the static analysis the CA is concerned about the possibility of a rebellion driven by high taxes. A successful rebellion overcomes the central and local authorities and eliminates all taxes. A rebellion can succeed only if rebels in the various districts unite. Each agent incurs cost ρ of rebelling; he is willing to rebel if and only if the gain from a successful rebellion exceeds the cost of rebellion:

$$G_i(\theta) \equiv \lambda_i + \min\{\max\{\theta, 0\}, \tau_i\} \geq \rho$$

Summing over all districts, the fraction of potential rebels is

$$\alpha \equiv \int_0^1 [1 - F(\rho - \lambda_i - r_i)] 1_{\{\tau_i > \rho - \lambda_i\}} di$$

(so $1 - \alpha$ is the fraction of docile agents). We assume that the probability of rebellion H smoothly and strictly increases with the number of potential rebels α.[67]

[67] This is consistent with many possible stories, with and without coordination failures. To give but one

Summing up, the CA's objective function W^r when more broadly rebellion is possible equals the previous expression times the probability $1-H$ that the tax system generates enough docile agents across the territory so as to prevent a rebellion:[68]

$$W^r \equiv [1-H(\alpha)]W$$

where, recall,

$$W = \int_0^1 [\lambda_i + R_i^c(\tau_i)]di.$$

Local authorities' objective functions do not reflect the threat of rebellion, as each district is infinitesimal and can free-ride, counting on the other districts to moderate their taxes so as to limit the threat of rebellion. Their objective function is unchanged.

Proposition 6 ***(cooptation)***
(i) Under the threat of rebellion, the discriminatory tax τ_i is still increasing in the LA's identity strength c_i.
(ii) Under a threat of rebellion the CA induces a reduction in the discriminatory tax rate of counterattitudinal LAs when the cost of rebellion is low, and a reduction in the discriminatory tax rate of soft or zealous LAs when the cost of rebellion is high. The discriminatory tax revenue (R_i) now is inverted-U shaped in the LA's identity strength (c_i), with a peak for a secular LA $(c_i = 0)$.
(iii) Assume that H is indexed by a parameter ν of threat of rebellion: $H(\alpha|\nu)$, with density $h(\alpha|\nu)$ satisfying the monotone hazard rate property (MLRP): $h(\alpha|\nu)/[1-H(\alpha|\nu)]$ is increasing in ν. Then a higher threat of rebellion (a higher ν) makes cooptation more likely.

Remark (agency benefits of counterattitudinal tax collectors). We have assumed that the tax collectors' identity is given by the available local competency pool (which as we will discuss is a reasonable assumption for early Islam in Egypt). Nonetheless, it is useful to examine whether the CA would be willing to incur costs to replace existing tax collectors. Here the predictions are drastically different depending on whether there is a threat of rebellion. In the absence of such a threat, counterattitudinal tax collectors are a nuisance to the CA as they do not convert enough agents. In contrast, counterattitudinal

example, we could assume that, after taxes are set but before deciding whether to rebel, the agents learn what it takes for making the rebellion successful; namely, the rebellion will be successful if and only if $\alpha \geq \varepsilon$, where $\varepsilon \geq 0$ is the CA's capacity to counter the rebellion. The parameter ε is ex-ante distributed according to smooth cdf $H(\varepsilon)$. In the absence of coordination failure, a rebellion occurs whenever all those who gain from it have a mass exceeding ε. If there are more agents willing to rebel than is needed for the rebellion to be successful, an arbitrary selection mechanism will do. Furthermore, exactly the same analysis holds even if there are coordination failures, under the reasonable assumption that the probability of rebellion strictly increases with the number of potential rebels.

[68] This version generalizes that considered in the dynamic framework. There $H(\alpha) = 1$ iff $\alpha > 1 - F(\hat{\theta})$.

tax collectors may help the CA avert a rebellion as their preferences make them committed to treating agents more leniently. The proof of Proposition 6 shows that the CA may not want to replace a tax collector with one whose preferences are more aligned with his objectives, even if it were costless to do so.

Proof of Proposition 6

The CA's welfare is $[1-H(\alpha)]W$, where $[1-H(\alpha)]$ is the probability of staying in power and W the welfare when in power. Let $w_i = \lambda_i + (\tau_i - c)[1-F(\tau_i - r_i)]$ denote the CA's welfare corresponding to district i in the absence of threat of rebellion (so τ_i is as in Proposition 1 and Figure 1), and $\widehat{w}_i(\widehat{\tau}_i) \equiv \lambda_i + (\widehat{\tau}_i - c)[1-F(\widehat{\tau}_i - r_i)]$ denote the CA's welfare when the transfer demand to district i (possibly) accounts for the threat of rebellion (that is, is chosen with an eye on maximizing $[1-H(\alpha)]W$ and not just W). Only values $\widehat{\tau}_i \leq \tau_i$ are relevant when adding a no-rebellion constraint.

If $\lambda_i + \tau_i \leq \rho$, agents in district i will not join a rebellion in the absence of policy change (such districts are "not rebellion-prone"), and so at the optimum $\widehat{\tau}_i = \tau_i$. A district is more likely not to be rebellion-prone, the lower λ_i, r_i, and c_i.

So, we will be interested only in districts such that $\lambda_i + \tau_i > \rho$. For these, either the CA induces $\lambda_i + \widehat{\tau}_i > \rho$; in this case the number of rebels in district i is an exogenous $1-F(\rho - \lambda_i - r_i)$ and the optimal discriminatory tax is still $\widehat{\tau}_i = \tau_i$. Or, by strict quasi-concavity and the property that the optimal rebellion-free tax satisfies $\tau_i \leq \tau_i^a(c)$ (Proposition 1), $\lambda_i + \widehat{\tau}_i = \rho$ and $\widehat{w}_i \equiv \widehat{w}_i(\rho - \lambda_i) = \lambda_i + (\rho - \lambda_i - c)[1-F(\rho - \lambda_i - r_i)] < w_i$. Let $x_i = 1$ if $\lambda_i + \widehat{\tau}_i > \rho$, and $x_i = 0$ otherwise. The probability of staying in power is then $1-H(\alpha)$, where

$$\alpha = \int_{i\in[0,1]} [1-F(\rho - \lambda_i - r_i)] 1_{\{\lambda_i + \tau_i > \rho\}} x_i di$$

The CA solves

$$\max_{\{x_i \in [0,1]\}_{i\in[0,1]}} \left[1 - H\left(\int_{i\in[0,1]} [1-F(\rho - \lambda_i - r_i)] 1_{\{\lambda_i + \tau_i > \rho\}} x_i di\right)\right] \times \left[\int_{i\in[0,1]} [w_i 1_{\{\lambda_i + r_i \leq \rho\}} + [x_i w_i + (1-x_i)\widehat{w}_i] 1_{\{\lambda_i + \tau_i > \rho\}}] di\right]$$

Solving this program (for districts that are rebellion-prone), there exists a parameter $\xi \equiv \left[\frac{h(\alpha)}{1-H(\alpha)} W\right]$ (determined country-wide, i.e. independant of i) such that for districts such that $\lambda_i + \tau_i > \rho$, then $\widehat{\tau}_i = \tau_i$ if and only if the cost of detering rebels in district i relative to the fraction of discouraged rebels in that district exceeeds the country-wide threshold:

$$\xi \leq \frac{w_i - \widehat{w}_i}{1-F(\rho - \lambda_i - r_i)}$$

- Suppose, first, that $\lambda_i + \tau_i^m > \rho$. The implementability condition implies that district i cannot be made rebellion free if $c_i \geq 0$. Let $c^* < 0$ be defined by $\tau_i^a(c^*) = \rho - \lambda_i$. Because of the implementability constraint, only districts satisfying $c_i \leq c^*$ can be made rebellion free, so that for $c^* < c_i \leq 0$, the optimal policy remains that under no rebellion threat. For $c_i \leq c^*$, strict quasi-concavity implies that the optimal policy is either $\tau_i^a(c^*)$, or the no-rebellion-threat policy τ_i^m, and (A) implies that the choice between the two is the same for all $c_i \leq c^*$.

- Second, assume that $\lambda_i + \tau_i^m \leq \rho$. Let $c^* \geq 0$ be defined by $\tau_i^a(c^*) = \rho - \lambda_i$. If $c^* > c$, then the optimum is either the same as in the absence of threat of rebellion, or the same for $c_i < c^*$ and $\tau_i^a(c^*)$ for all $c_i \geq c^*$. Next, suppose that $c^* \leq c$. Then, for $c_i \leq c^*$, the no-rebellion policy policy does not generate rebellion in district i ; so the discriminatory tax is unchanged. For $c_i > c^*$, the CA faces a choice between eliminating the threat of rebellion in the district (inducing $\widehat{\tau}_i = \tau_i^a(c^*) = \rho - \lambda_i$) and sticking to the no-rebellion policy τ_i. The CA optimally quells the rebellion in district i if $c_i \in [c^*, c^{**}]$ where $c^* < c^{**} \leq c$ or $c^{**} = +\infty$ (if not, $\widehat{\tau}_i = \tau_i$).

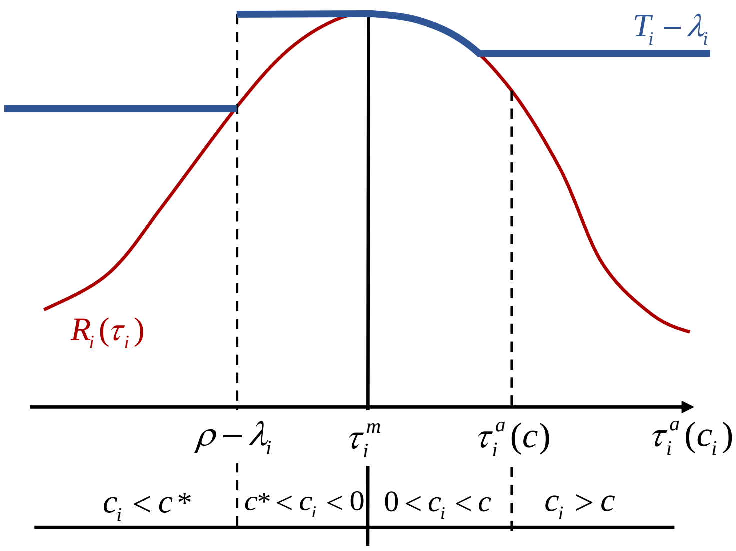

Figure A.1: **Revenue and LA identity strength under the threat of rebellion**

Finally, index the distribution $H(\alpha|\nu)$, where ν is an index of the threat of rebellion. Suppose that ξ does not increase as ν decreases. Then, from (A), the set of districts i such that $x_i = 1$ expands, increasing W. Furthermore $h(\alpha|\nu)/[1-H(\alpha|\nu)]$ increases as well, and so ξ increases, a contradiction. This implies that part (iii) of Proposition 6 holds.

C. *Proof of Proposition 3*

The agents' equilibrium strategy can be described by the following cutoff rule at date t:

$$\theta_t^* = \max\left\{\theta_{t-1}^*, \frac{\tau_t}{K_t}\right\}$$

and the discriminatory tax obeys:

$$\tau_t = K_t \max\{\theta^*, \theta_{t-1}^*\}.$$

To see that this is an equilibrium, note that the date-t cutoff, if interior $(\theta_t^* > \theta_{t-1}^*)$ satisfies $\left(1 + \frac{\beta}{1-\beta} x_{t+1}\right)\theta_t^* = \tau_t$: Either the ruler is removed at date $(t+1)$ and then the cutoff type enjoys θ_t^* forever; or the ruler remains in place and then type $\theta_t^* \le \theta_{t+1}^*$ prefers (weakly or strongly) to pay the tax τ_{t+1}.

As for the ruler, note that the equilibrium behaviors deliver the upper bound on his intertemporal payoff that would correspond to the no-external-challenge environment ($x_t \equiv 0$ for all t):

$$W_t^{\max}(\theta_{t-1}^*) = \begin{cases} \dfrac{1}{1-\beta}[\lambda + (\theta^* - c)[1 - F(\theta^*)]] & \text{if } \theta_{t-1}^* \le \theta^* \\ \dfrac{1}{1-\beta}[\lambda + (\theta_{t-1}^* - c)[1 - F(\theta_{t-1}^*)]] & \text{if } \theta_{t-1}^* > \theta^* \end{cases}$$

To see this, assume that $\theta_{t-1}^* \le \theta^*$, say (the proof is the same in the opposite case, due to strict quasi-concavity of the adjusted tax revenue). Let the ruler charge $K_t\theta^*$. Then

$$W_t = [\lambda + (K_t\theta^* - c)[1 - F(\theta^*)]] + \beta x_{x+1}\left[-\frac{c[1 - F(\theta^*)]}{1-\beta}\right] + \beta(1 - x_{t+1})W_{t+1}.$$

So W_{t+1} is equal to $\frac{1}{1-\beta}[\lambda + (\theta^* - c)[1 - F(\theta^*)]$, then W_t takes this value as well. The upper bound on the ruler's continuation payoff can be reached though a stationary policy $\theta_{t+k}^* = \max\{\theta^*, \theta_{t-1}^*\}$. So no deviation for any history can yield more that the equilibrium strategy. ||

D. *Proof of Proposition 4*

Assume that $\lambda + \min\{\hat{\theta}, \tau^c\} > \rho$, so there is a real threat of rebellion. The CA's optimization program is

$$\max_{\{\hat{\tau}, \hat{\lambda}\}} \hat{\lambda} + (\hat{\tau} - c)[1 - F(\hat{\tau})]$$

s.t.

$$\begin{cases} \hat{\lambda} + \min\{\hat{\theta}, \hat{\tau}\} \le \rho & \text{(no rebellion)} \\ \hat{\lambda} \le \lambda & \text{(uniform tax cannot exceed its extractive level)} \\ \hat{\tau} \in [\tau^m, \tau^c] & \end{cases}$$

Suppose that the CA chooses $\{\hat{\lambda}, \hat{\tau}\}$ such that $\hat{\lambda} + \hat{\tau} \le \rho$ (that is, $\hat{\tau} \le \hat{\theta}$ and so the marginal rebel is a non-convert). Then the CA has welfare $\hat{\lambda} + (\hat{\tau} - c)[1 - F(\hat{\tau})] = \rho - \hat{\tau} + R^c(\hat{\tau})$, which is decreasing in $\hat{\tau}$ for $\hat{\tau} \ge \tilde{\tau}$, where

$$\tilde{\tau} + \frac{F(\tilde{\tau})}{f(\tilde{\tau})} = c.$$

Let us restrict the consideration set for the discriminatory tax. First, $\hat{\tau} < \tau^m$ is not implementable. Next, $\hat{\tau} < \tilde{\tau}$ is always weakly dominated: Consider a small change $\delta\hat{\tau} = +\varepsilon$ and $\delta\hat{\lambda} = -\varepsilon$; then the no-rebellion constraint, $\hat{\lambda} + \min\{\hat{\tau}, \hat{\theta}\} \leq \rho$, remains satisfied and $\delta(\hat{\lambda} + R^c(\hat{\tau})) = \varepsilon((R^c)' - 1) > 0$ for $\hat{\tau} < \tilde{\tau}$. Finally, $\tau < \rho - \lambda$ is not feasible unless $\hat{\theta}$ is a convert, i.e. $\lambda + \hat{\theta} = \rho$ and $\hat{\theta} < \tau$.

Let $\tau^* \equiv \max\{\tau^m, \tilde{\tau}, \rho - \lambda\}$. Because we are interested only in the case of a rebellion threat $(\lambda + \tau^c > \rho)$, $\tau^* < \tau^c$.

We distinguish three regions:

Region 1: $\hat{\theta} < \tau^m$. Then $\hat{\theta}$ is a convert, $\lambda + \hat{\theta} = \rho$ and $\hat{\tau} = \tau^c$. Welfare is

$$W^1 \equiv \rho - \hat{\theta} + R^c(\tau^c).$$

Region 2: $\hat{\theta} > \tau^c$. Type $\hat{\theta}$ is then necessarily a non-convert, and

$$W^2 \equiv \rho - \tau^* + R^c(\tau^*).$$

Region 3: $\tau^m \leq \hat{\theta} \leq \tau^c$. Either $\hat{\tau} > \hat{\theta}$ (the marginal rebel is a convert) and then at the optimum $\hat{\tau} = \tau^c$. Welfare is then $W^3 = W^1$. Furthermore, welfare W^1 can be obtained for any $\hat{\theta} \in [\tau^m, \tau^c]$.

Or $\hat{\tau} \leq \hat{\theta}$ (the marginal rebel is a non-convert). Then $\hat{\tau} = \tau^*$, yielding welfare $W^3 = W^2$. But, unlike for W^1, W^2 is not feasible for any $\hat{\theta} \in [\tau^m, \tau^c]$: It requires that $\tau^* \leq \hat{\theta}$.

Optimal welfare is therefore W^1 for $\hat{\theta} \in [\tau^m, \tau^*]$. On $[\tau^*, \tau^c]$, note that $dW^1/d\hat{\theta} = -1$ while $dW^2/d\hat{\theta} = 0$. Furthermore

$$W^1(\tau^*) - W^2(\tau^*) = R^c(\tau^c) - R^c(\tau^*) > 0 > W^1(\tau^c) - W^2(\tau^c) = -(\tau^c - \tau^*).$$

Therefore in this interval $W^3 = W^2$ if and only if $\hat{\theta} \geq \theta^*$ where

$$\theta^* \equiv R^c(\tau^c) - R^c(\tau^*) + \tau^*.$$

Putting all three regions together

(1) For $\hat{\theta} < \theta^*$, $W = W^1$, $\hat{\lambda} = \rho - \hat{\theta}$, and $\hat{\tau} = \tau^c$.

(2) For $\hat{\theta} > \theta^*$, $W = W^2$, $\hat{\lambda} = \rho - \tau^*$, and $\hat{\tau} = \tau^*$. ||

E. *Proof of Proposition 5*

(i) For all t, the CA chooses $\{\lambda_t(\theta^*_{t-1}), \tau_t(\theta^*_{t-1})\}$ so as to maximize:

$$W_t = \sum_{k=0}^{+\infty} \beta^k [\lambda_{t+k} + (\tau_{t+k} - c)[1 - F(\theta^*_{t+k})].$$

s.t :

$$\lambda_t \leq \lambda \;\; \text{(uniform tax capped at extractive capacity)}$$

$$\sum_{k=0}^{+\infty} \beta^k [\lambda_{t+k} + (\tau_{t+k} - \hat{\theta}) 1_{\{\hat{\theta} > \theta^*_{t+k}\}}] \leq \frac{\rho - \hat{\theta} 1_{\{\hat{\theta} > \theta^*_{t-1}\}}}{1 - \beta} \;\; \text{(no-rebellion constraint)}$$

$$\tau_t \in [\tau^m, \tau^c] \;\; \text{(implementability).}$$

The agents' strategy can be described by $\theta_t^*(\tau_t, \lambda_t, \theta_{t-1}^*) \geq \theta_{t-1}^*$, the cutoff rule at date t (types $\theta \geq \theta_t^*$, and only them, keep their identity up to date t included). Type θ solves

$$U_t = \sum_{k=0}^{+\infty} \beta^k [-\lambda_{k+t} - (\tau_{k+t} - \theta) 1_{\{\theta > \theta_{t+k}^*\}}].$$

Lemma 1 *Suppose that type $\hat{\theta}$ converts at some date $T \in \{1, \dots, +\infty\}$. From date $T+1$ on, $\lambda_t = \min(\lambda, \rho)$ and $\tau_t = \tau^c$.*

Proof of Lemma 1. Consider an equilibrium path $\{\tau_t, \lambda_t\}_{t \geq 1}$ such that type $\widehat{\theta}$ converts at some date $T \in \{1, \dots, +\infty\}$. After that date, there is no threat of rebellion provided that for all $t \geq T+1$, $\sum_{k=0}^{+\infty} \beta^k \lambda_{t+k} \leq \frac{\rho}{1-\beta}$. Given that $\lambda_t \leq \lambda$ for all t, from date $T+1$ on, the CA optimally charges $\lambda_t = min(\lambda, \rho)$ and choose τ_t so as to maximize $R^c(\tau_t)$, so $\tau_t = \tau^c$.[69] ‖

Lemma 2 *For $t \leq T$, $\theta_t^* = \tau_t$.*

Proof of Lemma 2. The path of conversions if described by a sequence of cutoffs $\{\theta_t\}_{t \in \{1, \dots, +\infty\}}$, satisfying

$$\theta_1^* \leq \theta_2^* \leq \dots \leq \theta_{T-1}^* < \hat{\theta} \leq \theta_T^* \leq \theta_{T+1}^* \cdots.$$

Suppose that, for some t, $\tau_t > \theta_t^*$, implying that type θ_t^* loses utility at date t from not converting. This utility must be recouped in the future, and so there exists $k \geq 1$ (possibility infinite) such that

$$(\theta_t^* - \tau_t) + \beta(\theta_t^* - \tau_{t+1}) + \ldots + \beta^k(\theta_t^* - \tau_{t+k}) \geq 0.$$

This implies in particular that $\theta_{t+1}^* = \theta_t^*$. In an MPE, this implies that $\theta_{t+\ell}^* = \theta_t^*$, $\tau_{t+\ell} = \tau_{t+1} < \theta_t^*$, $\lambda_{t+\ell} = \lambda_t$ for all $\ell \geq 0$. Suppose, first, that $\tau_t < \theta_{t-1}^*$. Then, for any $\tau_t' \in [\tau_t, \theta_{t-1}^*)$, there is no new conversion at date t as any $\theta \geq \theta_{t-1}^*$ enjoys a current surplus, $\theta - \tau_t' > 0$, and keeps an option value. So for $\tau_t' \in [\tau_t, \theta_{t-1}^*)$, $\theta_t^* = \theta_{t-1}^*$ and the Markov property implies that the continuation equilibrium remains the same. But with $\tau_t' > \tau_t$, the revenue is higher for the CA. Therefore $\tau_t \geq \theta_{t-1}^*$ for all t. Because θ_t^* cannot recoup the loss in the future, $\tau_t \leq \theta_t^*$.

Next, let us show that $\tau_t = \theta_t^*$. Suppose, to the contrary, that $\tau_t < \theta_t^*$. Either $\theta_t^* > \theta_{t-1}^*$, but then type $\theta_t^* - \varepsilon$ should not convert, as $\theta_t^* - \varepsilon - \tau_t > 0$. Or, $\theta_t^* = \theta_{t-1}^*$ and then $\tau_t < \theta_{t-1}^*$, a contradiction. ‖

Now define $\iota \equiv \arg\max_{\tau \in [\tau^m, \tau^c]} \{-\min\{\tau, \hat{\theta}\} + (\tau - c)(1 - F(\tau))\}$. Also define $\tau^{**} \equiv \max(\tau^m, \tilde{\tau})$, and $\theta^{**} \equiv R^c(\tau^c) - R^c(\tau^{**}) + \tau^{**}$. With simple algebra, we have $\tau^c \geq \theta^{**} \geq \tau^{**}$, and

[69] Recall that $\theta_t^* > \tau^c$ is not implementable for any t because implementability requires that discriminatory taxes be below τ^c in each period and so no type above τ^c would ever convert.

Lemma 3 *$\iota = \tau^{**}$ if $\hat{\theta} > \theta^{**}$, and $\iota = \tau^c$ if $\hat{\theta} \le \theta^{**}$.*

Proof of Lemma 3. Either $\tau < \hat{\theta}$ and then the maximand, $(\tau - c)[1 - F(\tau)] - \tau$, is maximized at $\tilde{\tau}$ in the absence of the implementability constraint; so, $\tau = \max\{\tau^m, \tilde{\tau}\} = \tau^{**}$, yielding maximand $-\tau^{**} + R(\tau^{**})$. Or $\tau \ge \hat{\theta}$ and then the maximand, $R^c(\tau) - \hat{\theta}$, is maximized at $\tau = \tau^c$ and then equal to $R^c(\tau^c) - \hat{\theta}$. To see that $\theta^{**} \in (\tau^{**}, \tau^c)$, it suffices to observe that $(R^c)' < 1$ for $\tau > \tilde{\tau}$. ||

To find an upper bound for W_1, we first ignore the constraint that $\lambda_t \le \lambda$ for all λ. This constraint will be satisfied in two cases, and will need to be reintroduced in the third. Finally, consider the date-1 no-rebellion constraint. Rebelling at date 1 yields net cost $(\rho - \hat{\theta})/(1 - \beta)$ to the marginal rebel. Suppose that type $\hat{\theta}$ converts at some date $T \in \{1, \ldots, +\infty\}$. So, it must be the case that

$$\sum_{t=1}^{T-1} \beta^{t-1}[\lambda_t + \tau_t - \hat{\theta}] + \sum_{t \ge T} \beta^{t-1}\lambda_t \le \frac{\rho - \hat{\theta}}{1 - \beta},$$

where for $t \le T - 1$, $\tau_t \le \theta_t^* < \hat{\theta}$ from Lemma 2. So, the CA's date-1 welfare can be bounded above by using, successively, Lemma 2 and the date-1 no-rebellion constraint:

$$\begin{aligned} W_1 &= \sum_{t \ge 1} \beta^{t-1}[\lambda_t + (\tau_t - c)[1 - F(\theta_t^*)]] \\ &= \sum_{t \ge 1} \beta^{t-1}[\lambda_t + R^c(\tau_t)] \\ &\le \sum_{t=1}^{T-1} \beta^{t-1}[\rho + R^c(\tau_t) - \tau_t] + \sum_{t=T}^{+\infty} \beta^{t-1}[\rho + R^c(\tau^c) - \hat{\theta}]. \end{aligned}$$

Because $\tau_t \ge \tau^m$ and $R(\tau) - \tau$ is maximized at $\tilde{\tau}$, a new upper bound is $\frac{1}{1-\beta}\max\{R^c(\tau^{**}) - \tau^{**}, R^c(\tau^c) - \hat{\theta}\}$. And so:

$$W_1 \le \begin{cases} \dfrac{\rho - \hat{\theta} + R^c(\tau^c)}{1 - \beta} & \text{if } \hat{\theta} \le \theta^{**} \\ \dfrac{\rho - \tau^{**} + R^c(\tau^{**})}{1 - \beta} & \text{if } \hat{\theta} \ge \theta^{**}. \end{cases}$$

Next we show that this upper bound is reached for some MPE in the following cases.

(a) Suppose that $\hat{\theta} \le \theta^*$ and that the CA sets $\tau_t = \tau^c$ for all t, $\lambda_t = \min\{\lambda, \rho\}$ for $t \ge 2$ and $\lambda_1 < \min\{\lambda, \rho\}$ such that $\lambda_1 + \frac{\beta}{1-\beta}\min\{\lambda, \rho\} = \frac{\rho - \hat{\theta}}{1-\beta}$ (recall that the no-rebellion constraint is binding, so $\min\{\lambda, \rho\} > \rho - \hat{\theta}$). All agents $\theta \le \tau^c$ convert at date 1, and no conversion occurs later on. These strategies yields CA welfare equal to $[\rho - \hat{\theta} + R^c(\tau^c)]/(1 - \beta)$ and therefore are optimal for the CA; they also are optimal for the CA from date 2 on, and so form an MPE. If $\hat{\theta} > \tau^c$, then it can be shown that the upper bound on W_1 is the same as for $\hat{\theta} \ge \theta^{**}$ (and the strategies implementing this upper bound are the same as well).

(b) Suppose that $\hat{\theta} \ge \theta^*$ and that the CA set $\tau_t = \tau^{**}$ and $\lambda_t = \rho - \tau^{**} \le \lambda$ for all t and that $\hat{\theta} \ge \tau^{**}$. Then all conversions occur at date 1 and only types $\theta \le \tau^{**}$ convert. And the strategies yield the upper bound for W_1 in each period if a fortiori $\hat{\theta} \geq \theta^{**}$.

(c) If the maximum corresponds to τ^{**} and if $\tau^{**} < \rho - \lambda$, we will face a problem when implementing τ^{**} from 1 to $+\infty$, as the per-period uniform tax that would satisfy the no-rebellion constraint would exceed λ, which is impossible. Let us reintroduce the constraint that $\lambda_t \leq \lambda$ in a weaker form:

$$\sum_{t=1}^{+\infty} \beta^{t-1} \lambda_t \leq \frac{\lambda}{1-\beta}.$$

If this constraint is binding, the date-1 no-rebellion constraint becomes:

$$\sum_{t=1}^{T-1} \beta^{t-1}(\hat{\theta} - \tau_t) \geq \frac{\hat{\theta} + \lambda - \rho}{1-\beta}$$

(the RHS of this inequality is by assumption strictly positive). Letting $z_t \equiv 1_{\{\hat{\theta} > \theta_t^*\}} \in \{0,1\}$, and substituting the date-1 no-rebellion constraint,

$$\sum_{t=1}^{+\infty} \beta^{t-1}[\lambda_t + (\tau_t - \hat{\theta})z_t] \leq \frac{\rho - \hat{\theta}}{1-\beta},$$

$$W_1 \leq \sum_{t=1}^{+\infty} \beta^{t-1}[R^c(\tau_t) - (\tau_t - \hat{\theta})z_t + \rho - \hat{\theta}].$$

Maximize over z_t and τ_t the RHS of this inequality subject to the constraint coming from the upper bound on the uniform tax:

$$\sum_{t=1}^{+\infty} \beta^{t-1}(\hat{\theta} - \tau_t)z_t \geq \frac{\hat{\theta} + \lambda - \rho}{1-\beta} \qquad (\mu)$$

The period-by-period maximization amounts to solving

$$\max_{\{\tau_t, z_t\}} R^c(\tau_t) - (\tau_t - \hat{\theta})(1+\mu)z_t$$

and so τ_t and z_t are both constant over time (call these τ and z). Furthermore

$$z = 1 \iff (\hat{\theta} - \tau)(1+\mu) > R^c(\tau^c) - R^c(\tau).$$

When the constraint is non-binding ($\mu = 0$), then the solution is as in cases (a) and (b). When it is binding

$$\hat{\theta} - \tau = \hat{\theta} + \lambda - \rho \iff \tau = \rho - \lambda.$$

And so, letting $\tau^* \equiv \max\{\tau^m, \tilde{\tau}, \rho - \lambda\}$ and $\theta^* \equiv \tau^* + [R^c(\tau^c) - R^c(\tau^*)]$ the solution is the same as in Proposition 4, except for the sequencing of uniform taxes in case (a).

(ii) Suppose that $\hat{\theta} < \theta^{**}$. Could a coalition of size (at least) $1 - F(\widehat{\theta})$ coordinate and not convert at date 1, so that the rebellion constraint would remain at date 2? Could it do so repeatedly? Let the CA set $\tau_2 = \tau^c$ and $\lambda_2 = \lambda_1$ and continue doing so as long as type $\widehat{\theta}$ (and types below necessarily) has not converted[70]. From date 2 on, type $\widehat{\theta}$ in this deviation from equilibrium behavior obtains value function $V_1 \equiv \frac{\widehat{\theta} - \rho}{1-\beta}$, i.e. its date-1 value function. So, type $\widehat{\theta}$ does not want to deviate at date 1 if $-\lambda_1 - \tau^c + \beta V_1 \leq V_1$ or $\tau^c \geq \beta \frac{\widehat{\theta + \lambda - \rho}}{1-\beta}$. ‖

[70]Observing the discriminatory tax volume supplies this information; indeed, we have assumed that types are not observable.

Context	Tax	Low threat of rebellion	High threat of rebellion	Intensity of threat
		τ^m — τ^* — θ^*	θ^* — τ^c	$\hat{\theta}$
Statics	Uniform	$\hat{\lambda} = \rho - \hat{\theta}$	$\hat{\lambda} = \rho - \tau^*$	
	Discriminatory	$\hat{\lambda} = \tau^c$ (high)	$\hat{\tau} = \tau^*$ (low)	
Myopic (β = 0)	Uniform	$\lambda_1 = \rho - \hat{\theta} < \lambda_t = \min\{\rho, \lambda\}$ for $t \geq 2$	$\lambda_t = \rho - \tau^*$	
	Discriminatory	$\tau_t = \tau^c$	$\tau_t = \tau^*$	
Far-sighted (β > 0)	Uniform	$\lambda_1 = \rho - \hat{\theta} - \frac{\beta\lambda}{1-\beta} < \lambda_t = \min\{\rho, \lambda\}$ for $t \geq 2$	$\lambda_t = \rho - \tau^*$	
	Discriminatory	$\tau_t = \tau^c$	$\tau_t = \tau^*$	

Figure A.2: **Threat of rebellion (summary of Propositions 4 through 5)**

Proof of Corollary 8. For the sake of the argument, suppose that the LA is soft ($c_i > 0$) rather than counterattitudinal (the same reasoning works in the latter case). For a given p_i, the implementable set is $[\tau_i^m(p_i), \tau_i^a(c_i, p_i)]$, where $\tau_i^a(c_i, p_i)$ is the LA's preferred discriminatory tax, which solves: *max* $(\tau_i - c_i)[1 - F(\tau_i + p_i - r_i)]$ (and $\tau_i^m(p_i)$ solves the same program for $c_i = 0$). So $\tau_i^a(c_i, 0) = \tau_i^a(c_i)$, the discriminatory tax in the persecution-free environment. It is easily shown that persecutions reduce the discriminatory tax as the LA absorbs a fraction of its effect: $\frac{\partial \tau_i^a(c_i,p_i)}{\partial p_i} \in (-1, 0)$. The CA's payoff when the LA sets $\tau_i^a(c_i, p_i)$ (which is the tax in the implementable set that the CA prefers) is $W_i = \lambda_i + [\tau_i^a(c_i, p_i) - c][1 - F(\tau_i^a(c_i, p_i) + p_i - r_i)]$. Simple computations show that $\frac{\partial W_i}{\partial p_i}|_{p_i=0} \propto \frac{\partial \tau_i^a(c_i,p_i)}{\partial p_i}|_{p_i=0} [c - c_i] + c - \tau_i^a(c_i, 0)$. For $c = c_i$, the first term on the RHS is equal to 0 while the second term is strictly negative. The RHS is strictly increasing in c; and for c sufficiently large the CA can guarantee itself λ_i by choosing an infinite level of persecutions and gets strictly less than λ_i when choosing $p_i = 0$. Finally, the cutoff level c_i^* is defined by $\frac{\partial \tau_i^a(c_i,p_i)}{\partial p_i}|_{p_i=0} [c_i^* - c_i] + c_i^* - \tau_i^a(c_i) = 0$. $\|$

F. *Extensions*

(a) *Discrimination through access to public goods.* When direct discrimination is prohibited by the constitution or a higher-level polity (which was not the case for early Islam), we naturally observe more indirect forms of discrimination, such as neighborhood-based access to public goods, ethnicity-based patronage and incendiary rhetoric. Glaeser and Shleifer (2005) describe such forms of discrimination in 20th-century US, staging an Irish-catholic/Anglo-Saxon-protestant conflict in Boston and a black/white conflict in Detroit. In both examples, the mayor induced over the years substantial emigration of the minority out of the city, reinforcing the incumbent's political power;[71] Glaeser and Shleifer call this the "Curley effect," after the name of a Boston mayor who was in power for most of the 1913-1951 period. A direct, ethnic or race-based, tax discrimination being prohibited by the federal government, the ruler's hostility toward the minority shifted to presumably less efficient forms of utility extraction. Their paper also documents Robert Mugabe's tactic in Zimbabwe, which led to substantial migration by white farmers. In either case, more discrimination involved a revenue cost, in terms of either migration or incompetent management (patronage). And it increased the probability of a rebellion.

The trade-off between loss in revenue and preference alignment through emigration of members with a dissonant identity also arises in modern democracies when a ruler may also want to increase the cohesiveness of the polity. Democratic regimes and organizations sometimes function more efficiently when their membership is more homogeneous. For example, Hansmann (1996) argues that congruence in objectives facilitate both the flow of information and the fluidity of decision making in cooperatives. Besley et al. (2017) argue that districts with single party majority yield more cohesive policies, presumably because this cohesion facilitates agreement on the use of tax revenue and thereby raises incentives to collect tax revenue. Relatedly, Alesina et al. (1999) have shown that the provision of local public goods is facilitated by religious or ethnic homogeneity. Without applying a value judgment to such objectives, we can capture the ruler's demand for cohesiveness within the function.

(b) *Discriminatory empathy.* Suppose that LA i puts (positive or negative) weight $w_i(\theta)$ on type-θ agent's utility, where $w_i'(\theta) \leq 0$ and $\int_{-\infty}^{+\infty} w_i(\theta) dF(\theta - r_i) \equiv \bar{w}_i < 1$.[72] The LA's objective function is then

$$\begin{aligned} V_i \equiv\ & [\lambda_i + \tau_i[1 - F(\tau_i - r_i)] - T_i] + \Big[\int_{-\infty}^{\tau_i} w_i(\theta)(-\lambda_i) dF(\theta - r_i) \\ & + \int_{\tau_i}^{+\infty} w_i(\theta)[-\lambda_i + (\theta - \tau_i)] dF(\theta - r_i)\Big] = (1 - \bar{w}_i)\lambda_i + R_i^a(\tau_i) - T_i \end{aligned}$$

[71] Migration then reduces resistance to the ruler over time because of the majoritarian electoral system. By contrast, our time-decreasing resistance in Section 4.1 will be based on a reduced stake for the converts.

[72] The LA need not observe individual agents' types to form such preferences (and actually it does not).

where

$$R_i^a(\tau_i) \equiv \tau_i[1 - F(\tau_i - r_i)] + \int_{\tau_i}^{\infty} w_i(\theta)(\theta - \tau_i)dF(\theta - r_i).$$

Note that

$$(R_i^a)'(\tau_i) = [1 - w_i^+(\tau_i)][1 - F(\tau_i - r_i)] - \tau_i f(\tau_i - r_i)$$

where $w_i^+(\tau_i) \equiv E[w_i(\theta)|\theta \geq \tau_i] \leq \bar{w}_i$.

The difference with the model in the text is that $c_i \equiv \frac{-\tau_i w_i^+(\tau_i)}{1 - w_i^+(\tau_i)}$ depends on the discriminatory tax, which itself depends on the extent of discriminatory empathy. Substantial hostility to the high-identity agents is then required to be on the wrong side of the Laffer curve. More generally, lower empathy (w_i^+ falls) implies a higher discriminatory tax.

(c) *Social incentives: norms and network externalities.* When contemplating changing his identity, an agent may take into account not only his own preferences (θ) and the material incentive (τ_i), but also the resulting perception of his choice within his community. Letting $F_i(\theta) \equiv F(\theta - r_i)$, suppose that the potential convert has image concerns $\mu M_i^+(\theta_i^*) = \mu E_{F_i}[\theta|\theta \geq \theta_i^*]$ if he does not convert and $\mu M_i^-(\theta_i^*) = \mu E_{F_i}[\theta|\theta \leq \theta_i^*]$ if he does, where θ_i^* is the threshold type in district i and $\mu \geq 0$ is a parameter of intensity of image concerns. $M_i^+(\theta_i^*)$ and $M_i^-(\theta_i^*)$ are the upward and downward truncated means, respectively (i.e. the expectations of θ conditional on θ being above or below θ_i^*). The cutoff θ_i^* (or alternatively the tax $\tau(\theta_i^*)$ that induces θ_i^*) is then given by

$$\theta_i^* - \tau_i + \mu\Delta(\theta_i^* - r_i) = 0,$$

where $\Delta(\theta^*) \equiv E_F[\theta|\theta \geq \theta^*] - E_F[\theta|\theta < \theta^*]$. The co-variation of the threshold and the discriminatory tax is no longer 1 for 1 if $\mu>0$, and is given by:

$$\frac{d\tau_i}{d\theta_i^*} = 1 + \mu\Delta'(\theta_i^* - r_i).$$

Let us assume that image concerns are not too large, $1 + \mu\Delta'(\theta_i^* - r_i)>0$, and so the equilibrium threshold is unique and $\tau(\theta_i^*)$ well-defined. The new revenue function is $\hat{R}_i(\tau_i) \equiv \tau_i[1 - F(\theta_i^*(\tau_i) - r_i)]$. The analysis is unchanged, except that now LA i's objective function is:

$$V_i = \lambda_i + (\tau_i - c_i)[1 - F(\theta_i^*(\tau_i) - r_i)] - T_i.$$

Introducing social pressure adds a few interesting additional insights, though. If the distribution $f(\cdot)$ is unimodal, the function $\Delta(\theta^*)$ is U-shaped. When conversions are rare, the reputational concern is driven mainly by the strong stigma attached to conversions (and so $\Delta'(\theta_i^* - r_i)<0$). The discriminatory tax has a strong impact on the threshold because it not only provides a material incentive for conversion, but it also releases the social stigma attached to conversions. When in contrast there are few Copts remaining, reputational concerns are mainly driven by the social prestige attached to resistance (and so $\Delta'(\theta_i^* - r_i)>0$); the discriminatory tax impact on the threshold is then less than

1 for 1.[73] The model can also easily be extended to allow for *network externalities*.

(d) *Malthusian ruler.* Suppose now that agents care not only about consumption and identity, but also about the number of their children. We use a model à la Galor and Weil (2000) and enrich it through an identity decision. A district-i agent's utility is[74]

$$U(\theta) = \max_{z\in\{0,1\}} \frac{\rho^{1-\alpha}}{\alpha^{\alpha}(1-\alpha)^{1-\alpha}} a^{\alpha} n^{1-\alpha} + \theta z$$

s.t.

$$a + \rho n \leq y - \lambda_i - \tau_i z,$$

where z equals 1 if the agent maintains his identity and 0 if he converts, a is consumption, n the number of children, y the endowment, ρ the cost of a child's upbringing, and $\alpha \in (0,1)$. Hence

$$U(\theta) = y - \lambda_i + (\theta - \tau_i)z,$$

which yields, as in the model without fertility choice, cutoff

$$\theta^* = \tau_i.$$

Suppose now that LA i is motivated to reduce the number of non-converts:

$$V_i = \lambda_i + (\tau_i - c_i)[1 + \nu n_i][1 - F(\tau_i - r_i)] - T_i$$

where some weight $\nu > 0$ is put on the indirect conversions (of children). Let us show that n_i is a decreasing function of τ_i. A non-convert's number of children is given by $\rho n_i = (1-\alpha)(y - \lambda_i - \tau_i)$. So n_i is a decreasing function of τ_i. Note that the LA, when raising the poll tax, achieves double benefits: directly by inducing the adult generation to convert, and indirectly by making holdouts poorer and therefore reducing their reproductive rate. Appendix Section B.1 fails to find empirical support for this indirect mechanism in our historical context, but it might be relevant to other contexts.

B Empirics

B.1 Conversion or demographic Islamization?

An alternative theory of Egypt's, and the region's, Islamization traces the process to population replacement, in the sense that Arabs (Muslims) replaced the local non-Muslim populations of the region, rather than to conversions to Islam among the local populations. In the absence of Copts' conversion to Islam, five demographic processes could have driven the decline in Egypt's non-Muslim population share between 641 and 1200, and subsequently through 1848 (Fargues 2001):[75] Muslim immigration into

[73] One can go further in the elasticity analysis by assuming that $\Delta''(\theta_i^* - r_i) > 0$ (a hypothesis for which Jia and Persson (2017) find supporting evidence in a different context).

[74] In this version, the agent cares about his own identity or, alternatively, about the identity of his dynasty.

[75] This section draws on and expands the discussion in Saleh (2018, pp. 425-426).

Egypt, Coptic emigration, Muslims' higher fertility (net of child mortality), Muslims' lower adult mortality, and intermarriage between Coptic females and Muslim males (the opposite scenario is prohibited) without pre-marriage conversion, which results by law in a Muslim offspring.[76] These processes, we argue, are *not* the main causes of Islamization.

Muslim immigration. Arab immigration, the largest Muslim immigration wave in Egypt between 641 and 1200, was small compared to the Egyptian (Coptic) population. In 641, Egypt's population (2.7 million) was three times that of the Arab peninsula (1 million) (Russell 1958, p. 89). Russell (1966) estimates the number of Arab immigrants in 650 at 100,000. Furthermore, Arab immigration subsided after 833 with the shift to recruiting slave armies and the stoppage of state stipends to Arabs, which led Arabs to lose their military aristocratic position to Turks. It is also important to note that if Arab immigration were the sole driver of the decline in Egypt's non-Muslim population share between 641 and 1200, we would normally expect Arabs (Muslims) to be better off, on average, than Copts, because Arabs dominated by law the top white-collar positions in the military, judiciary, police, and the high-level bureaucracy, and because Copts were subject to a higher tax. This prediction contradicts though the papyrological evidence in 641-969 that shows that Copts were better off than Muslims; they were over-represented among white-collar workers and artisans and under-represented among farmers and unskilled non-agricultural workers (Saleh 2018).

Copt emigration. Copts rarely emigrated from Egypt, because of their unique Christian denomination that differed from both Catholics and Greek Orthodox Christians. Until today, Coptic Christianity has been considered a "heretical" "non-Chalcedonian" Oriental Orthodox Christian denomination, which split from the Roman/Byzantine Church at the Council of Chalcedon in 451. Egypt's Chalcedonian Christians, who remained loyal to the Roman/Byzantine Church, formed a small minority called the *Melkites*.

Coptic-Muslim fertility difference. Even if Arab immigration was small compared to Egypt's population, Muslims could have gradually replaced Copts over time if they had more children.[77] While this alternative hypothesis (which rules out Copt conversions to

[76] A marriage in which a Coptic male converts to Islam prior to marriage is excluded because the mechanism of converting the offspring in this case is paternal conversion, and not cross-marriage per se.

[77] In 641, Egypt's (coptic) population was about 2.5 million, and Arab immigrants were about 100,000 (4%). In 1200, Egypt's population was 2.3 million, with Muslims constituting 84% (1.9 million) and Copts 16% (0.3 million). This implies that over the course of 560 years (641 to 1200), Arabs grew by 18 times (0.53% annually), while Copts lost 88% of their population (declined by 0.38% annually). In order for Arab (Muslim) settlers to grow from a small minority (4%) in 641 to the majority (84%) by 1200 by fertility privilege alone (without Coptic conversions), total fertility rate must have been at least 2.3 child per woman for Muslims, and 1.8 for Copts, *assuming that there were neither (child) mortality*

Islam) still does not explain why Copts were better off than Muslims as early as in 641-969, we attempt to test it directly using the 1848 and 1868 census samples which were digitized by Saleh (2013). Because these censuses predate Egypt's demographic transition, which started in the second half of the twentieth century, they provide a glimpse of the demographics of medieval (Malthusian) Egypt. They also allow us to measure the number of *surviving* children, which is arguably a better measure of the *desired* number of children than the number of children *ever born*, which we do not observe. Specifically, our measure is fertility net of child mortality: the number of surviving children below 10 years and below 1 year. Measuring fertility from the population censuses is subject to two caveats, though: (1) We only observe children who reside with their parent(s) at the time of the census. But this is less of a concern for children below 10, who are more likely to live with their parent(s). (2) We do not observe the father and mother of every individual in the censuses (except for children of the household head), but we inferred the (potential) father and mother from the relationship to the household head (the household structure). The findings in Table B.1 reveal that Muslim males do *not* have more surviving children than Coptic males, whether we count the number of surviving children below 10 years of age or below 1. This null finding holds within each occupational group: unskilled non-agricultural workers, farmers, artisans, and white-collar workers. Furthermore, Muslim females have *fewer* children under 10 than their Coptic counterparts, especially in households headed by farmers and white-collar workers, but the difference is statistically insignificant if we measure fertility by the number of surviving children under 1 (except for females in households headed by white-collar workers).

Coptic-Muslim adult mortality difference. Measuring adult mortality from the population censuses is more challenging, because we do not observe deaths. Saleh (2018) measures adult life expectancy among Copts and Muslims by comparing the age distribution between 1848 and 1868. The findings in Table B.2 (taken from the Online Appendix of Saleh (2018)) show that Muslims had lower adult mortality (higher life expectancy) at younger ages (10-29 or 10-39), but higher adult mortality (lower life expectancy) at older ages (30-79 or 40-79). However, the differences are small in magnitude, and may be attributable to statistical caveats in the 1848 and 1868 censuses, namely, (1) the gap (20 years) that separates the two censuses is longer than ideal (5 or 10 years) as it increases the chance of population movement, and (2) age heaping (tendency to report age as a number ending in "0" or "5") and age exaggeration (for older individuals); since both phenomena are negatively correlated with socioeconomic

nor migration.

Table B.1: **Coptic-Muslim fertility difference in 1848 and 1868**

	Males				Females			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Children <10	Children <10	Children <1	Children <1	Children <10	Children <10	Children <1	Children <1
Copt	-0.050	-0.128	0.015	0.012	0.159	0.085	0.034	0.013
	(0.067)	(0.120)	(0.029)	(0.052)	(0.069)**	(0.075)	(0.023)	(0.022)
Farmer		0.067		0.050		0.243		0.073
		(0.062)		(0.015)***		(0.037)***		(0.012)***
Artisan		-0.070		-0.027		0.374		0.091
		(0.092)		(0.025)		(0.101)***		(0.029)***
White-collar		0.424		0.086		0.109		0.032
		(0.090)***		(0.030)***		(0.085)		(0.013)**
Copt * Farmer		0.261		-0.036		0.320		0.022
		(0.153)*		(0.048)		(0.121)***		(0.037)
Copt * Artisan		0.042		0.049		-0.223		-0.005
		(0.228)		(0.069)		(0.206)		(0.063)
Copt * White-collar		-0.118		0.012		0.373		0.147
		(0.188)		(0.066)		(0.213)*		(0.061)**
Constant	1.836	1.768	0.328	0.301	1.198	1.120	0.211	0.188
	(0.036)***	(0.059)***	(0.012)***	(0.011)***	(0.022)***	(0.027)***	(0.007)***	(0.006)***
Obs (individuals)	22119	22119	22119	22119	14780	14780	14780	14780
Clusters (districts)	106	106	106	106	98	98	98	98
R^2	0.00	0.01	0.00	0.00	0.00	0.01	0.00	0.01
Mean dep. var.	1.54	1.54	0.23	0.23	1.20	1.20	0.21	0.21

Notes: Robust standard errors clustered at the district level are in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$. The omitted group is unskilled non-agricultural Muslim workers.
Source: The 1848 and 1868 population census samples (Saleh 2013) and an over-sample of non-Muslims in Cairo in 1848 and 1868. Census samples are pooled and restricted to Copts and Muslims aged 15 to 60 years. Regressions are weighted by sample design. Because almost all females have missing occupations, we assigned the household head's occupational title to all household members with missing occupations, including females. Number of children is inferred from the relationship to the household head, and includes only surviving children residing with their parent(s) at the time of the census.

status, they are less prevalent among Copts.

Cross-marriages without pre-marriage conversion. Another way of replacing the Coptic population is by Arab (Muslim) males marrying (possibly more than one) Coptic females, as the off-spring in this case will be Muslim. Cross-marriages between Muslim males and Coptic females were rare as suggested by the dearth of cross-marriage contracts in the papyri in 641-969. The 1848 and 1868 population census samples record only two cross-marriages.

Table B.2: **Estimating adult life expectancy from the 1848 and 1868 population census samples**

	Copts				Muslims			
Age Group	Estimated size in 1848	Estimated size in 1868	Estimated life expectancy (method 1)	Estimated life expectancy (method 2)	Estimated size in 1848	Estimated size in 1868	Estimated life expectancy (method 1)	Estimated life expectancy (method 2)
0-9	90,740	117,801	NA	NA	1,148,827	1,458,614	NA	NA
10-19	32,981	51,600	41.45	42.9	377,685	603,264	43.44	44.82
20-29	33,290	52,466	44.59	44.59	406,293	622,071	49.08	48.73
30-39	40,100	36,657	30.44	32.2	457,208	481,535	32.97	32.65
40-49	27,031	26,187	25.46	24.72	348,101	360,926	25.9	23.79
50-59	15,325	25,345	22.61	21.02	243,063	288,588	21.83	19.98
60-69	11,406	12,595	17.67	16.1	171,180	195,387	16.88	13.53
70-79	7,849	10,899	11.52	9.03	99,442	111,561	12.26	8.68
80+	7,094	5,107	NA	NA	125,336	78,559	NA	NA

Notes: The handbook of the United Nations Population Division (2002, pp. 5-20) outlines a methodology for estimating adult mortality from any two consecutive censuses that are separated by an interval of x years, where x is a multiple of 5. The methodology uses the relative sizes of age cohorts, defined in groups of 5-year intervals, in the two censuses in order to estimate the probability of survival to an age $y+x$, conditional on being of age y in the first census. A slightly different methodology, the synthetic survival ratio, calculates the growth rate of each age cohort in order to make the methodology applicable to any census interval, i.e. not necessarily a multiple of 5. We applied the two methods to the census samples of 1848 and 1868, in order to estimate adult mortality by religious group. A few caveats arise though: (a) the time interval separating the two Egyptian censuses (20 years) is too long to apply the two methodologies; ideally, the interval should be around 5 or 10 years, (b) we do not have 100-percent samples of the two censuses and so there is a sampling error in estimating the size of each age cohort, and (c) there is a problem of age misreporting; in particular, age heaping and age exaggeration, which is typical in historical censuses and even contemporary censuses in developing countries. Age misreporting is likely correlated with socioeconomic status and may thus vary in a non-random way across religious groups, where Muslims are more likely than Copts to misreport their true age. In order to mitigate age misreporting, we defined age groups in intervals of 10 years instead of 5 years.
Source: The 1848 and 1868 population census samples. This table is reproduced from the Online Appendix of Saleh (2018).

B.2 Data sources

Cross-sectional evidence. We use the following sources of data:

- Identity strength of local authorities (c_i): We constructed a dummy variable at the *kura* level that takes value 1 if at least one Arab tribe settled permanently in *kura i* between 700 and 969, based on al-Barri (1992). This is a secondary source that draws on medieval Muslim narratives, and in particular, *al-Bayan wal-I'rab 'amman fi Ard Misr min al-A'rab* (Arab Tribes in Egypt) by al-Maqrizi (died in 1442).

- Identity strength of Copts (r_{ji}):

  1. We constructed a dummy variable that takes value 1 if it is believed, according to pre-641 local Coptic legends, that the village was visited by the Holy Family during its legendary biblical flight to Egypt. The list of villages is recorded in Anba Bishoy (1999) and Gabra (2001). They are both based on the apocryphal book *Vision of Theophilus* that was translated and published by Mingana in 1931 (Mingana 1931). The book is attributed to Theophilus, the patriarch of Alexandria from 384 to 412, but both the authorship and date are doubtful.
  2. We constructed a dummy variable that takes value 1 if a Coptic saint or martyr is believed to have lived in the village before 641. The list of Coptic saints and martyrs is based on the Coptic Synaxarium *Le Synaxaire arabe-jacobite* translated by R. Basset (Basset 1907).

- Proportion of converts (F_{ji}):

  1. We collected village-level data on Coptic churches and monasteries circa 1200 from the medieval chronicle *History of Churches and Monasteries* by the Coptic priest Abul-Makarim Sa'dullah ibn Jirjis ibn Mas'ud, who died circa 1208 (Abul-Makarim 1984). We use the two-volume version edited by Anba Samuel in 1984. The first volume of this version is a reprint of the 1895 book *The Churches and Monasteries of Egypt and Some Neighbouring Countries* that was translated and edited by B. Evetts and A. Butler, and was wrongly attributed to Abu-Saleh The Armenian.
  2. We constructed a second village-level dataset on Coptic churches and monasteries circa 1500 from the medieval chronicle *al-Mawa'iz wal-I'tibar fi Zhikr al-Khitat wal-'Athar* (Sermons and Considerations in Examining Plans and Monuments) by the medieval historian al-Maqrizi (2002).

3. The universe of Egyptian villages is based on the list of villages in the 1315 cadastre recorded in the medieval manuscript: *al-tuhfa al-saniyya bi asma' al-bilad al-misriya* (Book of Names of Egyptian Localities) by Ibn al-Ji'an (died circa 1250) Ibn al-Ji'an (1898).
4. The list of Egyptian villages mentioned in Byzantine sources is constructed from Amélineau (1893).
5. The number of Coptic households in the *kura* of *Fayum* is constructed from Rapoport (2018) based on the 1245 cadastre of *Fayum* in *Tarikh al-Fayum (History of Fayum)* by al-Nabulsi (died circa 1250).

- Discriminatory tax rate (τ_{hi}): We constructed an individual-level dataset on poll tax payments from papyrological tax registers and receipts in 641–1100. We employ Morimoto (1981, pp. 67-79, 85-87) for Greek papyri, and the Arabic Papyrology Database for Arabic papyri.[78]

- Total tax transfer ($\tilde{T}_{ji}$): We collected village-level data on total tax transfer (*'ibra*) per unit of land from the cadastral surveys of 1315 and 1375, based on Ibn al-Ji'an (1898).

- Byzantine-period *kura*-level controls: We employ the natural logarithm of urban population circa 300, based on Wilson (2011, pp. 185-187). Byzantine military garrisons circa 600 are constructed from Maspero (1912). *Autopract* estates circa 600 are constructed from Hardy (1931).

- Geographic village-level controls are from the Food and Agriculture Organization's Global Agro-Ecological Zones (FAO-GAEZ) Data Portal 3.0.1. Crop suitability indices are under irrigation and intermediate input level. FAO-GAEZ does not report crop suitability indices under irrigation and *low* input level. Population is from the 1897 population census (Ministère des finances 1898).

Time-series evidence.

- Identity strength of central authority (c_t):

[78] We do not employ two other sets of tax papyri. First, there are other Coptic and Greek poll tax registers and receipts in 641–800 that we do not use because they have not been digitized yet. These papyri are either from the same *kuras* as in our sample, and thus adding them will not augment the statistical power of our analysis, or from monasteries, and hence are not representative of the non-monastic population. Second, there are poll tax receipts from Nessana in Palestine, which we do not use because they do not vary within Palestine (they come from a single location).

1. Measure 1 is at the governor level. It takes value 1 if the governor is hostile toward non-convert Copts. It is based on two medieval Coptic chronicles: (1) *The Chronicle of John, the Bishop of Nikiu* for the Rashidun period (641–661): John of Nikiu (located in the Nile Delta) died circa 690 CE, and was thus an eyewitness of the Arab Conquest of Egypt. His chronicle was translated by R. H. Charles and published as John of Nikiu (1916). (2) *History of the Patriarchs of the Coptic Church of Alexandria* for 661–1170: This chronicle is the "Book of the Popes" of the Coptic Church of Alexandria. It covers the history of the Church from its customary foundation by Saint Mark circa 49 CE up until 1894 CE. The book was compiled by multiple authors, but the most important author is Severus ibn al-Muqaffa, the bishop of *Hermopolis*, who lived from 915 until the end of the 10th century. Severus compiled the biographies of popes from 49 up to circa 1000 CE. Volume I of the book, covering the period from 49 CE to 849 CE, was translated into English and published by B. Evetts in four parts as Ibn al-Muqaffa (1910). The following parts were translated into English and published from 1943 to 1974, of which we use Volume II: Part I-III and Volume III: Part I in Ibn al-Muqaffa (1943).
2. Measure 2 is at the Caliph level. It takes the value 1 if the Caliph is not known for drinking alcohol. We use Sirhan (1978) for the period 641–868, al-Dhahabi's *The Lives of Noble Figures* (al-Dhahabi 1982) for 868–969, and al-Maqrizi's *History of the Fatimid Caliphs* (al-Maqrizi 1996) for 969–1170.

- Budget needs (B_t): The yearly number of major military campaigns that were initiated by the Caliphate against foreign empires is constructed from Mikaberidze (2011).
- External threats (x_t): The yearly number of major military attacks that were initiated by foreign empires against the Caliphate is constructed from Mikaberidze (2011).
- Nile adverse shocks is constructed from Chaney (2013).

B.3 Figures and tables

This section presents additional figures and tables.

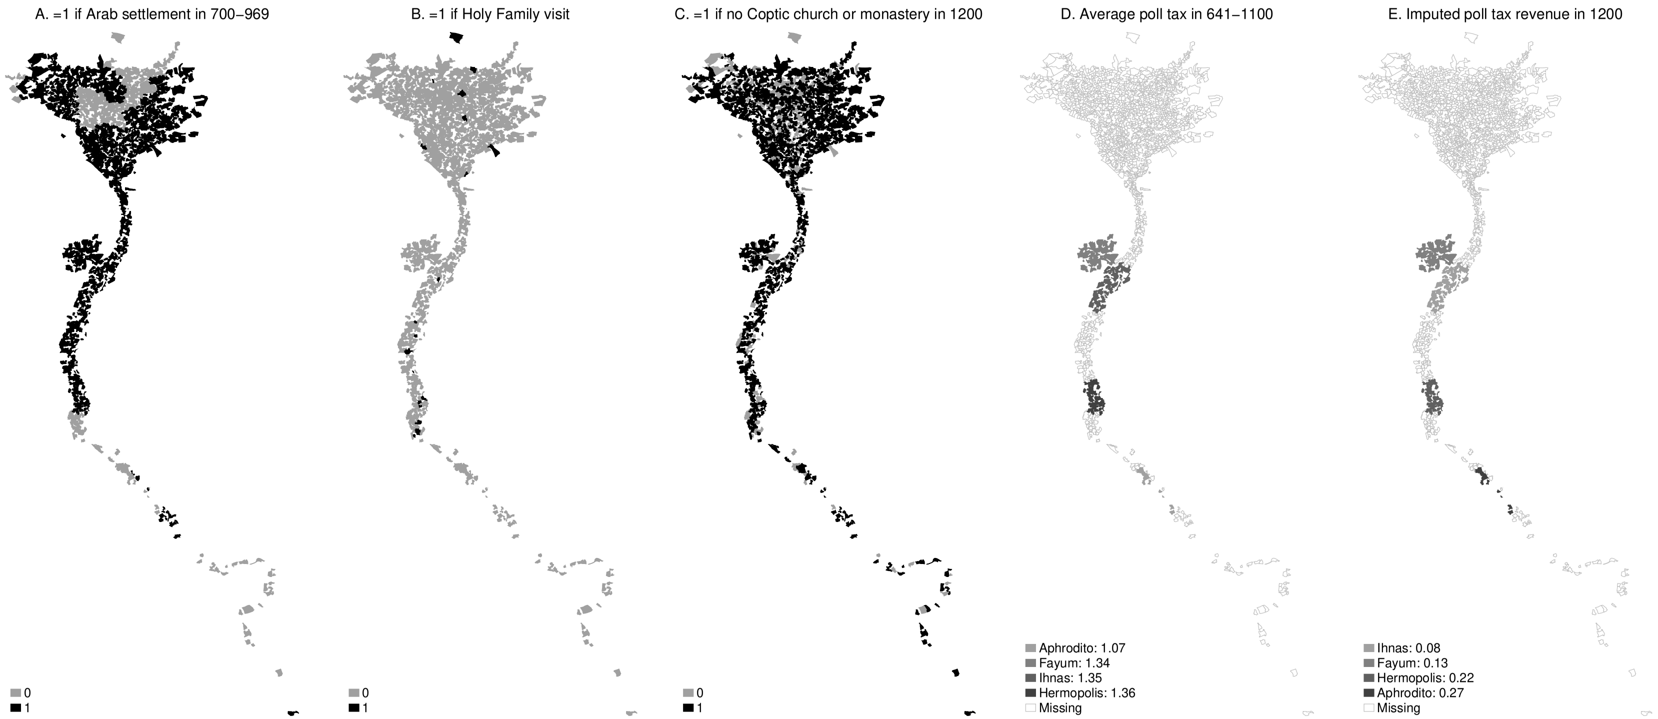

Figure B.1: **Cross-sectional spatial heterogeneity in determinants and outcomes**

Notes: (1) The map shows 1,782 villages that appear in the 1315 cadastre, which defines our universe of villages, using the boundaries of villages in the 2006 population census. (2) The Nile Delta refers to the northern triangle on the map. The Nile Valley covers the whole region to the south of the Delta.
Sources: See Appendix B.2.

Table I Register of Gold-Tax Assessment for "Five Fields",
The 3rd Indiction (704/05 : A.H. 85/86)

Taxpayers	Location of fields	Taxes: Land sol.	Taxes: Corn art.	Taxes: Land sol.	Taxes: Poll sol.	Total sol.	Corn tax art.
Mēnas Apollōs	Belekau			½	3	3½	½
Kaumas Antheria	Sarseltōh			2½	2½	5	3
Psoios Andreas	Pkathakē Pkarou	1 ⅙	1 ⅓	1⅙	1½	2⅔	1⅓
Horsenuphios Hermaōs	Ammōniu Pankul & others Piah Alau	8½ 1½ ½	10 1½ 1	10½	4	14½	12½
Abraham Theodosios	Piah Boōn Piah Kam Hagiu Biktōr	½ 1 1	½ 1½ 1	2½	4½	7	3
Bethanias Pkaloos	Pkarou			⅓	0	⅓	⅓
Taam, Johannes Th[ ]-liaie & Eudoxia	Pkarou & Belekau			2⅓	0	2⅓	2⅓

P.Cair.Arab. 202 : List of poll-tax payers. (unknown (Egypt); 1. 1. 801 - 31. 12. 900 CE)

Inv. No.: **P.Cair.EgLib.inv. 293 recto and verso** Material: **papyrus** Height: **19 cm.** Width: **12 cm.** Language(s): **Arabic** Kind: **List, account** Edition(s): **Grohmann, P.Cair.Arab. 202** Translation(s): **Grohmann, P.Cair.Arab. 202**

recto 1 *** سدس وثمن شنوده القرّاش نصف وربع

recto 2 *** وثمن بستّة كرسدوده ثلثة دنانير ونصف وثلثمن

recto 3 *** وسدس كيل بستّه ثلثي وسدس ثمن

recto 4 *** دينر أدّى سويرس أبلّوا ثمن دينر

Figure B.2: **Examples of the secondary sources of the poll tax sample**

Top: Morimoto (1981, p. 67): Register of "Five Fields" in *Aphrodito* in 704/05 CE. Bottom: Arabic Papyrology Database: List of poll-tax payers in 801-900.

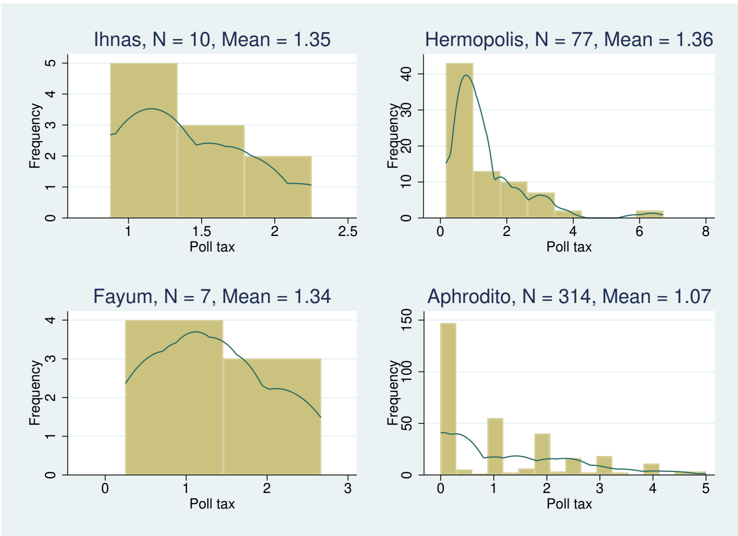

Figure B.3: **Histogram of individual poll tax payments by *kura* in 641–1100**

Notes: Arab settlement =1 in *Ihnas*, *Hermopolis*, and *Fayum*, and =0 in *Aphrodito*. Sources: Greek papyri in Morimoto (1981, pp. 67-79, 85-87) and Arabic papyri in the Arabic Papyrology Database. Sample is restricted to papyri from a known location.

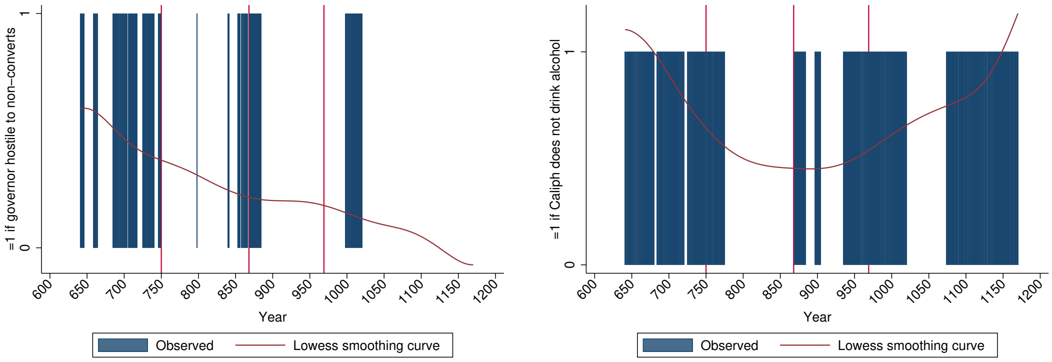

Figure B.4: **Measures of c_t in 641–1170**

The vertical red lines at years 750, 868, and 969, indicate major dynastic changes. 641–750: Rashidun and Umayyads; 750–868: First Abbasid Period; 868–969: Tulunids, Second Abbasid Period, Ikhshidids; 969–1170: Fatimids. Source: See Appendix Section B.2.

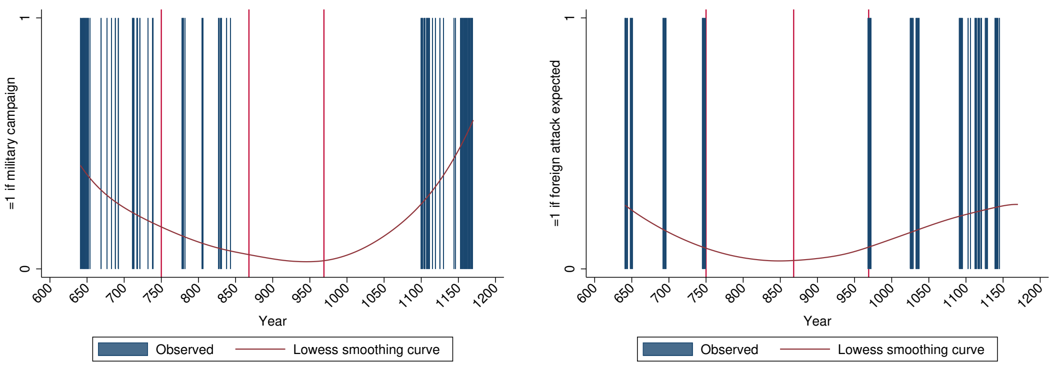

Figure B.5: **Measures of B_t and x_t in 641–1170**

Expected foreign attacks is a dummy variable =1 in the five years preceding a foreign attack. The vertical red lines at years 750, 868, and 969, indicate major dynastic changes. 641–750: Rashidun and Umayyads; 750–868: First Abbasid Period; 868–969: Tulunids, Second Abbasid Period, Ikhshidids; 969–1170: Fatimids. Source: See Appendix Section B.2.

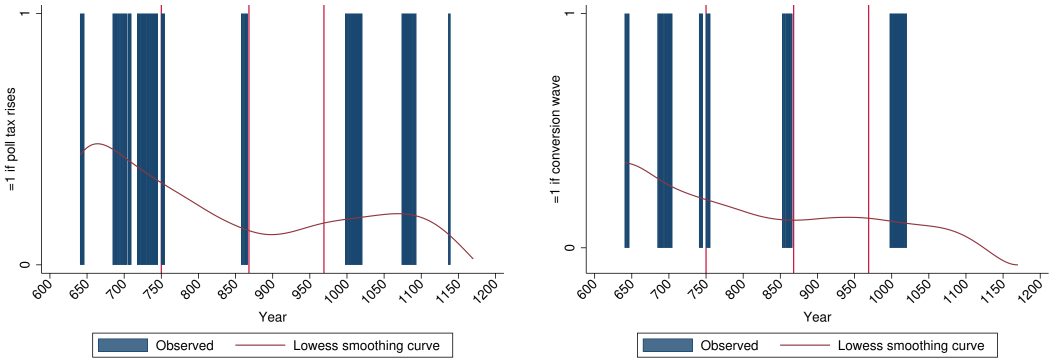

Figure B.6: **Poll tax rises ($\Delta\tau_t$) and conversion waves (ΔF_t) in 641–1170**

The vertical red lines at years 750, 868, and 969, indicate major dynastic changes. 641–750: Rashidun and Umayyads; 750–868: First Abbasid Period; 868–969: Tulunids, Second Abbasid Period, Ikhshidids; 969–1170: Fatimids. Source: See Appendix Section B.2.

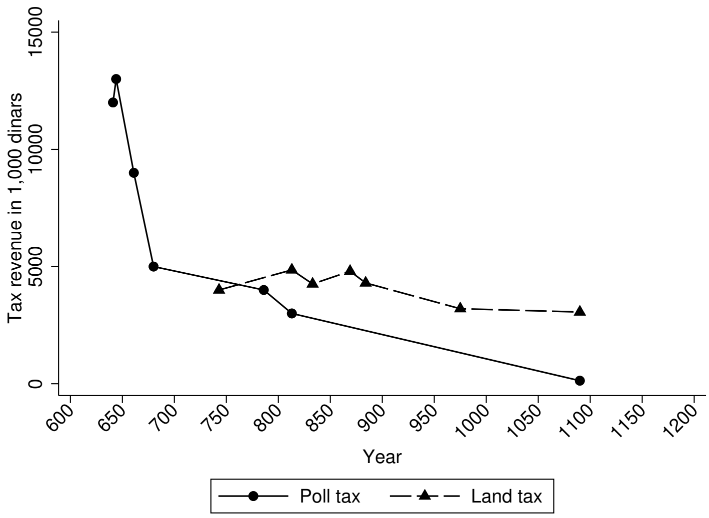

Figure B.7: **Total poll and land tax revenues in 641–1170**

Sources: Courbage and Fargues (1997); poll tax revenue in 1090: Mahmoud (2009).

Table B.3: **Local determinants of F_{ji} in 1200 - Alternative measure of r_{ji}**
Dependent variable: =1 if no Coptic church or monastery in village j circa 1200

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) 2SLS
=1 if Arab settlement (c_i)	0.083		0.083	0.078	0.074	0.111
	(0.033)**		(0.031)**	(0.029)**	(0.030)**	(0.050)**
=1 if pre-641 Coptic saint or martyr (r_{ji})		-0.503	-0.500	-0.502	-0.520	-0.517
		(0.093)***	(0.098)***	(0.097)***	(0.091)***	(0.091)***
Byzantine *kura*-level controls?	No	No	No	Yes	Yes	Yes
Geographic village-level controls?	No	No	No	No	Yes	Yes
Obs (villages)	1782	1778	1778	1778	1748	1748
Clusters (*kuras*)	42	42	42	42	42	42
R^2	0.01	0.03	0.04	0.04	0.05	0.05
Mean dep. var. in control	0.78	0.85	0.78	0.78	0.78	0.78
KP Wald F-stat						16.65

Notes: Standard errors clustered at the *kura* level are in parentheses. Byzantine-period controls are: (1) the logarithm of urban population in *kura* i circa 300, and (2) a dummy variable =1 if there was a Byzantine garrison in *kura* i circa 600. Geographic controls are: (3) FAO-GAEZ suitability index to the cultivation of barley, wheat, beans, and maize, under irrigation and intermediate input level, (4) mean temperature, (5) temperature range, (6) slope, and (7) rainfall. * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Sources: See Appendix Section B.2.

Table B.4: **Local determinants of F_{ji} in 1200 - Amélineau's villages only**

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) 2SLS
=1 if Arab settlement (c_i)	0.037		0.059	0.079	0.123	0.073
	(0.087)		(0.081)	(0.076)	(0.069)*	(0.124)
=1 if Holy Family visit (r_{ji})		-0.597	-0.607	-0.611	-0.581	-0.578
		(0.044)***	(0.049)***	(0.053)***	(0.069)***	(0.063)***
Byzantine *kura*-level controls?	No	No	No	Yes	Yes	Yes
Geographic village-level controls?	No	No	No	No	Yes	Yes
Obs (villages)	163	163	163	163	157	157
Clusters (*kuras*)	39	39	39	39	37	37
R^2	0.00	0.08	0.08	0.09	0.14	0.14
Mean dep. var. in control	0.54	0.60	0.54	0.54	0.54	0.54
KP Wald F-stat						15.71

Notes: Robust standard errors clustered at the *kura* level are in parentheses. * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Source: See Appendix Section B.2.

Table B.5: **Measuring F_{ji} in 1500 and in 1848–1868**

	=1 if no Coptic church or monastery in 1500				=1 if Copt in 1848–1868			
	(1) OLS	(2) OLS	(3) OLS	(4) 2SLS	(5) OLS	(6) OLS	(7) OLS	(8) 2SLS
=1 if Arab settlement (c_i)	0.035	0.035	0.039	0.040	-0.071	-0.075	-0.031	-0.079
	(0.025)	(0.024)	(0.011)***	(0.011)***	(0.040)*	(0.042)*	(0.020)	(0.031)**
=1 if Holy Family visit (r_j)		-0.280	-0.278	-0.278		0.171	0.137	0.129
		(0.068)***	(0.074)***	(0.073)***		(0.061)***	(0.045)***	(0.043)***
Byzantine controls?	No	No	Yes	Yes	No	No	Yes	Yes
Geographic controls?	No	No	Yes	Yes	No	No	Yes	Yes
Obs (villages)	1782	1782	1751	1751	16641	12935	12935	12935
Clusters (*kuras*)	42	42	42	42	42	41	41	41
R^2	0.01	0.05	0.12	0.12	0.02	0.03	0.07	0.07
Mean dep. var. in control	0.95	0.95	0.95	0.95	0.12	0.12	0.12	0.12
KP Wald F-stat				15.61				5.60

Notes: Robust standard errors clustered at the *kura* level are in parentheses. * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Source: Village-level data on Coptic churches and monasteries in 1500 constructed from al-Maqrizi (2002). Individual-level data on religious affiliation in 1848–1868 are from the 1848 and 1868 individual-level population census samples restricted to Egyptian local free Coptic and Muslim employed men of a rural district of origin who are at least 15 years of age and with non-missing information on age, religion, occupation, and district of origin.

Table B.6: **Local determinants of F_{ji} in 1200: spatial-autoregressive model with spatial-autoregressive errors**

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) 2SLS
=1 if no Coptic church or monastery in 1200						
=1 if Arab settlement (c_i)	0.081		0.078	0.076	0.072	0.103
	(0.020)***		(0.020)***	(0.021)***	(0.021)***	(0.026)***
=1 if Holy Family visit (r_{ji})		-0.621	-0.610	-0.612	-0.619	-0.616
		(0.080)***	(0.079)***	(0.079)***	(0.079)***	(0.079)***
Byzantine *kura*-level controls?	No	No	No	Yes	Yes	Yes
Geographic village-level controls?	No	No	No	No	Yes	Yes
Obs (villages)	1730	1730	1730	1730	1730	1730
Mean dep. var. in control	0.78	0.85	0.78	0.78	0.78	0.78

Notes: Columns (1)-(5) report the results of estimating a spatial autoregressive model with spatial standard errors with inverse distance weighting matrix estimated using generalized spatial two-stage least squares (GS2SLS) (STATA command spreg). Column (6) reports a spatial-autoregressive model with spatial-autoregressive errors and additional endogenous variables (STATA command spivreg). * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Source: See Appendix Section B.2.

Table B.7: **Relevance and exogeneity of the instrumental variables**

	(1) =1 if Arab settlement in 700–969	(2) =1 if Holy Family visit	(3) Log (urban population) in 300	(4) =1 if Byzantine garrison in 600	(5) FAO-GAEZ cereals suitability	(6) Mean temperature	(7) Mean Temperature range	(8) Mean slope	(9) Mean rainfall
Kura's distance to *Arish* (km)	0.011	0.005	-0.005	0.008	-0.001	0.000	0.008	0.002	-0.198
	(0.004)**	(0.006)	(0.001)***	(0.005)	(0.000)	(0.003)	(0.004)*	(0.003)	(0.287)
=1 if *Kura* borders desert	3.715	1.168	-1.004	2.223	-0.130	-1.982	0.529	0.093	-36.071
	(1.020)***	(1.419)	(0.492)**	(1.211)*	(0.112)	(0.668)***	(1.058)	(0.677)	(67.489)
Bordering desert × Dist. *Arish*	-0.012	-0.006	0.006	-0.008	0.000	0.009	0.001	0.000	0.053
	(0.004)***	(0.006)	(0.002)***	(0.005)	(0.000)	(0.003)***	(0.004)	(0.003)	(0.288)
Observations	42	42	42	42	42	42	42	42	42
R^2	0.376	0.134	0.109	0.061	0.377	0.828	0.776	0.383	0.379

Notes: White-Huber robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. A constant term is included in all regressions.
Sources: See Appendix Section B.2.

Table B.8: **Local determinants of F_{ji} in 1200: Alternative specifications**

	(1)	(2)
=1 if Arab settlement (c_i)	0.076	0.101
	(0.032)**	(0.068)
=1 if Holy Family visit (r_{ji})	-0.626	-0.640
	(0.181)***	(0.108)***
Arab settlement × Holy Family visit	0.003	
	(0.211)	
=1 if *Autopract* estates circa 600		-0.006
		(0.068)
Geographic village-level controls?	Yes	Yes
Byzantine *kura*-level controls?	Yes	Yes
Obs (villages)	1751	575
Clusters (*kuras*)	42	21
R^2	0.05	0.09
Mean dep. var.	0.78	0.78

Notes: Standard errors clustered at the *kura* level are in parentheses. * p <0.10, ** p <0.05, *** p <0.01. A constant is included in all regressions.
Sources: See Appendix Section B.2.

Table B.9: **Evaluating the national representativeness of the poll tax sub-sample**

(a) Village-level variables

	Villages out of poll-tax sample			Villages in poll-tax sample			
	N	Mean	SD	N	Mean	SD	Diff
=1 if no church or monastery in 1200	1589	0.83	0.38	195	0.89	0.31	0.061*
=1 if no church or monastery in 1500	1589	0.98	0.15	195	0.94	0.24	-0.038
ibra per *feddan* in 1375	1336	3.23	2.33	176	4.01	6.40	0.793
ibra per *feddan* in 1477	1336	2.78	2.03	176	3.51	6.45	0.749
=1 if on H. Family route	1589	0.01	0.11	195	0.03	0.16	0.014
=1 if pre-641 Coptic saint or martyr	1585	0.02	0.12	195	0.03	0.17	0.016
FAO-GAEZ cereals suitability index	1560	0.68	0.10	191	0.66	0.10	-0.024
Mean temperature	1560	20.98	0.82	191	21.88	0.30	0.899***
Mean temperature range	1560	14.17	1.04	191	16.34	0.23	2.167***
Mean slope	1560	3.43	0.61	191	3.90	0.63	0.467***
Mean rainfall	1560	50.26	33.27	191	6.43	3.31	-43.832***

(b) *Kura*-level variables

	Kuras out of poll-tax sample			Kuras in poll-tax sample			
	N	Mean	SD	N	Mean	SD	Diff
=1 if Arab settlement in 700–969	38	0.63	0.49	4	0.75	0.50	0.118
Log (urban population) in 300	38	10.00	0.73	4	10.57	0.72	0.570
=1 if Byzantine garrison in 600	38	0.42	0.50	4	1.00	0.00	0.579***
Kura's distance to *Arish* (km)	38	354.07	148.34	4	425.86	83.63	71.792
=1 if *Kura* borders desert	39	0.74	0.44	4	1.00	0.00	0.256***

Notes: The "Diff" column reports the coefficient of the following regression: $y = \alpha_1 + \alpha_2 polltaxsample_i + \varepsilon$, where y is the outcome of village j located in *kura* i in panel (a), or the outcome of *kura* i in panel (b), and $polltaxsample_i$ is a dummy variable =1 if *kura* i is in the papyrological poll tax sub-sample. Standard errors are clustered at the *kura* level in panel (a), and are White-Huber heteroskedasticity-robust in panel (b). $*p < 0.10, **p < 0.05, ***p < 0.01$.
Sources: See Appendix Section B.2.

Table B.10: **Local determinants of F_{ji} in 1200 - Poll tax sub-sample**
Dependent variable =1 if no Coptic church or monastery in village in 1200

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) 2SLS
=1 if Arab settlement (c_i)	0.152		0.166	0.237	0.643	0.240
	(0.558)		(0.585)	(0.162)	(0.348)	(0.124)
=1 if Holy Family visit (r_j)		-0.505	-0.516	-0.501	-0.663	-0.520
		(0.784)	(0.796)	(0.695)	(0.915)	(0.746)
Byzantine *kura*-level controls?	No	No	No	Yes	Yes	No
Geographic village-level controls?	No	No	No	No	Yes	No
Obs (villages)	195	195	195	195	191	195
Clusters (*kuras*)	4	4	4	4	4	4
R^2	0.01	0.07	0.08	0.08	0.12	0.08
Mean dep. var. in control	0.75	0.91	0.75	0.75	0.75	0.75
KP Wald F-stat						1.17

Notes: The first-stage regression for the 2SLS in column (6) is $ArabSettlement_i = \alpha_0 + \alpha_1 DistancetoArish_i + v_i$. P-values are in parentheses, estimated by clustering standard errors at the *kura* level, using Wild Cluster Restricted (WCR) bootstrap for OLS regressions, and Wild Restricted Efficient (WRE) clustered bootstrap for IV regressions, with Webb weights and 999,999 replications. A constant is included in all regressions. Byzantine-period controls are: (1) the logarithm of urban population in *kura* i circa 300. Geographic controls are: (2) FAO-GAEZ suitability index to the cultivation of barley, wheat, beans, and maize, under irrigation and intermediate input level, (3) mean temperature, (4) temperature range, (5) slope, and (6) rainfall.
Source: See Appendix Section B.2.

Table B.11: **Local determinants of tax transfer in 1375**

	Tax transfer per unit of land			=1 if LA is Mamluk		
	(1)	(2)	(3)	(4)	(5)	(6)
=1 if no Coptic church or monastery in 1200		-0.806	-0.731		-0.123	-0.118
		(0.264)***	(0.209)***		(0.033)***	(0.037)***
=1 if Arab settlement in 700–969	-0.142	-0.076		0.078	0.088	
	(0.293)	(0.277)		(0.070)	(0.069)	
=1 if HF visit	0.913	0.444	0.365	0.248	0.174	0.034
	(0.415)**	(0.458)	(0.526)	(0.112)**	(0.119)	(0.141)
Kura fixed effects?	No	No	Yes	No	No	Yes
Geographic village-level controls?	No	No	Yes	No	No	Yes
Obs (villages)	1485	1485	1460	1746	1746	1599
Clusters (*kuras*)	40	40	40	42	42	40
R^2	0.00	0.01	0.11	0.01	0.02	0.22
Mean dep. var.	3.30	3.30	3.30	0.42	0.42	0.42

Notes: Tax transfer (*'ibra*) is in army dinars (≈13.3/20 dinars) per *feddan* (= 6,368 square meters) of land. Standard errors clustered at the *kura* level are in parentheses. Geographic village-level controls are: (3) FAO-GAEZ suitability index to the cultivation of barley, wheat, beans, and maize, under irrigation and intermediate input level, (4) mean temperature, (5) temperature range, (6) slope, (7) rainfall, and (8) population in 1897 ÷ land area in 1315. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. A constant is included in all regressions.
Sources: See Appendix Section B.2.

Appendix References

Abul-Makarim, S. (1984). *Tarikh al-Kana'is wal-Adyera (History of Churches and Monasteries)*. Dar al-Na'am, Cairo. Anba Samuel (Ed.).

al-Barri, A. K. (1992). *al-Qaba'il al-'Arabiya fi Misr fil Qurun al-Thalatha al-Ula lil Hijra (Arab Tribes in Egypt in the First Three Hijri Centuries)*. Egyptian General Book Authority, Cairo.

al-Dhahabi, S. (1982). *Siyar A'lam al-Nubala' (The Lives of Noble Figures)*. al-Risala, Beirut.

al-Maqrizi, T. (1996). *Itti'az al-Ḥunafa' bi-Akhbar al-A'immah al-Fatimiyin al-Khulafa' (History of the Fatimid Caliphs)*. Islamic Affairs Supreme Council, Cairo. Jamal-el-Din al-Shayyal (Ed.).

al-Maqrizi, T. (2002). *al-Mawa'iz wal-I'tibar fi Zhikr al-Khitat wal-'Athar (Sermons and Considerations in Examining Plans and Monuments)*, volume 4.1. Al-Furqan Islamic Heritage Foundation, London. A. F. Sayyid (Ed.).

Alesina, A., Baqir, R., and Easterly, W. (1999). Public Goods and Ethnic Divisions. *Quarterly Journal of Economics*, 114:1243–1284.

Amélineau, E. (1893). *La Géographie de l'Egypte à l'Epoque Copte*. Imprimerie nationale, Paris.

Anba Bishoy, P. (1999). *The Flight of the Holy Family to Egypt*. Nubar Printing House, Cairo.

Basset, R. (1907). Le Synaxaire Arabe-Jacobite (Redaction Copte). In Blundell, R., Newey, W., and Persson, T., editors, *Patrologia Orientalis*. Firmin-Didot, Paris. 1907-1929.

Besley, T., Jensen, A., and Persson, T. (2017). Norms, Enforcement, and Tax Evasion. mimeo.

Chaney, E. (2013). Revolt on the Nile: Economic shocks, religion, and political power. *Econometrica*, 81:2033–2053.

Courbage, Y. and Fargues, P. (1997). *Christians and Jews under Islam*. I. B. Tauris Publishers, London-New York. J. Mabro (Tr.).

Fargues, P. (2001). Demographic Islamization: Non-Muslims in Muslim Countries. *SAIS Review*, 21(2):103–116.

Gabra, G. (2001). *Be Thou There: The Holy Family's Journey in Egypt*. American University in Cairo Press, Cairo.

Galor, O. and Weil, D. (2000). Population, Technology and Growth: From the Malthusian Regime to the Demographic Transition. *American Economic Review*, 90:806–828.

Glaeser, E. and Shleifer, A. (2005). The Curley Effect. *Journal of Law, Economics, and Organization*, 21:1–19.

Hansmann, H. (1996). *The Ownership of Enterprise*. Belknap Press, Cambridge, MA.

Hardy, E. R. (1931). *The Large Estates of Byzantine Egypt*. Columbia University Press, New York.

Ibn al-Ji'an (1898). *al-Tuhfa al-Saniyya bi Asma' al-Bilad al-Misriya (Book of Names of Egyptian Localities)*. al-Matba'a al-Ahliya, Cairo.

Ibn al-Muqaffa (1910). History of the Patriarchs of the Coptic Church of Alexandria, Volume I. In *Patrologia Orientalis*, volume 1,5,10. Firmin-Didot Et Co., Paris. B. Evetts (Tr.).

Ibn al-Muqaffa (1943). *History of the Patriarchs of the Coptic Church of Alexandria, Volume II: Part I-III, Volume III: Part I*. Publications de la Société d'Archéologie Copte, Cairo. Yassa Abd al-Masih, Aziz Suryal Atiya, O.H.E. Burmester, Antoine Khater (Trs. and Eds.), 1943, 1948, 1959, 1968.

Jia, R. and Persson, T. (2017). Ethnicity in Children and Mixed Marriages: Theory and Evidence from China. mimeo.

John of Nikiu (1916). *The Chronicle of John, Bishop of Nikiu*. Williams and Norgate, London-Oxford. R. H. Charles (Tr.).

Mahmoud, A. A. (2009). *al-jizya fi misr 1713-1856 (The Poll Tax in Egypt in 1713-1856)*. Supreme Council of Culture, Cairo.

Maspero, J. (1912). *Organisation Militaire de l'Égypte Byzantine*. Bibliothèque de l'École des Hautes-Études, Paris.

Mikaberidze, A. (2011). *Conflict and Conquest in the Islamic World: A Historical Encyclopedia [2 volumes]: A Historical Encyclopedia*. ABC-CLIO, Santa Barbara, CA.

Mingana, A. (1931). Vision of Theophilus. In *Woodbrooke Studies*, volume 3, pages 1–92. W. Heffer and Sons Limited, Cambridge.

Ministère des finances (1898). *Recensement Général de l'Egypte: 1 juin 1897*. Imprimerie nationale, Le Caire.

Morimoto, K. (1981). *The Fiscal Administration of Egypt in the Early Islamic Period*. Dohosha, Kyoto.

Rapoport, Y. (2018). *Rural Economy and Tribal Society in Islamic Egypt: A Study of al-Nabulusi's Villages of the Fayyum*. Brepols Publishers, Turnhout.

Russell, J. C. (1958). *Late Ancient and Medieval Population*. The American Philosophical Society, Philadelphia.

Russell, J. C. (1966). The Population of Medieval Egypt. *Journal of the American Research Center in Egypt*, 5:69–82.

Saleh, M. (2013). A Pre-Colonial Population Brought to Light: Digitization of the Nineteenth-Century Egyptian Censuses. *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 46(1):5–18.

Saleh, M. (2018). On the Road to Heaven: Taxation, Conversions, and the Coptic-Muslim Socioeconomic Gap in Medieval Egypt. *Journal of Economic History*, 78(2):394–434.

Sirhan, J. (1978). *al-Musamara wal-Munadama ‘ind al-‘Arab hatta al-Qarn al-Rabi‘ al-Hijri (The Boon-Companionship in the Arab World until the 4th Century A.H.)*. PhD thesis, American University of Beirut, Beirut.

United Nations Population Division (2002). Methods for estimating adult mortality. Technical report, United Nations, New York.

Wilson, A. (2011). City Sizes and Urbanization in the Roman Empire. In Bowman, A. and Wilson, A., editors, *Settlement, Urbanization, and Population*, pages 161–195. Oxford University Press, Oxford.