Salt Wars

Desiree A. Desierto, Patrick Fitzsimmons and Mark Koyama*

July 6, 2026

Abstract

Salt was a key fiscal resource in premodern Europe. The ability of states to appropriate salt shaped patterns of fragmentation and consolidation of territories. We propose a model demonstrating that when and where salt was *non-appropriable* by a central authority, it produced more fragmented partitions. Territories with non-appropriable resources could fund rebellion and resist incorporation, which made conquest of such territories unprofitable. In equilibrium, such territories were more likely to remain independent. To test this hypothesis, we construct a new dataset of 251 historical salt sites, classify them by extraction technology, and exploit temperature variation driven by the Little Ice Age as a source of exogenous identification. Colder temperatures reduced coastal/marine salt production, raising the relative value of inland salt. Where inland salt was appropriable by a central authority, specifically rock salt, we find this temperature shock generated more consolidation. Where it was dispersed and non-appropriable, specifically salt from brine springs, it produced greater fragmentation.

JEL: N43
Keywords: Salt; State Capacity; Fiscal Capacity; Warfare; Battles; Little Ice Age

*Desiree Desierto: ddesierto@ufl.edu. Patrick Fitzsimmons: prfitz@sas.upenn.edu. Mark Koyama: mkoyama@ufl.edu. We are grateful for comments from the audience at Markets and Society. We thank Andy Ferrara for detailed comments on an earlier draft. We are particularly grateful to Anton Howes whose substack posts on salt taxation inspired us to begin this project.

1 Introduction

The growing literature studying the foundations of modern states in preindustrial Europe highlights several related stylized facts. First, historically Europe was highly politically fragmented (Jones, 2003; Abramson, 2017; Ko et al., 2018; Fernández-Villaverde et al., 2023). Second, European states were frequently at war and the continent was extraordinarily conflict prone (Hoffman, 2015). And third, that these competitive pressures led to investment in fiscal capacity (Gennaioli and Voth, 2015; Cantoni et al., 2024; Becker et al., 2025; Cervellati et al., 2026). Hundreds of polities—kingdoms, duchies, bishoprics, city-states—competed over territory in a landscape where borders shifted constantly and frequent interstate warfare had to be funded by new sources of revenue. These factors are seen by many as crucial ingredients in the Great Divergence that subsequently opened up between Europe and the rest of the world.

Nonetheless, despite agreement on these facts, many aspects of European state-building are contested. For example, while political fragmentation is associated with innovation and institutional reforms (Mokyr, 2016), it also led to frequent destructive warfare (Martines, 2013). Moreover, the effects of natural resource revenues on state capacity are not clear-cut. Scholars have long argued that the inflow of American silver eroded fiscal capacity in Iberia, for example (Elliott, 1961; Drelichman, 2005; Henriques et al., 2026).[1]

We study the effects on European borders and battles of a highly taxable resource – salt. Salt was universally consumed, costly to substitute, and hence inelastically demanded, making it a natural fiscal resource for premodern states (Adshead, 1992). One prominent line of argument associated with Tilly (1985, 1990) associates taxable resources, such as salt, with the ability of modernizing states to build fiscal capacity and expand their borders. From this perspective, salt should be associated with greater political consolidation. In contrast, another line of reasoning highlights the *appropriability* of resources by the state as a critical determinant for whether these resources are associated with state building (Scott, 2017; Mayshar et al., 2017, 2022; Huning and Wahl, 2023; Desierto and Koyama, 2026, 2025).

[1] Conversely, deliberate fiscal-administrative reform could build state capacity: Chiovelli et al. (2024) show that the eighteenth-century Bourbon intendancy reforms raised crown revenue and strengthened the state across the Spanish empire.

To disentangle the relationship between salt and political consolidation, fragmentation, and conflict, a theoretical framework is required.

We propose a general model that demonstrates how the appropriability of resources generates different-sized polities in equilibrium. Consider a set of territories with an adjacency structure. At any time period, this set is partitioned into polities, with each polity controlling one or more territories. Each territory is endowed with resources, some of which it can secure, hide, or keep away from any polity, and are thus *non-appropriable*. At each period, each territory chooses whether to rebel from their polity, support a war of annexation against a neighboring polity, or maintain the status quo, thereby determining the partition in the next period.

The model identifies two channels through which resource appropriability affects polity size and the levels of conflict in equilibrium. First, a *rebellion channel*: territories rich with non-appropriable resources have higher fighting capacity precisely because they can fund rebellions from revenues that no polity can seize. They are thus able to resist incorporation. Second, a *war deterrence channel*: polities whose territories have large non-appropriable resources are also unattractive targets for annexation, because their territories will be more likely to rebel post-conquest, thereby dissipating the gains from war. Our main result thus establishes that territories with larger non-appropriable resources produce (weakly) more fragmented partitions than those with smaller non-appropriable resources. Three regimes arise in equilibrium. When non-appropriable resources are uniformly low across territories, a single polity (an Empire) controls all territories. When non-appropriable resources are low in some territories, but high in others, a number of polities emerge, where the size of each polity is bounded by the fighting capacity of its constituent territories. Lastly, when non-appropriable resources are uniformly high, territories remain independent and therefore constitute singleton polities.

A testable prediction immediately follows: there should be greater fragmentation whenever and wherever salt is non-appropriable.

The empirical challenge, however, is that salt endowments are fixed and geographically determined, and therefore difficult to disentangle from other drivers of political structure. Our identification strategy exploits the interaction between salt endowments and climate variation driven

by the Little Ice Age. We distinguish between three types of salt extraction: solar evaporation of water from coastal/marine salt works, the boiling of liquid from brine springs, and the mining of rock salt. Coastal salt production was therefore heavily dependent on temperature, and brine-salt production only marginally so, as the wood used for brine boiling competed for fuel-use in cold periods. In contrast, rock salt mining is climate-insensitive, as it took place underground. As temperatures fell after approximately 1300—the onset of the European Little Ice Age (Pfister, 1981; Campbell, 2016), the output from coastal/marine salt works declined, shifting production towards rock salt mines and brine springs. Temperature variation thus provides exogenous, time-varying activation of the value of different types of salt.

In turn, the type of salt provides variation in the size of non-appropriable resources. In particular, rock salt is appropriable by a central authority. Rock salt mines were underground operations that extracted solid halite deposits. They were geographically fixed and capital-intensive, and easily monitored as output came out through a single chokepoint (the mine entrance). These characteristics enabled the central authority to monopolize rock salt production. In contrast, brine springs were operated by many small producers who could easily boil water over wood fires, and were geographically dispersed and therefore harder to monitor. Coastal/marine salt works had even lower barriers to entry, as anyone along a suitable coastline could evaporate seawater.

Thus, salt from rock salt mines is appropriable, while salt from brine springs and coastal waters are non-appropriable. All else equal, a territory with brine springs or coastal waters has larger non-appropriable resources, while a territory with rock salt mines has smaller. If the model holds, then as temperatures fell in the late Middle Ages, which decreased salt production from the coast and increased production from brine springs and mines, we should observe greater fragmentation in regions of territories with brine springs, but more consolidation in coastal regions and regions with rock salt mines.

We conduct our analysis using both a panel of 476 fixed grid cells (200km $\times$ 200km) and a panel of over 600 historical polities, observed at five-year intervals from 1100 to 1790. Our results are as follows. First, without considering the extraction method, salt in colder temperatures appears to increase the number of polity borders within a grid cell – it therefore makes the grid cell more

fragmented. The coefficient on the interaction of salt and temperature is sizable: a 1°C cooling raises the border count by about 0.72 for *each* salt mine in a cell.

This result runs counter to the Tilly hypothesis. Increased revenues from salt appear to decrease, rather than increase, consolidation. We then provide evidence for our own hypothesis that the effect depends on the appropriability of salt. In particular, we distinguish between types of salt extraction, and find that the effect of the interaction between brine salt and temperature on the number of borders is twice as large in magnitude and *opposite* in sign to the effect of the interaction between rock salt and temperature, and of the interaction between coastal salt and temperature. Colder periods are associated with more borders in cells with brine salt, but fewer borders in cells with rock salt or coastal salt. Consistent with our model's prediction, an increase in the value of non-appropriable resources (more brine salt in colder weather) produces greater fragmentation, while an increase in the value of appropriable resources (more rock salt or less coastal salt in colder weather) generates more consolidation.

To address the possible endogeneity of salt production, we construct several instrumental variables that exploit independent sources of geological variation. We first instrument all salt mine locations with evaporite basin deposits interacted with terrain ruggedness. Then specifically for brine salt, we use the salt-bearing rock strata underlying central European brine springs, and specifically for coastal salt, sea-surface salinity. The 2SLS regression results confirm the OLS results. Moreover, the coefficient on the interaction between brine salt and temperature is approximately twice the OLS coefficient, which is likely due to attenuation of classical measurement error, as historical brine sites are undercounted in the data.

We then test the mechanism in our model. Regions of territories with large non-appropriable resources are more fragmented because territories are likely to rebel from the polity (rebellion channel), and are also unattractive to conquering polities (war deterrence channel). Conversely, regions of territories with large appropriable resources are more consolidated since the territories are not likely to rebel, and are thus attractive to conquer.

Since rebellions pertain to individual territories, while wars of annexation (conquest) are waged by polities, we consider each channel separately. In our panel of grid cells, we find that there are

more civil war battles during colder weather in brine-salt regions, and less in rock and coastal-salt regions. Consistent with the model, the results indicate that there are more rebellions when the value of non-appropriable resources increases (more brine salt in colder weather), and less when it decreases (less coastal salt in colder weather) or when the value of appropriable resources increases (more rock salt in colder weather).

The patterns should be opposite when considering wars of annexation, since rebellious territories are unattractive to conquering polities. We confirm this using polity-level data. In particular, we find that in colder weather, polities with brine springs fight less interstate wars, while polities with rock salt mines and coastal salt works fight more.

We contribute to several important literatures. First, we build on existing work on state size and borders. Within economics, the literature on the size of nations was pioneered by Alesina et al. (1999); Alesina and Spolaore (2003, 2005) and Bolton and Roland (1997). Relatedly, Desmet et al. (2025) examine factors that make secession, focusing ethnolinguistic differences. Several papers examine the determinants and consequences of historical borders. Kitamura and Lagerlof (2020) examine borders from 1500 to the present in Europe and the Middle East at the grid-cell level. They find geographic obstacles help to predict borders and that border density is positively associated with economic development today, effects that they attribute to the positive impact of interstate competition. Depetris-Chauvín and Özak (2024) relate borders directly to conflict, showing that violence in contemporary Africa concentrates along *de facto* historical ethnic boundaries. Acharya and Lee (2018) are interested in the emergence of territorial states in the Middle Ages. They develop a model in which economic development generates rents that lead to the formation of territorial states. Ko et al. (2018) examine the question of why Europe is more fragmented than China. Fernández-Villaverde et al. (2023) study the process of state and border formation globally using a simple conflict-based model of state expansion and contraction. Schönholzer and Weese (2026) examine how border changes affected city growth in premodern Europe, identifying a channel whereby rulers in more competitive environments improve institutions in response to a competitive threat. Our contribution is to show that a specific fiscal resource—salt—and the extent to which it is appropriable by the state, shaped both the geography of borders and the

nature of conflict.

Second, this paper contributes to a body of work on war and the rise of fiscal states in medieval and early modern Europe (Gennaioli and Voth, 2015; Cantoni et al., 2024; Chambru et al., 2024). The Tilly hypothesis—that war made the state, and the state made war (Tilly, 1985)—has been widely discussed but is difficult to test rigorously because warfare and state capacity are jointly determined. Dincecco and Onorato (2017) links European battles to urban growth. Becker, Ferrara, Melander, and Pascali (2025) show that wars in early modern Germany caused greater fiscal and state capacity as well as the representation of craft guilds.

A parallel literature on the origins of state capacity in Europe focuses on the persistence of political fragmentation in Europe. While Tilly (1990) associated the growth of more powerful states with a *reduction* in political fragmentation, the empirical work of Abramson (2017) establishes that in fact many smaller polities persisted until 1800.[2] Europe's persistent political fragmentation is widely seen as a factor in both its high levels of inter-state conflict (Hoffman, 2015; Ko et al., 2018) but also its high levels of innovation and prosperity (e.g. Mokyr, 2016; Fernández-Villaverde et al., 2023). Our analysis disrupts the Tilly paradigm by providing a key insight — when an exogenous fiscal resource (salt) is non-appropriable by the state, it increases civil/internal conflict (rebellion) and political fragmentation; if it is appropriable, it increases interstate conflict (wars of annexation) and consolidation. This finding, then, can reconcile the traditional bellicist with the more recent anti-bellicist findings in political economy. Salt revenues fund inter-state wars to make or enlarge the state, but only if they are appropriable; otherwise, non-appropriable salt revenues actually deter war and prevent state expansion, although at the expense of fueling internal conflict. Either way, salt revenues funded armies on all sides, creating a wasteful tournament that shaped Europe's polities (Hoffman, 2015). The geography of our results reinforces this reading. The brine-driven fragmentation effect is strongest in central Europe, within the lands of the Holy Roman Empire. The Empire is the standard counterexample to the bellicist thesis—a region of endemic warfare that nonetheless remained divided among hundreds of small polities into the nineteenth century (Whaley, 2012; Wilson, 2009; Abramson, 2017)—and it is where our mechanism

[2]Numerous recent works in historical political economy find evidence for and against Tilly's hypothesis. See *inter alia* Kaspersen et al. (2017), Gorski and Sharma (2017), Spruyt (2017) and Cederman et al. (2023).

should bite hardest: dispersed, non-appropriable salt wealth sustained the fiscal independence of small territories where Tilly's logic predicts they should have been absorbed (see Section 6 and Appendix Section J).

Finally, we are not the first authors to study salt. There have been several recent papers on salt taxes in ancien régime France. Giommoni and Loumeau (2026) show that the *gabelle*—the French salt tax—had persistent negative effects on local economic development, while Giommoni et al. (2026) provide systematic evidence that extractive indirect taxation, including the salt tax, was an important driver of revolutionary unrest in late eighteenth-century France. These studies focus on the consequences of salt taxation within a single country. Our focus is different, as we consider the role played by salt in the development of fiscal states across Europe as a whole. By assembling a comprehensive dataset of 251 salt production sites, including extraction methods, and exploiting the interaction between salt endowments and climate, we provide the first systematic evidence that salt resources shaped the political and military geography of medieval and early modern Europe.

The remainder of this paper is organized as follows. Section 2 discusses the role of salt and the Little Ice Age in European history. Section 3 presents a model of territorial consolidation and fragmentation based on resource appropriability. Section 4 describes our data sources, including the classification of salt sites by extraction technology. Section 5 presents the empirical analysis, Section 6 reports robustness checks, and Section 7 tests the conflict mechanism. Section 8 concludes.

2 Salt and State-Building

2.1 Salt in the Premodern Economy

Salt was critical for food preservation—meat, fish, butter, and cheese could not be stored for extended periods without salting—and for livestock feed, leather tanning, and textile production (Adshead, 1992; Kurlansky, 2002). Before the advent of modern refrigeration, salt was an indispensable household good with no close substitutes, and demand was correspondingly inelastic. Governments recognized this early: salt taxes were a feature of fiscal systems from Roman times

onwards.[3]

There were three principal methods of salt production in premodern Europe. First, *solar evaporation* of seawater in shallow coastal ponds was the most common technique around the Mediterranean and along the Atlantic coast (Multhauf, 1976). As M. R. Bloch noted, "the production of salt from concentrated brine is much easier than quarrying or mining. The method consists essentially of bringing the natural brine into shallow ponds enclosed by low earth walls and allowing the sun to evaporate the water. The deposited salt layers are then harvested." This technique was cheap but required "abundant sunshine and low atmospheric humidity," making it highly sensitive to climate (Bloch, 1976, 339).

Second, *brine springs*—natural outcrops of salt-laden water—could be boiled to produce salt. Moreover by creating a succession of ponds controlled by pumps and sluices, it was possible to increase salinity. The Mediterranean brines might rely on solar evaporation but in central and northern Europe it was common to heat brines artificially. This method was common in central Europe (e.g., the Salzkammergut in Austria, Halle in Saxony) and in Cheshire, England, and was less directly climate-dependent but more fuel-intensive.

Third, *rock salt mining* extracted solid halite deposits formed by the evaporation of ancient seas millions of years ago. Major mines included Wieliczka and Bochnia in Poland, Hallstatt in Austria, and several sites in Transylvania. Rock salt mining was capital-intensive but produced large volumes and was independent of contemporary climate.

There are several historical examples of salt mines funding state building. The rich salt deposits of Transylvania provided the fiscal foundations for the medieval Kingdom of Hungary (Draskóczy, 2018; Weisz, 2018). In early modern Poland, for example, the income from salt mines provided up to a third of royal revenue (Hanik, 1988, 8).

Salt taxation also funded war. There are many examples of this. In the Second Punic War, Rome imposed a tax on salt to pay for the conflict (Multhauf, 1976, 12). It was during the Hundred Years War that the French monarchy fully established and eventually made permanent, its various taxes on salt (Sands and Higby (1949), Multhauf (1976, 13)). Medieval Venice had a

[3]While we focus on the role of salt taxes in premodern Europe, they were of critical importance in Chinese and Japanese history too.

monopoly on local salt and even waged wars "to eliminate the competition from other northwest Adriatic cities . . . so as to seize or destroy their saltworks" (Laszlo, 1998, 66). Alfonso V of Aragon utilized royal revenues from state monopolies to fund his imperial ambitions in Southern Italy, with the salt monopoly being the most lucrative of these sources. Alfonso imposed strict controls over salt production, created direct taxes, and required households to purchase salt from the royal stocks (Aloisio, 2010). In areas of strong state control, like the Habsburg salt production sites, salt sources were liquid enough that they could be mortgaged and transferred between rulers to provide military relief and support (Adshead, 1992, 257). So important were these sources, that salt-producing areas were prone to conflict. The Este family fought for decades to control areas rich in salt and fish, with farms of these resources becoming the largest source of revenue for the family (Dean, 2002). In the French Wars of Religion, Huguenots repeatedly took control of salt producing regions, understanding that they represented a large portion of royal revenue (Hughes, 1925).

The French *gabelle* is the most extensive salt monopoly in European history, but the gabelle revenues could only be extracted through an enormous enforcement apparatus, as there was persistent regional variation and endemic smuggling. We discuss the gabelle in detail in Appendix 2.

2.2 The Little Ice Age as a Differential Shock

The Little Ice Age refers to a period of cooling that affected Europe from approximately 1300 to 1850, with particularly severe cold episodes in the 15th, 17th, and early 19th centuries (Pfister, 1981; Lamb, 1982; Campbell, 2016; Waldinger, 2022). Average temperatures during the coldest decades were approximately 1–2°C below the medieval warm period.[4]

The Little Ice Age had direct consequences for salt production. The physics of evaporation makes this relationship mechanical: the standard Penman equation used by hydrologists to model evaporation from open water shows that evaporation rate depends on the saturation vapor pressure e_s, which is an exponential function of temperature (Penman, 1948; Akridge, 2008). A

[4]Note there has been some skepticism about the Little Ice Age, see Kelly and Gráda (2014). In particular, the term may be misleading if it is taken to mean that temperatures were uniformly lower throughout this period. But lower, more volatile and wetter climate is attested to across multiple qualitative and quantitative sources. See Appendix B.

1–2°C reduction in growing-season temperature—the magnitude of the Little Ice Age—reduces the saturation vapor pressure and hence the rate at which brine concentrates in solar evaporation pans. Akridge (2008) validates this relationship quantitatively: numerical simulations of solar saltmaking, calibrated against ethnographic observations at a Mexican coastal saltworks, show that evaporation rates of 4–5 mm/day at 27°C would decline substantially with lower temperatures, longer cloud cover, and shorter production seasons. For brine evaporation by boiling, the direct climate effect is smaller—heat comes from fuel, not the sun—but cold periods increase competition for wood, the primary fuel source. Beck (1993) documents severe deforestation around the southwestern German salines from their fuel consumption, a problem that would have been exacerbated during cold periods when wood was also needed for heating.

The historical record confirms these effects. Coastal salt works that relied on solar evaporation became less productive as temperatures fell, cloud cover increased, and growing seasons shortened.[5] The great salt pans of the Bay of Bourgneuf, which had supplied much of northern Europe in the 14th century, declined markedly during the 15th century. Mediterranean salt works were less affected but still experienced reduced yields. As coastal/marine sources of salt became scarcer and more expensive, the relative value of inland mineral salt increased.

3 A Model

We construct a parsimonious model of how polities evolve. Polities grow as territories conquer other territories through wars of annexation; they shrink as territories secede through rebellion. The main factor that drives rebellion and conquest is not the size of resources per se, but the extent to which such resources are appropriable. A territory whose resources are difficult to appropriate by another territory – whose resources are largely *non-appropriable*, is likely to remain independent. It has more resources that it can keep for itself. Thus, it has greater (fighting) capacity to rebel against any polity, and is less strategically valuable to the polity. In contrast, a territory with more appropriable resources is less likely to rebel once conquered, and is therefore more likely to

[5]Kurlansky (2002) reports that even during the medieval warm period, Viking attempts to use solar evaporation failed: "Traces of such Viking operations in the seventh century have been found in Normandy. But the northern climate would have made these saltworks unproductive. The climate has too much rainfall and not enough sunlight."

be conquered. In equilibrium, territories that have sufficiently small non-appropriable resources tend to consolidate into a single polity, while territories with sufficiently large non-appropriable resources tend to fragment into independent realms.

The main results in 3.4 apply to *any* kind of resource, including salt. Additional results specific to salt are in 3.5.

3.1 Set up

There is a set $\mathcal{N}$ of N territories, with a generic territory indexed by i. At any time period t, each territory $i \in \mathcal{N}$ has resource endowment $e_i \in \mathbb{R}_+$, of which $n_i \in [0, e_i]$ are non-appropriable by any other territory.[6]

Also, at any t, the territories in $\mathcal{N}$ are partitioned into a set $\mathcal{P}_t$ of P_t polities.[7] Each polity $j \in \mathcal{P}_t$ is identified with the subset of territories it controls, so that $i \in j$ denotes that territory i belongs to polity j, with $|j|$ the number of such territories.[8] Polity j has total resources $E_j = \sum_{i \in j} e_i$, which is the sum of the resources of the territories it controls. Each territory $i \in j$ obtains a share of E_j — for simplicity, we assume equal sharing, which gives each territory $\frac{E_j}{|j|}$.

The partition $\mathcal{P}_t$ is endogenous and evolves over time: territories may gain independence through successful rebellion or be annexed through conquest, altering both the composition of polities and the total number P_t.

In particular, a territory $i \in j$ may rebel against j: if successful, it secedes from j and forms its own singleton polity $\{i\}$, retaining e_i for itself but forfeiting its share $\frac{E_j}{|j|}$ of j's resources. The outcome of the rebellion, however, depends on the relative fighting capacities of $i \in j$ and j. Specifically, $i \in j$'s rebellion is successful if $f_{i \in j} + \delta \geq F_j$, where $f_{\in j}$ denotes $i \in j$'s fighting capacity, F_j polity j's, and $\delta \in \mathcal{R}$ is a random variable with a density function that exists everywhere, and is differentiable, single-peaked, and symmetric around zero. Thus the probability that $i \in j$ successfully rebels from j is $\rho_{i \in j} = \Pr[\delta \geq (F_j - f_{i \in j})]$.

Fighting capacities are the amount of resources that the fighting party can muster. A rebel

[6] Resources simply replenish at each time period, to capture a Malthusian world of zero economic growth. Thus, e_i, n_i are fixed across time periods.

[7] Formally, $\{j\}_{j \in \mathcal{P}_t}$ constitutes a partition of $\mathcal{N}$: $\bigcup_{j \in \mathcal{P}_t} j = \mathcal{N}$ and $j \cap j' = \emptyset$ for all $j \neq j' \in \mathcal{P}_t$.

[8] Since $\mathcal{P}_t$ partitions $\mathcal{N}$, we have $\sum_{j \in \mathcal{P}_t} |j| = N$.

territory $i \in j$ can only use what the polity cannot appropriate, so $f_i = n_i$. During rebellion, i effectively withdraws n_i from j's resources. Let $\mathcal{R}_j \subseteq j$ denote the set of territories in j currently in rebellion; then polity j's fighting capacity is $F_j = E_j - \sum_{i \in \mathcal{R}_j} n_i$, which shrinks as more territories rebel simultaneously.

A polity can also expand by conquering a neighbor. Let $\mathcal{M}_t = \{\{j, -j\}\}$ denote a maximal matching on $\mathcal{P}_t$, where each matched pair $\{j, -j\}$ represents two polities that may contest annexation.[9] As $\mathcal{P}_t$ evolves through rebellions and annexations, the induced matching $\mathcal{M}_t$ updates accordingly. Polities not included in any matched pair — which arise whenever the matching is not perfect — have no annexation opportunity that period. The probability that j conquers $-j$ depends on their relative fighting capacities: $\omega_j = \Pr[\gamma \leq (F_j - F_{-j})]$, where γ is a random variable with a density function that exists everywhere, and is differentiable, single-peaked, and symmetric around zero. Since either j or $-j$ wins, the victor annexes all territories of the losing polity that have not successfully rebelled, along with their resources.

3.2 Sequence of Events

At $t = 0$, the initial partition $\mathcal{P}_0$ is any partition of $\mathcal{N}$. The initial matching $\mathcal{M}_0$ is the maximal matching on $\mathcal{P}_0$ induced by the adjacency structure, which may or may not be perfect depending on $\mathcal{P}_0$ and the adjacency structure.

At each time period thereafter, $t = 1, 2, ...\infty$, the following sequence of events occurs.

1. *Decisions.* Let $j \in \mathcal{P}_t$ denote the polity to which territory i belongs at the start of period t. Each territory $i \in j$ simultaneously chooses whether to rebel against j. Non-rebel territories in j then vote on whether to wage a war of annexation against $-j$; a declaration of war is made if all non-rebel territories in j vote in favor.[10]

[9] The matching is induced by a fixed adjacency structure on $\mathcal{N}$: two polities are matchable if they contain adjacent territories. A natural example is a line, where territories are ordered $1, 2, \ldots, N$ and adjacency means consecutive positions; polities then occupy contiguous intervals and each has at most two neighbors. A grid or hexagonal tiling are alternative examples. Nothing in the model depends on the specific structure.

[10] The unanimity rule assumption is inconsequential: since resources are shared equally, all non-rebel territories in j face identical payoffs and therefore vote identically. Hence, other aggregation rules, e.g. majority rule, can be adopted without altering the results.

2. *Outcomes.* The fighting capacities of territories and polities are determined jointly, and rebellion and war outcomes are realized simultaneously. Each rebel territory $i \in j$ succeeds with probability ρ_{ij} and forms singleton polity $\{i\}$; a failed rebel remains in j. If either j or $-j$ (or both) declared war, annexation is contested: j prevails with probability ω_j and $-j$ with probability $1 - \omega_j$.

3. *Updating.* The new partition $\mathcal{P}_{t+1}$ and matching $\mathcal{M}_{t+1}$ are formed from the outcomes of rebellions and annexations, as described in the setup.

4. *Payoffs.* Each territory obtains a share of its (new) polity's resources: e_i if it is a successful rebel or a singleton, $\frac{E_{j'}}{|j'|}$ otherwise, where $j' \in \mathcal{P}_{t+1}$ denotes the polity it belongs to after updating.

3.3 The Game

3.3.1 Players, Actions, and Strategies

The players are the territories $i \in \mathcal{N}$. At each period t, the state of the game is the partition $\mathcal{P}_t$ (with matching $\mathcal{M}_t$ induced by the adjacency structure). Each territory $i \in j \in \mathcal{P}_t$ chooses an action a_i^t from a state-dependent action space:

$$a_i^t \in \begin{cases} \{R, W, S\} & \text{if } j \in \mathcal{M}_t, \\ \{R, S\} & \text{if } j \notin \mathcal{M}_t, \end{cases}$$

where R (rebel) denotes declaring rebellion against j, W (war) denotes remaining in j and voting for annexation of $-j$, and S (stay) denotes remaining in j and voting against annexation (or simply remaining when j is unmatched). The action W is available only when j is matched; an unmatched endpoint polity has no opponent to contest, so its territories choose only between rebelling or staying. A war against $-j$ is declared if all non-rebel territories in j choose W. Let $\mathbf{a}^t = (a_i^t)_{i \in \mathcal{N}}$ denote the action profile at t.

A strategy for territory i is a mapping $\sigma_i : \mathcal{P}_t \mapsto a_i^t$ from the current state to an action. A strategy profile is $\boldsymbol{\sigma} = (\sigma_i)_{i \in \mathcal{N}}$.

3.3.2 Payoffs

Let $\beta \in (0,1)$ be a common discount factor. Territory i's per-period payoff is

$$u_i(\mathbf{a}, \mathcal{P}) = \begin{cases} e_i & \text{if } i \text{ successfully rebels or is a singleton,} \\ \dfrac{E_{j'}}{|j'|} & \text{otherwise,} \end{cases}$$

where $j' \in \mathcal{P}_{t+1}$ is the polity i belongs to after updating.

3.4 Equilibrium

Our solution concept is Markov Perfect Equilibrium (MPE). Let

$$V_i(\mathcal{P}; \boldsymbol{\sigma}) = \mathbb{E}\left[\sum_{s=0}^{\infty} \beta^s \, u_i(\boldsymbol{\sigma}(\mathcal{P}_s), \mathcal{P}_s)\right]$$

denote territory i's expected continuation payoff from current state $\mathcal{P}$, where $\mathcal{P}_s$ denotes the state s periods hence, under strategy profile $\boldsymbol{\sigma}$. A *Markov Perfect Equilibrium* (MPE) is a strategy profile $\boldsymbol{\sigma}^* = (\sigma_i^*)_{i \in \mathcal{N}}$ such that, for every territory i, every period t, and every state $\mathcal{P}_t$,

$$V_i(\mathcal{P}_t; \boldsymbol{\sigma}^*) \geq V_i(\mathcal{P}_t; \sigma_i, \boldsymbol{\sigma}^*_{-i})$$

for all strategies σ_i.

Since the game is time-homogeneous — payoffs and transition probabilities depend on t only through $\mathcal{P}_t$, and e_i, n_i, and β are all fixed across periods — we restrict attention to *stationary* MPE, in which strategies do not depend on t directly. Formally, a stationary MPE is a strategy profile $\boldsymbol{\sigma}^*$ in which each $\sigma_i^* : \mathcal{P} \mapsto a_i$ is a time-invariant mapping from states to actions, and $\boldsymbol{\sigma}^*$ constitutes a MPE. In a stationary MPE, territory i's value function satisfies the Bellman equation

$$V_i(\mathcal{P}; \boldsymbol{\sigma}^*) = \mathbb{E}\left[u_i(\mathcal{P}') + \beta V_i(\mathcal{P}'; \boldsymbol{\sigma}^*)\right],$$

where $\mathcal{P}'$ is the partition resulting from the equilibrium actions $\boldsymbol{\sigma}^*(\mathcal{P})$, and the expectation is taken over the stochastic outcomes of rebellion and annexation.

Since the state space (all partitions of $\mathcal{N}$) and the action space $\{R, W, S\}$ are both finite and $\beta \in (0,1)$, existence of a stationary MPE in mixed strategies is guaranteed by standard results for finite discounted stochastic games (Fink, 1964). We focus on stationary MPE in pure strategies.

In particular, we characterize such equilibria by the partition outcomes they produce. A partition outcome is either non-absorbing or absorbing.

Definition 1. *A partition $\mathcal{P}^*$ is absorbing under strategy profile $\boldsymbol{\sigma}^*$ if $\mathcal{P}_{t+k} = \mathcal{P}^*$ with probability one for all $k \geq 0$ whenever $\mathcal{P}_t = \mathcal{P}^*$.*

Since e_i and n_i are fixed parameters and the partition is the only state variable, $\mathcal{P}^*$ is absorbing if and only if equilibrium play at $\mathcal{P}^*$ prescribes no rebellion (NR) and no war (NW): (i) (NR condition) $\sigma_i^*(\mathcal{P}^*) \neq R$ for all $i \in j \in \mathcal{P}^*$, and (ii) (NW condition) $\sigma_i^*(\mathcal{P}^*) = S$ for all $i \in j \in \mathcal{P}^*$. A partition is non-absorbing when either condition fails at that partition: equilibrium play prescribes R for some territory (so the partition breaks apart), or W for some polity (so the partition consolidates).

In turn, whether or not the conditions hold depends on the size of non-appropriable resources:

Lemma 1.

1. (No-rebellion.) *Suppose $e_i + \beta V_i(\mathcal{P}^*_{-i}; \boldsymbol{\sigma}^*) > \frac{E_j/|j|}{1-\beta}$, so that rebellion is profitable if feasible. There exists a threshold $\bar{n}_i \geq 0$ such that condition* (NR) *is satisfied if and only if $n_i \leq \bar{n}_i$.*

2. (No-war.) *$V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is non-increasing in n_k for $k \in -j$. Consequently, there exists a threshold $\bar{N}_{-j} \geq 0$ on the aggregate non-appropriable resources of $-j$ such that condition* (NW) *is satisfied whenever $\sum_{k \in -j} n_k \geq \bar{N}_{-j}$.*

A territory is more emboldened to rebel the more non-appropriable resources it has since it has greater fighting capacity against the polity to which it belongs. At the same time, a polity is less likely to want to annex its neighbor the greater the non-appropriable resources of that neighbor, since the latter would be prone to rebellion and therefore less valuable to the polity.

The foregoing implies that when individual territories have large non-appropriable resources, one would expect more rebellions and fewer annexations, and therefore greater fragmentation. Equivalently, when individual territories have large appropriable resources, one would expect less rebellions and more annexations, and therefore greater consolidation.[11]

[11] By Lemma 1, the NR condition is satisfied if and only if $n_i \leq \bar{n}_i$, which is equivalent to $(e_i - n_i) \geq (e_i - \bar{n}_i)$.

To formally establish this, we first define what it means for one partition to be more fragmented than another.

Definition 2. *Partition $\mathcal{P}$ is a refinement of partition $\mathcal{P}'$, written $\mathcal{P} \preceq \mathcal{P}'$, if every $j \in \mathcal{P}$ is a subset of some $j' \in \mathcal{P}'$. We say $\mathcal{P}$ is strictly more fragmented than $\mathcal{P}'$ if $\mathcal{P} \preceq \mathcal{P}'$ and $\mathcal{P} \neq \mathcal{P}'$.*

The refinement order gives a partial order on partitions. We name three reference points on this order:

- *Independent Kingdoms*: the finest partition $\{\{i\} : i \in \mathcal{N}\}$, in which every territory is its own sovereign polity.
- *Empire*: the coarsest partition $\{\mathcal{N}\}$, in which all territories are unified under a single polity.
- *Principalities*: any partition strictly between these extremes, $\{\{i\} : i \in \mathcal{N}\} \prec \mathcal{P} \prec \{\mathcal{N}\}$, in which multiple polities of intermediate size coexist.

A cardinal measure of fragmentation is the number of polities $|\mathcal{P}|$: Independent Kingdoms maximizes it at $|\mathcal{P}| = N$, Empire minimizes it at $|\mathcal{P}| = 1$, and Principalities generate $1 < |\mathcal{P}| < N$.

The main results are the following.

Theorem 1. *Let $\mathbf{n} = (n_i)_{i\in\mathcal{N}}$ and $\mathbf{n}' = (n'_i)_{i\in\mathcal{N}}$ be two vectors of non-appropriable resources. If $n'_i \geq n_i$ for all $i \in \mathcal{N}$, with strict inequality for at least one i, then any absorbing partition under $\mathbf{n}'$ is weakly more fragmented than any absorbing partition under $\mathbf{n}$.*

Corollary 1. *There exist thresholds $\underline{\mathbf{n}} \leq \bar{\mathbf{n}}$ such that absorbing partitions become weakly more fragmented as $\mathbf{n}$ increases: specifically,*

1. *If $\mathbf{n} \leq \underline{\mathbf{n}}$, the unique absorbing partition is Empire.*
2. *If $\underline{\mathbf{n}} < \mathbf{n} < \bar{\mathbf{n}}$, the absorbing partitions are Principalities (strictly between Empire and IK).*

Thus, the larger i's appropriable resources $(e_i - n_i)$, the more likely that the NR condition is satisfied, i.e. i does not rebel. Similarly, the NW condition is satisfied if and only if $\sum_{k\in-j} n_k \geq \bar{N}_{-j}$, which is equivalent to $\sum_{k\in-j}(e_k - n_k) \leq \sum_{k\in-j} e_k - \bar{N}_{-j}$. Thus, the larger $\sum_{k\in-j}(e_k - n_k)$ is, the less likely that the NW condition is satisfied, which makes war of annexation more likely. Thus large appropriable resources, which discourage rebellion but encourage wars of annexation, lead to greater consolidation.

3. If $\bar{\mathbf{n}} \leq \mathbf{n}$, the unique absorbing partition is Independent Kingdoms.

When each territory has sufficiently large non-appropriable resources, $\bar{\mathbf{n}} \leq \mathbf{n}$, each remains independent. When each territory has sufficiently small non-appropriable resources, $\mathbf{n} \leq \underline{\mathbf{n}}$, they all combine into a single polity. The heterogeneity of $\mathbf{n}$ across territories is what generates Principalities as a stable intermediate outcome: territories with $n_i > \bar{n}_i$ successfully rebel out of any polity that absorbs them, while those with $n_i \leq \bar{n}_i$ lack the capacity to do so and remain, producing polities that are bounded in size by the rebellion capacity of their constituent territories.

The three named partitions – Independent Kingdoms, Empire, Principalities, exhaust all cases of absorbing partitions. The larger $\mathbf{n}$ is, the more fragmented the (absorbing) partition. Any other equilibrium outcome is non-absorbing. Figure 1 summarizes.

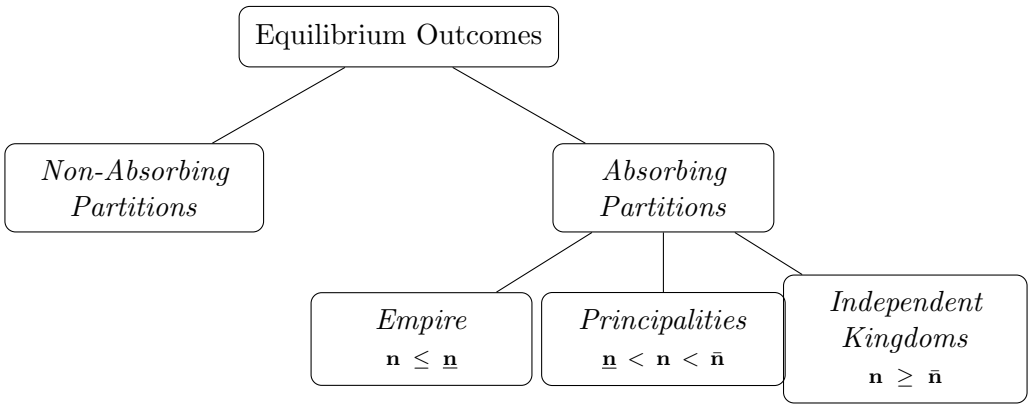

Figure 1: Types of equilibrium outcomes. Larger non-appropriable resources $\mathbf{n}$ generate greater fragmentation among absorbing partitions.

3.5 An Application to Salt

In the model, e_i includes all the resources that a territory has, some of which (n_i) are non-appropriable, while others ($e_i - n_i$) are appropriable. For some territories, e_i may indeed include salt, and the issue is whether or not it is non-appropriable. If it is, then it contributes to the size of n_i which, by Lemma 1, should make it more likely for the territory to rebel from its polity and less likely to be annexed by another. If most other territories have such salt, then by Theorem 1 and Corollary 1, the equilibrium partition should be more fragmented. However, if salt is an

appropriable resource, then the opposite pattern should hold – territories with salt are more likely to stay in their polity, and therefore more likely to be annexed by other polities. The equilibrium partition is more consolidated.

The appropriability of salt depends on extraction technology. Mining rock salt is capital intensive, and rock salt mines have a single entrance, which allows effective monitoring. A polity can therefore more easily monopolize or appropriate rock salt from any territory under its control. In contrast, evaporating liquid from brine springs or coastal seawater is a much less capital intensive way to produce salt, and thus poses low barriers to entry. Brine springs and the coast are also geographically dispersed, making it difficult for any central authority (polity) to monitor them. A territory with brine springs or near the coast has access to salt that is non-appropriable by the polity.

Empirically, salt production is potentially correlated with many other drivers of political structure. For identification, therefore, we need an exogenous shift in its relative value. Exploiting the temperature shock associated with the Little Ice Age, moreover provides a source of exogenous variation in the value of different types of salt.

A negative temperature shock has two effects on salt, as documented in Section 2. Lower temperatures reduce solar evaporative yield – by significantly decreasing salt production from all coastal salt works, it makes salt more scarce. This, then, raises its value, and triggers greater salt production from alternative sources, i.e. from brine springs and rock salt mines. In equilibrium, negative temperature shocks increase the value of non-appropriable resources of territories with brine springs, and increase the value of appropriable resources of territories with rock salt mines.

To formally demonstrate, let any salt producer be of type $i = \{R, B, C\}$, where R denotes rock-salt producer, B a brine-spring salt producer, and C a coastal-salt producer. In the short run, capital stock is fixed. Write the cost function of firm-type i as

$$\kappa_i(q;\theta) = F_i + VC_i(q;\theta),$$

where F_i is the fixed cost of the sunk capital stock, VC_i is variable cost (strictly convex in production quantity q), and $\theta \geq 0$ indexes the severity of a cold shock. Because rock-salt mining is capital-intensive, $F_R \gg F_B \approx F_C$; in particular, R's average cost (AC) only reaches its minimum

at high output – the fixed cost is spread over many units before declining noticeably, while B's and C's average costs reach their minima at low output.

Cold raises the cost of solar evaporation, which means it enters C's production intensively, but B's only marginally. In contrast, R uses no solar energy. We can thus make the following assumption on the marginal cost (MC) of i:

Assumption 1

$$\frac{\partial MC_C}{\partial \theta} \gg \frac{\partial MC_B}{\partial \theta} > 0 = \frac{\partial MC_R}{\partial \theta}.$$

Given the price P of salt, the first-order condition $P = MC_i(q_i^*; \theta)$ obtains i's equilibrium supply level $q_i^*(P; \theta)$. Aggregate supply is $Q^S(P; \theta) = r\, q_R^*(P) + b\, q_B^*(P; \theta) + c\, q_C^*(P; \theta)$, where r is the number of rock-salt producers, b the number of brine-salt producers, and c the number of coastal-salt producers. By the implicit function theorem, each producer's supply is non-increasing in the severity of the cold shock:

$$\frac{\partial q_i^*}{\partial \theta} = -\frac{\partial MC_i/\partial \theta}{\partial MC_i/\partial q} \leq 0,$$

with strict inequality for $i \in \{B, C\}$. Thus, $\frac{\partial Q^S}{\partial \theta} < 0$.

Denote market demand as $D(P)$, with $D' < 0$. Equilibrium price $P^*(\theta)$ therefore solves $D(P^*) = Q^S(P^*; \theta)$. One can then differentiate $D(P^*) = Q^S(P^*; \theta)$ with respect to θ to get:

$$\frac{dP^*}{d\theta} = \frac{\partial Q^S/\partial \theta}{D'(P^*) - \partial Q^S/\partial P} > 0.$$

Thus, a cold shock raises the price of salt. Assume that the shock is large enough such that the increase in price outweighs the marginal cost increase of B.[12] That is:

Assumption 2

$$q_B^* \frac{dP^*}{d\theta} > \frac{\partial C_B}{\partial \theta}.$$

Then a cold shock increases profits of rock salt and brine salt producers, but decreases profits of coastal salt producers:

[12]This is actually a natural assumption to make, since, by Assumption 1, $\partial C_B/\partial \theta$ is small, and a substantial contraction of coastal supply drives P^* up by a lot.

Proposition 1.

$$\frac{d\pi_R}{d\theta} > 0, \qquad \frac{d\pi_B}{d\theta} > 0, \qquad \frac{d\pi_C}{d\theta} < 0.$$

Figure 2 illustrates the mechanism.

To summarize, a cold shock contracts the supply of coastal salt, increases the market price of salt, and thereby *raises* the profit—and hence the resource value—of brine and rock producers while *lowering* that of coastal producers. Because brine and coastal salt are non-appropriable while rock salt is appropriable, the shock raises the non-appropriable resources of territories with brine springs, lowers them in territories with coastal salt works, and raises the appropriable resources of territories with rock salt mines. By Theorem 1 and Corollary 1, the following prediction is immediate.

Prediction 1. *All else equal, when the temperature drops, there should be greater fragmentation among territories with brine springs, and more consolidation among territories with rock salt mines and coastal salt works.*

4 Data

We conduct the main analysis at the grid-cell level. We construct a grid-cell panel that overlays 476 200km × 200km cells in Europe, observed at five-year intervals from 1100 to 1790 (135 periods). For each cell-period, we also observe: the number of distinct polities passing through the cell from Abramson 2017, 435 battles (1302–1700) from Kitamura 2021, and the mean area of polities overlapping the cell. We also construct a polity-level panel that tracks 631 unique polities across 30,712 polity-year observations.

We compile a new dataset of 251 historical salt production sites across Europe, recording each site's geographic coordinates, extraction technology, and period of activity. Our primary sources are Adshead (1992), Multhauf (1978), and Beck (1993). Of the 476 grid cells, 70 contain at least one salt production site. We classify each site by extraction type: rock salt (underground mining through a single chokepoint, N=27), brine evaporation (dispersed inland springs, boiled over wood or coal fires, N=75), and coastal/marine evaporation (coastal ponds requiring warmth and sunshine, N=91). The majority of salt was *non-appropriable*: the dispersed brine and coastal/marine

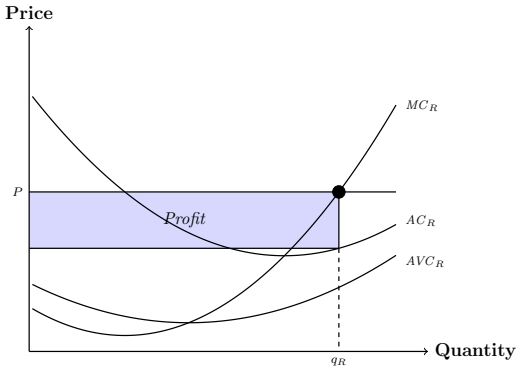

(a) Rock: pre-shock

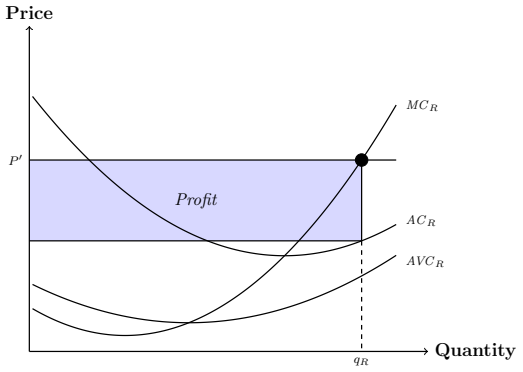

(b) Rock: at P'

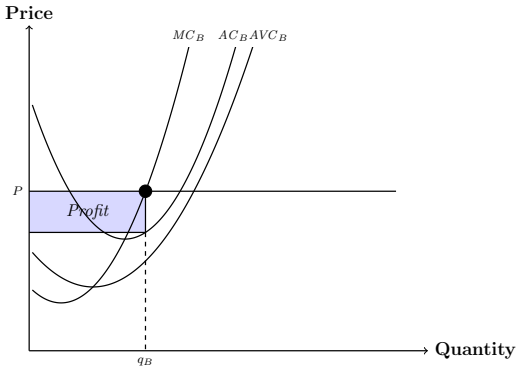

(c) Brine: pre-shock

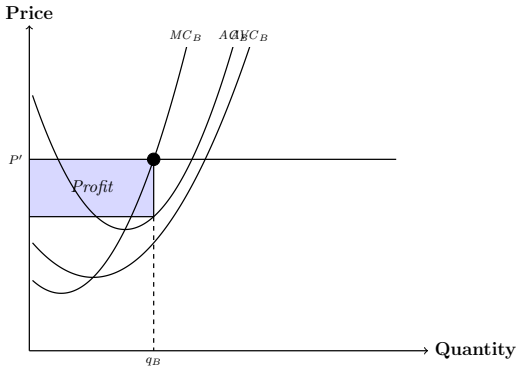

(d) Brine: at P'

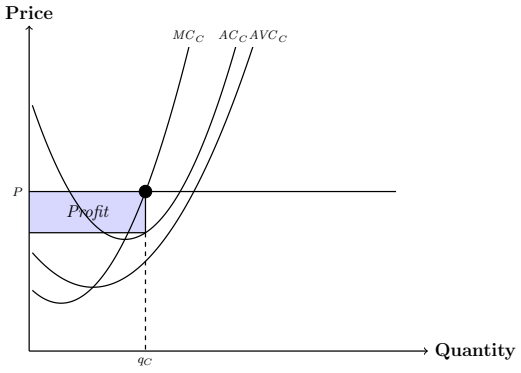

(e) Coastal: pre-shock

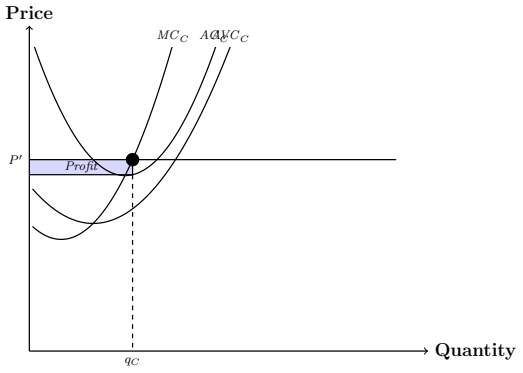

(f) Coastal: at P'

Figure 2: The cold shock and the value of salt production, by type. Each panel plots marginal cost (MC), average cost (AC), and average variable cost (AVC) for a representative firm against the market price; the shaded rectangle is profit, $(P - AC) \times q$. Each row is one salt type, shown before the cold shock (left) and at the new equilibrium price P' (right). Coastal production is climate-sensitive: the shock raises its costs, so coastal output and its profit rectangle shrink (panels e and f). Costs of rock-salt production are unchanged, while those from brine increase only slightly, so at the higher price P' their profit rectangles grow (panels a–d). The shock therefore raises the value of brine and rock salt production and lowers that of coastal.

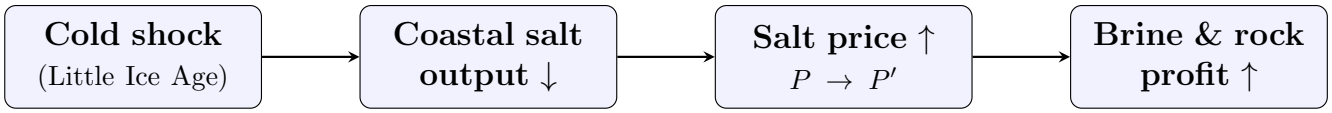

Figure 3: The logic of the empirical argument. A cold shock cuts climate-sensitive coastal salt output, raising the price of salt; with costs of rock salt production unchanged, and of brine only slightly higher, brine and rock salt producers earn higher profits.

types together account for 166 of the 193 classified sites (86%), leaving rock salt—the highly appropriable type—as the 14% exception. The same imbalance holds among salt cells: of the 70 grid cells with salt, 26 contain brine and 28 coastal/marine sites, but only 10 contain rock.

There are 61 unclassified sites due to having the general classification of "saline", we both show our results are robust to alternative classifications of these sites and to various methods of assigning likely extraction technology to sites (Appendix Sections O and O). Nine sites are recorded as having multiple technologies and contribute to multiple type counts. Full classification details, source documentation, and the rationale for each coding decision are in Appendix Section A.

To verify that cold temperatures decreased coastal salt production and raised the price of salt (Section 5.1), we draw on two sources on the quantity and price of salt. For *quantities*, we use the site-level production tables in Multhauf (1978), which report output for individual production regions at irregularly spaced dates. After converting all figures to millions of kilograms per year, this yields an unbalanced panel of 40 production regions observed between 1205 and 1796, matched to extraction type and local temperature (Appendix Section U).

For *prices*, we use the Allen–Unger commodity-price database, standardized to grams of silver per litre. Because silver standardization removes currency debasement but not the secular depreciation of silver during the Price Revolution, we deflate by Robert Allen's consumer price indices to obtain a real salt-price series spanning 1300–1799. Appendix Section V details the price data, the consumer price index, and an internal-deflator robustness check.

Our primary temperature source is Guiot and Corona (2010), which provides reconstructed growing-season (April-September) temperature anomalies at 5° resolution from 1100 to 1790 for 412 of 476 grid cells. We compute lagged rolling means at windows of 5, 10, 15, 20, 25, 30, and 50 years and interact these with salt endowments. We additionally employ two alternative datasets of

reconstructed temperature (CCSM4 and MPI-ESM) as robustness checks (Appendix Section B).

It is important to control for geography as certain geographic features correlate with border frequency (Kitamura and Lagerlof, 2020). Geographic controls include caloric suitability (Galor and Ömer Özak, 2016), soil nutrients and workability (Fischer et al., 2008), elevation and terrain ruggedness (Farr et al., 2007), and average length of navigable rivers. Urbanization data are from Buringh (2021). Since geographic controls are time-invariant at the grid-cell level, we interact them with period fixed effects. Appendix Section A provides full variable definitions and summary statistics.

5 Empirical Analysis

5.1 Cold Shocks to the Value of Salt

The main hypothesis we test is Prediction 1. The latter, in turn, relies on the following causal chain: cold temperature (Little Ice Age) decreases coastal salt production, and therefore total salt supply, which then raises the price of salt, and increases the profits of brine and rock-salt producers. (See Figure 3.)

While we do not have data on profits or revenues from salt production, we can establish the first, more crucial, part of the argument – that the Little Ice Age decreased coastal salt production and increased the price of salt.

Quantities. Using the Multhauf (1978) site-level production data, we estimate the effect of temperature on coastal salt output. Interacting each extraction type with the lagged temperature mean and absorbing site and year fixed effects (Appendix Table A22), we find that colder temperatures indeed reduce coastal/solar salt production. Figure A7 shows the same pattern in the raw, pooled data. Cold depresses output for salt whose technology depends on solar evaporation, as expected.

Prices. If cold temperatures decrease coastal salt production and make salt more scarce, they should also raise the price of salt. Using the real salt-price series, we regress the annual real price

on the ten-year temperature mean. The real salt price rises significantly as temperatures fall: a slope of -0.23 ($p < 0.001$), robust to detrending (-0.25 with a linear year trend, -0.13 with a quadratic). Figure 4 displays this negative real-price–temperature relationship.

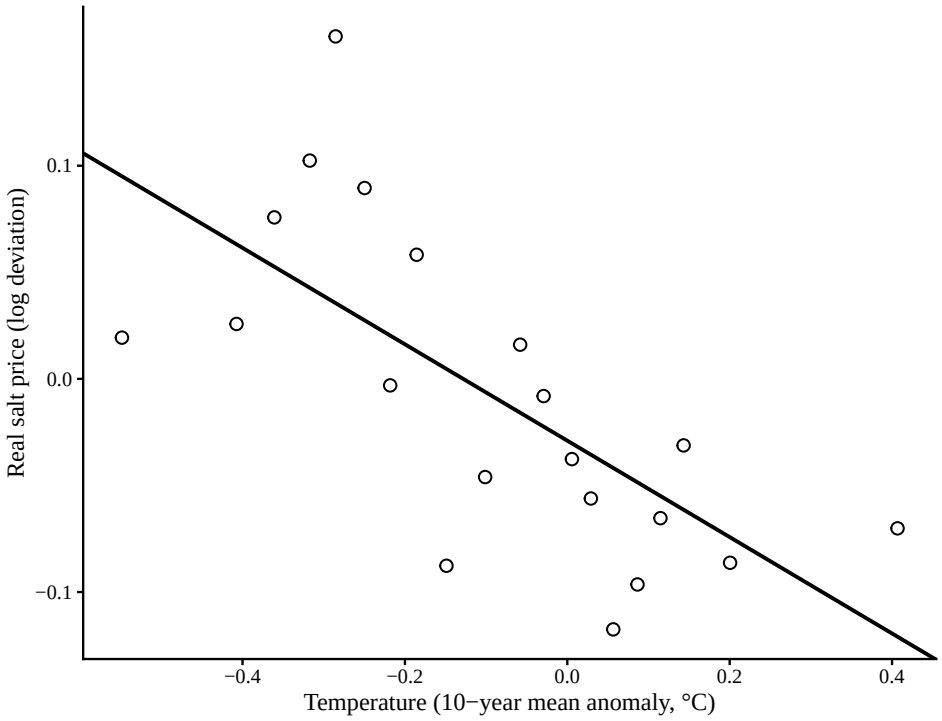

Figure 4: Real Salt Prices Rise as Temperatures Fall

Notes: Binned scatter (20 equal-count bins) of the annual European real salt price against the ten-year mean temperature anomaly, 1300–1799. The real salt price is the silver salt-price index from the Allen–Unger commodity-price database deflated by Robert Allen's consumer price indices (grams of silver, seven European cities); both in log deviations from series means. The line is the OLS fit; the slope is -0.23 ($p < 0.001$) and is robust to linear and quadratic year trends. Temperature from the Guiot and Corona (2010) (EJ-grid) reconstruction. See Appendix Section V.

The pattern suggests that the cold temperatures in the Little Ice Age made salt more valuable per unit. In turn, unless brine and rock salt production decreased dramatically – unlikely since they are climate-insensitive, any increase in the price of salt would have raised the profits of brine and rock-salt producers.

5.2 Grid Cell Analysis

To test Prediction 1, we study within-cell variation across 476 fixed 200km × 200km grid cells covering Europe, observed at 5-year intervals from 1100 to 1790. Grid cells are exogenous geographic units whose boundaries do not change over time. For each cell-period, we observe the number of polity borders passing through the cell.

The baseline specification is:

$$Y_{ct} = \alpha_c + \delta_t + \beta \cdot \text{Salt}_c \times \text{Temp}_{ct} + \gamma \cdot \text{Temp}_{ct} + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the border count in cell c at time t, α_c are cell fixed effects, δ_t are five-year period fixed effects (135 periods), Salt_c is the count of salt sites in the cell, and Temp_{ct} is the rolling mean of the Guiot and Corona (2010) temperature reconstruction at the indicated window (10, 25, or 50 years). Since salt-site counts are time-invariant, the direct effect of salt is absorbed by cell fixed effects; β is identified from within-cell variation in temperature. Throughout the grid-cell analysis, we report Conley (1999) spatial HAC standard errors with a 100km bandwidth to account for spatial correlation in both temperature shocks and political boundaries across neighboring cells. Appendix Table A6 reports standard errors at 200km and 500km bandwidths.

5.3 Aggregate Salt and Political Fragmentation

The Tilly hypothesis predicts that the presence of fiscal resources such as salt increases state capacity and should therefore lead to consolidation. If this hypothesis holds, we should observe a negative relationship between the value of (total) salt resources and the number of borders in a grid-cell.

Table 1 presents the results. Column (1) uses period fixed effects only: cells with more salt sites have significantly more polity borders. Columns (2)–(4) add cell fixed effects and interact salt with temperature to more adequately capture the value of salt sites; because salt-site count is time-invariant, its direct effect is absorbed by the cell fixed effects and identification comes from within-cell variation in temperature. The salt × temperature interaction is negative. The coefficient of -0.724 implies that a 1°C degree cooling raises the border count by about 0.72 for

each salt mine in a cell.[13] The effect is robust to a secular time trend (column 3, −0.462) and grows in magnitude at a 50-year window (column 4, −1.137), consistent with the argument that sustained cold was responsible for shifting the value of salt as a fiscal resource.

Table 1: Salt and Political Consolidation: Grid Cell Analysis

	Dependent Variable: Border Count			
	(1) Period FE	(2) Salt × Temp	(3) + Time Trend	(4) 50yr Window
Salt Count	1.809 ∗ ∗			
	(0.593)			
Temperature (GC)		0.163 ∗ ∗	0.190 ∗ ∗	0.333 ∗ ∗
		(0.075)	(0.088)	(0.146)
Salt × Temperature		−0.724∗	−0.462∗	−1.137
		(0.408)	(0.246)	(0.711)
Salt × Time Trend			0.149	
			(0.104)	
Cell FE	No	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes
Observations	64,260	54,796	54,796	51,500
R^2 (Within)	0.020	0.043	0.088	0.050

Notes: 476 grid cells observed at five-year intervals (1100–1790). All specifications include five-year period fixed effects (135 periods). Column (1) uses random effects (between-cell variation); salt count measures the number of salt sites within each 200km grid cell. Columns (2)–(4) include cell fixed effects; since salt count is time-invariant, its direct effect is absorbed. Temperature is the Guiot–Corona rolling mean (Guiot and Corona, 2010): 10-year window in columns (2)–(3), 50-year in column (4). The sample is smaller in columns (2)–(4) because the reconstruction covers 412 of 476 cells, and further reduced in column (4) by the longer lag. Time trend = (year − 1100)/100. Conley (1999) spatial HAC standard errors (100km bandwidth) in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Figure 5 visualizes the salt × temperature interaction. Both border count and temperature are residualized on cell and five-year period fixed effects, so the figure shows within-cell, within-period variation. There is a strong negative relationship: salt sites at higher temperatures (warmer) are associated with fewer borders. Equivalently, salt sites at colder temperatures are associated with more borders. This implies that as salt becomes more valuable, territories do not consolidate, but rather fragment. This directly runs counter to the Tilly hypothesis.

[13] As temperature is coded as a positive anomaly, a negative coefficient means that colder periods bring *more* borders to salt-rich cells.

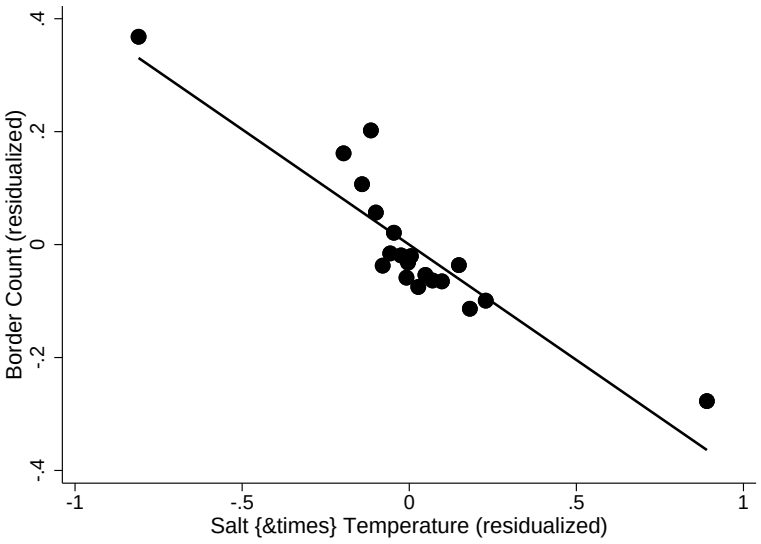

Figure 5: Salt × Temperature and Border Count: Binscatter by Salt Status (Grid Cell Level)

Notes: Both border count and temperature (Guiot–Corona 10-year rolling mean) are residualized on cell and five-year period fixed effects. Observations are split by salt status (salt count > 0 vs. $= 0$) and binned into 10 temperature quantiles within each group. Lines are linear fits through the bin means. The salt-cell line slopes downward: moving left (colder) is associated with more borders. The divergent slopes visualize the negative salt × temperature interaction from column (2) of Table 1. 476 cells, 5-year intervals (1100–1790).

5.4 Disaggregated Salt and Political Fragmentation

The foregoing results do not make sense from a Tilly perspective, as it treats all salt resources in the same manner. We contend that the effect of resources depends on whether it is appropriable or non-appropriable. Only when salt is appropriable can it fund greater consolidation.

To test our hypothesis, we distinguish between different sources of salt. The majority of salt sites in Europe are brine springs and coastal salt works which were relatively difficult to appropriate. There are a number of rock salt mines, which were the opposite—concentrated, observable, capital-intensive, and extracted through a single chokepoint, and therefore appropriable by a central authority.

Table 2 decomposes the aggregate salt × temperature interaction of Table 1 by extraction type. All specifications include cell and period fixed effects; the three salt-type × temperature interactions enter simultaneously.

The interaction of rock salt with temperature has a positive coefficient, implying that colder

periods generate less borders in rock-salt regions. The results on coastal/marine × temperature imply the same. In contrast, the coefficient on brine × temperature is negative, implying that colder periods generate more borders in brine regions.

These patterns are consistent with Prediction 1. Colder temperatures decrease the value of coastal salt—a non-appropriable resource—which should then increase consolidation in regions with coastal salt works. The values of both rock salt and brine salt increase, but the former should generate consolidation since rock salt is appropriable, and the latter fragmentation since brine salt is non-appropriable.

The signs of the estimated coefficients hold across all temperature windows, and the magnitudes of the coefficients increase monotonically with the window length. This suggests that the effects are being driven by sustained cold, and not simply short-run fluctuations.

Table 2: Type-Specific Salt × Temperature: Borders (Cell + Year FE)

	Dependent Variable: Border Count		
	(1)	(2)	(3)
	10yr	25yr	50yr
Rock × Temperature	+0.897∗	+1.319∗	+1.777∗
	(0.521)	(0.720)	(0.938)
Brine × Temperature	−1.902 ∗ ∗	−2.633 ∗ ∗	−3.352 ∗ ∗
	(0.923)	(1.214)	(1.488)
Coastal/marine × Temperature	+0.128	+0.205	+0.316
	(0.134)	(0.178)	(0.225)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes
Observations	54,796	53,560	51,500
R^2 (Within)	0.076	0.096	0.116

Notes: 476 grid cells, five-year intervals. All specifications include cell and period fixed effects. All three type × temperature interactions are entered simultaneously. Temperature is the Guiot–Corona rolling mean at the indicated window. Standard errors clustered by grid cell in parentheses; on this 200km grid these coincide with Conley (1999) spatial HAC standard errors at a 100km cutoff (uniform kernel, allowing unrestricted within-cell serial correlation). Wider spatial bandwidths are reported in Appendix Table A6. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Figure 6 traces the interactions between salt-type and temperature century by century. The three types follow strikingly different trajectories: rock-salt cells move from weak fragmentation

in the 14th century to strong consolidation by the 17th–18th, brine cells fragment monotonically, and coastal/marine cells consolidate in every century once the temperature interaction is controlled for (Appendix Table A17). Approximately, per century we observe −0.28 borders for each additional rock site (consolidation), +0.38 for each brine site (fragmentation), and −0.14 for each coastal/marine site (consolidation). The same climate shock pushed rock and coastal salt sites toward consolidation and brine salt sites toward fragmentation, with the cumulative effect growing as the Little Ice Age deepened. The effects survive the inclusion of century interactions alongside the type × temperature terms, ruling out a generic "salt × time" confounder (Appendix Table A17).

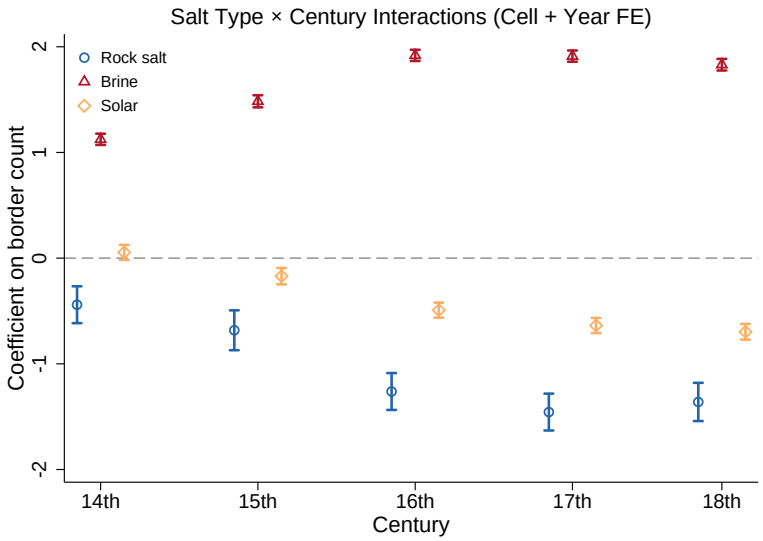

Figure 6: Century-Specific Salt–Border Coefficients by Extraction Type

Notes: Each point represents the coefficient on salt type count × century dummy from a cell fixed effects regression with border count as the dependent variable. All three type interactions entered simultaneously. Whiskers show 95% confidence intervals. Rock salt cells consolidate over time (negative trend); brine cells fragment monotonically (positive trend); coastal/marine cells show a sign reversal explained by the temperature mechanism (see Appendix Table A17).

A mineral placebo test confirms that the results are specific to salt rather than a generic effect of valuable underground resources in cold periods: deposits of iron, copper, silver, gold, and lead show no interaction with temperature on borders, while the salt coefficient retains its sign and monotonic window pattern when the two are entered together. Full results are in Appendix

Section C.

The pattern is also not an artifact of measuring fragmentation with border counts. Re-estimating using the log of mean polity area within each cell as an alternative, consolidation-oriented outcome, the coefficients on each interaction term reverse signs as expected: brine × temperature is positive and significant at every window—colder periods produce *smaller* polities in brine cells (fragmentation)—while rock × temperature and coastal/marine × temperature are negative—colder periods produce *larger* polities (consolidation)—though the rock arm is imprecise given the small number of rock cells (Appendix Table A19).

Lastly, we also conduct our analysis using a complementary measure of political consolidation and fragmentation, the Herfindahl index (HHI) of polity shares within each cell-year. Our results parallel the main specification. Colder temperatures in rock and coastal/marine salt regions are associated with consolidation, but with fragmentation in brine salt regions (Table 3).[14]

Table 3: Type-Specific Salt × Temperature: HHI Concentration (Cell + Year FE)

	Dependent Variable: HHI polity concentration		
	(1) 10yr	(2) 25yr	(3) 50yr
Rock × Temperature (pred. −)	−0.0749**	−0.0937**	−0.0989
	(0.0325)[0.0309]	(0.0469)[0.0451]	(0.0664)[0.0650]
Brine × Temperature (pred. +)	+0.0252***	+0.0321***	+0.0334***
	(0.0094)[0.0126]	(0.0112)[0.0147]	(0.0119)[0.0148]
Coastal/marine × Temperature (pred. −)	−0.0134	−0.0242**	−0.0367***
	(0.0085)[0.0080]	(0.0097)[0.0093]	(0.0102)[0.0107]
Cell FE, Period FE	Yes	Yes	Yes
Observations	46,718	46,718	46,718

Notes: Dependent variable is the (normalized) within-cell Herfindahl index of polity shares, bounded in $(0, 1]$ (higher = more consolidated). On HHI, consolidation is a *negative* salt × temperature coefficient and fragmentation *positive*; predicted signs after each row label. All three type × temperature interactions are entered simultaneously with cell and period fixed effects. Standard errors clustered by grid cell in (parentheses), with Conley (1999) spatial-HAC standard errors at 200km in [brackets] as a robustness check; stars are based on the clustered standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

[14]Robustness checks for the HHI results are in Appendix Table A33.

5.5 Instrumental Variables

A natural concern with the results so far is that the locations of salt sites may be endogenous: they may be correlated with other determinants of political fragmentation, such as the level of economic development required to exploit salt deposits. Fixed effects address this only partially. We therefore construct IVs and run 2SLS regressions.

We first very the results from 5.3. We use the geological variation in evaporite-basin deposits. These are sedimentary formations created by the evaporation of ancient seas, and are thus a geological prerequisite for salt. As an instrument for salt sites, we interact each cell's evaporate-basin coverage (share of cell area overlapping a basin) and the mean terrain ruggedness (TRI) within that overlap; basin geology determines where salt *can* exist, while ruggedness affects extraction costs and hence whether sites were actually opened. Both are fixed by deep geological processes and are plausibly exogenous to medieval political outcomes.

Panels A and B of Table 4 report the 2SLS results. Panel A instruments the salt *level* (number of salt sites); because basin $\times$ TRI is time-invariant it uses period fixed effects only.[15] The 2SLS estimate of 2.297 is positive and 27% larger than OLS (1.809). Because basin $\times$ TRI is cross-sectional, however, the first stage is weak once standard errors are clustered by cell (first-stage $F \approx 5$, far below the figure the homoskedastic formula returns); we therefore rely on the weak-instrument-robust Anderson–Rubin 95% set, [1.10, 6.60], which remains positive and excludes zero. Panel B instruments the salt $\times$ temperature interaction with (basin $\times$ TRI)$\times$ temperature and includes cell and year fixed effects; the 2SLS coefficient keeps its negative sign (-0.33 at the 10-year window) but is not statistically significant once standard errors are appropriately clustered. The estimated 2SLS coefficients have the same signs as those in 5.3. This confirms that the Tilly hypothesis is inconsistent with the data.

We next conduct 2SLS regressions on the effect of each salt type and temperature interaction. In this case, however, we cannot use the same IV because evaporite-basin geology is common to *both* rock and brine: Permian and Triassic evaporites generate bedded rock salt and the brine springs that dissolve out of those same strata alike. To test whether the opposing type-specific

[15] Cell fixed effects would absorb the instrument.

responses are themselves causally identified, we instrument each type's temperature interaction with an independent geological measure (Panels C and D). Panel C instruments brine × temperature with the fraction of each cell covered by Upper Triassic (Keuper) formations from the EGDI GeologicUnitView database—the salt-bearing strata underlying central European brine springs[16]—and Panel D instruments coastal/marine × temperature with sea-surface salinity, the key determinant of solar-evaporation productivity.[17] The 2SLS results for brine salt show negative and significant coefficients at every window, so colder periods are thus associated with fragmentation in brine cells. The coefficients are roughly twice the OLS magnitude (−3.90 vs. −1.86 at the 10-year window; −5.75 vs. −3.26 at 50 years). This is consistent with attenuation of classical measurement-error, as historical brine sites are likely undercounted in the data.[18] The 2SLS coefficients on coastal/marine salt × temperature are positive (+0.88 to +1.92) but, once standard errors are clustered, are not statistically significant at any window.[19]

6 Robustness

Our main results, specifically our finding that different types of salt had different effects on political fragmentation, are robust. Table 5 reports the three salt-type × temperature coefficients while controlling for other important drivers of border formation, including features emphasized in the existing literature (e.g. Kitamura and Lagerlof (2020)). We also demonstrate robustness to using alternative temperature reconstructions, sample-period splits, functional forms, and different spatial resolutions. The brine × temperature coefficient is negative and precisely estimated in

[16]The brine result does not hinge on the Keuper map. Instrumenting brine × temperature jointly with three near-orthogonal geological proxies—Keuper coverage, karst aquifer × evaporite lithology (IHME1500), and Mesozoic sedimentary-basin coverage (InDepth)—leaves the 2SLS essentially unchanged (−3.78, −4.79, and −5.55 at the 10-, 25-, and 50-year windows), and the over-identified specification permits a test of the exclusion restriction.

[17]We are not able to construct an instrument specific to *rock* salt. Because rock salt and brine arise from the same subsurface evaporite formations and differ only in extraction technology—mined bedded halite versus springs dissolved out of it—any geology-based instrument that predicts rock also predicts brine; since the brine effect is large and opposite in sign, it then dominates the 2SLS and the rock estimate collapses onto the brine one.

[18]There is also a LATE interpretation: the Keuper instrument moves brine endowment precisely in the central-European cells where dispersed Triassic brine springs make salt least appropriable (Angrist and Imbens, 1995), so the IV may estimate the effect for the cells where the appropriability mechanism is strongest while OLS dilutes this with brine of more mixed character elsewhere in Europe.

[19]This coefficient is identified off only about thirty coastal cells; the coastal IV should therefore be read as indicative rather than precisely estimated.

Table 4: Salt and Political Fragmentation: Instrumental-Variable Estimates

	(1)	(2)	(3)	(4)
Panel A: Aggregate salt level *(endogenous = salt count); period FE*				
	OLS	First Stage	Reduced Form	2SLS
Salt Count	1.809 ∗ ∗			2.297 ∗ ∗
	(0.594)			(0.627)
Basin × TRI		0.0192 ∗ ∗	0.0442 ∗ ∗	
		(0.0083)	(0.0176)	
Clustered first-stage $F = 5.4$; AR 95% set [1.10, 6.60] excludes zero; $N = 64{,}260$.				
Panel B: Aggregate salt × temperature; cell + year FE				
	OLS (10yr)	2SLS (10yr)	2SLS (+Trend)	
Salt × Temperature	−0.724∗	−0.328	−0.472∗	
	(0.408)	(0.250)	(0.278)	
Salt × Time Trend			−0.095	
			(0.106)	
Instrument: (basin × TRI)× temperature; clustered first-stage $F \approx 4$; $N = 54{,}796$.				
Panel C: Brine × temperature (border count); cell + year FE				
	OLS	Keuper IV		
10-year window	−1.862 ∗ ∗	−3.897 ∗ ∗∗		
	(0.910)	(1.141)		
25-year window	−2.572 ∗ ∗	−4.936 ∗ ∗∗		
	(1.199)	(1.450)		
50-year window	−3.264 ∗ ∗	−5.745 ∗ ∗∗		
	(1.474)	(1.715)		
Clustered first-stage $F \approx 44$. Instrument: Keuper (Upper Triassic) geology × temperature.				
Panel D: Coastal/marine × temperature (border count); cell + year FE				
	OLS	Salinity IV		
10-year window	+0.035	+0.885		
	(0.092)	(0.790)		
25-year window	+0.093	+1.360		
	(0.112)	(1.058)		
50-year window	+0.206	+1.924		
	(0.129)	(1.298)		
Clustered first-stage $F \approx 10$; Instrument: sea-surface salinity × temperature.				

Notes: 476 grid cells, five-year intervals (1100–1790); dependent variable is border count throughout. *Standard errors are clustered by grid cell* (equivalently, Conley spatial-HAC at the 100km bandwidth on this 200km grid); a homoskedastic formula substantially overstates precision for these panel and cross-sectional regressors and is not used. *Panels A–B* use the aggregate instrument basin × TRI (share of each cell overlapping an evaporite basin × mean terrain ruggedness within the overlap); because the instrument is cross-sectional the first stage is weak under clustering, so Panel A relies on the weak-instrument-robust Anderson–Rubin set (Andrews et al., 2006), which excludes zero, and the aggregate salt × temperature 2SLS (Panel B) is not significant under clustering. *Panels C–D* instrument the type-specific temperature interactions: brine × temperature with Upper Triassic (Keuper) geology (a strong instrument, $F \approx 44$; the brine 2SLS is also robust to instrumenting jointly with three near-orthogonal geological proxies, see text), and coastal/marine × temperature with sea-surface salinity (borderline, $F \approx 10$, identified off ∼31 coastal cells). The brine 2SLS is about twice OLS, consistent with measurement-error attenuation, and is the only arm with both a strong instrument and a significant 2SLS; the coastal 2SLS is not significant under clustering. No instrument is available for *rock* salt (rock and brine share the same evaporite geology; see text); the rock arm rests on the OLS decomposition of Table 2. Standard errors in parentheses. ∗∗∗ $p < 0.01$, ∗∗ $p < 0.05$, ∗ $p < 0.1$.

every specification, and coefficients of rock × temperature and coastal/marine × temperature are positive throughout (though they lose precision in some specifications). Dropping any one macro-region also does not overturn any of our results.

A related question is *where* in Europe the results come from. The leave-one-out exercises (Appendix Tables A4 and A5) show that the brine effect is largest in central Europe and the lands of the Holy Roman Empire (Appendix Table A10). This makes sense historically. Central Europe is an important counterexample to the bellicist account associated with Tilly. It was a region of near-continuous warfare that nonetheless remained divided among hundreds of small territories until the nineteenth century (Whaley, 2012; Wilson, 2016, 2023). A theory of European state formation must explain why war failed to make large states precisely there. Our mechanism may provide part of the answer the Empire sat atop the Triassic strata that fed Europe's dispersed brine springs, so the region's salt wealth was structurally difficult for any central authority to appropriate, and cold shocks that raised its value strengthened exactly the small territories that Tilly's logic predicts should have been swallowed.[20]

A full set of robustness checks, with different approaches to modeling standard errors and additional checks—Conley spatial-HAC bandwidths, spatial unit-root corrections, leave-one-cluster-out, randomization inference, non-overlapping time units, Holy Roman Empire institutional interactions, alternative battle data, sensitivity to unclassified sites, and type-specific instruments for the brine and coastal/marine interactions—are reported in the Appendix.

We also rule out other explanations. Two alternative channels deserve direct attention. First, the salt–temperature interaction could reflect an economic development channel rather than a political mechanism: if temperature shocks affect urbanization in salt-rich cells, and urbanization in turn shapes political boundaries, our results might capture indirect economic effects. Appendix Table A26 addresses this: the salt × temperature interaction has *no* significant effect on urban population at any temperature window (coefficients -0.015 to $+0.023$, all $p > 0.2$). We do not include urbanization as a control in our main specifications because it is potentially a "bad

[20]Note that the mechanism is not merely a proxy for HRE institutions: the brine × temperature coefficient remains negative outside the Empire (Appendix Table A10), and within the Empire salt-poor cells do not fragment the way brine-rich cells do.

Table 5: Robustness of the Type-Specific Salt × Temperature Result

	Rock × Temp	Brine × Temp	Coastal/marine × Temp
Confounder controls (full sample)			
Baseline (Table 2)	$+1.319^{*}$	-2.633^{**}	$+0.205$
+ Black Death mortality	$+1.385^{*}$	-2.570^{**}	$+0.175$
+ Other minerals	$+1.296^{*}$	-2.691^{**}	$+0.199$
+ Geography (elevation, ruggedness)	$+1.418^{*}$	-2.652^{**}	$+0.221$
+ Coastal proximity, sea salinity	$+1.373^{*}$	-2.626^{**}	$+0.144$
+ Urbanization (benchmark years)	$+0.119$	-1.172^{**}	-0.003
Alternative temperature series and sample period			
CCSM4 reconstruction	$+0.842$	-1.649^{*}	$+0.358^{*}$
MPI-ESM reconstruction	$+0.251$	-0.474^{**}	$+0.187^{**}$
Functional form and spatial resolution			
+ Type × century dummies	$+0.662$	-1.684^{***}	-0.333^{**}
Poisson fixed effects	$+0.097$	-0.120^{***}	$+0.015$
150km grid	$+0.457$	-0.967^{**}	$+0.067$
Leave-one-region-out (most influential)	$+0.827^{***}$	-1.285^{**}	-0.108

Notes: Each row re-estimates the border-count specification of Table 2 at the 25-year temperature window, with all three type × temperature interactions entered simultaneously and cell and period fixed effects. Coefficients shown; Conley (1999) spatial HAC standard errors (100km) underlie the significance stars. The Poisson coefficients are count-model semi-elasticities and are not directly comparable in magnitude to the linear estimates; signs and significance are. The urbanization row is estimated on the ten Buringh (2021) benchmark years ($N \approx 3{,}708$); within that sample the type pattern is unchanged with versus without urban controls. Full results with standard errors, and additional checks, are in Appendix Tables A27 (Black Death), A17 (century interactions), A5 (Lorraine exclusion), and 4 (instrumental variables, Section 5.5), among others. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

control"—downstream of salt endowments and subject to post-treatment bias—but the null effect on urbanization rules out the economic-development channel as a confound.

Second, Black Death mortality (1347–1351) could confound the results through a demand-side channel: higher mortality raised per capita incomes and hence demand for meat and salt preservation. Black Death mortality was lower in central Europe (Jedwab et al., 2022, 2024), where brine salt was disproportionately located, so it could differentially affect our estimates. Controlling for Black Death mortality interacted with temperature and time trends leaves the salt × temperature coefficient essentially unchanged, and the type-specific opposing signs are fully preserved (Appendix Table A27). Black Death mortality does not confound the salt–temperature mechanism.

7 Mechanisms:

The mechanism by which non-appropriable resources generate fragmentation is given by Lemma 1 of our model. Non-appropriable resources increase the fighting capacity of a territory and thus emboldens it to rebel (rebellion channel). At the same time, territories with large non-appropriable resources are therefore unattractive targets for annexation by a polity, precisely because they are more likely to rebel post-conquest (war deterrence channel).

Since the rebellion channel relates to an individual territory's decision, while the war deterrence channel pertains to a polity's decision, we use different units of observation to test each channel separately.

For the rebellion channel, we can simply count the number of rebellions in a given region, i.e. a grid cell, without specifying the particular polity in the grid cell in which the rebellion is located. We can do so because the calculus is the same for any territory in any polity – the territory is more likely to rebel when it has large non-appropriable resources. In contrast, wars of annexation are between polities, which necessitates using polities as units of observation when testing the war deterrence channel.

Of course, the main data hurdle is in distinguishing between types of battles, to identify incidents of rebellion and wars of annexation. We elaborate on our strategy below.

7.1 The Rebellion Channel

All else equal, regions with larger non-appropriable resources should experience more rebellions than regions with less.

There are several challenges to testing this prediction. First, comprehensive and comparable data on rebellions does not exist for all Europe for the period of analysis. Second, there was often no clear delineation between rebellions/civil wars and interstate wars in the medieval period as these were frequently intertwined.[21] We instead proceed by using battles, a direct, reduced-form, indicates of conflict following Dincecco and Onorato (2016), Leeson and Russ (2018) and Dincecco

[21]For example, the Hundred Years War between England and France was also a time of French civil war, during which Burgundy established itself as a semi-independent polity.

et al. (2022). We then focus on civil wars – as internal wars they are a sharper proxy for rebellion, and show that these are responsible for the empirical patterns that we document.

Panel A of Table 6 uses all battles. The estimated coefficients on rock × temperature is positive, which implies that in colder weather (lower temperature), regions with more rock salt mines experience less rebellions. Now since colder weather increase the value of rock salt, and rock salt is appropriable, the results suggest that less rebellions occur when appropriable resources are relatively larger or, equivalently, when non-appropriable resources are relatively smaller.

The results on brine × temperature provide even more direct evidence for the rebellion channel. Since the estimated coefficients are negative, the results imply that when the value of non-appropriable resources increases (brine salt in lower temp/colder weather), rebellions also increase.

The negative coefficients on coastal/marine × temperature, however, do not appear to validate the rebellion channel. In colder weather, the value of coastal/marine salt—a non-appropriable resource—is lower, but it appears to be associated with more rebellions. Nonetheless, note that the magnitude of the coefficients are smaller than those for brine × temperature.

The patterns are confirmed when we restrict battles to civil war battles (Panel B).

These conflict results are robust. For example, at a coarser unit—the polity level—the same patterns appear and, as in Table 6, run through civil conflict (see Appendix T).

7.2 The War Deterrence Channel

While we can test the rebellion channel at the cell-level, this analysis is not well suited to exploring the relationship between salt and external wars. This is because the cell-level analysis does not permit us to assign a salt site to a given polity, and will thus fail to capture instances in which a polity uses salt rent revenues earned in one place to fund conquest somewhere else.

We thus test the war deterrence channel by moving from the grid cell to the polity. We hand-assign each salt site to the polity that initially held it — this fixes each rock, brine, and coastal salt endowment to the polity, thereby preventing any subsequent conquest from contaminating the treatment. When a new polity appears in the data, we are able to identify it as the union of previous polities, e.g. Poland-Lithuania from 1385, Castile and Aragon into Spain from 1479. We count, for each polity and five-year period, the *interstate* battles it fights outside the territories

Table 6: Salt and Conflict: All Battles and Civil Wars

Temperature window	10-year	25-year	50-year
Panel A. All battles			
Salt × Temp	-0.0127^{***}	-0.0222^{***}	-0.0373^{***}
	(0.0030)	(0.0050)	(0.0102)
Rock × Temp	$+0.0112^{*}$	$+0.0244^{**}$	$+0.0272$
	(0.0061)	(0.0104)	(0.0329)
Brine × Temp	-0.0214^{***}	-0.0439^{***}	-0.0834^{***}
	(0.0071)	(0.0133)	(0.0318)
Coastal/marine × Temp	-0.0107^{***}	-0.0143^{***}	-0.0142^{***}
	(0.0040)	(0.0040)	(0.0051)
Panel B. Civil wars battles only			
Salt × Temp	-0.0070^{***}	-0.0109^{***}	-0.0202^{**}
	(0.0019)	(0.0039)	(0.0089)
Rock × Temp	$+0.0045$	$+0.0151^{*}$	$+0.0362^{*}$
	(0.0061)	(0.0086)	(0.0194)
Brine × Temp	-0.0134^{***}	-0.0255^{**}	-0.0540^{*}
	(0.0044)	(0.0102)	(0.0278)
Coastal/marine × Temp	-0.0063^{***}	-0.0076^{**}	-0.0083^{**}
	(0.0024)	(0.0030)	(0.0040)
Cell FE, Period FE	Yes	Yes	Yes
Observations	53,560	53,560	53,560

Notes: 476 grid cells × five-year periods. The dependent variable is the number of battles per cell-period from the 538 georeferenced battles (1302–1700); Panel A counts all battles, Panel B only those coded as civil wars battles. The aggregate row interacts the total salt count with temperature; the type rows enter the rock, brine, and coastal/marine counts × temperature simultaneously. Cell and period fixed effects throughout; Conley (1999) spatial-HAC standard errors (100km) in parentheses. The opposing-sign signature (brine and coastal/marine −, rock +) is carried by civil conflict. Results are similar under the narrow civil definition. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

it controls. This proxies for the number of external, or annexationist, wars. Finally, we interact each polity's salt endowment with the temperature in its lands.

Table 7 reports the results. Overall, the signs on the estimated coefficients of the interaction between salt and temperature are *opposite* of those from Table 6 (civil wars/rebellion). This is consistent with our model, as Lemma 1 implies that rebellion should make annexation less attractive. Non-appropriable resources which enable greater rebellion should also be associated with fewer wars of annexation.

The positive coefficients on brine × temperature reveal that in colder temperatures, in which

Table 7: Salt-Type Endowment × Temperature and External (Annexation) Battles

	Dep. var.: external battles per polity-period		
Temperature window	10-year	25-year	50-year
Rock × Temp	-0.080^{***}	-0.069^{***}	-0.103^{**}
	(0.019)	(0.022)	(0.050)
Brine × Temp	$+0.031^{***}$	$+0.042^{***}$	$+0.076^{***}$
	(0.005)	(0.007)	(0.020)
Solar × Temp	+0.001	−0.005	−0.014
	(0.006)	(0.013)	(0.032)
Polity FE, Period FE	Yes	Yes	Yes
Polities	56	56	56
of which rock-endowed	4	4	4
Observations	1,891	1,891	1,891

Notes: Polity × five-year-period panel, 1300–1700. A polity's baseline rock/brine/solar endowment (hand-coded deposit→polity, fixed at its core holdings) is interacted with the temperature in its territory. The dependent variable counts the interstate battles the polity fights in cells it does *not* control ("external"/annexation wars). Polities are chained over time (Poland–Lithuania from 1385; Castile+Aragon=Spain from 1479); the Habsburg hereditary, Bohemian, and (post-1526) Hungarian lands are treated as a single entity (see Appendix T.3). Polity and period fixed effects throughout; standard errors clustered by polity in parentheses. Cold (lower temperature) raises external warfare for *rock*-endowed (appropriable) polities and lowers it for *brine*-endowed (non-appropriable) polities. A specification treating the Habsburg lands as separate states yields the same pattern and is available on request. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

(non-appropriable) brine salt is more valuable, there are fewer external wars fought by polities with more brine springs. In the baseline specification, a one-unit fall in temperature generates 0.042–0.076 *less* external battles for a polity having one more brine-spring site. In contrast, an equivalent shock generates 0.07-0.10 *more* external battles for a polity having one more rock salt mine. The results thus suggest that larger non-appropriable (appropriable) resources deter (enable) wars of annexation.

The role of resource-appropriability is clearly evident in the data. When non-appropriable (salt) resources become more valuable, they can fund more rebellion, and at the same time more effectively deter annexation. This explains why regions with large non-appropriable resources are more fragmented.

8 Concluding Comments

We develop a theory of how the appropriability of a fiscal resource affects the size of polities in equilibrium. Territories are grouped into polities; each period, a territory can secede through rebellion, engage in a war of annexation against a neighbor, or maintain the status quo. The decisive variable is the appropriability of a territory's resources—how much of their value a central authority can capture.

When resources are non-appropriable, a territory can both fund resistance to incorporation and credibly threaten revolt if conquered, so it is neither easily absorbed nor worth absorbing; when they are appropriable, the reverse holds and a ruler who controls them is strengthened. The model generates a clear prediction: as non-appropriable resources rise, equilibrium partitions become more fragmented.

Salt was universally consumed and difficult to substitute. But how appropriable the rents were from salt depends on how it was extracted. Rock salt was easy for the state to appropriate, brine springs, and coastal salt works were hard to appropriate.

Our identification strategy exploits the interaction between salt endowments and temperature variation associated with the Little Ice Age. Cooler temperatures reduced the outcome of climate-sensitive coastal sat. This raised the value of rock salt and brine salt. The political response divided sharply by extraction type. In rock-salt regions—concentrated, capital-intensive, monopolizable—cold periods brought consolidation; in brine regions—dispersed and hard to monitor—they brought fragmentation. The opposing signs hold across temperature windows, on a finer grid, under spatial-HAC standard errors, and on a concentration-based measure as well as on border counts, and the fragmenting effect of non-appropriable brine is confirmed when salt is instrumented by independent geological variation.

Consistent with our theory, the mechanism we identify as responsible for these results is conflict. A single mechanism explains the contrast. Where salt was concentrated and observable—the rock-salt mine, with its single chokepoint—rulers could tax it, and a more valuable deposit strengthened whichever state held it. Where salt was dispersed—brine springs scattered across a landscape, each worked by a small producer—the same rise in value flowed to many hands, financing rebellion and

resistance rather than the center. The conflict record shows both sides of this: non-appropriable salt is associated with internal conflicts and battles, while appropriable salt funded the wars of annexation by which states enlarged themselves.

These results recast a classic claim. Tilly held that war made the state. We find that war made the state only where the fiscal sinews of war were appropriable. This provides novel insights into Europe's persistent political fragmentation. It is fitting that our results are strongest in central Europe: the Holy Roman Empire is where the Tilly thesis fails most conspicuously—endemic warfare there produced not consolidation but a mosaic of hundreds of small polities that survived into the nineteenth century. The geography of non-appropriable salt wealth helps explain why the region at the heart of Europe's wars remained its most fragmented.

References

Abramson, S. (2017). The economic origins of the territorial state. *International Organization 71*(1), 97–130.

Acharya, A. and A. Lee (2018). Economic foundations of the territorial state system. *American Journal of Political Science 62*(4), 954–966.

Adshead, S. A. M. (1992). *Salt and Civilization*. London: Macmillan.

Akridge, D. G. (2008). Methods for calculating brine evaporation rates during salt production. *Journal of Archaeological Science 35*, 1453–1462.

Alesina, A., R. Baqir, and W. Easterly (1999). Public goods and ethnic divisions. *The Quarterly Journal of Economics 114*(4), 1243–1284.

Alesina, A. and E. Spolaore (2003). *The Size of Nations*. Cambridge, MA: MIT Press.

Alesina, A. and E. Spolaore (2005). War, peace, and the size of countries. *Journal of Public Economics 89*(7), 1333–1354.

Aloisio, M. (2010). Salt and royal finance in the kingdom of Naples under Alfonso the Magnanimous. *Medioevo adriatico: ricerche della Società internazionale per lo studio dell'Adriatico nell'età medievale, SISAEM: 3, 2010*, 9–28.

Andrews, D. W. K., M. Moreira, and J. H. Stock (2006). Optimal two-sided invariant similar tests for instrumental variables regression. *Econometrica 74*(3), 715–752.

Angrist, J. D. and G. W. Imbens (1995). Two-stage least squares estimation of average causal effects in models with variable treatment intensity. *Journal of the American Statistical Association 90*(430), 431–442.

Beck, G. (1993). The economic exploitation of man and nature in the production of salt: The case of the southwest German salines in the eighteenth and nineteenth centuries. In *Bad Nauheimer Beiträge zur Salzgeschichte*. University of Göttingen.

Becker, S., A. Ferrara, E. Melander, and L. Pascali (2025). Wars, taxation and representation: Evidence from five centuries of German history. *Journal of the European Economic Association 23*(6), 2241–2288.

Becker, S. O., P. D. Boll, and H.-J. Voth (2026). Testing and correcting for spatial unit roots in regression analysis. *Stata Journal 26*(2), 177–202.

Bloch, M. R. (1976). Salt in human history. *Interdisciplinary Science Reviews 1*(4), 336–352.

Bolton, P. and G. Roland (1997, 11). The Breakup of Nations: A Political Economy Analysis*. *The Quarterly Journal of Economics 112*(4), 1057–1090.

Buringh, E. (2021). The Population of European Cities from 700 to 2000. *Research Data Journal for the Humanities and Social Sciences 6*(1), 1–18.

Campbell, B. M. (2016). *The Great Transition: Climate, Disease and Society in the Late-Medieval World*. Cambridge: Cambridge University Press.

Cantoni, D., C. Mohr, and M. Weigand (2024). The rise of fiscal capacity: Administration and state consolidation in the Holy Roman Empire. *Econometrica 92*(5), 1439–1472.

Cederman, L.-E., P. Galano Toro, L. Girardin, and G. Schvitz (2023). War did make states: Revisiting the bellicist paradigm in early modern Europe. *International Organization 77*, 324–362.

Cervellati, M., S. Lazzaroni, G. Prarolo, and P. Vanin (2026). The state system and growth in Europe 1000-1800 AD. *Journal of the European Economic Association Forthcoming*.

Chambru, C., E. Henry, and B. Marx (2024, November). The dynamic consequences of state building: Evidence from the French Revolution. *American Economic Review 114*(11), 3578–3622.

Chiovelli, G., L. Fergusson, L. R. Martinez, J. D. Torres, and F. Valencia Caicedo (2024). Bourbon reforms and state capacity in the Spanish Empire. Discussion Paper 19042, Centre for Economic Policy Research. Also BFI Working Paper 2024-36.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics 92*(1), 1–45.

Conley, T. G. and M. Kelly (2025). The standard errors of persistence. *Journal of International Economics 153*, 104027.

Dean, T. (2002). *Land and power in late medieval Ferrara: the rule of the Este, 1350-1450*. Number 7. Cambridge University Press.

Depetris-Chauvín, E. and Ö. Özak (2024). Borderline disorder: (de facto) historical ethnic borders and contemporary conflict in Africa. Discussion Paper 13736, Centre for Economic Policy Research. Revised May 2024; also IZA Discussion Paper 13736.

Desierto, D. and M. Koyama (2025). Feudal political economy. *Economic Theory*.

Desierto, D. and M. Koyama (2026). Castles. *European Economic Review 184*, 105281.

Desmet, K., I. O. Ortín, and Ö. Özak (2025, 05). Is secessionism mostly about income or identity? a global analysis of 3,153 subnational regions. *The Economic Journal 135*(668), 1261–1299.

Dincecco, M., J. Fenske, A. Menon, and S. Mukherjee (2022, 04). Pre-colonial warfare and long-run development in India. *The Economic Journal 132*(643), 981–1010.

Dincecco, M. and M. G. Onorato (2016). Military conflict and the rise of urban Europe. *Journal of Economic Growth 21*(3), 259–282.

Dincecco, M. and M. G. Onorato (2017). *From Warfare to Welfare*. Cambridge: Cambridge University Press.

Draskóczy, I. (2018). Salt mining and trade in Hungary from the mid-thirteenth century until the end of the middle ages. In *The Economy of Medieval Hungary*, Leiden, The Netherlands, pp. 205 – 218. Brill.

Drelichman, M. (2005, December). All that glitters: Precious metals, rent seeking and the decline of Spain. *European Review of Economic History 9*(03), 313–336.

Elliott, J. H. (1961). The decline of Spain. *Past and Present 20*(1), 52–75.

Farr, T. G., P. A. Rosen, E. Caro, R. Crippen, R. Duren, S. Hensley, M. Kobrick, M. Paller, E. Rodriguez,

L. Roth, D. Seal, S. Shaffer, J. Shimada, J. Umland, M. Werner, M. Oskin, D. Burbank, and D. Alsdorf (2007). The shuttle radar topography mission. *Reviews of Geophysics 45*(2), RG2004.

Fernández-Villaverde, J., M. Koyama, Y. Lin, and T.-H. Sng (2023, 01). The Fractured-Land Hypothesis*. *The Quarterly Journal of Economics 138*(2), 1173–1231.

Fink, A. M. (1964). Equilibrium in a stochastic n-person game. *Journal of Science of the Hiroshima University, Series A-I (Mathematics) 28*(1), 89–93.

Fischer, G., F. Nachtergaele, S. Prieler, H. T. van Velthuizen, L. Verelst, and D. Wiberg (2008). Global agro-ecological zones assessment for agriculture (GAEZ 2008). Technical report, IIASA and FAO, Laxenburg, Austria and Rome, Italy.

Galor, O. and Ömer Özak (2016, October). The Agricultural Origins of Time Preference. *American Economic Review 106*(10), 3064–3103.

Gennaioli, N. and H.-J. Voth (2015). State capacity and military conflict. *Review of Economic Studies 82*(4), 1409–1448.

Giommoni, T. and G. Loumeau (2026). Taxation with a grain of salt: The long-term effect of extractive fiscal institutions on development. *Review of Economics and Statistics*. Published online ahead of print.

Giommoni, T., G. Loumeau, and M. Tabellini (2026, February). Extractive taxation and the French Revolution. Working Paper 34816, National Bureau of Economic Research.

Gorski, P. and V. S. Sharma (2017). Beyond the Tilly thesis: "family values" and state formation in Latin Christendom. In L. B. Kaspersen and J. Strandsbjerg (Eds.), *Does War Make States? Investigations of Charles Tilly's Historical Sociology*, Chapter 4, pp. 98–124. Cambridge University Press.

Guiot, J. and C. Corona (2010, April). Growing season temperature in Europe and climate forcings over the past 1400 years. *Plos One 5*(4), 1–15.

Hanik, M. (1988). *Wieliczka: Seven Centuries of Polish Salt*. Warsaw: Interpress.

Henriques, A., K. K. Karaman, and N. Palma (2026). State capacity and executive constraints in early modern Europe. *Public Choice*. Published online ahead of print.

Hoffman, P. T. (2015). *Why Did Europe Conquer the World?* Princeton, NJ.: Princeton University Press.

Hughes, E. (1925). The English monopoly of salt in the years 1563-71. *The English Historical Review 40*(159), 334–350.

Huning, T. R. and F. Wahl (2023). You reap what you know: Appropriability and the origin of European states. *European Journal of Political Economy 79*, 102432.

Jaques, T. (2007). *Dictionary of Battles and Sieges: A Guide to 8,500 Battles from Antiquity through the Twenty-first Century*. Westport, CT: Greenwood Press. 3 volumes.

Jedwab, R., N. D. Johnson, and M. Koyama (2022, March). The economic impact of the Black Death. *Journal of Economic Literature 60*(1), 132–78.

Jedwab, R., N. D. Johnson, and M. Koyama (2024). Pandemics and cities: Evidence from the Black Death and the long-run. *Journal of Urban Economics 139*(C).

Jones, E. L. (1981, 2003). *The European Miracle* (3rd ed.). Cambridge: Cambridge University Press.

Kaspersen, L. B., J. Strandsbjerg, and B. Teschke (2017). Introduction: State formation theory: Status, problems, and prospects. In L. B. Kaspersen and J. Strandsbjerg (Eds.), *Does War Make States? Investigations of Charles Tilly's Historical Sociology*, pp. 1–22. Cambridge University Press.

Kelly, M. (2019). The standard errors of persistence. Working paper, University College Dublin.

Kelly, M. and C. Ó. Gráda (2014). The waning of the little ice age: climate change in early modern Europe. *Journal of Interdisciplinary History 44*(3), 301–325.

Kitamura, S. (2021). World historical battles database. Open Science Framework. Accessed: April 2026.

Kitamura, S. and N.-P. Lagerlof (2020). Geography and state fragmentation. *Journal of the European Economic Association 18*(4), 1726–1769.

Ko, C. Y., M. Koyama, and T.-H. Sng (2018). Unified China and Divided Europe. *International Economic Review 59*(1), 285–327.

Kurlansky, M. (2002). *Salt: A World History*. New York: Walker and Company.

Kwass, M. (2000). *Privilege and the politics of taxation in eighteenth-century France: liberté, égalité, fiscalité*. Cambridge: Cambridge University Press.

Kwass, M. (2014). *Contraband: Louis Mandrin and the Making of a Global Underground*. Cambridge, MA: Harvard University Press.

Lamb, H. (1982). *Climate, History, and the Modern World*. London: Methuen.

Laszlo, P. (1998). *Salt: Grain of Life*. New York: Columbia University Press. Translated by Mary Beth Mader.

Leeson, P. T. and J. W. Russ (2018). Witch trials. *The Economic Journal 128*(613), 2066–2105.

Martines, L. (2013). *Furies: War in Europe, 1450-1700*. London: Bloomsbury Press.

Mayshar, J., O. Moav, and Z. Neeman (2017). Geography, transparency and institutions. *American Political Science Review 111*(3), 622–636.

Mayshar, J., O. Moav, Z. Neeman, and L. Pascali (2022). The origin of the state: Land productivity or appropriability? *Journal of Political Economy 130*(4), 1091–1144.

Mokyr, J. (2016). *Culture of Growth*. Princeton, NJ.: Princeton University Press.

Müller, U. K. and M. W. Watson (2024). Spatial unit roots and spurious regression. *Econometrica 92*(5), 1661–1695.

Multhauf, R. P. (1976). Geology, chemistry, and the production of common salt. *Technology and Culture 17*(4), 634–645.

Multhauf, R. P. (1978). *Neptune's Gift: A History of Common Salt*. Baltimore: Johns Hopkins University Press.

Penman, H. L. (1948). Natural evaporation from open water, bare soil, and grass. *Proceedings of the Royal Society of London. Series A 193*, 120–145.

Pfister, C. (1981). The little ice age: Thermal and wetness indices for Central Europe. In R. I. Rotberg and T. K. Rabb (Eds.), *Climate and History*, Princeton, N.J., pp. 85–116. Princeton University Press.

Sands, T. and C. P. Higby (1949). France and the salt tax. *The Historian 11*(2), 145–165.

Schönholzer, D. and E. Weese (2026). Creative destruction in the European state system: 1000-2000.

Scott, J. C. (2017). *Against the Grain*. Princeton, NJ.: Princeton University Press.

Spruyt, H. (2017). War and state formation: Amending the bellicist theory of state making. In L. B. Kaspersen and J. Strandsbjerg (Eds.), *Does War Make States? Investigations of Charles Tilly's Historical Sociology*, Chapter 3, pp. 73–97. Cambridge University Press.

Stock, J. H. and M. Yogo (2002). Testing for weak instruments in linear iv regression. working paper t0284, NBER.

Tilly, C. (1985). Warmaking and statemaking as organized crime. In P. Evans, D. Rueschemeyer, and T. Skocpol (Eds.), *Bringing the State Back In*, Cambridge, UK, pp. 169–192. Cambridge University Press.

Tilly, C. (1990). *Coercion, Capital, and European States, AD 990-1990*. Oxford: Blackwell.

Waldinger, M. (2022). The economic effects of long-term climate change: Evidence from the Little Ice Age. *Journal of Political Economy 130*(9), 2275–2314.

Weisz, B. (2018). Royal revenues in the Árpádian age. In J. Laszlovszky, B. Nagy, P. Szabó, and A. Vadas

(Eds.), *The Economy of Medieval Hungary*, Leiden, The Netherlands, pp. 253–264. Brill.

Whaley, J. (2012). *Germany and the Holy Roman Empire 1493–1806*, Volume Vol 1. Oxford: Oxford University Press.

Wilson, P. H. (2009). *The Thirty Years War: Europe's Tragedy*. Cambridge MA: Harvard University Press.

Wilson, P. H. (2016). *The Holy Roman Empire: A Thousand Years of Europe's History*. London: Allen Lane.

Wilson, P. H. (2023). *Iron and Blood: A Military History of the German-Speaking Peoples since 1500*. Cambridge, MA: Belknap Press of Harvard University Press.

Woods, F. (1913). *The Influence of Monarchs: Steps in a New Science of History*. Macmillan.

Woods, F. A. (1906). *Mental and moral heredity in royalty: A statistical study in history and psychology*. H. Holt.

Appendices

Appendix Table of Contents

1 Theory Appendix

Proof of Lemma 1. The absorbing partition conditions are derived from the one-shot deviation principle. Since $\mathcal{P}^*$ is absorbing, territory $i \in j$ receives $\frac{E_j}{|j|}$ in every period, so its continuation value simplifies to $V_i(\mathcal{P}^*; \boldsymbol{\sigma}^*) = \frac{1}{1-\beta}\frac{E_j}{|j|}$ for all $i \in j \in \mathcal{P}^*$.

No-rebellion. Territory $i \in j$ does not rebel at $\mathcal{P}^*$ if and only if

$$\rho_{ij}(n_i)\left[\frac{E_j/|j|}{1-\beta} - e_i - \beta V_i(\mathcal{P}^*_{-i}; \boldsymbol{\sigma}^*)\right] \geq 0, \tag{NR}$$

where $\mathcal{P}^*_{-i}$ denotes the partition in which i has successfully seceded from j, and $\rho_{ij}(n_i) = \Pr[\delta \geq E_j - 2n_i]$ is the probability of successful rebellion, increasing in n_i.

No-war. Under equal sharing, the per-territory payoff of the merged polity $j \cup -j$ is the same regardless of which polity prevails, so ω_j drops out. Polity j does not initiate war against $-j$ if and only if

$$\frac{E_j/|j|}{1-\beta} \geq V_i(j \cup -j; \boldsymbol{\sigma}^*), \tag{NW}$$

where $V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is territory i's continuation value in the merged polity, which depends on its subsequent stability.

Part 1. Since $\Delta_i \equiv \frac{E_j/|j|}{1-\beta} - e_i - \beta V_i(\mathcal{P}^*_{-i}; \boldsymbol{\sigma}^*) < 0$ by assumption, condition (NR) reduces to $\rho_{ij}(n_i) \cdot \Delta_i \geq 0$, which holds if and only if $\rho_{ij}(n_i) = 0$. Now, $\rho_{ij}(n_i) = \Pr[\delta \geq E_j - 2n_i]$, where the argument follows from $f_i = n_i$ and $F_j = E_j - n_i$ at the absorbing partition. Since $E_j - 2n_i$ is strictly decreasing in n_i, $\rho_{ij}(n_i)$ is non-decreasing in n_i. Define

$$\bar{n}_i = \max\left\{n \geq 0 : \Pr[\delta \geq E_j - 2n] = 0\right\}.$$

The maximum is attained because $\Pr[\delta \geq E_j - 2n]$ is continuous in n (since δ has a density), so the set is closed. Then $\rho_{ij}(n_i) = 0$ if and only if $n_i \leq \bar{n}_i$, so (NR) is satisfied if and only if $n_i \leq \bar{n}_i$.

Part 2. Let $j' = j \cup -j$ denote the merged polity. For any territory $k \in -j$, the rebellion success probability in j' is $\rho_{kj'}(n_k) = \Pr[\delta \geq E_{j'} - 2n_k]$, which is non-decreasing in n_k by the same argument as Part 1. A successful rebellion by k removes n_k from j''s resource pool, reducing the

per-territory share available to $i \in j$. Since $\rho_{kj'}(n_k)$ is non-decreasing in n_k, the expected resource loss from post-annexation rebellions is non-decreasing in n_k, and therefore $V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is non-increasing in n_k for each $k \in -j$. It follows that $V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is non-increasing in $\sum_{k \in -j} n_k$. Define

$$\bar{N}_{-j} = \min \left\{ \sum_{k \in -j} n_k : V_i(j \cup -j; \boldsymbol{\sigma}^*) \leq \tfrac{E_j/|j|}{1-\beta} \right\}.$$

The minimum is attained because $V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is continuous in $\sum_{k \in -j} n_k$ (since δ has a density), so the set is closed. Then (NW) is satisfied whenever $\sum_{k \in -j} n_k \geq \bar{N}_{-j}$. □

Proof of Theorem 1. The result follows from two channels, each established in Lemma 1.

Rebellion channel. By Lemma 1(1), $\rho_{ij}(n_i)$ is non-decreasing in n_i. For any territory $i \in j$ with $\Delta_i < 0$, condition (NR) requires $\rho_{ij}(n_i) = 0$, i.e., $n_i \leq \bar{n}_i$. Since $n'_i \geq n_i$, any territory for which (NR) is binding under $\mathbf{n}$ will have its condition weakly tightened under $\mathbf{n}'$. Territories for which $n'_i > \bar{n}_i$ now strictly violate (NR), so any partition containing them in a multi-territory polity is non-absorbing under $\mathbf{n}'$.

War channel. By Lemma 1(2), $V_i(j \cup -j; \boldsymbol{\sigma}^*)$ is non-increasing in n_k for $k \in -j$. Since $n'_k \geq n_k$ for all k, condition (NW) is weakly easier to satisfy under $\mathbf{n}'$: any partition that was non-absorbing due to a profitable war under $\mathbf{n}$ may become absorbing under $\mathbf{n}'$, but only at a weakly more fragmented partition where war is no longer profitable.

Since both channels rule out absorbing partitions that are weakly less fragmented than those sustainable under $\mathbf{n}$, any absorbing partition under $\mathbf{n}'$ is weakly more fragmented than any absorbing partition under $\mathbf{n}$. □

Proof of Corollary 1. When $\mathbf{n} \geq \bar{\mathbf{n}}$, (NR) fails for every territory in any multi-territory polity: no polity with more than one territory can be absorbing. Since IK satisfies (NW) trivially — there are no resources to gain from annexing a singleton that will immediately rebel — IK is the unique absorbing partition. When $\mathbf{n} \leq \underline{\mathbf{n}}$, rebellion is infeasible everywhere ($\rho_{ij}(n_i) = 0$ for all i), so (NR) holds at any partition. (NW) fails at every partition coarser than Empire, so Empire is the only partition where no polity prefers to expand, making it the unique absorbing partition. For intermediate $\mathbf{n}$, both (NR) and (NW) can be jointly satisfied only at partitions strictly between IK and Empire. □

Proof of Proposition 1. Short-run profit is $\pi_i = P\,q_i^* - C_i(q_i^*;\theta)$. By the envelope theorem,

$$\frac{d\pi_i}{d\theta} = \underbrace{q_i^*\,\frac{dP^*}{d\theta}}_{\text{price gain}} - \underbrace{\frac{\partial C_i}{\partial \theta}}_{\text{cost increase}}.$$

Rock salt (R): $\partial C_R/\partial\theta = 0$ by Assumption 1, so $d\pi_R/d\theta = q_R^*(dP^*/d\theta) > 0$.

Brine springs (B): $d\pi_B/d\theta > 0$ follows directly from Assumption 2.

Coastal salt (C): Since B and C are both small-scale, low-capital operations, $q_C^* \approx q_B^*$. Yet $\partial C_C/\partial\theta \gg \partial C_B/\partial\theta$ (Assumption 1). The same price gain that covers B's small cost increase cannot cover C's much larger one:

$$\frac{\partial C_C}{\partial\theta} \gg \frac{\partial C_B}{\partial\theta} \approx q_C^*\,\frac{dP^*}{d\theta} \quad\Longrightarrow\quad \frac{d\pi_C}{d\theta} < 0.$$

□

2 History Appendix

2.1 The French Gabelle

The French *gabelle*, one of the most extensive salt monopolies in European history, might seem to challenge the claim that dispersed salt production is non-appropriable. France taxed salt heavily from the fourteenth century onward, relying on both brine and coastal/marine salt (Sands and Higby, 1949). If dispersed salt is so difficult to monopolize, how did the French crown manage it?

The answer is revealing: France tried harder than any other European state to appropriate dispersed salt revenues, and the result was enormous enforcement costs, endemic smuggling, and persistent regional variation reflecting the limits of state reach. Initially, there was no salt monopoly and salt production was widely dispersed. Only in mid-16th century did the sale of salt become a monopoly (Sands and Higby, 1949, 148). The gabelle as it was instituted by the 17th century was a tax on *consumption* at the point of sale—households were required to purchase a minimum quantity (*sel de devoir*) from royal warehouses (*greniers à sel*)—not a tax on production, because monitoring dispersed production sites was infeasible even for the most powerful monarchy in Europe (Adshead, 1992; Kwass, 2000).

Salt smuggling (*faux-saunage*) was endemic throughout the ancien régime. Salt was among the four most highly trafficked goods (Kwass, 2014, 88). This was due to the huge differences in the price of salt between say Brittany, where salt marshes were abundant and inland Maine, where salt cost 20-40 times more. Salt gangs routinely transported salt from where it was relatively untaxed and affordable to where it was scarce and highly taxed.

As a result, the gabelle required a vast enforcement apparatus: the *fermiers généraux* who purchased tax-farming contracts, armed guards (*gabelous*) who patrolled borders between tax zones, and a system of informers and inspections (Kwass, 2014).[22] The regional variation in the gabelle itself—the *grandes gabelles* of northern France, the *petites gabelles* of the south, the *pays rédimés* that had bought exemptions, and the *provinces franches* that paid no salt tax at all—reflects

[22]For example, "Often regular troops were used with the object of stopping smuggling, and frequently pitched battles took place between the smugglers and the salt guards. In the winter of 1708 it is reported that 27,000 or more royal troops were used in an effort to stop the smuggling of salt." (Sands and Higby, 1949, 156).

the state's inability to impose uniform monopolization across regions with different production technologies and political traditions. The contrast with rock salt is instructive: at Wieliczka, the Polish crown could station guards at the single mine entrance and capture revenue with minimal enforcement cost. The gabelle, by contrast, required an entire bureaucratic apparatus to approximate the same result—and even then, smuggling was pervasive and the tax was a major source of popular grievance that contributed to revolutionary unrest (Giommoni and Loumeau, 2026; Giommoni et al., 2026).

The gabelle therefore does not contradict the appropriability argument. We acknowledge that France did extract substantial gabelle revenue—roughly 6% of royal income for four centuries—so the effective n_i was not infinite. The model's n_i should be understood as *relative* to rock salt, not as a binary distinction between "appropriable" and "non-appropriable." The Polish crown could capture revenue with minimal enforcement cost (n_i close to zero). The gabelle, by contrast, required an entire bureaucratic apparatus—tax farmers, armed guards, border patrols, mandatory purchase quotas—to approximate the same result, and even then smuggling was pervasive and the tax was a major source of popular grievance (Giommoni and Loumeau, 2026; Giommoni et al., 2026). The relevant comparison is not whether France *could* tax dispersed salt, but the *cost* of doing so: the enormous enforcement costs, persistent smuggling, and regional variation are precisely what one would expect when a state attempts to appropriate resources with high n_i. It is this difference in appropriation costs—low for rock salt, high for dispersed brine and coastal/marine salt—that our model captures through the parameter n_i and that our empirical results confirm.

3 Empirical Appendix

A Data Description

A.1 Salt Mine Locations

We compile a dataset of 251 historical salt production sites across Europe, recording the name, geographic coordinates, extraction type, and period of activity. Our main sources include Adshead (1992) (*Salt and Civilization*), Multhauf (1978) (*Neptune's Gift*), and Beck (1993).[23] For the period a site is active, we draw on the periodization in Adshead (1992): *Dark and Light Ages* (500–1000 CE), *Middle Ages* (1000–1500 CE), and *Early Modernity* (1500–1800 CE). Each site is geocoded to latitude–longitude coordinates.

Type classification. We classify each site by extraction technology:

- **Rock salt** (27 sites): Underground mining of solid halite deposits. Capital-intensive, requiring organized labor and infrastructure, with output through a single observable chokepoint (the mine entrance). Examples: Wieliczka-Bochnia (Poland), Hallstatt and Hallein (Austria), Cardona (Catalonia).

- **Brine evaporation** (75 sites): Extraction from inland springs or wells, concentrated by boiling over wood or coal fires. Geographically dispersed, with many small springs each operated by a local producer. Examples: Schwäbisch Hall and the Muschelkalk salines (Germany), Droitwich and the Cheshire wiches (England), Halle (Saxony).

- **Coastal/marine evaporation** (91 sites): Evaporation of seawater or lake brine in shallow coastal ponds. Low capital requirements, low barriers to entry, diffuse coastline production. Examples: Ibiza, Trapani, Peccais, Bay of Bourgneuf, Crimean salt flats.

[23] The dataset was compiled and coded by Patrick Fitzsimmons, Mark Koyama, and Kyle Raines, with type classifications verified against the primary sources. A detailed data appendix documenting each site, its classification, and the rationale for inclusion or exclusion is available from the authors. The dataset evolved from an initial 270 locations to 251 after removing duplicates (e.g., multiple entries for the Salins-les-Bains complex), consolidating satellite operations (e.g., Weissbach merged into Niedernhall), and dropping sites with no confirmed production during the 1100–1790 sample period.

An additional 61 sites remain unclassified (recorded only as "Saline"). Nine sites have multiple classifications and contribute to both type-specific counts. Appendix Section O tests sensitivity to alternative assignments of unclassified sites.

Salt count measures. We construct three measures: (i) the *static* measure counts all sites ever active within each unit; (ii) the *temporal* measure restricts to sites with documented start dates on or before the observation year; (iii) the *inclusive temporal* measure treats missing start dates as active from 1100. Because 87% of salt variance is between units, the temporal measure is our preferred polity-level specification. In the grid-cell analysis, we construct type-specific counts (rock, brine, coastal/marine) with corresponding interaction terms.

A.2 Political Boundaries

We use Abramson (2017) for political boundaries at five-year intervals from 1100 to 1790 (135 periods)[24], based on the *Centennia Historical Atlas* and *Euratlas*. We standardize polity names across periods following Woods (1906, 1913).[25] The final panel contains 30,712 polity-year observations spanning 631 unique polities. For the grid-cell analysis, we overlay a regular grid of 200km $\times$ 200km cells on Europe in an Albers equal-area conic projection and count the number of distinct polity borders passing through each cell in each period. We additionally compute mean polity area and HHI concentration for each cell-year from the cell–polity crosswalk.

Figure A1 illustrates the political geography of Europe at four dates spanning the sample period. The number of polities roughly doubles from 112 in 1100 to 253 by 1300, coinciding with the onset of European cooling. The dense cluster of borders in central Europe—corresponding to the Holy Roman Empire, where brine salt is concentrated—persists across all four centuries even as peripheral regions (France, England, Iberia, Russia) consolidate into larger states.

[24]Four five-year cross-sections (1450, 1455, 1460, and 1465) are missing from the source data, so the panel contains 135 rather than 139 periods between 1100 and 1790 ($476 \times 135 = 64{,}260$ cell-period observations).

[25]This standardization fixes variations (Castille/Castile) and connects historical polities to successor states (Muscovy $\rightarrow$ Russia, Habsburgs $\rightarrow$ Austria).

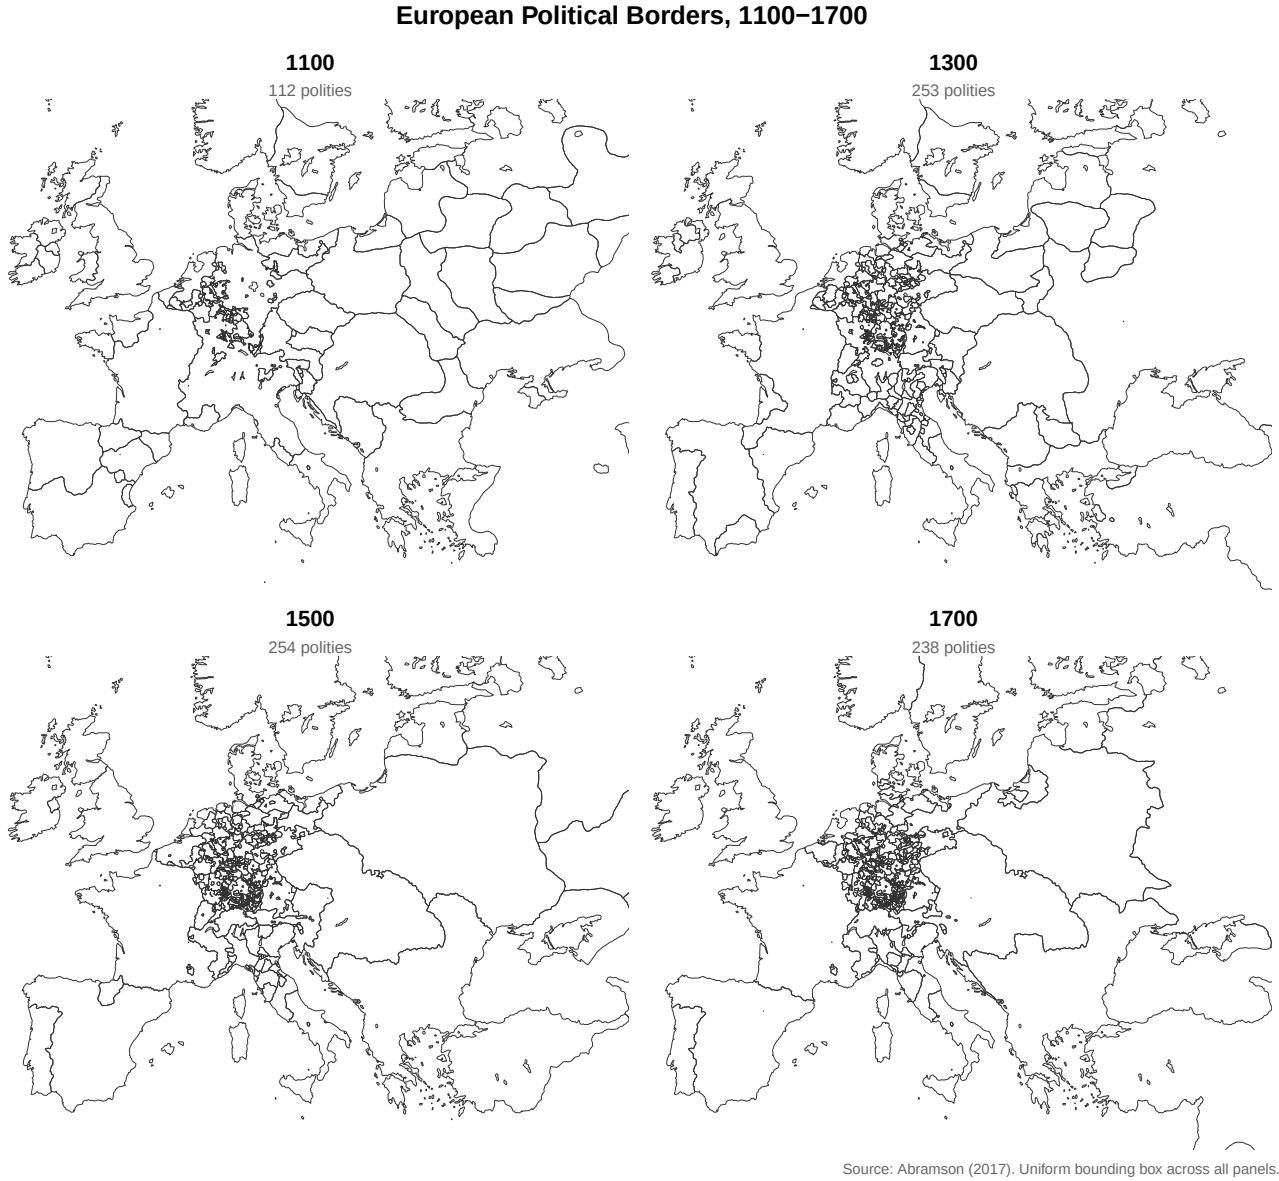

Figure A1: European Political Borders, 1100–1700

Notes: Political boundaries from Abramson (2017) at four dates spanning the sample period. All panels use the same bounding box for visual comparison. The number of polities roughly doubles from 1100 (112 polities) to 1300 (253 polities), then remains elevated through 1700 (238 polities). The persistent concentration of small polities in central Europe corresponds to the Holy Roman Empire, where dispersed brine salt production sustained political fragmentation.

A.3 Battle Data

Battle data come from Kitamura (2021) and Wikipedia, comprising 538 geocoded battles from 1302 to 1700, matched to grid cells by point-in-polygon; 435 are traditional (interstate) battles falling within the grid. For the polity panel, we link battles to polities through belligerent names, standardized using the same conventions as the boundary data. Battles are classified as traditional (interstate) or civil (same polity on both sides). As a robustness check, we use the Historical Conflict Event Dataset (HCED) compiled by Miller (Harvard Dataverse), which contains 892 battles over 1300–1700 (Appendix Section P; cell-year correlation $r = 0.68$).

A.4 Temperature

Our primary source is the gridded reconstruction of Guiot and Corona (2010), providing growing-season (April–September) temperature anomalies from 600 to 2007 at 5° resolution, covering 412 of 476 grid cells[26]. The reconstruction draws on tree rings, historical documents, pollen series, and ice-core isotopes. We compute the spatial mean temperature within each grid cell or polity boundary and construct lagged rolling means at 5, 10, 15, 20, 25, 30, and 50 years. For robustness, we employ CCSM4 and MPI-ESM paleoclimate simulations (Appendix Section B).

A.5 Geographic Controls and Urbanization

Geographic controls include: caloric suitability (Galor and Ömer Özak, 2016); soil nutrient availability and workability (Fischer et al., 2008); elevation and terrain ruggedness index (Farr et al., 2007); and distance to navigable rivers. Each control is computed as the spatial mean over the polity or cell area. Since these are time-invariant at the grid-cell level, we interact them with period fixed effects. Urbanization data are city-level population figures for 2,262 European settlements from Buringh (2021), aggregated to grid cells by summing populations of all cities within each cell.

A.6 Summary Statistics

Tables A1 and A2 report summary statistics for the grid-cell and polity-level analyses, respectively.

[26] Cells outside the reconstruction's spatial domain drop from the temperature specifications.

Table A1: Summary Statistics: Grid-Cell Panel

Variable	Obs	Mean	SD	Min	Max
Panel A: Outcomes and salt					
Border count	64,260	1.882	4.418	0	76
Battle count (interstate)	64,260	0.007	0.119	0	13
Civil war count	64,260	0.000	0.010	0	1
Salt mine count	64,260	0.334	1.075	0	8
Rock salt count	64,260	0.027	0.198	0	2
Brine salt count	64,260	0.122	0.653	0	7
Coastal/marine salt count	64,260	0.099	0.477	0	6
Panel B: Temperature (5-year rolling mean)					
Guiot–Corona (5°)	39,150	−0.066	0.198	−0.696	0.878
Guiot–Corona (interpolated)	55,208	−0.061	0.384	−2.948	1.979
CCSM4	39,555	−0.459	0.577	−4.168	1.691
MPI-ESM	39,555	−0.493	0.503	−3.425	1.660
Panel C: Controls					
Urban population (000s)	4,760	20.704	60.545	0	865

Notes: 476 grid cells (200km) observed at five-year intervals, 1100–1790 (135 periods). Guiot–Corona temperature covers 412 of 476 cells; CCSM4 and MPI-ESM cover a similar subset (39,555 obs). Guiot–Corona (interpolated) covers 411 cells (55,208 obs). Battle data available 1302–1700. Urban population available for 35 cells with nonzero Buringh data. 70 cells contain at least one salt mine; 406 do not.

B Temperature Data

We use three independent temperature reconstructions to ensure that our results are not driven by the choice of climate data. Figure A2 displays the mean temperature anomaly across all grid cells for each source over the sample period.

- **Guiot–Corona** (Guiot and Corona, 2010): A multi-proxy reconstruction of European growing-season temperatures based on pollen, tree-ring, and documentary evidence. This is our primary temperature source, available for 412 of our 476 grid cells. The series clearly displays the Medieval Climate Anomaly (warmer conditions through circa 1300) followed by the Little Ice Age (cooler conditions, roughly 1400–1850).

- **CCSM4**: Output from the Community Climate System Model 4 (NCAR), a coupled general circulation model simulation of the last millennium. Covers all 476 cells.

Table A2: Summary Statistics: Polity-Level Panel

Variable	Obs	Mean	SD	Min	Max
Panel A: Outcomes and salt					
Battle count (interstate)	30,712	0.022	0.256	0	18
Civil war count	30,712	0.002	0.051	0	3
log(Salt mines + 1), temporal	30,712	0.201	0.608	0	4.522
log(Salt mines + 1), static	30,712	0.258	0.693	0	4.564
Polity area (km^2)	30,712	45,681	329,000	0	7,250,000
Panel B: Temperature and geography					
Temperature (Guiot–Corona, 5yr)	30,561	−0.183	0.173	−0.656	0.603
Caloric suitability	30,712	3,139	536	0	4,235
Mean elevation (m)	30,712	352	287	2	2,053
Terrain ruggedness	30,712	28.79	32.53	0.06	244
Soil workability	30,712	1.813	0.644	0	4.500
Soil nutrient content	30,712	1.578	0.609	0	4.500
Distance to rivers (km)	30,712	23.04	113.5	0	1,546

Notes: Polity-year panel observed at five-year intervals, 1100–1790 (30,712 observations after removing 28 duplicate polity-year records). The temporal salt measure activates mines at their known production start dates (34 mines turn on circa 1500). Temperature from the Guiot–Corona reconstruction (Guiot and Corona, 2010). Geographic controls aggregated from grid-cell data to polity boundaries using area-weighted means. Battle data available 1302–1700.

- **MPI-ESM**: Output from the Max Planck Institute Earth System Model, another coupled general circulation model run over the last millennium. Covers all 476 cells.

All three series display the characteristic cooling trend associated with the Little Ice Age, though they differ in amplitude and timing of specific fluctuations.

C Mineral Placebo Test

A concern is that the salt × temperature interaction captures a generic effect of valuable underground resources during cold periods, rather than anything specific to salt. We test this using data on other mineral deposits—iron, copper, silver, gold, and lead—from the USGS Mineral Resources Data System (MRDS). These minerals were economically important in premodern Europe, are extracted from underground deposits, and—crucially—their production is *not* climate-sensitive. If our identification is valid, mineral × temperature should be null.

We identify 798 mineral deposit sites across 172 European grid cells (compared with 70 salt

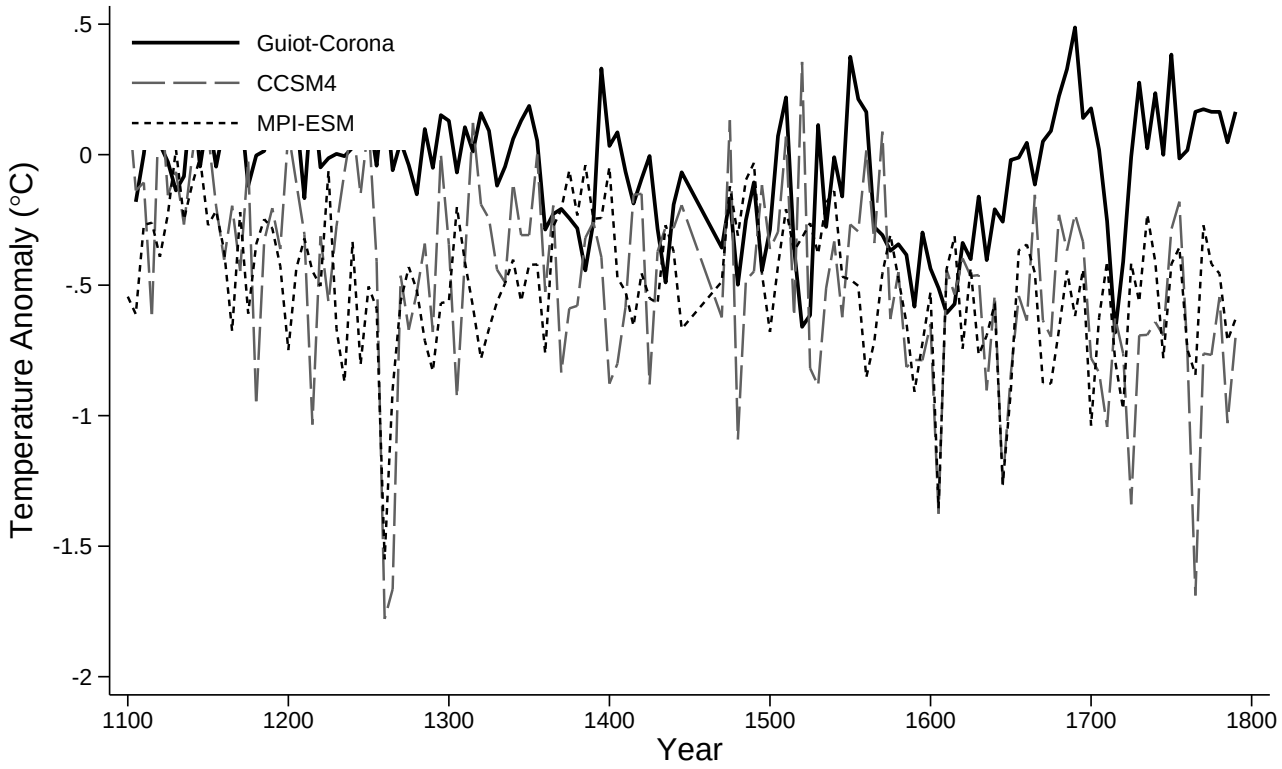

Figure A2: Temperature Anomalies Over the Sample Period, by Source

Notes: Each line shows the mean 5-year rolling temperature anomaly (°C) across all available grid cells for each reconstruction. Guiot–Corona (Guiot and Corona, 2010) covers 412 cells; CCSM4 and MPI-ESM cover all 476 cells.

cells). The salt–mineral correlation at the cell level is moderate ($r = 0.31$): 24 of 70 salt cells also contain iron deposits, concentrated in the Lorraine/Saar/Rhine region of Western Europe.

Table A3 presents the placebo results. Columns (1)–(2) show the salt × temperature and mineral × temperature interactions entered separately. When run alone (column 2), the mineral × temperature coefficient is negative (−0.059) but statistically insignificant ($p = 0.10$) with Conley standard errors. Columns (3)–(5) enter both interactions simultaneously at increasing temperature windows. The mineral × temperature coefficient is economically zero and statistically insignificant at all windows in the horse race (−0.011 to +0.011, all $p > 0.77$), while the salt coefficient retains its sign, magnitude, and monotonic window pattern (−0.712 at 10 years to −1.149 at 50 years). The salt coefficient is not individually significant in the horse race with spatial standard errors,

reflecting limited statistical power when two collinear regressors (24 of 70 salt cells also contain mineral deposits) are entered simultaneously.

The key diagnostic is not whether salt retains significance conditional on minerals, but whether minerals show *any* independent temperature channel. They do not. For salt, longer windows produce larger coefficients—consistent with sustained climate shifts driving the mechanism. For minerals, the coefficient shrinks toward zero and eventually turns positive—inconsistent with any structural climate channel and consistent instead with the standalone result being driven by geographic overlap between mineral and salt cells. Once salt is controlled for directly, minerals contribute no independent temperature-driven fragmentation effect.

Table A3: Mineral Placebo: Salt vs. Other Minerals × Temperature

	Dependent Variable: Border Count				
	(1) Salt 10yr	(2) Mineral 10yr	(3) Horse Race 10yr	(4) Horse Race 25yr	(5) Horse Race 50yr
Salt × Temp	−0.724∗		−0.712	−0.948	−1.149
	(0.408)		(0.440)	(0.602)	(0.769)
Mineral × Temp		−0.059	−0.011	−0.006	+0.011
		(0.036)	(0.039)	(0.053)	(0.069)
Cell FE	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes

Notes: 476 grid cells, five-year intervals. Mineral count is the number of USGS MRDS mineral deposits (iron, copper, silver, gold, lead) within each cell; 172 cells contain at least one mineral deposit. Columns (1)–(2) enter salt × temperature and mineral × temperature separately. Columns (3)–(5) enter both simultaneously at different temperature windows (10, 25, 50 years). Temperature is the Guiot–Corona rolling mean. All specifications include cell and period fixed effects. Conley (1999) spatial HAC standard errors (100km bandwidth) in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

D Leave-One-Out Robustness

Table A4 reports the *type-specific* salt × temperature coefficients (rock, brine, and coastal/marine; 25-year Guiot–Corona window, cell and period fixed effects) when major salt-producing clusters are dropped one at a time. The opposing signs are stable: dropping any single cluster leaves the brine coefficient negative and the rock coefficient positive. The Lorraine cluster (cells 265–266, 7 brine sites each, in the heartland of the Holy Roman Empire) is the most influential—excluding

it attenuates the brine coefficient from −2.63 to −0.71 (still significant at $p < 0.1$) and the rock coefficient from +1.32 to +0.35, consistent with Lorraine's position as the densest brine region in the most politically fragmented part of Europe. Even without Lorraine, the opposing-sign pattern survives; its type-specific consequences are examined in detail in Appendix Section E.

Table A4: Leave-One-Cluster-Out: Type-Specific Salt × Temperature (25yr, Cell + Year FE)

Dropped cluster	Rock × Temp	Brine × Temp	Coastal/marine × Temp
None (baseline)	+1.319* (0.720)	−2.633** (1.214)	+0.205 (0.178)
Adriatic	+1.335* (0.730)	−2.660** (1.223)	+0.208 (0.181)
Alpine (Hallstatt)	+1.318* (0.799)	−2.634** (1.214)	+0.205 (0.178)
Atlantic France (Bourgneuf)	+1.353* (0.724)	−2.664** (1.219)	+0.102 (0.131)
Cheshire	+1.436* (0.761)	−2.853** (1.262)	+0.222 (0.187)
Lorraine	+0.353 (0.386)	**−0.712*** (0.402)	+0.052 (0.098)
Lüneburg	+1.220* (0.699)	−2.643** (1.215)	+0.196 (0.177)
Mediterranean	+1.316* (0.725)	−2.620** (1.224)	+0.209 (0.182)
SW Germany (Muschelkalk)	+1.269** (0.566)	−2.690** (1.352)	+0.188 (0.187)
Transylvania	+1.380* (0.804)	−2.637** (1.217)	+0.204 (0.177)
Wieliczka	+1.308 (0.885)	−2.649** (1.222)	+0.209 (0.178)

Notes: Each row drops all grid cells in the named salt cluster and re-estimates the type-specific specification of Table 2 (all three type × temperature interactions entered simultaneously, 25-year window). Clusters are defined by geographic proximity of salt sites (cell IDs per the replication code). Lorraine (cells 265–266) is the densest brine region in the sample. Standard errors clustered by grid cell in parentheses (equivalently, Conley spatial-HAC at the 100km bandwidth on this 200km grid), consistent with Table 2. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E Type-Specific Robustness to Lorraine Exclusion

The leave-one-out analysis in Table A4 shows that the Lorraine cluster (cells 265–266, containing 7 salt sites each—all brine) is the most influential region for the aggregate salt × temperature coefficient. This is unsurprising: Lorraine contains 12 of the 75 brine sites in the dataset and sits at the junction of France, the Holy Roman Empire, and Burgundy—the most politically fragmented area of medieval Europe. The question is whether the *type-specific* opposing signs, the paper's central finding, survive its exclusion.

Table A5 reports the full type-specific decomposition with and without Lorraine. Two findings are noteworthy. First, the opposing signs survive at all temperature windows: rock × temperature remains positive and brine × temperature remains negative. Under our preferred cluster-by-cell

standard errors, the brine coefficient remains significant at the ten percent level at every window; the rock coefficient—identified off only ten rock cells—retains its sign and monotonicity but loses statistical significance. Second, the monotonic increase with window length is preserved: brine coefficients grow from −0.49 (10yr) to −0.71 (25yr) to −0.92 (50yr), and rock from +0.21 to +0.35 to +0.52.

Magnitudes attenuate by approximately 70%—as expected when the densest cluster is removed from any regression. The relevant test is not whether the magnitude is preserved but whether the *pattern* survives: the opposing signs by type and the monotonic window effects both do, and the brine coefficient remains statistically distinguishable from zero. As we discuss in Section 6, the concentration of the brine effect in Lorraine and the surrounding lands of the Holy Roman Empire is consistent with the mechanism: this was the region where dispersed, non-appropriable salt wealth met the political conditions under which it could sustain fragmentation.

Table A5: Type-Specific Salt × Temperature: Baseline vs. Excluding Lorraine

	Baseline (Full Sample)			Excluding Lorraine		
	10yr	25yr	50yr	10yr	25yr	50yr
Panel A: Dependent Variable = Border Count						
Rock × Temp	+0.897*	+1.319*	+1.777*	+0.214	+0.353	+0.523
	(0.521)	(0.720)	(0.938)	(0.295)	(0.386)	(0.469)
Brine × Temp	−1.902**	−2.633**	−3.352**	−0.493*	−0.712*	−0.923*
	(0.923)	(1.214)	(1.488)	(0.280)	(0.402)	(0.556)
Coastal/marine × Temp	+0.128	+0.205	+0.316	+0.005	+0.052	+0.150
	(0.134)	(0.178)	(0.225)	(0.080)	(0.098)	(0.117)
Panel B: Aggregate Salt × Temp → Borders						
Borders	−0.724*		−1.137	−0.191**		−0.180
	(0.408)		(0.710)	(0.091)		(0.146)
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Cells	476	476	476	474	474	474

Notes: Panel A enters all three type × temperature interactions simultaneously with cell and period fixed effects. Panel B reports the aggregate salt × temperature interaction. The Lorraine cluster comprises cells 265 and 266, each containing 7 salt deposits (all brine). These two cells account for 12 of the 75 brine sites in the dataset. Standard errors clustered by grid cell in parentheses (equivalently, Conley spatial-HAC at the 100km bandwidth on this 200km grid), consistent with Table 2. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

F Spatial Standard Errors

The main tables report standard errors clustered by grid cell; on this 200km grid these coincide with Conley (1999) spatial HAC standard errors at a 100km cutoff (uniform kernel, allowing unrestricted serial correlation within cell), since no two cell centroids lie within 100km of one another. Table A6 shows how inference changes at wider spatial bandwidths (200km, 500km), which admit correlation across neighboring cells.

Panel A shows the aggregate salt × temperature specification on borders. Conley standard errors move only modestly with the bandwidth (e.g., from 0.556 at 100km to 0.578 at 500km for the 25-year window), and the coefficient (-0.955) remains significant at the ten percent level out to 500km; the 50-year coefficient is marginal ($p = 0.11$) at the widest bandwidth.

Panel B shows the type-specific temperature interactions on borders: even at 500km—a bandwidth that encompasses most of central Europe—the brine × temperature coefficient (-2.633) remains significant at the five percent level (Conley SE = 1.15, $p = 0.022$), and rock at the ten percent level. Panel C repeats the exercise for the complementary HHI concentration measure (on which consolidation is a *negative* coefficient and fragmentation positive). The opposing-sign pattern—rock and coastal/marine consolidate, brine fragments—survives the spatial-HAC correction at every bandwidth, with all three coefficients remaining significant even at 500km.

G Spatial Autocorrelation and Spatial Unit Roots

A growing literature warns that spatial autocorrelation in regression residuals can severely distort inference in cross-sectional persistence regressions (Kelly, 2019; Conley and Kelly, 2025). Müller and Watson (2024) develop a formal test for spatial unit roots and propose a GLS transformation that removes spatial non-stationarity before estimation; Becker et al. (2026) provide a Stata implementation (the SPUR package). We apply both diagnostics to our grid-cell panel.

Spatial autocorrelation. We compute Moran's I on cell-and-year-FE residuals at each of the 135 five-year cross-sections using $k = 8$ nearest neighbors. For border counts, the average Moran's I is 0.378 (mean $z = 18.9$, significant in all 135 periods). For battle counts, the average is 0.208

Table A6: Conley Spatial HAC Standard Errors

	Coefficient	Standard Errors			
		IID	Conley 100km	Conley 200km	Conley 500km
Panel A: Salt × Temperature → Borders					
10-year window	-0.724^{*}	0.018	0.407	0.427	0.434
25-year window	-0.955^{*}	0.022	0.556	0.584	0.578
50-year window	-1.137	0.025	0.710	0.749	0.713
Panel B: Type × Temperature → Borders (25yr)					
Rock × Temp	$+1.319^{*}$	0.126	0.720	0.835	0.768
Brine × Temp	-2.633^{**}	0.038	1.214	1.368	1.147
Coastal/marine × Temp	$+0.205$	0.039	0.178	0.196	0.177
Panel C: Type × Temperature → HHI Concentration (25yr)					
Rock × Temp	-0.094^{*}	0.012	0.047	0.045	0.050
Brine × Temp	$+0.032^{**}$	0.004	0.011	0.015	0.014
Coastal/marine × Temp	-0.024^{**}	0.004	0.010	0.009	0.012

Notes: All specifications include cell and period fixed effects (implemented via residualization for the Conley estimator). Conley (1999) spatial HAC standard errors with a uniform kernel at the indicated distance cutoff in kilometers, allowing unrestricted serial correlation within cell. On this 200km grid the 100km column coincides with clustering by grid cell and hence with the standard errors of the main tables. Panel A reports the aggregate salt × temperature interaction on borders at three windows. Panel B reports the type-specific interactions on borders (25yr), entered simultaneously. Panel C reports the type-specific interactions on HHI polity concentration (25yr); on HHI, consolidation is a *negative* coefficient and fragmentation positive (rock and coastal/marine consolidate, brine fragments). At the most conservative 500km bandwidth, the aggregate interaction is significant at 10% at the 10- and 25-year windows, brine at 5%, rock at 10%, and coastal/marine is insignificant; the HHI coefficients remain significant at the 500km bandwidth (rock at 10%, brine and coastal/marine at 5%). Stars on each coefficient reflect the most conservative (500km) Conley SE.

(mean $z = 10.5$, significant in 96 of 135 periods). By the benchmarks of Kelly (2019), both are non-trivial, confirming the importance of spatial standard error corrections—which we address throughout the paper with Conley (1999) spatial HAC standard errors.

However, the concerns raised by Kelly (2019) target *cross-sectional* persistence regressions in which a time-invariant historical variable predicts a time-invariant contemporary outcome. In such settings, any two spatially smooth variables will correlate, and high Moran statistics indicate that the effective sample size is far smaller than the nominal one. Our setting differs in three respects. First, we estimate a *panel* with cell fixed effects, so the spatial levels are already absorbed; the Moran statistics reported above reflect autocorrelation in within-cell deviations, not in cross-sectional levels. Second, our coefficient of interest is identified from a *time-varying* interaction (salt

× temperature), not from a time-invariant regressor. Spatial noise cannot generate a spurious salt × temperature interaction the way it can generate a spurious correlation between two spatially smooth cross-sectional variables. Third, our type-specific results produce *opposing signs* from the same temperature shock—rock × temperature is positive while brine × temperature is negative—which a spatially smooth confounder would be unable to generate.

SPUR correction. We apply the Müller and Watson (2024) spatial GLS (SPUR) transformation, implemented via Becker et al. (2026), to all main specifications. The transformation is applied year-by-year to each cross-section of ∼476 cells, projecting out the leading eigenvectors of the spatial correlation matrix before estimation. Table A7 reports the results.

For *borders*, the SPUR correction strengthens the results. All type-specific temperature interactions retain their opposing signs at $p < 0.001$: rock × temperature is positive (+0.54, +0.77, +1.01 at 10, 25, 50-year windows) and brine × temperature is negative (−0.96, −1.35, −1.76). Coefficients are approximately half the OLS magnitude, reflecting the removal of spatially persistent variation, but standard errors shrink proportionally. The aggregate salt × temperature coefficient is also highly significant at all windows ($p < 0.001$). The monotonic increase with window length is preserved.

H Poisson Estimation

Border counts and battle counts are non-negative integers. While OLS is standard in the literature, we verify that our results are robust to Poisson estimation, which respects the count nature of the data. Table A8 reports conditional fixed-effects Poisson estimates alongside the baseline OLS for comparison. Because Poisson coefficients are semi-elasticities (the conditional mean is $\exp(X\beta)$), magnitudes differ from OLS, but signs and significance are the key comparison.

All results are qualitatively unchanged. For the aggregate border specification (Panel A), the salt × temperature interaction is negative and highly significant under Poisson at all three windows (−0.061 to −0.098, all $p < 0.001$ with standard errors clustered by cell). The type-specific border results (Panel B) preserve the pattern: brine is strongly negative (−0.120, $p < 0.001$), while rock is positive but—consistent with the small number of rock cells—not significant under clustering

Table A7: SPUR Spatial Unit Root Correction (Müller and Watson 2024; Becker et al. 2026)

	Borders		
	10yr	25yr	50yr
Panel A: Aggregate Salt × Temperature			
Salt × Temp	−0.240 ∗ ∗		−0.371 ∗ ∗
	(0.046)		(0.074)
Panel B: Type-Specific			
Rock × Temp	+0.540 ∗ ∗	+0.768 ∗ ∗	+1.010 ∗ ∗
	(0.099)	(0.126)	(0.156)
Brine × Temp	−0.957 ∗ ∗	−1.347 ∗ ∗	−1.756 ∗ ∗
	(0.154)	(0.198)	(0.249)
Coastal/marine × Temp	+0.062 ∗ ∗	+0.082 ∗ ∗	+0.114 ∗ ∗
	(0.016)	(0.021)	(0.025)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes

Notes: All specifications residualize on cell and five-year period fixed effects, then apply the Müller and Watson (2024) spatial GLS transformation (Becker et al. 2026 implementation) year-by-year to each cross-section of ∼476 cells. Heteroskedasticity-robust standard errors in parentheses (post-transformation). The transformation removes spatial non-stationarity by projecting out leading eigenvectors of the spatial correlation matrix. Coefficients are smaller in magnitude than OLS because the transformation absorbs spatially persistent variation, but all border results retain significance. ∗ ∗ ∗ $p < 0.01$, ∗∗ $p < 0.05$, ∗ $p < 0.1$.

(+0.097, $p \approx 0.23$), and coastal/marine is positive and insignificant (+0.015). The complementary HHI concentration results (Panel C) likewise survive Poisson estimation, with the opposing signs preserved: rock and coastal/marine consolidate (negative on the HHI scale) and brine fragments (positive).

I Spatial Spillovers

A cell's border count depends not only on its own resources but also on those of its neighbors: if a neighboring cell contains contested salt deposits, political boundaries may spill into adjacent cells. To test whether the *type-specific* results are driven by own-cell salt or by neighboring cells' salt, we construct a spatial lag of each type count—the sum of rock, brine, and coastal/marine counts in all cells whose centroids lie within 300km (queen contiguity on the 200km grid)—and enter own-cell and neighbor-cell type × temperature interactions in a horse race.

Table A8: OLS vs. Poisson Fixed Effects

	OLS (Cell FE)		Poisson (Cond. FE)	
	Coeff	SE	Coeff	SE
Panel A: Salt × Temperature → Borders				
10-year window	-0.724^{*}	(0.408)	-0.061^{***}	(0.017)
25-year window	-0.955^{*}	(0.556)	-0.080^{***}	(0.022)
50-year window	-1.137	(0.710)	-0.098^{***}	(0.026)
Panel B: Type × Temperature → Borders (25yr)				
Rock × Temp	$+1.319^{*}$	(0.720)	$+0.097$	(0.082)
Brine × Temp	-2.633^{**}	(1.214)	-0.120^{***}	(0.011)
Coastal/marine × Temp	$+0.205$	(0.178)	$+0.015$	(0.030)
Panel C: Type × Temperature → HHI Concentration (25yr)				
Rock × Temp	$-0.094**$	(0.047)	$-0.143*$	(0.076)
Brine × Temp	$+0.032***$	(0.011)	$+0.054*$	(0.029)
Coastal/marine × Temp	$-0.024**$	(0.010)	$-0.034**$	(0.016)

Notes: OLS columns reproduce the baseline cell + period FE specification. Poisson columns report conditional fixed-effects Poisson (Hausman, Hall, and Griliches 1984) estimates; cells with all-zero outcomes are dropped by the conditional likelihood. Poisson coefficients are semi-elasticities: a unit increase in salt × temperature changes $\log E[Y]$ by the reported amount. All specifications include period fixed effects. Panel C uses HHI polity concentration as the outcome (on which consolidation is a negative coefficient and fragmentation positive); its Poisson column models 100×HHI as a quasi-count, and Panel C standard errors are clustered by grid cell. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Table A9 reports the results (border count, 25-year window). The own-cell brine coefficient (-1.96) is roughly eight times the neighbor-cell brine coefficient (-0.24), and the own-cell rock coefficient remains positive, confirming that the opposing-sign mechanism operates primarily through a cell's own salt resources rather than through proximity to contested neighbors. With spatial-HAC standard errors the own-cell brine coefficient is estimated less precisely in the horse race, but its magnitude continues to dominate the neighbor effect.

J Holy Roman Empire Institutional Interaction

As discussed in Section 6, the brine fragmentation effect is concentrated in central Europe, particularly in the decentralized Holy Roman Empire. Table A10 formalizes this by interacting salt × temperature with an HRE indicator. We classify cells as HRE if their dominant polity in 1400 was an HRE member state (Austria, Bavaria, Saxony, Brandenburg, Bohemia, Hesse, Thuringia,

Table A9: Spatial Spillovers: Own vs. Neighbor Type-Specific Salt × Temperature

	Rock × Temp	Brine × Temp	Coastal/marine × Temp
(1) Own-cell salt only			
Own salt × Temp	+1.319*	−2.633**	+0.205
	(0.720)	(1.214)	(0.178)
(2) Own + neighbor salt (horse race)			
Own salt × Temp	+1.517*	−1.961	+0.201
	(0.907)	(1.297)	(0.172)
Neighbor salt × Temp	−0.267	−0.244*	+0.057
	(0.347)	(0.142)	(0.055)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes

Notes: 476 grid cells, five-year intervals; border count, 25-year temperature window. Neighbor type counts are the sum of each type's count in all cells with centroid within 300km (queen contiguity). Panel (1) enters the three own-cell type × temperature interactions; panel (2) adds the three neighbor-cell interactions in a horse race. All specifications include cell and period fixed effects. Conley (1999) spatial HAC standard errors (100km) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Württemberg, Baden, the Palatinate, the ecclesiastical territories of Cologne, Mainz, Trier, and Münster, Lüneburg, Mecklenburg, Friesland, Burgundy, and the HRE itself, among others; the full list is in the replication code). This yields 20 HRE cells out of 476.

Column (1) reproduces the baseline. Column (2) splits the aggregate interaction by institutional context: the point estimate is −2.133 in HRE cells versus −0.096 outside—a 22-fold difference in magnitude—though with only 20 HRE cells neither arm is precisely estimated under clustering. The brine decomposition is sharper. Columns (3)–(4) show that brine × temperature is −3.848 ($p = 0.026$) in the HRE versus −0.285 (not significant) outside, and the sample split confirms this: −3.478 ($p = 0.045$) within the Empire against −0.277 (not significant) elsewhere.

The time trend results (columns 5–6) point the same way. Salt cells *fragment* over time within the HRE (+0.488 borders per century, $p = 0.023$) but *consolidate* outside it (−0.045, $p = 0.001$). The brine trend within the Empire is similarly positive (+0.773, $p = 0.027$); the non-HRE brine trend is negative but imprecise.

These findings are consistent with the model's prediction that the political consequences of non-appropriable resources depend on institutional context. In the decentralized HRE, where

$E_j/|j|$ was low (many small polities), dispersed brine revenues sustained local lords' independence and the rebellion threshold $\bar{n}_i$ was easily exceeded. In centralized monarchies (France, England, Aragon), even dispersed salt revenues could be captured by a strong crown, so the appropriability distinction mattered less. We view this concentration as the mechanism operating where it should: it is within the decentralized Empire that non-appropriable salt wealth had the greatest scope to sustain small polities (see the discussion in Section 6).

Table A10: Holy Roman Empire Institutional Interaction

	Temperature Interaction (25yr)				Time Trend	
	(1) Baseline	(2) Salt×HRE	(3) Brine×HRE	(4) Splits	(5) Salt×HRE	(6) Brine×HRE
Salt×T	−0.955*					
	(0.556)					
Salt×T×HRE		−2.133				
		(1.404)				
Salt×T×non-HRE		−0.096				
		(0.080)				
Brine×T×HRE			−3.848**			
			(1.725)			
Brine×T×non-HRE			−0.285			
			(0.205)			
Brine×T (HRE)				−3.478**		
				(1.732)		
Brine×T (non-HRE)				−0.277		
				(0.213)		
Salt×Trend×HRE					+0.488**	
					(0.214)	
Salt×Trend×non					−0.045***	
					(0.014)	
Brine×Trend×HRE						+0.773**
						(0.350)
Brine×Trend×non						−0.035
						(0.024)
Cell FE	Yes	Yes	Yes	—	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable: border count. HRE cells are those whose dominant polity in 1400 was a Holy Roman Empire member state (20 of 476 cells). Column (4) reports separate sample splits (HRE: $N = 2{,}080$; non-HRE: $N = 51{,}480$). All other columns use the full sample with cell and period fixed effects. T = temperature (25-year Guiot–Corona rolling mean). Trend = (year − 1100)/100. "non" = non-HRE. Standard errors clustered by grid cell in parentheses, consistent with Table 2; the column (4) HRE split clusters on only 20 cells, so its p-value should be read with caution. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

K Alternative Grid Size

The 200km grid is arbitrary. Table A11 re-estimates the main specifications using a 150km grid (787 cells, 108 with salt). The results are qualitatively unchanged: the aggregate salt × temperature interaction is negative and significant at the five percent level at all windows on the finer grid (indeed stronger than at 200km, where the 50-year window is imprecise), and the type-specific pattern is preserved—brine remains negative and significant, while rock and coastal/marine keep their signs but, resting on fewer sites per cell, lose statistical significance. Magnitudes are smaller because each finer cell contains fewer salt sites and borders.

Table A11: Alternative Grid Size: 150km vs. 200km

	200km (baseline) 476 cells, 70 salt	150km 787 cells, 108 salt
Panel A: Salt × Temperature → Borders		
10-year window	-0.724^{*} (0.408)	-0.256^{**} (0.118)
25-year window	-0.955^{*} (0.556)	-0.361^{**} (0.162)
50-year window	-1.137 (0.710)	-0.454^{**} (0.205)
Salt count (RE)	1.809^{***} (0.593)	1.280^{***} (0.339)
Panel B: Type × Temperature → Borders (25yr)		
Rock × Temp	$+1.319^{*}$ (0.720)	+0.458 (0.330)
Brine × Temp	-2.633^{**} (1.214)	-0.967^{**} (0.408)
Coastal/marine × Temp	+0.205 (0.178)	+0.067 (0.048)

Notes: Panel A reports the salt × Guiot–Corona temperature interaction at three windows with cell and period fixed effects. Panel B reports type-specific temperature interactions entered simultaneously. The 150km grid produces more cells (787 vs. 476) and more salt cells (108 vs. 70); coefficients are smaller in magnitude because each finer cell contains fewer salt sites and borders. Standard errors clustered by grid cell in parentheses, consistent with Tables 1 and 2; the 150km column matches the 150km row of Table 5. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$.

L Temperature Window Sensitivity

A potential concern is that the choice of temperature averaging window (10, 25, or 50 years) is *ad hoc*. Figure A3 plots the salt × temperature coefficient across all seven available windows (5, 10, 15, 20, 25, 30, and 50 years). The coefficient increases monotonically in magnitude from -0.57 (5-year) to -1.14 (50-year), with no sign changes, no anomalous windows, and tightly estimated confidence intervals that exclude zero at every point. This smooth pattern—consistent

with sustained climate shifts rather than short-run weather noise—is itself strong evidence that the result is not an artifact of a particular window choice.

Figure A4 shows the same exercise for the type-specific temperature interactions (all three types entered simultaneously). The rock and brine series diverge monotonically: rock increases from +0.71 to +1.78 and brine decreases from −1.51 to −3.60. The fan widens with longer windows, confirming that the opposing climate responses of appropriable and non-appropriable salt types are a robust structural feature of the data, not a statistical artifact.

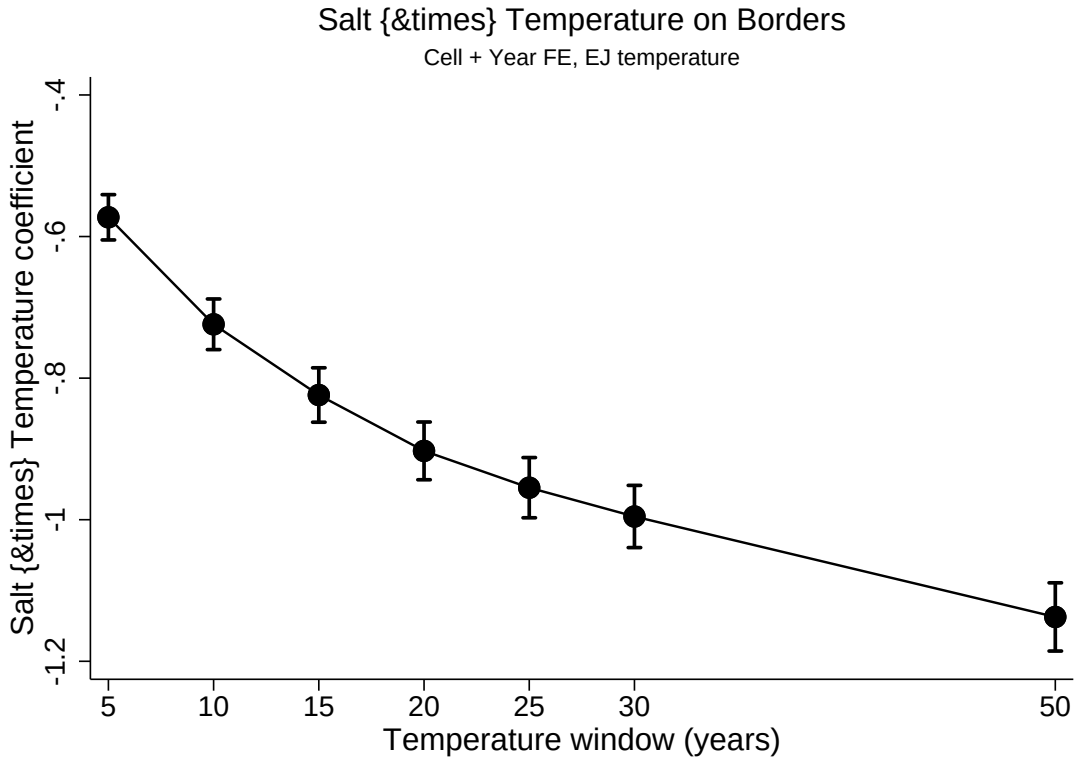

Figure A3: Salt by Temperature Coefficient Across Windows

Notes: Each point is the salt × Guiot–Corona temperature interaction coefficient from a separate cell + period FE regression at the indicated window length. Whiskers show 95% confidence intervals. The coefficient increases monotonically in magnitude from −0.57 (5-year) to −1.14 (50-year).

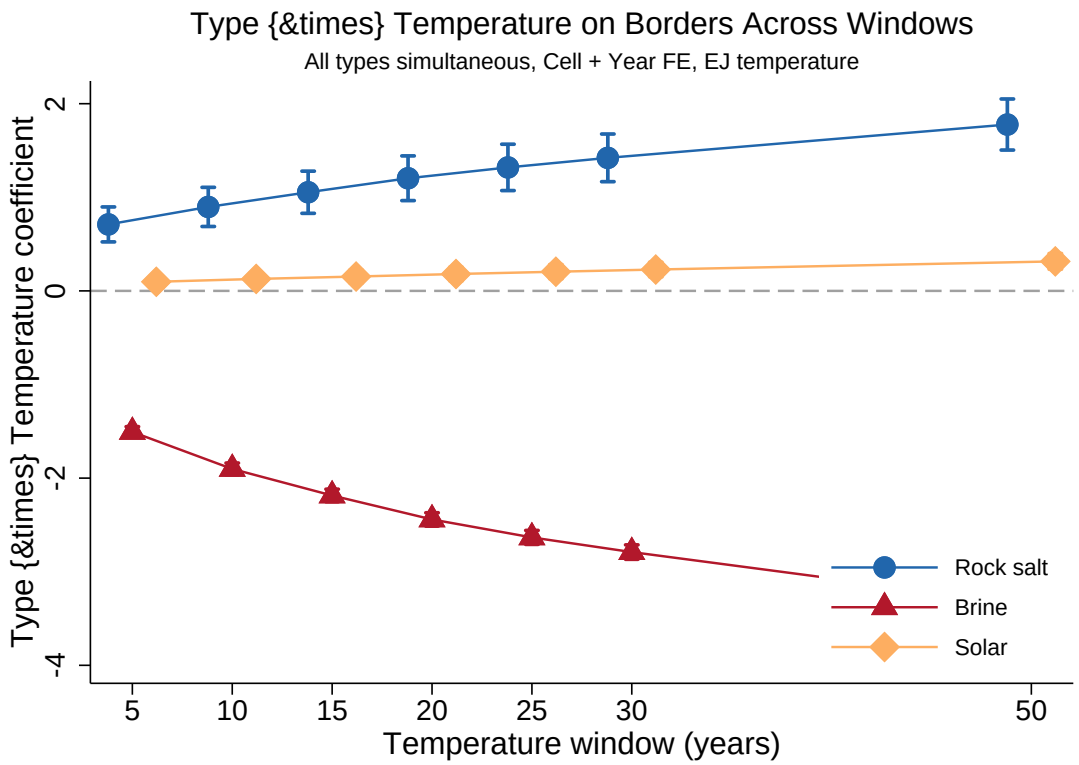

Figure A4: Type × Temperature Coefficients Across Windows

Notes: Each point is the type-specific temperature interaction coefficient from a regression with all three types entered simultaneously (cell + period FE). Rock (blue circles) is positive and increasing; brine (red triangles) is negative and decreasing; coastal/marine (orange diamonds) is small and positive. The divergence widens with longer windows, consistent with sustained climate shifts driving type-specific political responses.

M Sub-Period Analysis

Table A12 estimates the salt × temperature interaction separately for three sub-periods: 1100–1400 (transition from the Medieval Warm Period to early Little Ice Age), 1400–1600 (peak Little Ice Age onset), and 1600–1790 (mature Little Ice Age with more stable borders).

The effect is strongest in the earliest period: the salt × temperature coefficient is −1.620 ($p < 0.001$) in 1100–1400, −0.307 ($p < 0.001$) in 1400–1600, and small and of the opposite sign (+0.037, significant at the five percent level) in 1600–1790, by which point the European state system had largely consolidated. The triple interaction (column 4) confirms: the pre-1400 coefficient (−2.123) is four times larger than the post-1400 coefficient (−0.516).

This declining pattern is consistent with the identification strategy. During 1100–1400, temperature variation was large (the transition from medieval warmth to early cooling), and political borders were fluid—salt-rich areas were actively contested. By 1600–1790, borders had largely stabilized (the Peace of Westphalia in 1648 froze many territorial boundaries), and within-period temperature variation was smaller. The type-specific results reinforce this: brine × temperature is −4.301 in 1100–1400 but only −1.032 in 1400–1600 and null thereafter.

Table A12: Sub-Period Analysis: Salt × Temperature (25yr, Cell + Year FE)

	(1) 1100–1400	(2) 1400–1600	(3) 1600–1790	(4) Pre/Post 1400
Panel A: Aggregate				
Salt × Temp	−1.620 ∗ ∗	−0.307 ∗ ∗	+0.037	
	(0.036)	(0.024)	(0.015)	
Salt × Temp × Pre-1400				−2.123 ∗ ∗
				(0.044)
Salt × Temp × Post-1400				−0.516 ∗ ∗
				(0.026)
Panel B: By type				
Rock × Temp	+2.331 ∗ ∗	+0.693 ∗ ∗	−0.459 ∗ ∗	
	(0.187)	(0.150)	(0.093)	
Brine × Temp	−4.301 ∗ ∗	−1.032 ∗ ∗	+0.008	
	(0.072)	(0.042)	(0.026)	
Coastal/marine × Temp	−0.175 ∗ ∗	+0.256 ∗ ∗	−0.086 ∗	
	(0.055)	(0.051)	(0.034)	
Observations	23,072	15,244	16,068	53,560
Cell FE	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes

Notes: Columns (1)–(3) estimate the specification separately on each sub-period. Column (4) uses the full sample with interactions split by pre/post 1400. Temperature is the 25-year Guiot–Corona rolling mean. All specifications include cell and period fixed effects. Standard errors in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$.

N Time-Unit Robustness

With a 10-year rolling temperature mean observed at five-year intervals, consecutive observations share five years of temperature data, introducing mechanical serial correlation in the regressor. While this does not bias coefficients, it could inflate t-statistics. Table A13 re-estimates the

type-specific specification on non-overlapping panels.

Panel A uses a 5-year temperature mean at five-year intervals, which eliminates temporal overlap (at the cost of a shorter averaging horizon). Panel B aggregates the panel to decadal intervals, preserving the 10- and 50-year averaging horizons while ensuring that consecutive observations no longer share temperature data. In every case the opposing signs survive: the brine × temperature coefficient is negative and significant (ranging from −1.51 at the non-overlapping 5-year window to −3.53 at the decadal 50-year window), and rock × temperature is positive throughout. Temporal overlap does not drive the type-specific results.

Table A13: Time-Unit Robustness: Type-Specific Salt × Temperature, Non-Overlapping Windows

	Rock × Temp	Brine × Temp	Coastal/marine × Temp
Panel A: 5-year intervals, non-overlapping 5yr window			
Type × Temp (5yr)	+0.710*	−1.507**	+0.098
	(0.416)	(0.756)	(0.108)
Panel B: Decadal intervals, non-overlapping window			
Type × Temp (10yr)	+0.910*	−1.951**	+0.130
	(0.532)	(0.948)	(0.137)
Type × Temp (50yr)	+1.800*	−3.530**	+0.326
	(0.950)	(1.509)	(0.228)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes

Notes: Dependent variable is border count; all three type × temperature interactions enter simultaneously. Panel A uses the baseline five-year grid-cell panel (476 cells, 1100–1790) with a 5-year temperature window that has zero overlap between consecutive observations. Panel B aggregates the panel to decadal intervals (years divisible by ten), at which the 10- and 50-year rolling means have zero overlap. All specifications include cell and period fixed effects. Temperature is the lagged rolling mean from Guiot and Corona (2010). Conley (1999) spatial HAC standard errors (100km) in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

O Sensitivity to Unclassified Salt Sites

Approximately 24% of salt sites remain unclassified—recorded only as "Saline" in the historical sources without specifying extraction method. These sites contribute to the aggregate `salt_count` but not to any type-specific count, creating a gap between the total and the sum of types. Figure A5 shows that unclassified sites (gray circles) are geographically dispersed across Europe—they are not concentrated in any single region and overlap spatially with classified sites of all three types.

This suggests the unclassified sites are not systematically different from the classified ones. Table A14 bounds the type-specific results under extreme assumptions about what the unclassified sites actually are.

Column (1) reproduces the baseline. Columns (2)–(4) assign all unclassified sites to a single type: brine, coastal/marine, or rock respectively. Column (5) allocates them proportionally to the observed type shares (14% rock, 39% brine, 47% coastal/marine).

The brine × temperature coefficient is negative and highly significant in every scenario, ranging from −1.618 (when all unclassified are added to brine, diluting the measure) to −2.663 (when all unclassified are assigned to coastal/marine, leaving brine unchanged). The rock coefficient is always positive, ranging from +0.255 (worst case: all unclassified are rock) to +1.747 (proportional allocation). All signs are preserved in every scenario, and all coefficients remain significant at $p < 0.01$.

Table A14: Sensitivity to Unclassified Salt Sites

	(1) Baseline	(2) All → Brine	(3) All → Coastal/marine	(4) All → Rock	(5) Proportional
Rock × Temp	+1.319 ∗ ∗∗	+1.242 ∗ ∗∗	+1.300 ∗ ∗∗	+0.255 ∗ ∗∗	+1.747 ∗ ∗∗
	(0.126)	(0.128)	(0.126)	(0.049)	(0.121)
Brine × Temp	−2.633 ∗ ∗∗	−1.618 ∗ ∗∗	−2.663 ∗ ∗∗	−2.647 ∗ ∗∗	−2.437 ∗ ∗∗
	(0.038)	(0.028)	(0.039)	(0.040)	(0.036)
Coastal/marine × Temp	+0.205 ∗ ∗∗	+0.204 ∗ ∗∗	+0.148 ∗ ∗∗	+0.203 ∗ ∗∗	+0.337 ∗ ∗∗
	(0.039)	(0.039)	(0.031)	(0.039)	(0.037)
Cell FE	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable: border count. Temperature is the 25-year Guiot–Corona rolling mean. All three type × temperature interactions are entered simultaneously. Column (1) uses classified sites only. Columns (2)–(4) add all unclassified sites to brine, coastal/marine, or rock respectively. Column (5) allocates unclassified sites proportionally to observed type shares (14% rock, 39% brine, 47% coastal/marine). The aggregate salt × temperature coefficient (−0.955) is unchanged across scenarios since total salt count does not depend on type classification. Standard errors in parentheses. $***\ p < 0.01$.

A Geography-Based Assignment of Unclassified Sites

The bounds in Table A14 are deliberately agnostic: they ask what happens if every unclassified site were of a single type, or if the unclassified sites split according to the aggregate type shares. A

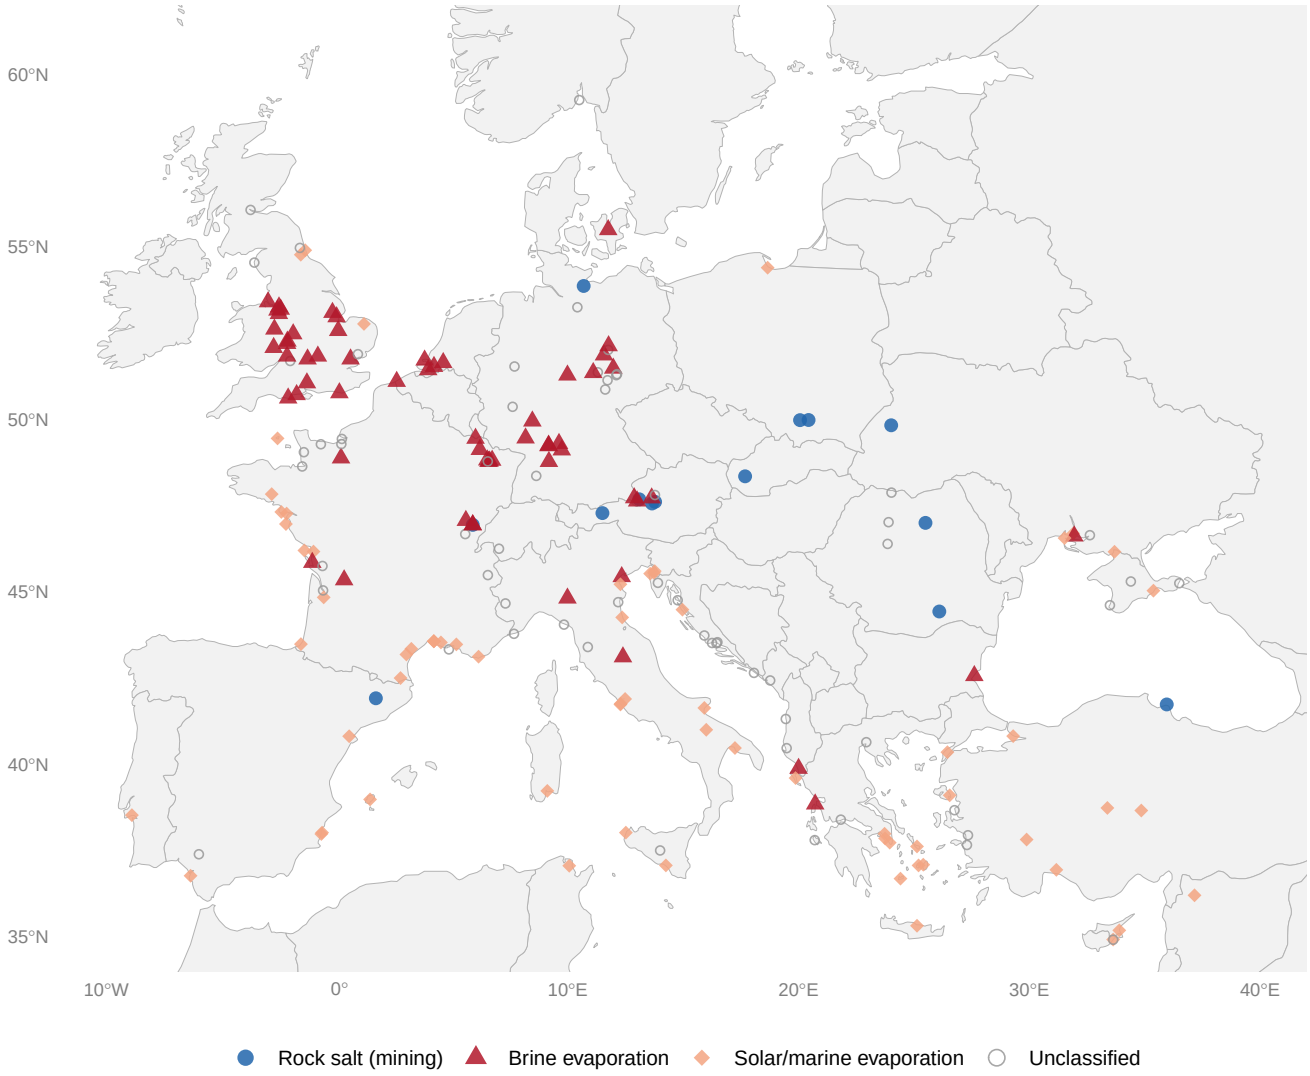

Figure A5: Salt Sites by Extraction Type, Including Unclassified

Notes: Blue circles = rock salt (mining); red triangles = brine evaporation; orange diamonds = coastal/marine evaporation; gray open circles = unclassified ("Saline" only). Of 202 European salt sites, 67 are unclassified. Unclassified sites are geographically dispersed, appearing in central Europe, France, the Alpine region, Eastern Europe, and the Baltic—broadly similar to classified sites. For multi-type sites, the most appropriable type is shown.

natural refinement asks instead *where* each unclassified site sits and assigns it a type accordingly. Extraction technology is closely tied to local geography—solar evaporation requires warm, saline coastlines; brine boiling exploits inland saline springs in karstic and Triassic formations; rock salt requires a mineable halite deposit. We therefore use the classified sites to learn the mapping from location to extraction type and apply it to the unclassified sites.

Concretely, we fit a penalized multinomial logit of extraction type on a parsimonious set of geographic predictors: point-sampled sea-surface salinity and a coastal indicator (which signal solar evaporation), latitude and elevation, proximity to mapped salt-deposit basins, and the karst–evaporite and Triassic (Keuper) characteristics of the surrounding grid cell (which signal brine). The model is trained on the classified sites and assessed by leave-one-out cross-validation. It achieves 82% out-of-sample accuracy against a 48% majority-class base rate, with a multiclass Brier score of 0.32 (versus 0.67 under random guessing). Brine and coastal/marine sites are recovered well (recall 0.89 and 0.86); rock is the difficult class (recall 0.39), for a concrete reason: no continental rock-salt geology is observable. Halite rarely outcrops, and the available evaporite-basin layers cover the central-European Zechstein but not the Alpine, Carpathian, or Iberian rock salt. We therefore treat rock with particular caution.

We use the fitted model in four ways, reported in Table A15. Column (2) is a *soft* assignment: each unclassified site is given fractional membership in each type equal to its predicted probabilities, and the per-cell unclassified mass is split accordingly. This is the geography-based analogue of the proportional column of Table A14, with location-specific rather than global shares. Column (3) imposes a substantively motivated bound: because rock-salt mines are capital-intensive, famous, and almost always documented as such, a site recorded merely as "Saline" is *a priori* unlikely to be rock, so we assign the unclassified mass only to brine and coastal/marine (the rock prior is set to zero), leaving the rock count—and hence the rock coefficient—at its classified baseline. Column (4) hand-verifies the handful of sites the model is most likely to misread: several sites it cannot place geologically are in fact rock salt—Bex (Saline de Bex, Alpine), Rónászék, Szamos Ujvár, and Maros Ujvár (Transylvanian salt towns), Volterra (Tuscan *salgemma*), and the rock component of Solikamsk on the Kama—while Po Valley and Moûtiers are coastal/marine; these

are fixed by hand and the remainder geography-imputed. Column (5) propagates the classification uncertainty into the standard errors: we draw fifty hard type assignments from the predicted probabilities, re-estimate on each, and combine by Rubin's rules.

Across all five treatments the qualitative pattern is unchanged. The brine–temperature interaction is negative and tightly estimated everywhere (between -1.95 and -2.63), and the rock–temperature interaction is positive and significant in every column. This includes the inference-honest multiple-imputation estimate, which remains significant at the five-percent level despite a fraction of missing information near 0.9, and the hand-verified override, which—by adding six geologically-invisible rock sites whose cells do not display the consolidation pattern—if anything understates the effect. The geography-based assignment thus corroborates the bounds in Table A14: the type-specific temperature results are not an artifact of how the unclassified sites are treated. Figure A6 maps the predicted types and shows that the assignment follows geography—coastal sites to coastal/marine, the central-German Triassic belt to brine, and the Transylvanian salt towns to rock.

P Alternative Battle Data

Our baseline conflict analysis uses the Kitamura (2021) World Historical Battles Database (442 georeferenced European battles within our grid, 1302–1700). To check that the rebellion-channel mechanism is not an artifact of one battle source, we re-estimate it on the Historical Conflict Event Dataset (HCED) compiled by Miller, which draws primarily on Jaques (2007) and cross-references Harbottle, Clodfelter, and other encyclopedias. The HCED contains 892 geocoded European battles in the same window—more than double the Kitamura count—and is only moderately correlated with it (cell-year $r = 0.68$), so it provides substantial independent variation.

Table A16 reproduces the type-specific decomposition of Table 6 on each dataset. The aggregate salt $\times$ temperature interaction is negative and highly significant in both, grows monotonically with window length, and is roughly twice as large in the larger HCED. The type decomposition replicates the key signature: brine $\times$ temperature is strongly negative in both datasets—about -0.044 (Kitamura) and -0.085 (HCED) at the 25-year window, both $p < 0.001$—so colder periods raise conflict precisely where salt is non-appropriable, and the effect grows with window length.

Table A15: Geography-Based Assignment of Unclassified Salt Sites

	(1) Baseline	(2) Geography (soft)	(3) Brine/Coastal bound	(4) Manual override	(5) Multiple imputation
Rock × Temp	$+1.319^{***}$	$+1.206^{***}$	$+1.288^{***}$	$+0.835^{***}$	$+0.882^{**}$
	(0.126)	(0.105)	(0.127)	(0.094)	(0.408)
Brine × Temp	-2.633^{***}	-2.044^{***}	-1.954^{***}	-2.036^{***}	-1.993^{***}
	(0.038)	(0.031)	(0.030)	(0.031)	(0.097)
Coastal/marine × Temp	$+0.205^{***}$	$+0.216^{***}$	$+0.246^{***}$	$+0.186^{***}$	$+0.188^{***}$
	(0.039)	(0.036)	(0.035)	(0.036)	(0.052)
Cell FE	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable: border count; all three type × temperature interactions entered simultaneously, with temperature the 25-year Guiot–Corona rolling mean. Column (1) reproduces the classified-only baseline. Column (2) splits the per-cell unclassified mass by the location-specific membership probabilities of a penalized multinomial logit of extraction type on geography (sea-surface salinity, coastal status, latitude, elevation, salt-deposit-basin proximity, and cell karst–evaporite and Triassic characteristics), trained on classified sites; leave-one-out accuracy is 82% (48% base rate). Column (3) assigns the unclassified mass only to brine and coastal/marine (rock prior set to zero), motivated by the fact that rock-salt mines are rarely recorded merely as "Saline." Column (4) hand-fixes eight verified sites (Bex, Rónászék, Szamos Ujvár, Maros Ujvár, Volterra, and Solikamsk to rock; Po Valley and Moûtiers to coastal/marine) and geography-imputes the rest. Column (5) is a multiple-imputation estimate over fifty draws from the predicted probabilities, combined by Rubin's rules; the larger rock standard error reflects the limited geological identification of rock. Standard errors in parentheses. ** $p < 0.05$, *** $p < 0.01$.

Coastal/marine × temperature is negative throughout. Rock × temperature is positive in both datasets but significant only in Kitamura, and there only at the 10- and 25-year windows; with just ten rock-salt cells it is the fragile arm, and the noisier HCED renders it insignificant throughout—a difference in precision, not in sign. Civil-war coding is available only for the Kitamura battles (Table 6, Panel B), so we report all battles for both datasets here.

Q Type-Specific Temperature Interactions: Full Specification

Table A17 presents the full specification with century interactions for all three types plus type-specific temperature interactions (25-year Guiot–Corona window, cell fixed effects). The type × temperature interactions retain their opposing signs after controlling for century-specific effects: brine × temperature remains strongly negative (-1.68, $p < 0.01$) and coastal/marine × temperature is negative (-0.33, $p < 0.05$). The rock × temperature coefficient remains positive ($+0.66$)

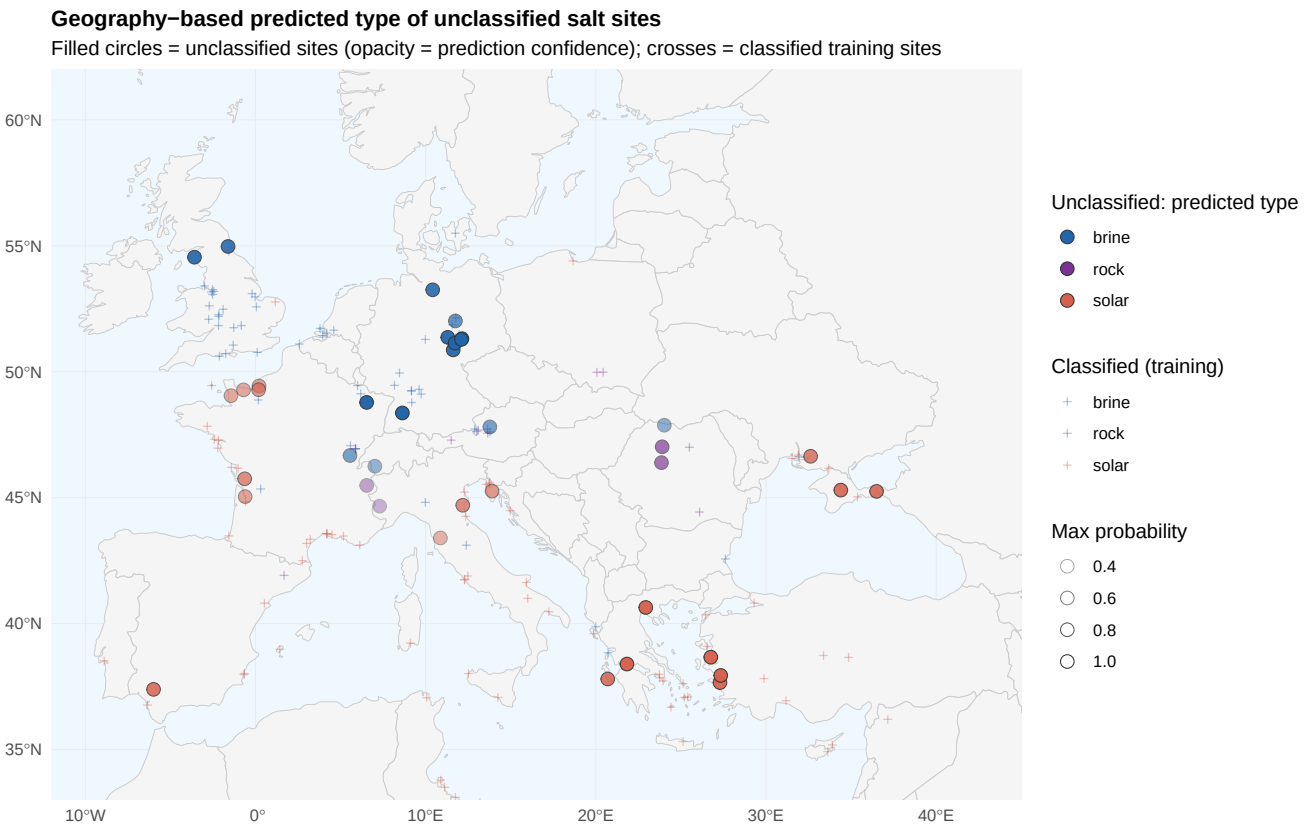

Figure A6: Geography-Based Predicted Type of Unclassified Salt Sites

Notes: Filled circles are the in-grid unclassified ("Saline") sites, colored by predicted extraction type, with opacity proportional to the prediction's maximum class probability; crosses are the classified training sites. Predictions follow geography: coastal Mediterranean, Aegean, and Black Sea sites are assigned coastal/marine; the central-German Triassic belt is assigned brine; and the Transylvanian salt towns are assigned rock.

but is not individually significant in this saturated specification that also includes twelve century interactions. Once the temperature channel is absorbed, the coastal/marine century coefficients show consolidation in every century (-0.07 to -0.74), resolving the apparent sign reversal visible in Figure 6.

R Cross-Sectional Horse Race

Table A18 enters all three type-specific salt counts simultaneously in a cross-sectional horse race regression (period fixed effects only, exploiting between-cell variation). Brine salt is the sole

Table A16: Alternative Battle Data: Kitamura vs. HCED (1300–1700)

Temperature window	10-year	25-year	50-year
Panel A: Kitamura (2021) — 442 battles			
Salt × Temp	−0.0132***	−0.0230***	−0.0388***
	(0.0030)	(0.0052)	(0.0106)
Rock × Temp	+0.0104*	+0.0236**	+0.0255
	(0.0062)	(0.0104)	(0.0328)
Brine × Temp	−0.0217***	−0.0442***	−0.0839***
	(0.0071)	(0.0134)	(0.0319)
Coastal/marine × Temp	−0.0113***	−0.0150***	−0.0154***
	(0.0041)	(0.0040)	(0.0050)
Panel B: HCED / Miller — 892 battles			
Salt × Temp	−0.0248***	−0.0402***	−0.0706***
	(0.0059)	(0.0080)	(0.0154)
Rock × Temp	+0.0041	+0.0052	+0.0086
	(0.0158)	(0.0189)	(0.0361)
Brine × Temp	−0.0487***	−0.0845***	−0.1636***
	(0.0101)	(0.0115)	(0.0338)
Coastal/marine × Temp	−0.0089	−0.0109	−0.0100
	(0.0069)	(0.0085)	(0.0097)
Cell FE, Period FE	Yes	Yes	Yes
Correlation (cell-year)		$r = 0.68$	

Notes: Each panel reproduces the rebellion-channel decomposition of Table 6 (aggregate salt × temperature and the three type × temperature interactions entered simultaneously) on a different battle dataset, restricted to 1300–1700. Kitamura battles from the World Historical Battles Database (Kitamura, 2021); HCED battles from the Historical Conflict Event Dataset (Miller, Harvard Dataverse), which draws primarily on Jaques (2007). Cell and period fixed effects throughout; standard errors clustered by grid cell in parentheses. On battles, brine and coastal/marine × temperature are negative (cold → more conflict where salt is non-appropriable) and rock is positive. Civil-war coding is available only for the Kitamura battles (Table 6, Panel B), so all battles are reported for both datasets here. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

significant predictor of political fragmentation ($\hat{\beta} = 2.97$, $p = 0.027$ with Conley standard errors). Rock salt is positive but insignificant (0.83, $p = 0.34$); coastal/marine salt is essentially zero (-0.08, $p = 0.81$). When each type is entered alone (columns 2–4), rock salt is significant (2.49, $p = 0.009$), but this significance disappears in the horse race—suggesting that rock salt cells overlap spatially with brine cells, and brine is the active ingredient.

Table A17: Type-Specific Temperature Interactions: Full Specification

	Dependent Variable: Border Count
Panel A: Century interactions (controlling for type × temp)	
Rock × 1301–1400	−0.330
Rock × 1401–1500	−0.641
Rock × 1501–1600	−0.969∗
Rock × 1601–1790	−1.190∗
Brine × 1301–1400	+0.742∗
Brine × 1401–1500	+1.094∗
Brine × 1501–1600	+1.404∗
Brine × 1601–1790	+1.530
Coastal/marine × 1301–1400	−0.071
Coastal/marine × 1401–1500	−0.314
Coastal/marine × 1501–1600	−0.655 ∗ ∗
Coastal/marine × 1601–1790	−0.741 ∗ ∗
Panel B: Type × temperature (25yr Guiot–Corona window)	
Rock × Temperature	+0.662
	(0.560)
Brine × Temperature	−1.684 ∗ ∗∗
	(0.629)
Coastal/marine × Temperature	−0.333 ∗ ∗
	(0.161)
Cell FE	Yes
Period FE	Yes
Observations	∼54,000
R^2 (Within)	0.176

Notes: 476 grid cells, five-year intervals. Temperature is the 25-year rolling mean of the Guiot–Corona reconstruction. All specifications include cell and period fixed effects. Century interactions use 1100–1300 as the omitted category. Conley (1999) spatial HAC standard errors (100km bandwidth) in parentheses for temperature interactions. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$.

S Mean Polity Size as an Alternative Outcome

Border count is the paper's primary measure of fragmentation, but it is not the only one. Table A19 re-estimates the type-specific decomposition with the log of mean polity area within each cell-year as the dependent variable—a direct, consolidation-oriented measure, since a larger mean polity indicates fewer and larger states. Because larger polities correspond to consolidation, the type × temperature signs flip relative to the border-count results of Table 2: brine × temperature is posi-

Table A18: Salt Type and Political Fragmentation: Horse Race

	Dependent Variable: Border Count			
	(1) All Types	(2) Rock Only	(3) Brine Only	(4) Coastal/marine Only
Rock salt count	0.829	2.492 ∗ ∗		
	(0.931)	(0.958)		
Brine salt count	2.975 ∗ ∗		3.013 ∗ ∗	
	(1.346)		(0.258)	
Coastal/marine salt count	−0.084			0.223
	(0.335)			(0.400)
Period FE	Yes	Yes	Yes	Yes
Observations	64,260	64,260	64,260	64,260

Notes: 476 grid cells, five-year intervals (1100–1790). All specifications use period fixed effects only, exploiting between-cell variation. Column (1) enters all three types simultaneously. Columns (2)–(4) enter each type separately. Conley (1999) spatial HAC standard errors (100km bandwidth) in parentheses for column (1); columns (2)–(4) retain standard errors from single-type regressions. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

tive and significant at every window (colder periods shrink polities in brine cells—fragmentation), while rock × temperature is negative (colder periods enlarge polities in rock cells—consolidation), consistent with the model. The rock coefficient is imprecise because few cells contain only rock salt.

T Polity-Level Analysis

This appendix reports the polity-level analysis, which provides external validity for the grid-cell results: it uses a different unit of analysis (the historical polity) and a different outcome (interstate warfare rather than borders). We present the aggregate salt × temperature result first, then its decomposition by extraction type.

T.1 Aggregate Salt × Temperature

The polity-level analysis uses the temporal salt measure, log(salt_mines_count_temporal+1), which exploits the limited within-polity variation arising from the documented start dates of salt production. The dependent variables are the count of all wars and, separately, civil wars and rebellions fought in each polity's territory in each five-year period—the 538 georeferenced battles attributed

Table A19: Type-Specific Salt × Temperature: Mean Polity Size (Cell + Year FE)

	Dependent Variable: Log Mean Polity Area		
	10-year	25-year	50-year
Rock × Temperature	−0.108	−0.176	−0.169
	(0.149)	(0.205)	(0.261)
Brine × Temperature	+0.154 ∗ ∗∗	+0.204 ∗ ∗∗	+0.246 ∗ ∗∗
	(0.044)	(0.047)	(0.051)
Coastal/marine × Temperature	+0.041	+0.019	−0.022
	(0.043)	(0.068)	(0.106)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes
Observations	47,603	46,718	45,238

Notes: Dependent variable is the log of mean polity area (km^2) within each cell-year, computed from the cell–polity crosswalk. All three type × temperature interactions enter simultaneously with cell and period fixed effects. Temperature is the Guiot–Corona rolling mean at the indicated window. A larger mean polity area indicates consolidation, so the signs are reversed relative to the border-count specification (Table 2). Standard errors clustered by grid cell in parentheses. $***\ p < 0.01$, $**\ p < 0.05$, $*\ p < 0.1$.

to the polity controlling the cell in which they were fought, with the civil/rebellion subset coded under the generous definition. All specifications include five-year period fixed effects and cluster standard errors at the polity level.

If salt's value as a fiscal resource increased during colder periods—because coastal salt production declined and mineral salt became relatively more important—then salt-rich polities should have experienced more conflict during cold episodes. The negative salt × temperature interaction predicted by this mechanism is exactly what we find. We use the Guiot–Corona reconstruction (Guiot and Corona, 2010), consistent with the grid-cell analysis.

Table A20 presents the results across three temperature windows (10, 20, and 50 years), with an area control and standard errors clustered by polity. Panel A uses all wars: the salt × temperature interaction is negative at every window and grows in magnitude with window length (−0.095 at 10 years, −0.129 at 20, and −0.261 at 50; significant at the 10%, 5%, and 1% levels). Panel B restricts the outcome to civil wars and rebellions (generous definition), and the same negative interaction is present at roughly half the magnitude (−0.050, −0.046, −0.101), significant at the 10- and 50-year windows. Mirroring the grid-cell conflict results (Table 6), salt-rich polities fight

more in colder periods and the effect operates through civil conflict; the estimates are noisier at this coarser, partly endogenous unit of analysis, and the type-specific decomposition does not carry to the heterogeneous polity (Table A21).

Table A20: Salt and Conflict at the Polity Level: All Wars and Civil Wars

Temperature window	10-year	20-year	50-year
Panel A. All wars			
Salt × Temp	−0.095*	−0.129**	−0.261***
	(0.051)	(0.055)	(0.078)
Panel B. Civil wars / rebellions only			
Salt × Temp	−0.050*	−0.046	−0.101***
	(0.030)	(0.029)	(0.033)
Period FE	Yes	Yes	Yes
Area control	Yes	Yes	Yes
Cluster SE (polity)	Yes	Yes	Yes
Observations	26,229	26,229	26,229

Notes: Polity × five-year-period panel. The dependent variable counts the battles fought in each polity's territory from the 538 georeferenced battles (1302–1700); Panel A counts all such battles, Panel B only those coded as civil wars or rebellions (generous definition). Each cell reports the log(salt, temporal) × temperature interaction from a regression that also includes log salt, temperature, and log polity area, with five-year period fixed effects and standard errors clustered by polity. Salt-rich polities fight more in colder periods, and the effect is present for both all wars and civil wars. The type-specific decomposition does not carry to the polity level, where heterogeneous territories aggregate cells of different salt types (Table A21). $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

The panel is observed at five-year intervals, and all specifications include five-year period fixed effects. Because we use 10-year rolling temperature means, consecutive observations at five-year intervals share five years of temperature data, introducing mechanical serial correlation in the regressor. While this does not bias coefficients, it could affect inference. Using a 5-year temperature mean, which eliminates the overlap, preserves the negative salt × temperature relationship (−0.057 for all wars, −0.028 for civil wars) but with smaller and less precise coefficients than the 10-year window—consistent with the effect strengthening as the temperature horizon lengthens, since the longer 20- and 50-year windows are the largest and most significant.

A concern with polity-level analysis is that polity boundaries are themselves an outcome of the fragmentation process. When a polity fragments, successor states enter the panel and may both inherit salt sites, mechanically amplifying the salt–war relationship; conversely, consolidation

removes polities from the panel. The grid-cell analysis avoids this concern entirely because cells are fixed geographic units. At the polity level, we note that the panel is dominated by long-lived polities: 75% of observations come from polities observed for at least 250 years, and only 4% come from polities lasting less than 100 years. To verify that endogenous entry and exit do not drive the results, we restrict the sample to the 255 polities observed for at least 50 periods ($\sim$250 years). The all-war salt $\times$ temperature coefficients are virtually unchanged, if anything slightly larger—-0.098 (10yr), -0.135 (20yr), and -0.276 (50yr), compared with -0.095, -0.129, and -0.261 in the full sample—and retain the same significance pattern under polity clustering. Restricting further to the 156 polities observed for at least 400 years preserves the pattern (-0.096, -0.135, -0.284), and the civil-war subset behaves the same way. Moreover, to the extent that salt-rich regions sustain more polities precisely because non-appropriable resources enable rebellion and prevent consolidation, endogenous entry is a *consequence* of the mechanism, not a confound.

T.2 Type Decomposition

Table A21 decomposes the polity-level salt $\times$ temperature interaction by extraction type, mirroring the grid-cell analysis in Table 2. All three type $\times$ temperature interactions are entered simultaneously with period fixed effects and polity-clustered standard errors.

The polity-level type decomposition does *not* reproduce the opposing signs found at the grid-cell level: all three type coefficients are negative, with coastal/marine being the most consistently significant ($p < 0.01$ at 25 and 50-year windows). Rock is marginally significant at 10 and 50-year windows but not at 25 years; brine is insignificant throughout.

This pattern reflects a fundamental difference between grid-cell and polity-level units of analysis. Grid cells are fixed geographic units, each containing a single salt type (or none); the type $\times$ temperature interaction isolates the response of a specific extraction technology to climate variation. Polities, by contrast, are large and heterogeneous: France contains rock, brine, and coastal/marine salt simultaneously, and a polity's type-specific count reflects the aggregation of many cells with different technologies. The opposing-sign mechanism—in which brine cells fragment while rock cells consolidate—operates *within* a polity's territory, producing cross-cutting effects that wash out when aggregated to the polity level. The grid-cell analysis, which isolates in-

dividual cells by extraction technology, is the more appropriate setting for testing the type-specific predictions of the model.

Table A21: Polity-Level Type-Specific Salt × Temperature

	Dependent Variable: Traditional War Count		
	(1) 10yr	(2) 25yr	(3) 50yr
Rock × Temperature	−0.148∗	−0.152	−0.248∗
	(0.081)	(0.099)	(0.141)
Brine × Temperature	−0.013	−0.051	−0.095
	(0.064)	(0.060)	(0.081)
Coastal/marine × Temperature	−0.110 ∗ ∗	−0.144 ∗ ∗∗	−0.218 ∗ ∗∗
	(0.047)	(0.054)	(0.069)
Period FE	Yes	Yes	Yes
Cluster SE (polity)	Yes	Yes	Yes
Observations	26,410	26,096	25,507

Notes: Polity-year panel observed at five-year intervals. All three type × temperature interactions enter simultaneously alongside log type-specific salt counts, temperature, and log polity area. Temperature is the Guiot–Corona rolling mean at the indicated window. Standard errors clustered by polity in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

T.3 Constructing the Polity-Level Annexation Dataset

The polity-level analysis above measures total warfare. A sharper test of the appropriability mechanism asks whether salt-rich states—especially those holding appropriable rock salt—fought more wars of annexation *outside* their own borders when salt became more valuable in cold periods. The cell-level results show *where* salt-driven conflict erupts, but they cannot capture a polity that converts salt rents earned in one place into conquest somewhere else. Building the dataset for this test proceeds in four steps: (i) mapping battle belligerents to polities; (ii) defining polity entities and chaining them through unions and renamings; (iii) classifying salt deposits and attributing them to polities at their core holdings; and (iv) defining the external-war outcome. Throughout, the goal is to attribute each polity a *fixed* salt endowment that is not contaminated by its later conquests, and to measure how its warfare beyond its borders responds to climate conditional on that endowment.

Mapping battle belligerents to polities. We begin from the catalogue of 538 georeferenced battles, 1300–1700, that underlies the cell-level conflict analysis. Each battle records the participants on each side as free text (e.g. "Kingdom of Castile, Order of Santiago"). We parse these strings into individual actors, handling nested coalitions, and map each actor to a canonical polity name using a hand-built crosswalk of roughly two hundred and fifty rules (e.g. "Kingdom of France" → France, "Spanish Empire" → Aragon, "Muscovy" → Russia), supplemented by a smaller set of additional corrections that we reviewed by hand (for example "Dutch republic" → Netherlands, "Duchy of Milan" → Milan). After these steps, ninety-four percent of belligerent mentions resolve to a polity present in the boundary data.

Because the polity-level test concerns *interstate* wars of annexation, we drop belligerents that are not territorial states. These fall into three groups: (a) non-territorial military actors (mercenary companies such as the *Tard-Venus* or the White Company, and free-floating fragments such as "supported by" clauses and named individual commanders); (b) civil-war and factional combatants—Scottish and Irish clans, the Covenanters, Jacobites, the Irish Brigade, and the Wars-of-the-Roses, Hussite, and Guelph/Ghibelline factions—which belong to internal rather than interstate conflict; and (c) actors outside the geographic coverage of the boundary maps (e.g. the North-African sultanates). The remaining battles are restricted to those coded as interstate (rather than civil war or rebellion) under our conflict classification, so that the dependent variable isolates inter-polity warfare.

Polity entities and temporal chaining. Polity territory comes from the gridded European boundary maps at five-year intervals. Because the same state can appear under different labels over time—and because the raw maps split some polities into numbered polygon fragments—we standardize the boundary names to the same canonical set used for the belligerents and merge polygon fragments to their parent polity. We then chain polities that are a single evolving state: the Kingdom of Poland and the Grand Duchy of Lithuania are treated as one entity, Poland–Lithuania, from the dynastic union of 1385; and the Crowns of Castile and Aragon are treated as Spain from 1479. Each polity is assigned a single, stable identity for the panel.

The Habsburg lands require special handling and are the one case we cannot resolve cleanly, so

we report two specifications. In the first, the Austrian hereditary lands, Bohemia, and Hungary are treated as *separate* states (each independent until its Habsburg acquisition). In the second, they are *bundled* into a single Habsburg entity, reflecting the dynastic union of the hereditary lands with Bohemia and royal Hungary after 1526. We report the bundled specification in the main text; the separate-states specification yields the same pattern and is retained as a robustness check. We do not bundle the Holy Roman Empire itself: in the boundary data the "Holy Roman Empire" is a small residual polygon (on the order of ten to twenty-five thousand square kilometres), not the imperial umbrella, and the Emperor in any case lacked the fiscal authority to tax the Empire at large—so folding it in would overstate the entity's fiscal reach.

One further correction is required for the reign of Charles V. Between roughly 1520 and 1556 the boundary data assign the Austrian hereditary lands—including the Salzkammergut and Tyrolean salt—to the Crown of Aragon, because Charles V simultaneously held the Spanish crown and the Empire and the "Austria" polygon disappears from the maps for these years. Left uncorrected, this would credit the Austrian salt and the Habsburg wars of this period to Spain. We therefore reassign the nine core Habsburg cells (Tyrol, Styria/Salzkammergut, Carniola, Bohemia, and royal Hungary), identified as cells that are Austrian before 1520 and after 1560 but Aragonese in between, from Aragon to Austria for the 1520–1556 window. This correction is applied in all specifications.

Salt deposits: classification and attribution. Each salt site in our deposit database carries a description of its extraction technology, which we classify into three appropriability types: *rock* (capital-intensive underground rock-salt and brine mines), *brine* (dispersed boiling of natural or pumped brine), and *solar* (coastal solar evaporation). A site may carry more than one type.

We attribute each deposit to the polity that held it, taking the controlling polity at a fixed reference year (1500)—a date at which the chained entities already exist, the boundary names are stable, and which predates the Charles V relabelling. For the twenty-seven rock deposits, which carry the analysis, we verified every attribution by hand against the historical record. The European rock deposits map to: Wieliczka, Bochnia, and the Galician works → Poland–Lithuania; Cardona → Spain; Aussee, Hallstatt, Hall (Tyrol), and Salins (Franche-Comté) → Austria; Hallein,

which belonged to the sovereign Archbishopric of Salzburg rather than the Habsburgs, → Salzburg; the Carpathian and Šoŏvar works → Hungary; and Wallachia → the Ottoman Empire. Fifteen nominally rock records were dropped: fourteen lie outside Europe (the Sahara, Arabia, Persia, Anatolia), and one, listed as "Lübeck," is a duplicate of the Lüneburg saltworks recorded at its Hanseatic export hub—Lübeck itself held no deposit—and is reassigned to Lüneburg's owner, Brunswick-Lüneburg, as brine.

Crucially, a polity's endowment is held *fixed* at these core holdings; it does not grow as the polity later conquers additional salt-bearing land. This is essential for identification: if conquest expanded the measured endowment, the salt variable would be partly an outcome of the very warfare we seek to explain. Fixing endowments at core holdings breaks that loop.

The external-war outcome. For each interstate battle we ask, of every belligerent polity, whether the battle was fought in a cell that polity controlled. A polity fighting in a cell it does not control is fighting *externally*—projecting force beyond its borders—which is our empirical proxy for a war of annexation or expansion. We aggregate these to counts of external wars per polity per five-year period.

We use *all* belligerents rather than only the attacking side because the attacker cannot be reliably identified for a majority of battles. The all-belligerent measure consequently captures power projection (fighting away from home) rather than strictly offensive initiation; we view this as the more conservative and better-measured object. "Control" is defined at the level of the two-hundred-kilometre grid cell, so the measure is coarse at the exact border; a finer point-in-polygon definition yields very similar results.

U Salt Production and Temperature

If the Little Ice Age caused a drop in temperatures around Europe, then there would be a heterogeneous effect on production depending on the country and site of production. Changes in temperature should affect production sites that rely purely on solar production methods, as both rock and brine salt extraction methods do not rely on sunlight. In order to test the temperature mechanism in our theory, we use site production data from Multhauf (1978). The appendix of

Multhauf (1978) contains production data tables both at the national and site level. The tables on national production contain observations overwhelmingly from the late 19th century and beyond, which is outside the scope of this paper. The site production data is broken down by individual sites within a country and contains observations at uneven temporal spacing, which results in an unbalanced panel of site, production amount, and year of observation. The sources for the observations are from a multitude of sources that Multhauf (1978) has pulled in. Production is converted to millions of kilograms per year across all tables, giving us consistent observations without a need for additional conversions for comparison. As Multhauf (1978) uses multiple sources for his observations, some sites have multiple observations per year with different sources. We use all observations recorded. If there is a range of years or production, we take the mean of the range.

We regress production on temperature, presenting the relationship both as a raw pooled scatter (Figure A7) and, with site and year fixed effects, as type-specific interaction regressions (Table A22). It can be seen that temperature has little discernible effect on both brine and rock salt, while having an obvious correlation with coastal/marine salt. A decrease in temperature shows that production at sites using coastal/marine extraction methods decrease, which is consistent with the idea that these sites require warmer weather for salt extraction. These results illustrate that the impact of the Little Ice Age on political fragmentation works through a differential supply shock. The decrease in temperature only affects production at coastal/marine sites. The decrease in coastal/marine site production raises the relative importance of inland salt sites that do not require solar power. It is through this shift in relative importance that we observe our political effect. Production decreases at coastal/marine sites in colder weather, raising the benefit of having inland salt sites for polities. Polities utilize and compete for inland salt sources, resulting in the difference in political equilibrium during the Little Ice Age.

The same sign appears in a within-site specification (Table A22), which interacts each extraction type with the lagged temperature mean and enters all three interactions simultaneously, absorbing time-invariant site characteristics and common European shocks: rock $\times$ temperature and brine $\times$ temperature are small and insignificant, while coastal/marine $\times$ temperature is positive. We treat this regression as *illustrative* rather than as a precise estimate. The coastal/marine

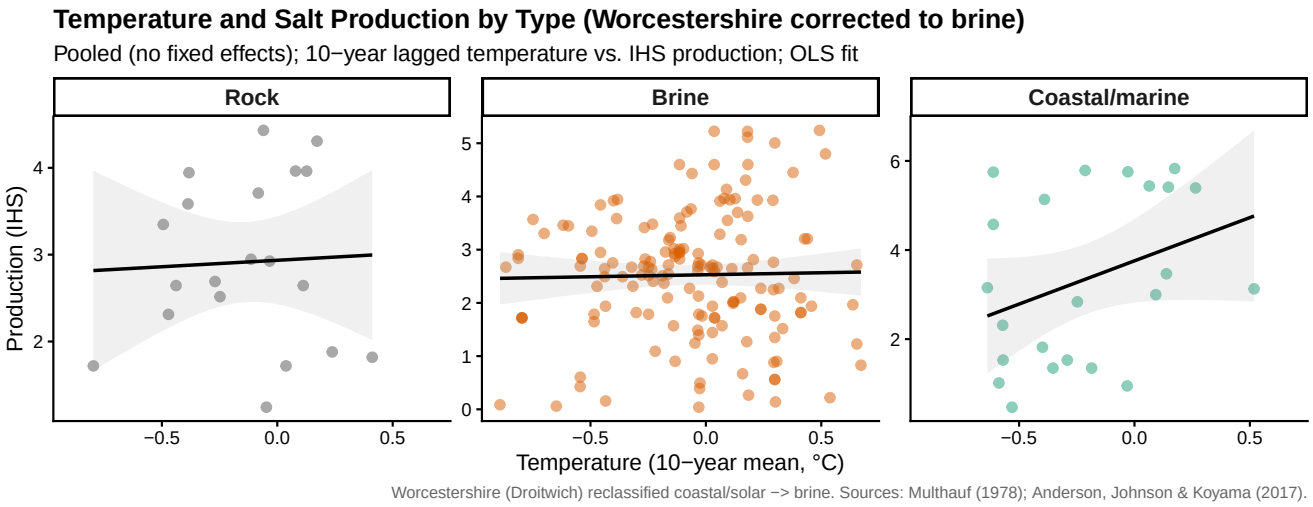

Figure A7: Temperature and Salt Production by Type (Pooled)

Notes: Each panel plots 10-year lagged mean temperature against salt production (inverse hyperbolic sine) for a different salt type, with no fixed effects. Production data from Multhauf (1978); temperature from Guiot and Corona (2010). The OLS fit (with 95% confidence interval) shows a positive relationship for coastal/marine salt—consistent with evaporation depending on warmth—but no relationship for rock or brine production. Worcestershire (Droitwich) is classified as brine rather than coastal/marine, correcting an ambiguous "salt pan" coding.

response is identified off only six regions—in practice the two with a genuine time series, Sardinia (1288–1410) and Brouage (1644–1780)—so the coefficient is implausibly large and unstable across temperature windows (+24 to +139), is essentially unchanged when the start year is varied because the same handful of observations always drive it, and cannot be estimated at all (collinearity) in narrower sub-periods. We therefore read the production data only for the *sign* of the type contrast—coastal sensitive, inland flat—which is visible in the raw scatter; the quantitative weight of the supply shock rests on the price evidence of Section 5.1.

V Salt Prices and Temperature

This appendix constructs the aggregate real salt-price series used in Section 5.1 and documents its robustness.

Data and the inflation problem. Salt prices are drawn from the Allen–Unger commodity-price database (`Data/Salt Data/Commodities.csv`), standardized to grams of silver per litre

Table A22: Temperature and Salt Production by Type (Site + Year FE; illustrative)

	Dependent Variable: Production (IHS)		
	(1) 5yr	(2) 10yr	(3) 25yr
Rock × Temperature	+0.068	−0.017	−0.077
	(0.282)	(0.323)	(0.462)
Brine × Temperature	+0.190	−0.156	−0.113
	(0.967)	(0.750)	(1.314)
Coastal/marine × Temperature	+24.07 ∗ ∗∗	+139.3 ∗ ∗∗	+24.24 ∗ ∗∗
	(2.224)	(5.720)	(2.421)
Site FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	179	179	179
R^2	0.983	0.983	0.983

Notes: 40 production regions, 1205–1796; unbalanced panel from Multhauf (1978). The outcome is the inverse hyperbolic sine of annual production (millions of kg). All three type × temperature interactions are entered simultaneously; "Temperature" is the lagged rolling mean of the Guiot and Corona (2010) (EJ-grid) anomaly at the indicated window. All specifications include site and year fixed effects. Standard errors clustered by region in parentheses. The coastal/marine interaction is identified off only six regions (chiefly Sardinia and Brouage); its magnitude is implausibly large and unstable across windows and the coefficient cannot be estimated in narrower sub-periods. The table is illustrative of the type *contrast* only; it should not be read as a stable elasticity. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

and aggregated to an annual European index (the mean of within-series log deviations, so the index captures common price movements rather than which markets happen to be observed in a given year). Standardizing to silver removes currency debasement but not the depreciation of silver itself during the Price Revolution: the influx of American bullion raised the silver price of essentially all goods between roughly 1500 and 1650. Salt is no exception—its nominal silver price trends strongly upward and is mechanically, and spuriously, *positively* associated with temperature (+0.58, $p < 0.001$; Table A23, row 1). Any price test must therefore deflate.

An external deflator. We deflate by Robert Allen's consumer price indices, constructed in grams of silver for seven European cities (Augsburg, Vienna, Paris, London, Antwerp, Gdansk, and Leipzig) that span and overlap our salt markets. The resulting real salt price—the salt index minus the Allen CPI—falls significantly with temperature: a slope of −0.23 ($p < 0.001$) on the ten-year temperature mean, robust to a linear year trend (−0.25) and a quadratic trend (−0.13),

both significant (Table A23). Cold decades raised the real value of salt. As a check on the deflator, we also build an internal deflator from the 136 other commodities in the Allen–Unger database; it correlates with the Allen CPI at 0.94. That internal deflator yields a smaller, insignificant slope (-0.07), because the commodity basket is itself weighted toward goods (fuel among them) whose real prices also rose in cold periods; the consumption-weighted Allen CPI is the appropriate, and the more conservative, choice. Figure A8 plots the nominal silver series, the Allen CPI, and the deflated real salt price over time.

Why not a market-level analysis. One might hope to sharpen the test by comparing coastal- and inland-supplied markets directly. We examined this but do not rely on it. The market-level series are few (eleven usable markets) and unbalanced; they carry no year fixed effects, so they do not difference out the Price Revolution; salt was the most heavily and variably taxed commodity in early-modern Europe, so quotations—French ones in particular, which reflect the *gabelle*—confound price with fiscal policy; and coastal "bay salt" differed in grade from refined inland salt. A coastal–inland contrast that appears significant in a yearly-means specification does not survive the inclusion of market and year fixed effects. We therefore rest the price evidence on the aggregate real price, which the external deflator identifies cleanly.

Table A23: Real Salt Price and Temperature

	Slope on temp.	SE	p
Nominal silver price	$+0.58$	0.11	< 0.001
Real (internal commodity-basket)	-0.07	0.05	0.150
Real (external Allen CPI)	$\mathbf{-0.23}$	**0.04**	$\mathbf{< 0.001}$
+ linear year trend	-0.25	—	< 0.001
+ quadratic year trend	-0.13	—	0.006

Notes: Dependent variable is the annual European salt-price index (log deviation, grams of silver); "real" rows deflate by the indicated consumer price index. "Temperature" is the Guiot and Corona (2010) ten-year lagged European-mean anomaly. 500 annual observations, 1300–1799. The external deflator is Robert Allen's seven-city silver consumer price index; the internal deflator is built from the 136 other commodities in the Allen–Unger database (correlation 0.94 with the Allen CPI). The nominal silver row reflects the Price Revolution rather than salt scarcity. Heteroskedasticity-robust standard errors; trend rows add a linear / quadratic year polynomial. A negative slope means salt is dearer in real terms when the climate is colder. $***\ p < 0.01$.

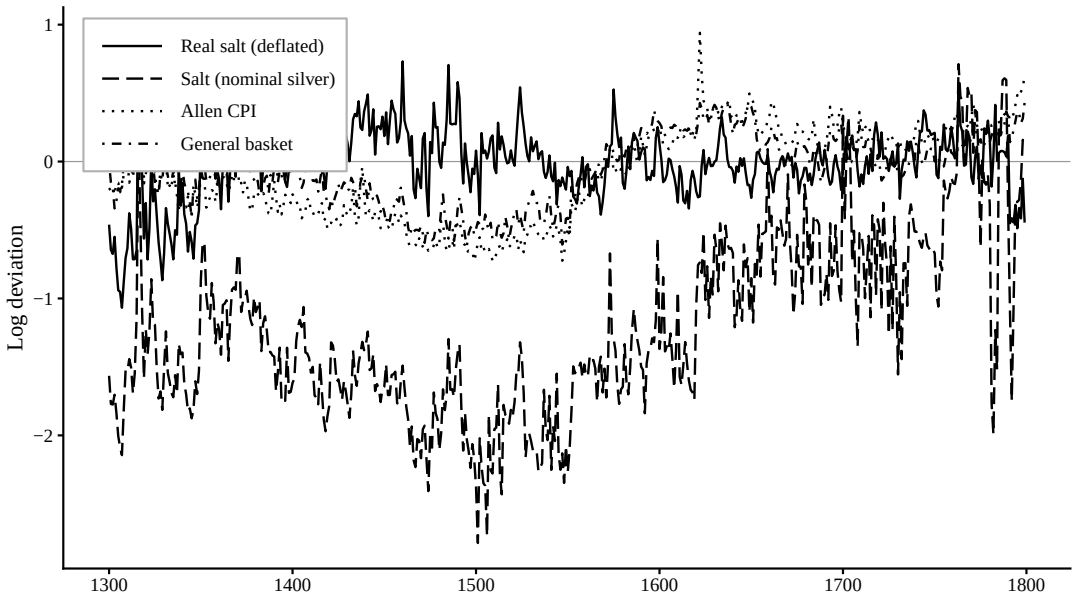

Figure A8: Nominal Silver, the Allen CPI, and the Real Salt Price, 1300–1799

Notes: Annual log deviations from series means. Solid line = real salt price (Allen-CPI deflated); long-dash = nominal silver salt price (dominated by the upward Price-Revolution trend); dotted = Allen consumer price index; dot-dash = internal commodity-basket deflator. The nominal silver price drifts up with the Price Revolution, while the deflated real salt price falls in cold periods. Sources: Allen–Unger commodity prices; R. Allen consumer price indices; Guiot and Corona (2010) temperature.

W Coastal Salt as a Substitute: Sea Salinity Analysis

Our identification strategy rests on the claim that colder temperatures increased the relative value of inland mineral salt by reducing coastal salt production. A direct implication is that the salt × temperature mechanism should be weaker in areas where productive coastal salt works provided a viable substitute for inland deposits. We test this prediction using data on sea surface salinity—a key determinant of the productivity of coastal salt works via solar evaporation.

We compute mean sea surface salinity (PSU) within a 5km coastal buffer of each grid cell, drawn from oceanographic raster data. Of our 476 grid cells, 184 are coastal (salinity > 0) and 292 are landlocked. We define "high salinity" cells as those above the coastal median (33.88 PSU), corresponding to Mediterranean and Atlantic coastal areas where solar evaporation was historically productive.

Heterogeneity by coastal substitute availability. Table A24 tests whether the salt × tem-

perature effect is attenuated in cells near productive coastal salt sources, using a triple interaction: salt × temperature × high-salinity neighbor. The results strongly confirm the substitute hypothesis. At the 10-year window, the main salt × temperature interaction is −1.655 ($p < 0.001$), but the triple interaction is +1.567 ($p < 0.001$)—nearly offsetting. The net effect near a high-salinity coast is approximately zero. This pattern is robust across temperature windows.

Table A24: Salt × Temperature: Heterogeneity by Coastal Salt Availability

	Borders			Battles	
	(1) 10yr	(2) 25yr	(3) 50yr	(4) Year FE	(5) Cell FE
Salt × Temperature	−1.655 ∗ ∗∗	−2.185 ∗ ∗∗	−2.664 ∗ ∗∗	−0.0125 ∗ ∗∗	−0.0130 ∗ ∗∗
	(0.027)	(0.032)	(0.035)	(0.002)	(0.002)
Salt × Temp × High-Sal. Neighbor	+1.567 ∗ ∗∗	+2.115 ∗ ∗∗	+2.623 ∗ ∗∗	+0.0024	+0.0030
	(0.035)	(0.041)	(0.045)	(0.003)	(0.003)
Cell FE	Yes	Yes	Yes	No	Yes
Period FE	Yes	Yes	Yes	Yes	Yes
Observations	54,796	53,560	51,500	54,796	54,796
R^2 (Within)	0.082	0.101	0.120	0.012	0.012

Notes: High-salinity neighbor equals one if the cell or any queen-adjacent cell has above-median sea surface salinity (> 33.88 PSU) within a 5km coastal buffer; 138 of 476 cells meet this criterion. Temperature is the Guiot–Corona rolling mean at the indicated window. Columns (1)–(3) and (5) include cell fixed effects; column (4) uses year fixed effects only. Standard errors in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Robustness to coastal controls. Table A25 adds progressively stringent coastal controls. The baseline salt × temperature coefficient of −0.724 is essentially unchanged when controlling for coastal × temperature (−0.736), continuous sea salinity × temperature (−0.753), high salinity × temperature (−0.749), or a coastal × time trend (−0.735). If anything, controlling for coastal trends makes the inland salt coefficient slightly *larger* in magnitude.

X Confound Controls: Urbanization and Black Death

X.1 Urbanization

Table A26 examines whether the salt × temperature mechanism operates through urbanization. Panel A regresses log urban population (from Buringh 2021) on the salt × temperature interaction. The coefficient is null at every window (−0.015 at 10 years, +0.023 at 25 years, −0.024 at 50 years;

Table A25: Robustness of Salt × Temperature to Coastal Controls

	Dependent Variable: Border Count				
	(1) Baseline	(2) + Coastal×T	(3) + Salinity×T	(4) + HighSal×T	(5) + Coast Trend
Salt × Temperature	−0.724 ∗ ∗∗	−0.736 ∗ ∗∗	−0.753 ∗ ∗∗	−0.749 ∗ ∗∗	−0.735 ∗ ∗∗
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Coastal × Temperature		0.376 ∗ ∗∗			
		(0.040)			
Salinity × Temperature			0.014 ∗ ∗∗		
			(0.001)		
High Salinity × Temperature				0.466 ∗ ∗∗	
				(0.048)	
Coastal × Trend					−0.160 ∗ ∗∗
					(0.006)
Cell FE	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes
Observations	55,860	55,860	55,860	55,860	55,860
R^2 (Within)	0.047	0.049	0.050	0.049	0.058

Notes: All specifications use cell and period fixed effects with the Guiot–Corona 10-year temperature window. Column (1) reproduces the baseline from Table 1, column (2). Coastal = 1 if sea salinity > 0 within 5km buffer (184 of 476 cells). Salinity is continuous sea surface salinity (PSU). High salinity = 1 if salinity exceeds the coastal median (33.88 PSU; 92 cells). Coastal trend = coastal × (year −1100)/100. Standard errors in parentheses. ∗ ∗ ∗ $p < 0.01$, ∗∗ $p < 0.05$, ∗ $p < 0.1$.

none significant), confirming that the mechanism is political, not economic. Panel B decomposes by salt type: no type × temperature interaction significantly predicts urbanization.

X.2 Black Death Mortality Control

The Black Death (1347–1351) raised per capita incomes through the Malthusian mechanism, potentially increasing demand for salt (for meat preservation) and affecting state consolidation through the tax revenue channel (Jedwab et al., 2022). Since Black Death mortality was lower in central Europe—where brine salt is concentrated (correlation with brine count = −0.124, $p = 0.007$)—this demand-side variation could confound the brine × temperature results.

We measure Black Death mortality using IDW spatial interpolation (power = 1.76) from 274 European cities with recorded mortality rates, using data from Jedwab et al. (2024) and originally created by Noel Johnson. The interpolated measure ranges from 4 to 93 percent across grid cells

Table A26: Salt × Temperature and Urbanization

	Dependent Variable: log(Urban Population)		
	(1) 10yr	(2) 25yr	(3) 50yr
Panel A: Aggregate salt × temperature			
Salt × Temperature	−0.015	+0.023	−0.024
	(0.018)	(0.019)	(0.023)
Panel B: Type × temperature			
Rock × Temperature	+0.050	+0.020	−0.200
	(0.114)	(0.117)	(0.134)
Brine × Temperature	−0.039	+0.047	+0.008
	(0.032)	(0.032)	(0.040)
Coastal/marine × Temperature	+0.027	+0.007	−0.059
	(0.032)	(0.040)	(0.044)
Cell FE	Yes	Yes	Yes
Period FE	Yes	Yes	Yes
Observations	3,708	3,708	3,708

Notes: Dependent variable is log urban population from Buringh (2021). Sample restricted to cell-years with nonzero Buringh data (3,708 observations across 35 cells). All specifications include cell and period fixed effects. Temperature is the Guiot–Corona rolling mean at the indicated window. Panel B enters all three type × temperature interactions simultaneously. No coefficient is significant at the 10% level. Urbanization is not included as a control in main specifications because it is potentially downstream of salt endowments (post-treatment bias). Standard errors in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

(mean = 43.4, SD = 4.0), with full coverage of all 476 cells.

Table A27 shows that the salt × temperature results are robust to controlling for Black Death mortality. Panel A presents the aggregate specification: the baseline coefficient of −0.724 attenuates by only 5% to −0.689 when controlling for both BD mortality × temperature and BD mortality × time trend (column 4). Panel B presents the type × temperature specification: the opposing signs are fully preserved (rock +0.889, brine −1.903, coastal/marine +0.120), and the BD × temperature coefficient is small and insignificant (+0.002) once salt types are controlled for. The salt–temperature mechanism is not an artifact of differential Black Death exposure.

Table A27: Salt × Temperature: Controlling for Black Death Mortality

	Panel A: Aggregate				Panel B: By Type	
	(1) Baseline	(2) + BD×T	(3) + BD×Trend	(4) + Both	(5) Baseline	(6) + BD×T
Salt × Temp	−0.724 ∗ ∗∗	−0.699 ∗ ∗∗	−0.688 ∗ ∗∗	−0.689 ∗ ∗∗		
	(0.018)	(0.018)	(0.018)	(0.018)		
Rock × Temp					+0.910 ∗ ∗∗	+0.889 ∗ ∗∗
					(0.106)	(0.107)
Brine × Temp					−1.898 ∗ ∗∗	−1.903 ∗ ∗∗
					(0.032)	(0.032)
Coastal/marine × Temp					+0.138 ∗ ∗∗	+0.120 ∗ ∗∗
					(0.033)	(0.033)
BD Mort. × Temp		+0.067 ∗ ∗∗		−0.006		+0.002
		(0.005)		(0.005)		(0.001)
BD Mort. × Trend			−0.033 ∗ ∗∗	−0.034 ∗ ∗∗		
			(0.001)	(0.001)		
Cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	54,796	54,796	54,796	54,796	54,796	54,796
R^2 (Within)	0.043	0.047	0.077	0.077	0.076	0.076

Notes: Dependent variable: border count. All specifications include cell and period fixed effects. Temperature is the 10-year Guiot–Corona rolling mean. BD mortality is the IDW-interpolated Black Death mortality rate (0–100 scale) from 274 European cities (Jedwab et al., 2024). BD mortality is time-invariant and absorbed by cell FE; its interactions with temperature and time trends are time-varying and identifiable. Brine cells have significantly lower BD mortality ($r = -0.124$, $p = 0.007$), reflecting the concentration of brine in central Europe where plague mortality was lower. Despite this correlation, the salt × temperature coefficient is essentially unchanged. Standard errors in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Y The Salt–Fragmentation Relationship Intensifies as Temperatures Decline

Table A28 presents century-specific salt coefficients from the cell fixed effects specification, with the 12th century as the reference period. The salt–fragmentation relationship is present throughout the sample period, consistent with the model's prediction that salt resources generate political competition. However, the relationship intensifies markedly as European temperatures decline: the aggregate coefficient nearly doubles from the 13th century (+0.55) to the 14th century (+1.01), coinciding with the transition from the Medieval Climate Anomaly to the cooler, more volatile conditions of the Little Ice Age documented in our temperature reconstructions (Figure A2). The

coefficient then plateaus at approximately 1.1 from the 15th through 18th centuries.

The type decomposition reveals that this intensification is driven by brine salt, whose coefficient rises from +0.76 in the 13th century to a peak of +2.22 in the 16th century and remains near that level through the 18th. Rock salt cells show the opposite trajectory: initially positive (+0.42 in the 14th century), the rock coefficient declines to near zero by the 16th century and turns slightly negative thereafter—consistent with rock salt areas consolidating as the fiscal advantage of appropriable resources grows during cold periods. This progressive divergence between brine and rock is inconsistent with any time-invariant confounder and instead reflects the gradual intensification of the appropriability mechanism as the Little Ice Age deepened.

Table A28: Century-Specific Salt Coefficients on Border Count (Cell + Year FE)

	(1) All Salt	(2) Rock	(3) Brine	(4) Coastal/marine
Reference: 12th century (1100–1199) = 0				
Salt × 13th century	+0.545 ∗ ∗∗	+0.413 ∗ ∗∗	+0.764 ∗ ∗∗	+0.190 ∗ ∗∗
	(0.019)	(0.109)	(0.031)	(0.045)
Salt × 14th century	+1.006 ∗ ∗∗	+0.422 ∗ ∗∗	+1.493 ∗ ∗∗	+0.269 ∗ ∗∗
	(0.019)	(0.109)	(0.031)	(0.045)
Salt × 15th century	+1.083 ∗ ∗∗	+0.362 ∗ ∗∗	+1.819 ∗ ∗∗	+0.099 ∗ ∗
	(0.021)	(0.116)	(0.033)	(0.048)
Salt × 16th century	+1.155 ∗ ∗∗	+0.048	+2.223 ∗ ∗∗	−0.180 ∗ ∗∗
	(0.019)	(0.109)	(0.031)	(0.045)
Salt × 17th century	+1.106 ∗ ∗∗	−0.149	+2.208 ∗ ∗∗	−0.335 ∗ ∗∗
	(0.019)	(0.109)	(0.031)	(0.045)
Salt × 18th century	+1.060 ∗ ∗∗	−0.116	+2.136 ∗ ∗∗	−0.404 ∗ ∗∗
	(0.020)	(0.110)	(0.031)	(0.046)
Cell FE	Yes	Yes	Yes	Yes
Period FE	Yes	Yes	Yes	Yes
Observations	64,260	64,260	64,260	64,260

Notes: Each column reports coefficients from a separate regression of border count on salt count × century dummies, with the 12th century (1100–1199) as the omitted reference. Column (1) uses total salt count; columns (2)–(4) use type-specific counts (each type estimated separately). All specifications include cell and period fixed effects. The aggregate coefficient (column 1) nearly doubles from the 13th to the 14th century, coinciding with the onset of European cooling. The brine coefficient (column 3) rises steadily to its 16th-century peak and stays near that level thereafter, while the rock coefficient (column 2) reverses from positive to near-zero, consistent with the appropriability mechanism intensifying as the Little Ice Age deepened. Standard errors in parentheses. $***$ $p < 0.01$, $**$ $p < 0.05$, $*$ $p < 0.1$.

Z Randomization Inference

With 70 salt cells out of 476, inference could be sensitive to the specific spatial distribution of salt deposits. We assess this with a permutation test: we randomly reassign salt counts across cells (preserving the panel structure and the total amount of salt in Europe), re-estimate the salt × temperature interaction, and repeat 1,000 times. If the true coefficient lies in the far tail of the permutation distribution, the result is unlikely to arise from any arbitrary set of 70 cells interacting with temperature.

Table A29 reports the results. The true aggregate salt × temperature coefficient on borders (−0.724) exceeds all but one of 1,000 permuted coefficients ($p = 0.001$). For battles, zero permutations produce a coefficient as extreme as the true value ($p < 0.001$). The type-specific results are equally robust: the brine × temperature coefficient (−1.898) exceeds every permutation ($p < 0.001$), and the rock × temperature coefficient (+0.910) exceeds all but two permutations on a one-sided test ($p = 0.002$). The salt–temperature mechanism cannot be generated by randomly assigning salt to any 70 cells in Europe.

Table A29: Randomization Inference: Permutation Test (1,000 Permutations)

	True Coefficient	Permutation p (two-sided)	Permutation p (one-sided)
Panel A: Aggregate salt × temperature			
Borders	−0.724	0.001	0.001
Battles	−0.013	< 0.001	< 0.001
Panel B: Type × temperature on borders			
Brine × Temp	−1.898	< 0.001	< 0.001
Rock × Temp	+0.910	0.041	0.002

Notes: Each permutation randomly reassigns salt counts (total, rock, brine, coastal/marine) across the 476 grid cells, preserving the panel structure and re-estimating the salt × temperature interaction with cell and period fixed effects (10-year Guiot–Corona temperature window). The two-sided p-value is the share of permutations where $|\hat{\beta}_{\text{perm}}| \geq |\hat{\beta}_{\text{true}}|$. The one-sided p-value tests the directional hypothesis (negative for borders/battles/brine, positive for rock). 1,000 permutations, seed = 20260428.

4 HHI Appendix: Concentration IV and Robustness

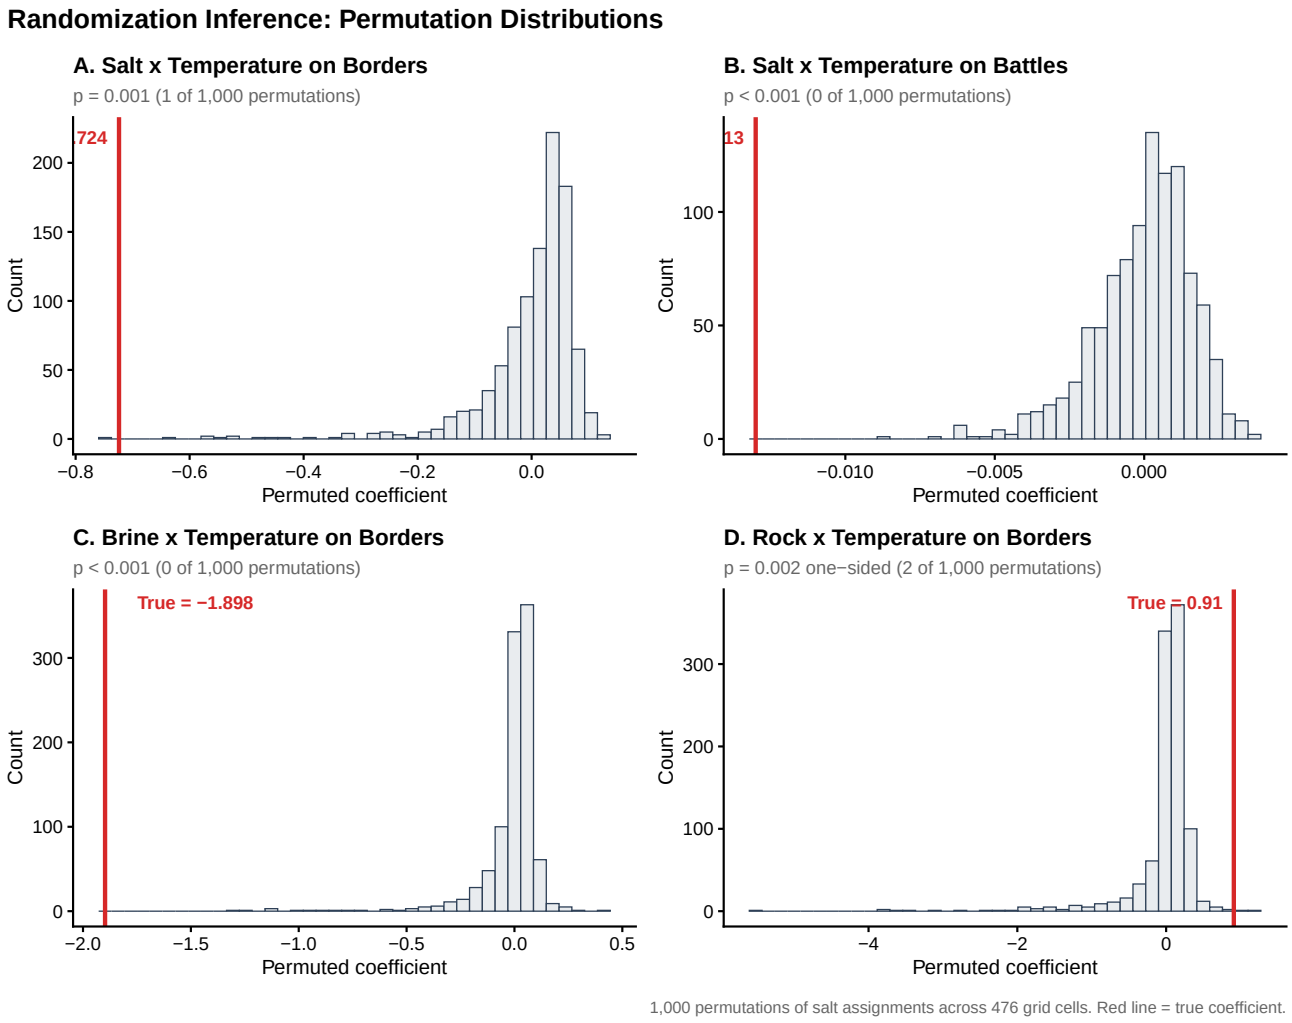

Figure A9: Permutation Distributions: True Coefficients vs. Random Salt Assignments

Notes: Each histogram shows the distribution of coefficients from 1,000 random reassignments of salt counts across 476 grid cells. The red line marks the true coefficient. Panel A: aggregate salt × temperature on borders ($p = 0.001$). Panel B: aggregate salt × temperature on battles ($p < 0.001$). Panel C: brine × temperature on borders ($p < 0.001$). Panel D: rock × temperature on borders ($p = 0.002$, one-sided). All specifications include cell and period fixed effects with the 10-year Guiot–Corona temperature window.

This appendix collects the instrumental-variables and robustness evidence for the concentration-based (HHI) results. Throughout, the outcome is the within-cell Herfindahl index of polity shares, normalized among the polities present so that it is bounded in $(0, 1]$; consolidation is a negative salt × temperature coefficient and fragmentation positive; all three type × temperature interactions are entered simultaneously with cell and period fixed effects. Standard errors are clustered by grid cell, with Conley (1999) spatial-HAC errors reported as a robustness check.

Grid resolution. Table A30 re-estimates the type table on a finer 150km grid. This serves two purposes. First, it is a spatial-resolution robustness check. Second, on the 150km grid cardinally-adjacent cell centroids are 150km apart, so the Conley bandwidth can capture genuine cross-cell spatial correlation at a sub-200km radius—something the 200km grid cannot, since there its nearest neighbours sit at roughly 170–215km and any smaller radius collapses toward clustering by cell. Brine fragmentation and rock consolidation both reappear with the predicted sign and significance on the finer grid; the coastal arm does not.

Table A30: Type-Specific Salt × Temperature on HHI: 150km Grid

	Dependent Variable: HHI polity concentration		
	(1) 10yr	(2) 25yr	(3) 50yr
Rock × Temperature (pred. −)	−0.0428	−0.0608*	−0.0765*
	(0.0262)[0.0241]	(0.0342)[0.0312]	(0.0437)[0.0398]
Brine × Temperature (pred. +)	+0.0181**	+0.0250**	+0.0322**
	(0.0088)[0.0108]	(0.0123)[0.0149]	(0.0152)[0.0180]
Coastal/marine × Temperature (pred. −)	+0.0080	+0.0088	+0.0081
	(0.0053)[0.0053]	(0.0062)[0.0060]	(0.0067)[0.0060]
Cell FE, Period FE	Yes	Yes	Yes
Observations	80,601	80,601	80,601

Notes: Dependent variable is the (normalized) within-cell Herfindahl index of polity shares, bounded in $(0, 1]$ (higher = more consolidated). On HHI, consolidation is a *negative* salt × temperature coefficient and fragmentation *positive*; predicted signs after each row label. All three type × temperature interactions are entered simultaneously with cell and period fixed effects (150km grid). Standard errors clustered by grid cell in (parentheses), with Conley (1999) spatial-HAC standard errors at 150km in [brackets] as a robustness check; stars are based on the clustered standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Instrumental variables. Table A31 re-estimates the type-specific 2SLS of Section 5.5 with HHI as the outcome, instrumenting brine × temperature with Keuper (Upper Triassic) geology × temperature and coastal × temperature with sea-surface salinity × temperature. The brine instrument is strong (cluster-robust first-stage $F \approx 44$, matching Table 4, since the first stage is the same) and the 2SLS brine coefficient is positive and significant across windows, confirming that the fragmenting effect of brine on concentration is causally identified. The coastal instrument is borderline ($F \approx 10$): the coastal 2SLS is negative (consolidation) at the 25- and 50-year windows but not statistically significant. The OLS rows enter each type's interaction on its own (with the

own-type count), matching the IV specification; they therefore differ slightly from the simultaneous three-type OLS of Table 3.

Table A31: Type-Specific IV: Salt × Temperature on HHI Concentration

	Dependent Variable: HHI concentration		
	(1) 10yr	(2) 25yr	(3) 50yr
Panel A: Brine × Temp (IV = Keuper × temp; pred. +)			
OLS	+0.0219**	+0.0275**	+0.0281**
	(0.0092)	(0.0112)	(0.0128)
2SLS	+0.0283***	+0.0360***	+0.0430***
	(0.0089)	(0.0107)	(0.0115)
First-stage F	44.0	44.8	44.4
Panel B: Coastal × Temp (IV = sea-salinity × temp; pred. −)			
OLS	−0.0125	−0.0231**	−0.0359***
	(0.0091)	(0.0102)	(0.0100)
2SLS	+0.0068	−0.0188	−0.0438
	(0.0382)	(0.0494)	(0.0619)
First-stage F	9.8	10.2	10.0
Own-type count, Temp, Cell FE, Period FE	Yes	Yes	Yes

Notes: 2SLS of HHI polity concentration on a type × temperature interaction, full 476-cell grid, cell and period fixed effects. The brine interaction is instrumented by Keuper (Upper Triassic) geology × temperature; the coastal interaction by sea-surface salinity × temperature. Higher HHI = more consolidated, so fragmentation is a *positive* coefficient and consolidation *negative*. Cluster-robust (by cell) first-stage F on the excluded instrument, matching Table 4; the Stock and Yogo (2002) 10% critical value of 16.38 assumes homoskedasticity, and by that benchmark the brine instrument is strong and the coastal instrument borderline. Standard errors clustered by grid cell in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Salt-specificity placebos. Because the rock and coastal consolidation effects operate at the secular frequency of the Little Ice Age cooling, they weaken under aggressive detrending; we defend them instead by showing they are specific to salt. Table A32 re-runs each arm against a geographic/geological placebo (generic coastal geography for the coastal arm, evaporite-basin coverage for the rock arm, and non-salt mineral deposits for brine): the salt interactions retain their sign and significance in every specification but one, and the placebos are null throughout with one exception: at the 50-year window the rock arm loses significance and the brine mineral placebo becomes marginally significant.

Table A32: Salt-Specificity Placebo Tests (HHI Concentration)

	Dependent Variable: HHI Concentration	
	(1) 25yr	(2) 50yr
A. Coastal arm		
Coastal salt × Temp	−0.0230**	−0.0344***
	(0.0099)	(0.0112)
Generic coastal × Temp (placebo)	−0.0092	−0.0169
	(0.0142)	(0.0194)
B. Rock arm		
Rock mine × Temp	−0.0918**	−0.0947
	(0.0462)	(0.0671)
Evaporite basin × Temp (placebo)	−0.0156	−0.0350
	(0.0352)	(0.0515)
C. Mineral placebo		
Brine × Temp (retained)	+0.0373**	+0.0449**
	(0.0174)	(0.0199)
Mineral deposits × Temp (placebo)	−0.0025	−0.0055*
	(0.0023)	(0.0030)
Cell FE	Yes	Yes
Period FE	Yes	Yes
Observations	46,718	45,238

Notes: HHI polity concentration; consolidation is a *negative* coefficient. Each panel is one regression with all three type × temperature interactions, cell and period fixed effects, plus the placebo interaction. The salt term is the effect; the placebo term is the same interaction for comparable non-salt cells (generic coastal geography in A, evaporite-basin geology in B, other mineral deposits in C). Salt effects are significant and their placebos are not, showing the response is salt-specific rather than a generic secular trend. Conley (1999) spatial HAC standard errors (200km bandwidth) in parentheses; t-statistics are ∼2–3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Robustness battery. Table A33 reports the three type × temperature coefficients on HHI across the same battery used for borders. Brine fragmentation is positive throughout and significant against the full confounder set (Black Death, other minerals, geography, coastal proximity/salinity) and under century dummies; it weakens under the CCSM4 and MPI-ESM reconstructions and when Western Europe is dropped. Rock and coastal consolidation hold against the confounder controls but, consistent with their secular-frequency nature, weaken under detrending and alternative temperature series. We were unable to compute HHI on the 150km grid, which would require rebuilding the cell–polity crosswalk at that resolution; the 150km robustness check

is therefore available for borders only.

Table A33: Robustness of the Type-Specific Salt × Temperature Result (HHI Concentration)

	Rock × Temp (pred. −)	Brine × Temp (pred. +)	Coastal × Temp (pred. −)
Confounder controls (full sample)			
Baseline	-0.0937^{**}	$+0.0321^{***}$	-0.0242^{**}
+ Black Death mortality	-0.0926^{**}	$+0.0332^{***}$	-0.0246^{**}
+ Other minerals	-0.0921^{**}	$+0.0368^{***}$	-0.0239^{**}
+ Geography (elevation, ruggedness)	-0.0988^{**}	$+0.0340^{***}$	-0.0273^{***}
+ Coastal proximity, sea salinity	-0.0957^{**}	$+0.0316^{***}$	-0.0251^{**}
+ Urbanization (benchmark years)	$+0.0214$	$+0.0203^{**}$	-0.0130
Alternative temperature series			
CCSM4 reconstruction	-0.0371	$+0.0158$	-0.0243^{**}
MPI-ESM reconstruction	$+0.0095$	$+0.0079$	-0.0074
Functional form and spatial resolution			
+ Type × century dummies	-0.0234	$+0.0309^{***}$	$+0.0046$
Poisson (100×HHI)	-0.1430^{*}	$+0.0537^{*}$	-0.0341^{**}
150km grid	-0.0608^{*}	$+0.0250^{**}$	$+0.0088$
Leave-one-region-out (drop Western)	-0.1066^{**}	$+0.0073$	-0.0191

Notes: Each row re-estimates the type-specific specification with *HHI polity concentration* as the outcome at the 25-year temperature window, all three type × temperature interactions entered simultaneously, cell and period fixed effects. On HHI, consolidation is a *negative* coefficient and fragmentation *positive*; predicted signs in the column heads. Standard errors clustered by grid cell underlie the stars. The Poisson row models 100×HHI as a quasi-count (semi-elasticities; signs and significance comparable, magnitudes not). Confounder controls enter as the named variable × temperature. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.