Doomsday to Today: 1,000 Years of Spatial Inequality

Luke Heath Milsom*

15th November, 2025

Abstract

To what extent are modern within-state spatial inequalities related to those in the distant past? Using granular data from the Doomsday Book, I show that areas of England that were 10% richer in 1086 are, on average, still between 1% and 2% richer almost 1,000 years later in 2020. This relationship is not mediated by physical characteristics of place, or pre-existing infrastructure, but rather by time-invariant measures of local market access. Using a historical natural experiment, the Harrying of the North, and a dynamic quantitative spatial economics model, I find suggestive evidence that correlated fundamentals rather than path-dependence can explain the observed persistence.

*KU Leuven and the Institute for Fiscal Studies, luke.heathmilsom@KULeuven.be. With thanks to Stephen Broadberry, Erik Buyst, Yannis Ioannidis, Frederic Vermeulen, Sebastiaan Maes, Frank Verboven, Frederik Maertens, Walker Hanlon, Sebastian Ellingsen, Stephan Heblich, Stephan Heblich, Hannah Zillessen, and seminar participants in Oxford, the IFS, KU Leuven, and the Urban Economics Association 2025 Berlin meeting, for helpful comments and discussions.

Large within-state spatial income inequalities are a pervasive feature of today's economies.[1] In this paper, I ask to what extent patterns of geographic inequality were already in place long ago. I focus on one particularly egregious example: England, and compare two points in time almost 1,000 years apart.[2] It is possible to answer these questions in this setting because of an astonishingly detailed historical source, the Doomsday Book [Delabastita and Maes, 2023]. This document was commissioned by William the Conqueror after his successful invasion of England in order to survey the value of his new possessions and raise tax revenue, and as such, it contains information on incomes at a fine geography from 1086.

Using this uniquely granular and long-run data from the Doomsday Book, I show that areas that are richer today were also richer 1,000 years ago. In particular, I find that areas that were 10% richer in 1086 remain, on average, between 1% and 2% richer today. This relationship is robust to a large number of potential specifications and a barrage of corrections for the possibility of spatial correlation causing a spurious relationship. In particular, I fit a spatial autoregressive model, adjust for Conley [Conley, 1999] standard errors, use the approach suggested by Müller and Watson [2022], and finally perform the semi-parametric thin plate spline correction due to Kelly et al. [2023].

I next turn to investigate potential mediating characteristics of place. To do this, I build a dataset of location-specific time-invariant characteristics encompassing: terrain ruggedness, pre-existing infrastructure in the form of mills or fisheries, Roman roads, Roman towns, proximity to the coast and London, soil fertility, and measures capturing a location's market access. I find that variables capturing local geographic features such as ruggedness and soil fertility, as well as those measuring pre-existing infrastructure, cannot explain the observed long-run persistence. However, measures of market access can explain about half of the persistence.

Having established long-run persistence and highlighted some potential mediating factors, I now turn to the challenging question of more systematically disentangling the channel through which this persistence operates. There are two general competing explanations as to why areas richer in 1086 could also be richer today. It may be due to *path-dependence* "we are richer today because we were richer yesterday" or, it could be due to *correlated fundamentals* "we are richer today because of local characteristic X, and this was also why we were rich yesterday" [Lin and Rauch, 2022].[3] Distinguishing between these competing

[1] For example: incomes in Massachusetts are on average almost double those in Mississippi, those in Shanghai are over three times those in Gansu province, and residents of South Tyrol in Italy enjoy incomes over twice those of their compatriots in Campania. Regional income data is from the St Louis Fed, the Chinese National Bureau of Statistics, and Eurostat for the US, China, and Italy, respectively.

[2] Today England is arguably one of the most regionally unequal of all industrialised economies [McCann, 2020], with the highest coefficient of variation in regional (NUTS1) gross value added in Western Europe [Stansbury et al., 2023] and a Northeast-Southeast healthy life gap of over 6 years [Bambra, 2016].

[3] Note that this dichotomisation applies to the relationship between any two periods and does not disallow

explanations is difficult, and particularly so over such an exceptionally long time frame. To lend credibility, I combine evidence from two orthogonal approaches, but the evidence presented remains suggestive in nature.[4]

First, I use evidence from a dynamic quantitative spatial economics model following Allen and Donaldson [2022]. This framework endogenises the spatial distribution of economic activity and embeds the two possible channels by allowing for correlated local fundamentals and path-dependence. Data from the Doomsday Book is sufficiently rich to allow me to invert the model and back out location-specific fundamental productivities and amenities in 1086, and modern data allows a similar procedure in 2020. However, historical data is not sufficiently rich to allow for the estimation of parameters in the economy under consideration, nor for such parameters to be allowed to vary over time. Instead, I calibrate the model using parameters from Allen and Donaldson [2022] and show robustness to parameters used by Ellingsen [2025] and random parameter perturbations. Using this framework, I perform two experiments. First, I shut down the potential role of correlated fundamentals by constructing counterfactual modern-day income distributions solving forward from the observed 1086 distribution and fundamentals but using random 2020 fundamentals. These counterfactual income distributions preserve the role of path dependence but disallow fundamentals from explaining any persistence. Regressing 1086 values on these counterfactual 2020 values, I find no correlation. Second, I perform the symmetric exercise, shutting down path-dependence by randomising the initial population distribution and solving the model forward but retaining correlated fundamentals. In this exercise, I find the same long-run correlation as with the actual data.

Secondly, I combine this evidence with that from an orthogonal approach: leveraging a natural experiment in the form of a large-scale localised spatial shock: The Harrying of The North. If the impact of such a shock can not be felt today, it provides evidence against path dependence and for local fundamentals.[5] In the early years of the conquest, many areas of England rebelled against William's rule. However, only one of these rebellions was brutally put down (that is, "Harried"[6]), a rebellion in the north of England in 1069/70. The reaction

different mechanisms affecting alternate or intermediate comparisons.

[4]Note that mediation analysis, controlling for some historical characteristics and observing whether the relationship attenuates, does not, in general, allow one to separately identify path dependence from fundamentals. Controlling for a characteristic may attenuate the relationship because in both 1086 and today, that characteristic was/is predictive of local incomes. However, it could also be that in 1086 the characteristic was important (or in some later period prior to 2020), and this led to the formation of agglomerations and infrastructure that persists today. In this second state of the world, the characteristic would still attenuate the long-run relationship without being a correlated fundamental. This identification issue will plague the use of any observed historical variable. The logic of this argument is similar to that employed by Casey and Klemp [2021] in their critique of persistence studies that take an instrumental variables approach.

[5]This is a common strategy employed in the literature, for example, see Davis and Weinstein [2002].

[6]The word harried comes from the verb to harry, which itself is from the Middle English word herigan, meaning to pillage, to plunder, or to make war.

from William was catastrophic; some commentators estimate that as much as three-quarters of the population of Yorkshire was killed or fled as a result [Strickland, 1998]. Comparing regions affected by the Harrying to other rebellious locations that didn't experience such reprisals, I find that the impact can still very much be felt some 16 years later, using data from the Doomsday Book. However, the impact can not be detected using modern data, implying that the Harrying did not result in a long-run change. This gives empirical evidence against path dependence and for fundamentals in this context.

Taken together, these approaches provide some suggestive evidence in favor of a large role for correlated fundamentals and against path-dependence. However, this evidence remains suggestive; both approaches have limitations. The model exercise relies on a framework calibrated to a different economy, with time-invariant parameters, that cannot capture structural changes in the English economy, such as industrialization, but rather takes them as given. The Harrying exercise relies on the quasi-randomness of William's decision to Harry only the northern rebellion, evidence for which comes only from historical discourse, and an almost 1,000-year-long parallel trends assumption.

This paper contributes to three main strands of the literature. First, it contributes to the literature on long-run persistence (see, for example, Cioni et al. [2022] for an overview). Following [Acemoglu et al., 2001], various papers have linked historic variables to long-run outcomes. See, for example, Alesina et al. [2013] on the plow and gender roles, Dell [2010] on Peru's "mita" and local development, Galor and Özak [2016] on crop suitability and time preference, Nunn and Wantchekon [2011] on the slave trade and mistrust, or Guiso et al. [2016] on city self-governance and civic capital. Closer to this paper are those that compare, in the language of Voth [2021], apples-to-apples, i.e., the same variable over time, such as Comin et al. [2010], Chen et al. [2020], Valencia Caicedo [2019] or Voigtländer and Voth [2012]. I contribute to this literature by considering very long-run persistence in a variable of central importance to economics that has thus received relatively little attention in this literature due primarily to a lack of data: income.[7] In addition, I take very seriously the issue of spatial correlation and consider the mechanisms through which the estimated persistence operates.

Second, I contribute to the literature on the determinants of the spatial distribution of economic activity, and in particular on the path dependence vs fundamentals debate [Lin and Rauch, 2022]. For example, using shocks to city population Davis and Weinstein [2002], Miguel and Roland [2011], Jedwab et al. [2024] find evidence of quick rebounding, suggesting a strong role for fundamentals. Turning to path-dependence, Bleakley and Lin

[7] Perhaps closest to this, Lee and Lin [2018] considers persistence in the spatial distribution of incomes in the US from 1880 to 2010. However, as well as considering a much shorter time span, Lee and Lin [2018] are mainly concerned with explaining persistence in the within-city distribution of income, whereas this paper is concerned with the cross-city or cross-area income distribution.

[2012], Heblich et al. [2021] and Jedwab et al. [2017] find evidence of persistence in the face of obsoletism, which is evidence for path-dependence. However, Michaels and Rauch [2018], Redding and Sturm [2008], Gibbons et al. [2024] and Dell [2010] find evidence that large shocks to local fundamentals can change long-run outcomes, providing some evidence for the fundamentals channel. In this paper, I attempt to disentangle these two possibilities in my setting by combining evidence from a historical natural experiment with that from a dynamic quantitative spatial economics model.

Finally, this paper contributes to the large literature on the economic history of England since the Norman Conquest, a body of work that spans history, economic history, and modern quantitative economics, and of which a comprehensive review is far beyond the scope of this paper. Foundational scholarship has documented the broad evolution of the English economy—from medieval agrarian structures and feudal institutions [Darby et al., 1979, Daunton, 2001] to the emergence of regional specialization, and industrialization [Crafts, 1985, Hunt, 1986, Broadberry et al., 2015, Crafts, 2018]. Other strands of this literature examine long-run regional inequalities and the historical roots of England's persistent North–South divide [Williamson, 1965, Hechter, 1971, Clark, 2007]. I build on and extend this literature by providing, to my knowledge, the first systematic quantification of millennium-scale spatial persistence in incomes within England. While prior work has richly described regional disparities and their evolution, the absence of directly comparable income measures over very long horizons has limited the ability to assess whether today's inequalities reflect deep historical patterns or more modern developments.

The rest of this paper proceeds as follows. Section 1 describes the Doomsday Book and gives historical context. Section 2 details the main long-run spatial persistence result, and section 3 discusses mediators of the uncovered persistence. Section 4 then analyses the mechanisms for this persistence, distinguishing between the role of path dependence vs that of fundamentals. Finally, section 5 concludes.

1 The Doomsday book

In 1066, William the Conqueror successfully invaded England from Normandy, defeating the Anglo-Saxon forces of Harold Godwinson at the Battle of Hastings and therefore settling the succession dispute that followed the death of Edward the Confessor. Some twenty years after his victory, and unsatisfied with the relatively low tax revenue from his new land, William commissioned a census of the value of everything in England.

The resulting book came to be known in subsequent centuries as the Domesday[8] Book, reflecting its comprehensive nature. It represents a fantastically detailed and complete ac-

[8] In this paper, I will use the modern spelling, Doomsday.

count of the economy of Medieval England. Only relatively recently, with the painstaking work of the Hull digitisation project [Palmer, 2016] has the Doomsday Book been digitised in a way suitable for quantitative analysis. For a more detailed discussion of the Doomsday Book for economic research, see Walker [2015] and Delabastita and Maes [2023]. Data is arranged at the level of the manor and covers ownership, population by type, ploughs, mills, land area, churches, and value. Manors were the level at which the agricultural feudal economy was organised and typically corresponded to a small settlement or part of a larger village.

The Doomsday Book is a unique quantitative historical source in Europe. One of the first ever effective censuses it pays particular detail not only to who (and where) but also to the productive capacity of manors. For this reason, as discussed in detail below, the Doomsday Book uniquely allows for the analysis of production, or incomes, at the individual or granular geographic level. This is not the first economics paper to leverage this data (see, for example, Heblich et al. [2023], Angelucci et al. [2022], and Rosenberg and Curci [2024]), but it is (to the best of the authors' knowledge) the first to leverage its fine geographic information to consider long-run spatial persistence.

The act of collecting data for the Doomsday Book was itself a remarkable exercise in central administration, unheard of in Medieval Europe. William sent commissioners to each corner of his new kingdom and interrogated local rulers on their holdings. Boards of English and Norman jurors were then tasked with verifying the validity of answers given by local lords. Although it was probably collected, no data for major cities survives, and therefore some large agglomerations in 1086 are not included in this analysis. In 1086, however, it is estimated that only eight agglomerations exceeded 2,000 in population [Russell, 1948].[9] The Doomsday Book suggests a population of England of around 1.5 million, implying that roughly 3% of the population is missing due to the omission of some major cities.

Crucial to my analysis is the manor-level recorded "valets" i.e. values. As with all variables in the Doomsday Book, and those in any data source from almost a thousand years ago, the exact interpretation is itself an area of scholarship. However, almost all modern historians and economic historians agree that "valets" should be interpreted as a measure of manor-level income [Delabastita and Maes, 2023, Walker, 2015, McDonald and Snooks, 1985]. Evidence in favour of this interpretation comes from various sources, for example, Galbraith [1929] shows that a specific manor was rented out at a rate exactly that of its value. This is also the conclusion reached by Britnell and Campbell [1995] and Roffe [2015], the latter of whom states, as noted by Delabastita and Maes [2023], that: "Doomsday values are a more or less accurate index of the productive capacity of estates". For this reason, I

[9] In descending order: London (17,850-10,000), Winchester (6,000-6,750), Norwich (4,444-4,750), York (4,134-5,000), Lincoln (3,560-4,500), Thetford (2,681-4,000), Bristol (2,310), Gloucester (2,146-2,750). Population estimates are from Russell [1948] and Darby and Darby [1986].

shall refer to value per capita as a measure of local income and discuss inequality in terms of income throughout this paper. I will also take as holistic a view of "per capita" as possible, and include all souls recorded as living in a locality in my denominator. This includes those denoted as slaves, but unfortunately is unlikely to include many women or children, as they were in general not recorded [Stafford, 1989]. It should also be noted that in this feudal economy, within-location inequality is likely to be extremely high, with local rulers enjoying most of any surplus generated.

Manors can be precisely geolocated, and thanks to the work of the Hull digitisation project, we know the location of manors across England and therefore could analyse spatial inequality in 1086 at this highly granular level. However, modern data doesn't live up to this degree of granularity and spatial precision. Instead, I aggregate Norman variables to modern local authorities. Throughout this paper, my main unit of analysis will therefore be these modern local authorities. As mentioned above, the Doomsday Book for some major agglomerations doesn't survive, in addition, some areas in the North West of England are unaccounted for. In total, I have data for 283 modern local authorities across England.

1.1 Other data sources

I use modern local authority level regional gross domestic product data from the office for national statistics, as well as local authority level wage estimates from the annual survey of hours and earnings. I augment this data with data on modern population estimates (at the local authority level) from the 2021 census, historical data on Roman towns [Hanson, 2016a,b] and Roman roads [McCormick, 2013], data on terrain ruggedness [Nunn and Puga, 2012], and data on local attainable yields over various common crops [FAO and IIASA, 2021]. I also draw on various historical sources cited throughout, in particular, to gather data on rebellions to William's rule. Finally, I use data on income proxies in intermediate periods in some parts of the analysis. In particular, I use data from Heldring et al. [2022] on the lay subsidies (in 1332 and 1525) and agricultural employment shares (in 1831), details on these data sources and construction are given in the appendix section A.1.

2 Persistence across space from Doomsday to today

Using data from the Doomsday Book I first turn to the empirical question of the degree of spatial income persistence over 1,000 years. Figure 1 shows the spatial distribution of incomes in 1086 and 2020 measured in period-specific standard deviations, over modern local authorities. It is immediately clear from these figures that the series depicted on the two maps is correlated. In 1086, the dominance of the south was perhaps even more stark than it is today, and in both periods, the relative poverty of the extremities of the country is evident.

Figure 1 Spatial distribution of incomes

(a) 1086 (b) 2020

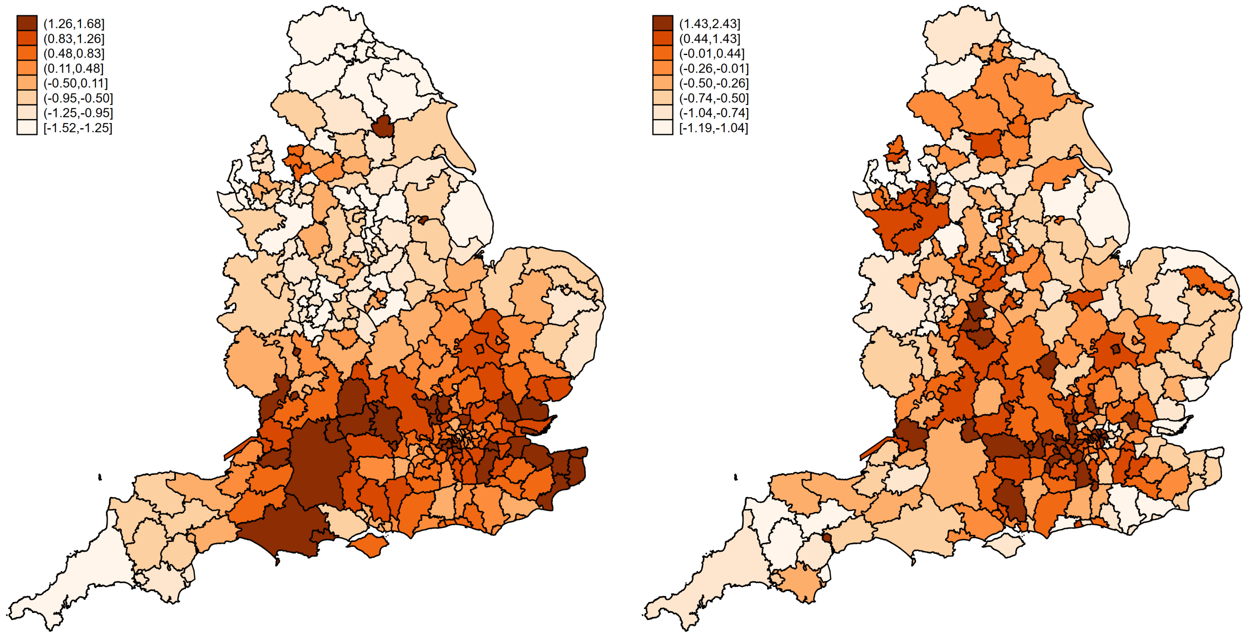

Notes: This figure shows the spatial distribution of (standardised) value per capita in panel (a) and (standardised) GDP per capita in panel (b). In both cases, the distribution is displayed over modern local authorities. Both series have been winsorised at the 5 and 95 percentile level. Both figures show the local authorities that appear in my main sample, i.e. those in England for which data from the Doomsday Book is available.

Although the eyeball econometrics encouraged by figure 1 are already compelling, 1 formalises the degree of persistence by regressing modern values on those from 1086 and show the results in table 1. Column one of 1 shows that areas that were 10% richer in 1086 remain, on average, almost 2% richer in 2020. In standard deviation terms, areas that were one standard deviation richer in 1086 are on average 0.3 standard deviations richer today. The relationship appears fairly linear and tight, with the 95% confidence interval indicating an elasticity between 0.135 and 0.257. In column two, I weigh by 1086 population and recover a slightly higher estimate; not shown is that weighting by the 2020 population also recovers a similar estimate. In column three, I include a second-order polynomial in local authority centroid latitude and longitude coordinates. In a crude manner, this specification controls somewhat for spatial dependencies, which are discussed in detail below. Anticipating the results from section 2.1, in column three the correlation is somewhat attenuated, but remains significant both statistically and economically. In column four, I include fixed effects for the nine English regions. Controlling for such cross-regional variation attenuates the estimated coefficient by 27%, indicating that the broad correlation is not driven by large-scale differences such as north vs south, but is also present in within-regional variation.

Finally, column five presents a rank-rank regression and recovers an even stronger correlation in ranks. Figure 3 in the appendix allows for non-linearities by considering the binscatter relationship.

Table 1 Spatial persistence over the very long run

	Log GDP per-capita 2020	Log GDP per-capita 2020	Log GDP per-capita 2020	Log GDP per-capita 2020	LA rank 2020
Log values per capita 1086	0.195***	0.231***	0.143***	0.136***	
	(0.0311)	(0.0356)	(0.0450)	(0.0468)	
Rank values per capita 1086					0.300***
					(0.0523)
Weighting	None	1086 Pop	None	None	None
Lat-lng polynomial			Yes		
Region FE				Yes	
Observations	283	283	283	283	283
R2	0.0971	0.135	0.194	0.131	0.0900

Notes: This table shows the results from regressing the spatial distribution of income in 2020 against that in 1086 across various specifications. In the first column, I show the raw results in a log-log specification that corresponds to figure 3. Column two weights the regression by the 1086 population. Column three includes a second-order interacted polynomial in local-authority centroids. Column four includes fixed effects for the 9 high-level regions of England. Finally, column five performs a rank-rank regression.

Table 6 in the appendix considers the robustness of the relationship to further specifications. It shows that the result of significant long-run persistence is robust to: not winsorising, excluding Harried local authorities, excluding modern London, using average or median wages from the 2020 ASHE data, or performing the analysis in levels (rather than logs).

Overall income inequality across space is similar in the two periods, albeit somewhat lower in 2020, with a coefficient of variation over geographies in 1086 of 0.44 and in 2020 of 0.34. As discussed above, it is likely that within-location inequality was significantly higher in 1086 under the feudal system; indeed, the Doomsday Book also records who owns many of the manors, and one individual stands out as owning around 20% of the entire wealth of the country: William.

Taken together, these results reveal a degree of spatial income persistence over 1,000 years that is striking. While the literature has documented meaningful income persistence over several decades or even centuries, evidence linking local economic advantage over a full millennium is virtually nonexistent. The elasticity estimates I recover, which imply that places advantaged at the time of the Norman Conquest remain noticeably richer today, underscore that long-run spatial inequality in England is extraordinarily slow-moving. I find that even in a setting marked by profound institutional, technological, and demographic upheaval, the relative ranking of places proves remarkably durable. As such, the results highlight that contemporary spatial disparities are not solely the product of recent shocks or policies but reflect deeply rooted historical forces that have persisted over centuries. Identifying what these sources may be, and whether they reflect some fundamental characteristic of place or simply path-dependence, is the question to which I now turn after discussing spatial autocorrelation.

2.1 Spatial correlation

Correlating two spatial variables can lead to a spurious relationship due to spatial autocorrelation, in much the same way that correlating two time series can also lead to nonsensical relationships [Conley and Kelly, 2025].

Spatial data is naturally ordered (A is close to B) in a similar way to that which time series data is naturally ordered (A is soon after B). However, spatial data operates in a two (or three, but I will not explicitly consider elevation) dimensional space, whereas time-series data operates over only one dimension. Spatial units also tend not to be uniformly spread over the topology they operate within and are measured on a continuous scale, whereas time series data is often uniform and measured on a discrete scale. These complications mean that to the authors' reading, the literature has not settled on one single approach to accounting for the possibility of spatial correlation. Instead, here I present results from four leading approaches, each of which either accounts for spatial dependence directly or adjusts standard errors ex-post.

Table 2 shows the results from each approach in the upper panel and the baseline results in the lower panel. Coefficient estimates are indicated above standard errors, which are given in brackets. Stars appended to standard errors indicate the usual significant levels. In column one, I present results using a spatial autoregressive model. This approach is analogous to using an autoregressive model to account for autocorrelation in time series data. When using an AR model in time series data, one must specify the number of lags to be used. Similarly, when using a spatial autoregressive model, one must specify the spatial weighting or decay matrix. One can make various choices here, but for the results presented in table 2, I've used the simplest approach, which is a linear inverse-distance weighted spatial decay matrix. By controlling for spatial correlation, much like when using an AR model on time series data, the estimated coefficient of interest can change as well as its standard errors. The estimated elasticity is attenuated, which suggests that spatial dependencies are causing some of the initially estimated relationship. However, the coefficient remains statistically and economically significant.

The second approach I adopt is that of Conley [1999], the results of which are presented in columns two, three, and four of table 2. Conley standard errors use a kernel to give more weight to observations close together when calculating standard errors, in a manner analogous to how White heteroskedasticity robust standard errors are calculated. Following the standard specification, I use a uniform kernel and vary the cutoff. In column two, I use a 10km cutoff, in column three a 50km cutoff, and in column four a 100km cutoff. In appendix figure 4, I show the estimated t-statistic on the coefficient of interest at cutoff intervals between 10 and 500km. As the cutoff is increased, the standard errors also increase up to 100km, whereafter they slowly decrease back to baseline. The smallest t-statistic remains above 2.7. Note that Conley standard errors are an ex-post correction (much like robust standard errors), so this procedure does not alter the estimated coefficient.

Thirdly, I control for spatial correlation by performing the procedure suggested in Müller and Watson [2022]. This paper proposes a procedure for constructing confidence intervals that account for many forms of spatial correlation by adjusting both the standard errors and critical values. As this procedure does not affect the point estimate, I only report the estimated confidence interval in table 2. The standard error in the Müller and Watson [2022] approach is calculated as the principal component from a given worst-case spatial correlation model; the critical value is then chosen to ensure correct coverage in some benchmark parametric model. This approach has the added advantage over Conley standard errors or Spatial autoregressive models, that it allows for the unequal placement of geographic units. In table 2, I display the estimated 95% confidence intervals for three such worst-case spatial correlation matrices. Although the confidence intervals sometimes include 0, in all cases, the vast majority of the mass is in the positive range.

In the final column, I correct for spatial autocorrelation using the semi-parametric thin plate spline methodology proposed in Kelly et al. [2023]. This process fits a two-dimensional non-parametric spline in latitude and longitude to directly and flexibly control for spatial dependencies. As discussed in Kelly et al. [2023], this approach allows me to separate out the spatial structure of the regression as a nuisance variable and carry out standard inference on the remaining parameters. This procedure attempts to directly control for the spatial dependencies and so will alter the coefficient estimates as well as the standard errors. This can be seen in column 8 of table 2 where the coefficient estimate is smaller, and the standard errors larger. Despite these corrections, the resulting point estimate remains economically significant and statistically significant at the 10% level (p-value: 0.079).

In sum, the results presented in Table 2 show that naive regressions likely overestimate the relationship between incomes across space in 1086 and 2020. However, the relationship on the whole survives and remains economically meaningful. In addition, procedures that attempt to directly control for the structure of spatial dependence by conditioning on some function of geography may inadvertently capture part of the mechanism behind the observed persistence. To see this, consider an example: suppose some geographic characteristic "X" causes areas to be richer in both 1086 and 2020, and that "X" displays some spatial structure, for example, soil fertility. Then, controlling flexibly for topology through a spatial autoregressive model or non-parametric thin plate splines may capture both spurious spatial correlations and this genuine dependency. Indeed, this dependency is an example of a local "fundamental" driving long-run persistence, the concept of which will be discussed in detail in section 4. For these reasons, Table 2, along with the naive results reported in Table 1, lead me to conclude that areas that were 10% richer in 1086 are on average somewhere between 1% and 2% richer today.

Table 2 Accounting for spatial correlation

	Spatial Autoregressive Model	Conley Standard Errors			Müller and Watson (2022)			Kelly, Mokyr, and, Ó Gráda (2023)
		d=10km	d=50km	d=100km	sc=0.03	sc=0.02	sc=0.01	
GDP per-capita 1086 (SD)	0.0902 (0.0448)**	0.1945 (0.0331)***	0.1945 (0.0520)***	0.1945 (0.0637)**	[-0.0579, 0.4470]	[-0.0081,0.3972]	[0.0467,0.3424]	0.0934 (0.0530)*
Observations	283	283	283	283	283	283	283	283
R2	0.1817	0.0971	0.0971	0.0971	0.0971	0.0971	0.0971	0.307

Results from the baseline model	
Coefficient	0.1945
Standard Error	0.0311***
Observations	283
R2	0.0971

Notes: This table presents the results of performing various corrections for the presence of spatial correlation. In the lower panel, I re-print the results from performing the baseline regression as shown in column one of table 1. In the upper panel, I perform 8 additional regressions adjusting for spatial correlation via four different approaches. In the first column, I fit a spatial autoregressive model with inverse-distance weighted spatial lags. In columns two, three, and four, I calculate Conley standard errors using a 10km, 50km, and 100km cutoff. In columns five, six, seven, and eight, I perform the spatial correlation robust inference procedure due to Müller and Watson [2022]. I specify a worst-case spatial correlation matrix with correlation values 0.03, 0.02, and 0.01 respectively. Finally, in column eight, I perform the semi-parametric thin plate spline correction due to Kelly et al. [2023].

3 What location characteristics mediate persistence

Having established persistence over an exceptionally long time horizon, the natural next question is what local characteristics might mediate this relationship? To answer this question I build a local-authority level database of time-invariant (since 1086) characteristics.

I consider four categories of local characteristics. First, ruggedness. Ruggedness captures how flat the local terrain is; flatter areas have lower ruggedness values. These areas are more amenable to agriculture, easier to build transport infrastructure to/from, and indeed to build on in general [Nunn and Puga, 2012, Henderson et al., 2018].

Second, I consider the role of the infrastructure that already existed in 1086. From the Doomsday Book itself, there is data on local mills and fisheries, which I aggregate into the raw number in each modern local authority. I also include a dummy variable that takes the value one if a major Roman town lies inside the local authority and another dummy variable that takes the value one if a major Roman road runs through the local authority. One can clearly see how previously existing infrastructure may have increased incomes in 1086, but it could also increase incomes in 2020. By the very nature of having more infrastructure these areas were relatively more developed in 1086, and this development could persist through the channel of path dependence or due to fundamental characteristics of these locations. For example it could be the case that the following characteristics remain predictive of higher incomes today: running water (mills), coastal access (fisheries), natural thoroughfares (Roman roads), or preexisting urban centres (Roman cities) [Dalgaard et al., 2022, Michaels and Rauch, 2018].

Third, I consider variables that capture the geographic centrality, or market access, of a location within England. The variables I use are: distance to the coast, distance to London[10], a measure of market potential based on Roman settlements, and a measure of market potential based on the 1086 population. Measures capturing some notion of local market access have been shown to have been associated with various local outcomes, including local wages [Donaldson and Hornbeck, 2016, Redding and Sturm, 2008, Asher and Novosad, 2020].

Finally, I consider variables that capture measures of local agricultural productivity. Using data on attainable yields from the FAO's global agro-ecological zones database, I construct measures of attainable yields for four key staple crops: wheat, oats, rye, and barley. FAO's measures use data on local climatic and geological features from the present day and so will not perfectly reflect conditions in 1086. However, I use their measure with no modern inputs, and consider it unlikely that local changes would be endogenous to local outcomes i.e. that such a measure would be a "bad control".

Table 3 provides the results from regressing log GDP per capita in 2020 against log value

[10] London was already the most populous city in 1086.

per capita in 1086, controlling for various combinations of the characteristics described above. In these multivariate regressions, I only include covariates that capture significantly difference sources of variation to avoid collinearity. In column one, I replicate the raw correlation. In column two I control for ruggedness, in column three for historic infrastructure, in column four for market potential, in column five for soil quality, in column six for soil quality and market potential, and finally in column seven for all of the above. In all regressions, I account for spatial autocorrelation using the approach due to Conley with a 50km cut-off. I take this approach rather than those that directly control for spatial dependence as partly what we want to capture in this exercise is some fundamental effect of geography that may be inadvertently controlled for in for example, in the thin-plate-spline approach.

Table 3 Accounting for long-run spatial persistence

	(1)	(2)	(3)	(4)	(5)	(6)
Log value pc 1086	0.195***	0.182***	0.195***	0.098*	0.194***	0.089*
	(0.052)	(0.050)	(0.052)	(0.052)	(0.055)	(0.050)
Log Ruggedness (mean)		-0.052*				-0.003
		(0.027)				(0.024)
Mills (SD)			-0.030*			-0.023
			(0.01)			(0.014)
Fisheries (SD)			-0.033***			-0.028**
			(0.013)			(0.012)
Roman Road			0.092**			0.063*
			(0.037)			(0.037)
Roman town			-0.009			0.047
			(0.054)			(0.055)
Log market potential				0.232***		0.195***
				(0.058)		(0.063)
Barley Attainable Yield (SD)					0.001	0.017
					(0.026)	(0.017)
Constant	10.51***	10.63***	10.45***	8.80***	10.05***	9.07***
	(0.098)	(0.131)	(0.110)	(0.409)	(0.102)	(0.470)
Observations	283	282	283	283	283	282

Notes: This table reports results from regressing log GDP per capita in 2020 against Log value per capita in 1086, controlling for various combinations of covariates. In column one, I report the raw relationship. In column two, I control for the ruggedness of a location. In column three, I control for measures of historic infrastructure. In column four, I control for log market potential. In column five, I control for the soil quality of a location proxied by the attainable yield of barley. Finally, in column six, I control for both log market potential and soil quality. Standard errors are calculated using the Conley approach with a 50km cut-off. Stars indicate usual significant levels.

Table 3 shows that the coefficient on log value per capita in 1086 is stable across specifications, with the exception of columns (4) and (6), the columns which include the market access covariate, where it decreases substantially.[11] This provides evidence to suggest that geographic characteristics of place and pre-existing infrastructure are not important determinants of the long-run persistence in income, but market access terms may be. However, it does not allow me to distinguish the channel through which time-invariant market access may mediate the relationship.

4 Path-dependence or Correlated Fundamentals: Suggestive Evidence

Section 2 shows that areas that were richer 1,000 years ago remain, on average, richer today, and section 3 shows how some variables may mediate this relationship. However, section 3 does not allow me to disentangle two possible channels through which each characteristic could mediate the observed persistence. In this section I turn to distinguishing between these channels and document more systematically the forces that may be causing the persistence. The two potential mechanisms are: *Path dependence* (We are rich today because we were rich yesterday) and *Correlated fundamentals* (We are rich today because of local characteristics X, this was also why we were rich yesterday).

In this section, I propose two approaches to provide evidence on which of the two channels is driving the observed long-run persistence. First, I turn to theory and impose structure on the data using a dynamic quantitative spatial economics model following Allen and Donaldson [2022]. This framework embeds both possible mechanisms, allowing me to estimate counterfactual distributions in the absence of each force and so discern which drives the observed persistence. I augment this approach with empirical evidence from a natural experiment that decimated the population in 1069/70 in part of the country and test whether this transient quasi-exogenous local shock led to a permanent change in local economic outcomes. The first approach represents a novel application of the Britton et al. [2021] framework, whereas the second approach is common in the literature, which has previously looked to leverage variation from transient exogenous local shocks (e.g. Davis and Weinstein [2002], Miguel and Roland [2011], Jedwab et al. [2024]).

Note that these exercises aim to understand why two periods in time might be related; they do not aim to shed light on the entire dynamic process linking these periods. England faced many shocks between 1086 and 2020, including, for example the Industrial Revolution and subsequent deindustrialisation. These shocks undoubtedly shaped the spatial distribu-

[11] The t-statistic for the test of equality of coefficients on the log value per capita in 1086 coefficient between columns (1) and (4) is 1.31, and between columns (1) and (6) is 1.47.

tion of incomes; indeed, there is evidence to suggest that local income in the early medieval period was negatively correlated with factories during the Industrial Revolution [Rosenberg and Curci, 2024], and the negative impacts of deindustrialisation are well documented [Franck and Galor, 2021, Timmer et al., 2015]. The presence of these convulsions does not invalidate exercises aiming to distinguish between path dependence and fundamentals, but interpretation must be made conditional upon them. If the uncovered persistence is due to path dependence or correlated fundamentals, this is in spite of intermediate shocks.

Each approach also has important limitations. As described below, the model exercise relies on a framework calibrated to a different economy, with time-invariant parameters, that cannot capture structural changes in the English economy, such as industrialisation, but rather takes them as given. The Harrying exercise, on the other hand, relies on the quasi-randomness of William's decision to Harry only the northern rebellion evidence for which comes only from historical discourse, and an almost 1,000-year-long parallel trends assumption.

4.1 Accounting directly for local fundamentals using a dynamic quantitative spatial economics model

The strategy in this section is to add structure to the data that allows for both correlated fundamentals and path-dependence, and then construct counterfactual 2020 distributions of economic activity solving the model forward from 1086, shutting down each possible channel in turn. This model, although quite general (for example, see Allen et al. [2020] for a discussion of generality in this type of model), does introduce considerable structure. The fundamentals uncovered are effectively structural residuals that allow us to rationalise the observed distribution of economic activity conditional on the model structure. The model necessarily cannot capture all potential sources of path dependence, and in particular, cultural or political mechanisms are not accounted for, nor is the possibility of endogenous growth or structural transformation.

Dynamic quantitative spatial economics models capture general mechanisms and economic interactions and have been successfully employed in a wide variety of settings. However, one may wonder whether they accurately reflect the medieval economy of Doomsday England. There are three main economic forces crucial for the model: Agglomeration economies, costly trade over space, and costly migration over space. There is no direct quantitative evidence on trade or migration in the Doomsday Book from which we can draw. However, there is considerable archaeological and anecdotal evidence of both trade and migration on a large scale. It is beyond the scope of this paper to provide a comprehensive assessment of sources here; however, Nightingale [2023] and McClure [1979], Smith [2014] provide an instructive introduction. Turning to agglomeration forces. Using Doomsday data, I regress values per capita on the local population and find a coefficient of 0.12 (0.027). This positive correlation implies that more populous areas were richer in 1086. Although purely correlational, this evidence remains indicative. Other work has also found a relationship between manor size and prosperity. For example, McDonald [2010] and McDonald and Snooks [1985] find evidence of size effects, and Delabastita and Maes [2023] shows that manors benefited from economies of scale across their often England-wide ownership network.

4.1.1 A dynamic quantitative spatial economics model

Here I follow Allen and Donaldson [2022] closely, and present a simplified version of their model where agents are myopic.[12] Within this framework, cities are connected, and goods and individuals are allowed to move (with some cost) between them. History impacts the future through dynamic agglomeration effects in productivity and amenities. Intuitively infrastructure built some time ago might enhance (or decrease) productivity today. The model also admits the potential for long-run persistence and multiple long-run spatial equilibria, whereby shocks can cause permanent changes to the distribution of economic activity. Here, I briefly describe the model as presented in Allen and Donaldson [2022].

There are arbitrarily many locations $i \in N$ and $t \in \mathcal{T}$ time periods. Each location i emits a unique good in an Armington fashion. A continuum of firms ω in i produce this homogeneous good ($q_{it}(\omega)$) under perfect competition and CRTS using labor ($l_i(\omega)$) as the only factor of production.

$$q_{it}(\omega) = A_{it} l_{it}(\omega), \qquad A_{it} = \bar{A}_{it} L_{it}^{\alpha_1} L_{it-1}^{\alpha_2} \tag{1}$$

Where $\bar{A}_{it}$ is exogenous productivity and L_{it} is the total number of workers. α_1 captures aggregate contemporaneous spillovers, α_2 captures aggregate historical productivity spillovers. Intuitively α_1 captures what is more traditionally thought of as agglomeration forces, whereas α_2 captures factors like historical infrastructure which remain productive in the next period.

Individuals have CES preferences over differentiated location-specific goods with the elasticity of substitution σ, therefore consumption welfare is captured by local real wages (w_{it}/P_{it}). A location also generates utility for individuals in the form of local amenities (u_{it}), and therefore location-time specific welfare is given by W_{it} in equation 2.

$$W_{it} = u_{it} \frac{w_{it}}{P_{it}}, \qquad u_{it} = \bar{u}_{it} L_{it}^{\beta_1} L_{it-1}^{\beta_2} \tag{2}$$

[12] Relative to Allen and Donaldson [2022] I set the intertemporal discount factor to zero, $\delta = 0$, rather than $\delta = 0.0535$. I choose to enforce myopia as I am not able to calibrate this parameter to the Doomsday economy, so instead assume a low-information world where expectations over future realisations are assumed to be equal to today's values.

Where $\bar{u}_{it}$ is exogenous productivity and β_1, β_2 are analogous to α_1, α_2. β_1 captures contemporaneous congestion forces i.e. from non-tradeables or land, and β_2 captures the impact of durable infrastructure on amenities.

Bilateral trade from locations i to j incurs exogenous, symmetric, iceberg trade costs denoted by τ_{ijt}. Iceberg trade costs and CES demand generate the familiar gravity equation in trade [Allen et al., 2020].

$$X_{ijt} = \tau_{ijt}^{1-\sigma} \left(\frac{w_{it}}{P_{it}} \right)^{1-\sigma} P_{jt}^{\sigma-1} w_{jt} L_{jt}, \qquad P_{it} = \left(\sum_{k=1}^{N} \tau_{kit} \left(\frac{w_{kt}}{A_{kt}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{3}$$

Individuals decide where to move to maximize utility given in 2 subject to iceberg moving costs μ_{ijt} and some idiosyncratic preference draw ε_{jt} which is drawn from a Frechet distribution with dispersion parameter θ. Given this distributional assumption, migration will also follow a gravity structure, and the number of people moving from i to j in period t will be given by equation 4.

$$L_{ijt} = \mu_{ijt}^{-\theta} \Pi_{it}^{-\theta} W_{jt}^{\theta} L_{it-1} \tag{4}$$

Where $\Pi_{it} = \left(\sum_k \mu_{ikt}^{-\theta} W_{kt}^{\theta} \right)^{1/\theta}$, is a measure of labor market access.

The dynamic equilibrium of this model is described in the appendix section B. The model can be solved via a simple iterative algorithm [Donaldson and Hornbeck, 2016]. For details and an in-depth discussion of the equilibrium properties of this model, and models in this class, see Allen and Donaldson [2022] and Allen et al. [2024]. It should be noted that the solution to the model is scale-invariant, and therefore although the total population is arbitrarily fixed, Malthusian dynamics at a UK level would not alter counterfactuals. Malthusian dynamics that vary at a local level would therefore simply magnify existing migration dynamics.

Persistence.

The dynamic quantitative spatial economics model described above has the attractive property that, depending on the parameter values, it can exhibit within-period equilibrium uniqueness but long-run equilibrium multiplicity [Allen and Donaldson, 2022], and shock persistence. Intuitively, α_1 and β_1 govern the within-period equilibrium properties of the model, if they are sufficiently small the dynamic equilibrium will exist and be unique. More specifically, if $\alpha_1 + \beta_1 < 1/\theta$, that is contemporaneous agglomeration forces are greater than the contemporaneous dispersion forces. If this condition holds the model will have a unique transition path.

Turning to the long-run equilibrium, that is the equilibrium to which the economy's transition path is converging. As discussed by Allen and Donaldson [2022], intuitively if agglomeration forces are strong enough location decisions may become self-reinforcing, and thus multiple equilibria could arise. The intuition here translates into a very similar condition as that discussed above, only now total agglomeration forces are what is important. That is multiple long-run equilibria will arise if $\alpha_1 + \alpha_2$ and $\beta_1 + \beta_2$ are sufficiently large. If this is the case the economy has the potential to exhibit path dependence where different initial population distributions may cause the economy to converge to different long-run equilibria. Indeed, each potential long-run equilibrium will be associated with a "basin of attraction", that is a set of possible initial population distributions that converge to the specific equilibrium. Therefore, although multiple equilibria may exist, it remains an empirical question to ask whether any one specific shock would be sufficient to cause a shift from one equilibrium to the next.

Estimation and calibration.

Although the Doomsday Book presents astonishingly detailed data from around a thousand years ago, it does not present any information on flows (of individuals or goods) over space, nor any changes over time in local populations. Without such information, it is difficult to credibly identify parameters in the above-described dynamic quantitative spatial economics model. Rather than attempt to do so in-credibly, in this paper, I take the approach of using previously-estimated parameters from the literature. However, no estimates exist that pertain to a Medieval economy. Instead, I compare results from two parameterisations and interrogate robustness to reasonable parameter deviations.

The model requires parameter estimates for contemporaneous and historical agglomeration economics (α_1, α_2), contemporaneous and historical amenity spillovers (β_1, β_2), the elasticity of substitution (σ), the amenity dispersion parameter (θ), and trade and migration costs over space τ_{ijt}, μ_{ijt}. These transport costs will depend on travel times between pairs of locations denoted as t_{ij}. I take these as time-invariant, measured using centroid-to-centroid distances, assuming a constant travel speed of 5km/h.

First, I consider values used in Allen and Donaldson [2022] who study the setting of the United States between 1800 and 2000. Allen and Donaldson [2022] finds $\{\sigma = 9,\ \theta = 4,\ \alpha_1 = 0.19,\ \alpha_2 = -0.041,\ \beta_1 = -0.26,\ \beta_2 = 0.31\}$ and parameterise iceberg migration costs as depending on travel times and iceberg trade costs as depending on the user cost of shipping using an exponential parameterisation: $\tau_{ij} = \exp(\kappa_\tau \cdot f_{ij}),\ \mu_{ij} = \exp(\kappa_\mu \cdot t_{ij})$, with $\kappa_\tau = 0.7\ \kappa_\mu = 0.02$ for 1850. Where f_{ij} is the user cost of shipping between i and j. I don't have information on bilateral user costs of shipping and instead use $f_{ij} = a \cdot t_{ij}$. I fit a using average travel times, $\bar{t}_{ij}$, and the relationship between trade and migration costs in Ellingsen [2025] such that $a = \phi \cdot \frac{\ln(\bar{t})}{\bar{t}} = 0.56 \times 0.103 = 0.058$. Therefore, for this parameterisation I have $\tau_{ij} = \exp(0.058 \cdot t_{ij})$ and $\mu_{ij} = \exp(0.02 \cdot t_{ij})$.

Second, as a robustness exercise, I use the parameterisation from Ellingsen [2025] who considers pre-industrial cities in the Spanish Empire in South America, and finds $\{\sigma = 5,\ \theta = 3.18,\ \alpha_1 = 0.055,\ \alpha_2 = 0.063\}$. However, he doesn't recover estimates for β_1, β_2 and so instead I use the values from Allen and Donaldson [2022] for these variables. Ellingsen [2025] uses the transport costs parameterisation $\tau_{ij} = t_{ij}^{\phi},\ \mu_{ij} = t_{ij}^{\lambda}$ with $\phi = 0.56$, and $\lambda = 0.363$.

Finally, I also consider robustness to reasonable variation across parameters (keeping transport costs fixed, and centring parameters at the Allen and Donaldson [2022] values), see figures 7 and 8 in the appendix, and find qualitatively similar results.

4.1.2 Model exercise

The model described above embeds path dependence and fundamentals as two possible channels explaining any long-run persistence. To leverage its structure to distinguish between these possibilities, I perform the following exercise.

First, after solving the model (see appendix section B) and using the parameter estimates discussed above, I can invert the model to back out the location fundamental productivities, $\bar{A}_{it}$, and amenities, $\bar{u}_{it}$, in 1086 and 2020 using data on population and wages in each period. Note that I constrain parameter values to be the same in each period. A necessary, but not sufficient, condition for correlated fundamentals to be driving the observed correlation between incomes in 1086 and 2020 is that these estimated fundamentals are also related over time. Note that these estimated fundamentals will capture structural economic features of a local economy, as well as temporary shocks that happen to be relevant at the time of measurement.

Table 4 summarises the correlation between estimated fundamentals using both the Allen and Donaldson [2022] and Ellingsen [2025] calibrations. Across time, productivities and amenities show a clear correlation between 0.13 and 0.22 — areas that were productive in 1086 are likely to also be productive in 2020 (and similarly for amenities). These correlations have been corrected for spatial autocorrelation using the parametric spatial autoregressive model approach described in subsection 2.1. Table 4 also shows that amenities and productivities are negatively correlated in the cross-section. This is intuitive; often, the nicest areas to live in are not the most productive, but rather those with the most beautiful nature, best climate, or best schools/ public infrastructure. Note that a correlation between fundamentals in 1086 and 2020 does not imply a correlation between fundamentals at intermediate periods. There is a good reason to suppose, for example, that during the Industrial Revolution, proximity to coal was an important fundamental that is less relevant both in 2020 and in 1086. Figures 11 and 12 in the appendix show the whole distribution of the estimated productivity and amenity fundamentals using the Allen and Donaldson [2022] calibration. Figures 10 and 9 in the appendix validate the estimated productivities and amenities, respectively, by showing that they correlate with the local characteristics discussed in section 3 as expected.

Table 4 Correlations between estimated fundamentals

Panel (A): Allen and Donaldson (2022)

	C Productivity 1086	Productivity 2020	Amenities 1086	Amenities 2020
Productivity 1086	1.00	0.13	-0.82	-0.22
Productivity 2020	0.13	1.00	-0.16	-0.84
Amenities 1086	-0.82	-0.16	1.00	0.22
Amenities 2020	-0.22	-0.84	0.22	1.00

Panel (B): Ellingsen (2025)

	C Productivity 1086	Productivity 2020	Amenities 1086	Amenities 2020
Productivity 1086	1.00	0.20	-0.70	-0.24
Productivity 2020	0.20	1.00	-0.13	-0.88
Amenities 1086	-0.70	-0.13	1.00	0.21
Amenities 2020	-0.24	-0.88	0.21	1.00

Notes: The table reports the pairwise correlations between the four estimated location fundamentals. In each case spatial autocorrelation has been accounted for by first residualising each variable on a parametric spatial autocorrelation specification as described in section 2.1. Panel (A) shows results using the parameterisation from Allen and Donaldson [2022], whereas panel (B) shows the results using the parameterisation from Ellingsen [2025].

This observed long-run correlation between fundamentals in 1086 and 2020 opens the door for fundamentals more generally to be driving the long-run correlation in incomes. However, observing some correlation in fundamentals is not informative about the degree to which fundamentals, rather than path dependence, may explain long-run persistence.

I can leverage the full structure of the model to shed further light on these questions by shutting down one of the two potential channels. In particular, I can use the model to ask: what is the correlation between 1086 and a counterfactual modern distribution with fundamentals uncorrelated to those in 1086? That is, if we remove the correlation in fundamentals, but maintain the possibility for path dependence — do we still recover a significant (statistically or economically) relationship between actual 1086 values and counterfactual 2020 values?

To operationalise this, I perform 250 experiments. For each experiment, m, I generate random vectors of amenity and productivity fundamentals v_a^m and v_p^m by randomly re-shuffling the actual amenity and productivity vectors estimated in 2020. I then start the model using the actual fundamental distributions in an initial period and propagate it forward using the random vectors of fundamentals over periods of 50 years until we have reached (just surpassed) 2020. This will generate a counterfactual income distribution for each experiment. I then regress this counterfactual income distribution on actual 1086 incomes, correcting for spatial correlation using the semi-parametric thin plate spline approach of Kelly et al. [2023]. This experiment preserves the possible path-dependence between 1086 and 2020 but removes any correlation between fundamentals.

Using the Allen and Donaldson [2022] calibration, I find that the average of the 250 recovered regression coefficients from the above exercise is 0.027, and that 7.6% of the coefficients are significantly different from zero at the 5% level. Similarly, using the Ellingsen [2025] calibration, I find an average across coefficients of -0.035, and that only 4.4% of the coefficients are significantly different from 0 at the 5% level. These results show that once the correlation between fundamentals has been removed, there is almost no meaningful relationship between 1086 values and counterfactual modern-day values. This is taken as additional evidence that the observed long-run persistence is mainly driven by fundamentals as opposed to path dependence.

These results, however, could be sensitive to the model parameters, which were taken from different studies in different time periods and locations. To investigate this, I perform the above analysis on 100 different parameter combinations. In each iteration, I randomly permute the parameter vector $\gamma = (\alpha_1, \alpha_2, \beta_1, \beta_2, \sigma, \theta)$. I do this while anchoring values to those in Allen and Donaldson [2022], which I denote by γ^{AD}, by multiplying γ^{AD} by a random vector R^{sim} whose elements are given by $(r_1^{sim} + 0.5)^{0.5}$ where $r_1^{sim} \sim U[0, 1]$. Thus, for each iteration, I use the random parameter vector given by $\gamma^{sim} = \gamma^{AD} \odot R^{sim}$, where $\odot$ denotes element-wise multiplication.[13] A graphical depiction of the parameter space I use is given in figure 8 in the appendix. I then perform the above analysis for each sim (I only consider 100 random counterfactual modern income distributions within each permutation due to computational constraints) and calculate the proportion of counterfactual income distributions that are significantly related to that in 1086 once controlling for spatial correlation via the thin plate spline method. Figure 7 in the appendix shows the results from this analysis. I find that on average 7.9% of relationships are significant at the 5% level, and on average 16.1% are significant at the 10% level. This indicates that some reasonable parameter values could have resulted in more than expected significant relationships, although the extent to which this occurs is minor. This tempers the conclusion from this section slightly; the weight of evidence remains that fundamentals are the main driving force, but there is some possibility that path dependence plays a more minor role.

Finally, the symmetric exercise to that above would be to instead ask: what is the correla-

[13] I also only retain parameter setups that allow for within-period equilibrium uniqueness, long-run multiplicity, and that converge relatively quickly.

tion between 1086 and a counterfactual modern distribution that starts from a random initial population distribution in 1086? That is, if we remove the possibility for path-dependence, but maintain the possibility for correlated fundamentals — do we still recover the same relationship between 1086 values and the counterfactual modern-day income distribution? I operationalise this by again performing 250 experiments, randomly re-shuffling the initial population distribution in each, but maintaining the actual fundamentals. I perform these experiments using the Allen and Donaldson [2022] parameterisation. Regressing counterfactual 2020 distributions on the actual 1086 distribution, controlling for spatial correlation using the thin plate spline approach, I find that none of the 250 coefficients are significantly different from those found using the actual 2020 distribution at the 1% level. Indeed, the average correlation between counterfactual and actual 2020 distribution across these experiments is 0.983.

Taken as a whole, these sets of model-based experiments provide evidence to suggest that correlated fundamentals rather than path-dependence are driving observed long-run persistence. However, these exercises suffer from various drawbacks. Most notably, I fix parameter values, including trade costs, over time and cannot estimate model parameters within the economy I am studying. Additionally, I take the Allen and Donaldson [2022] model "off-the-shelf" which does not allow for forces such as structural transformation or cultural persistence. These drawbacks motivate the inclusion of additional evidence from an orthogonal source.

4.2 The Harrying of the North as a "natural experiment"

A typical approach in the literature to distinguishing between path dependence and fundamentals is to track the relative performance of local outcomes after a local shock. Intuitively, if after a shock, cities recover back to their pre-shock trend, this is evidence of fundamentals playing a key role. On the other hand, if after a shock, the shocked area continues above or below the trend, this is taken as evidence of path dependence. For example, Davis and Weinstein [2002] finds that Hiroshima and Nagasaki returned to their pre-war position in the urban hierarchy soon after the horrifying and devastating effects of nuclear warfare. Here, I will take a similar approach, using a historical "natural experiment" and tracing the impact on affected local authorities over time.

The natural experiment I employ is the "Harrying of the North" — the brutal reprisals visited on the North in response to a rebellion against William's rule in 1069/70 (see, for example, Strickland [1998], or Dalton [2002] for more details on the Harrying or Vitalis [1854] for an almost first-hand account). In 1069, the last Wessex claimant, Edgar Ætheling, incited a rebellion in the North of England, centred in York. William raised an army and marched on York, but the rebels refused to meet him in open battle. William then decided to punish the northern shires and prevent future rebellions by using scorched earth tactics. He burned fields, destroyed food stores, and slaughtered livestock. This destruction was particularly effective as it came during winter, when acquiring food was already challenging, and many common people relied on their now destroyed winter stores.

Contemporary accounts, archaeological evidence, and evidence from the Doomsday Book attest to the scale of the resulting destruction. By some accounts, around 3/4 of the population of Yorkshire was killed or fled. For example, the Anglo-Norman chronicler Orderic Vitalis wrote:

> The King stopped at nothing to hunt his enemies. He cut down many people and destroyed homes and land. Nowhere else had he shown such cruelty. This made a real change. To his shame, William made no effort to control his fury, punishing the innocent with the guilty. He ordered that crops and herds, tools, and food be burned to ashes. More than 100,000 people perished from starvation.

It should be noted that not all historians are in agreement as to the exact scale of destruction. Dalton [2002] suggests that, given limited time and troops, destruction on the scale suggested by Vitalis would not have been possible. Similarly, Hagger [2021] and Horspool [2009] suggest that the scale of destruction was not so abnormal relative to other contemporary conflicts. However, even if exact numbers are disputed, all authors maintain that the scale was such as to shock the North and create scars for generations.

Of course, one may be concerned that the North's rebellion was not a random event. Perhaps this area chose to rebel because it was more negatively affected by the conquest or was more remote relative to the centres of Norman power, and so considered a rebellion to be more likely to be successful. To combat these concerns, I leverage a control group of local authorities where other rebellions against William's rule occurred. Within a similar time scale, Kent, Northumbria, the Welsh borders, and East Anglia all rebelled against Williams's rule. See appendix section A.3 for details on these other rebellions.[14] Figure 5 in the appendix shows on a map of England the Harried and other rebellious control modern local authorities. The first stage then relies on the fact that the Harrying of the North was qualitatively and quantitatively different from how William reacted to any other rebellions which were fairly peacefully dispersed perhaps following localised fighting.[15]

To investigate the initial impact of the Harrying, I use data available in the Doomsday Book on local values in 1066. These values were recorded via recall in the census of 1086.

[14] I don't include the rebellion in Exeter as it was confined to the city; however, results are robust to its inclusion.

[15] According to the chronicler Orderic Vitalis William regretted his actions in the north and confessed on his deathbed to having treated "the native inhabitants of the kingdom with unreasonable severity, cruelly oppressed high and low, unjustly disinherited many, and caused the death of thousands by starvation and war, especially in Yorkshire".

However, population estimates are not available in 1066, so instead I consider a measure of value per area. This measures the value created in 1066 or 1086 in each modern local authority by square kilometre. In figure 2, I then plot each local authority's 1066 value density rank (inverse rank — higher is better) against its 1086 value density rank, where ranks run over all 283 local authorities. Focusing first on the non-Harried locations, we see a fairly precisely estimated coefficient of unity, implying that otherwise rebellious, but not harried, locations maintained their position in the overall value density distribution.[16] For Harried locations, the picture could not look more different; these locations are positioned much lower in the 1086 value-density distribution than in the 1066 distribution. Harried locations that had among the highest value density in 1066, ranked around 250, were among the lowest in 1086, ranked around 50.

One may be concerned about selective recall driving these results, for example, those in Harried areas in 1086 may want to exaggerate the scale of the destruction by underestimating holdings in 1066. I can provide some evidence that this is not the case by considering another variable noted in the Doomsday Book, that of "waste". Manors were noted as "waste" in the Doomsday Book if they did not pay taxes. The most likely explanation for this tax status is that such manors effectively no longer existed and were not populated, that is, they had been laid to waste. However, it is possible that the term instead reflected manorial re-organisation, or some kind of tax break, see Strickland [1998] and Palliser [1993]. Figure 6 in the appendix shows the change in the proportion of manors noted as "waste" in 1066 and 1086 in both Harried and non-Harried control local authorities. The difference is startling, with the incidence of waste increasing by 45 percentage points in Harried locations relative to non-Harried control locations.[17]

[16] Note that as Harried locations fell down the rank distribution, one would mechanically expect other locations to move up the distribution. The fact that the control locations didn't move up the distribution implies that there may have been some, relatively small, penalty associated with rebelling. To be precise, as on average, Harried locations moved 100 ranks down, one would expect the remaining 260 locations to, on average, move 10 ranks up. If this average was maintained across the non-Harried control locations, one would therefore expect a slope coefficient of 1.1, which does lie within the 95% confidence interval of the estimated coefficient.

[17] Note that a comparison of levels across time is potentially misleading as manors in 1066 are only noted as waste conditional on them being inhabited in 1086 or on there remaining some indication, in 1086, that they once existed that was brought to the attention of the commissioners.

Figure 2 Short(er) run impact of the Harrying of the North

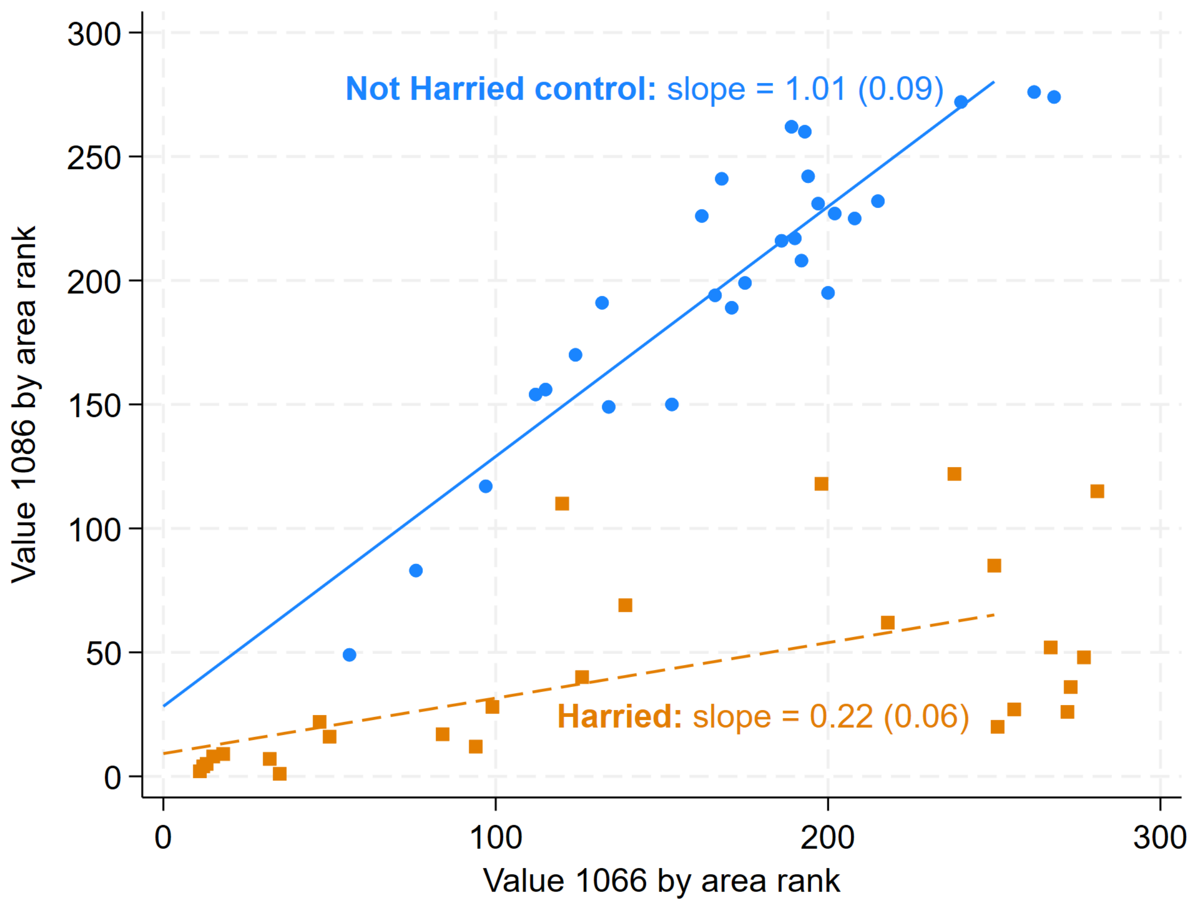

Notes: This figure shows the impact of the Harrying of the North on modern local authorities that were "harried" in 1086, some 16 or so years after the event. The figure plots, in dashed orange, Harried local authorities and in blue non-harried control local authorities that also rebelled against William's rule. On the x-axis is the 1066 value rank (scaled by the size of the local authority), and on the y-axis, the same variable for 1086. The slope and associated robust standard errors are given in the figure for each line.

Figure 2 shows a large impact of the Harrying of the North, some 16 years or so after it occurred. In table 5, I formalise the graphical results shown in figure 2 by performing the following regression on a sample of Harried treatment locations and rebellious but not Harried control locations

$$\text{Value}_{it} = \alpha_t + \beta_t \cdot \text{Harried}_i + \varepsilon_{it} \tag{5}$$

and reporting the β_t coefficients. Table 5 shows the results from estimating this regression on outcomes period by period. In column one, I consider the impact in 1066, before the Harrying occurred, and find no effect. This can be considered as a placebo exercise, or as a rather crude test for parallel pre-trends in a DiD analysis. In column two, I then confirm the large negative impact of the Harrying on the harried locations shown in figure 2, both using values by area and also using values per capita. Finally, in column three, I consider whether the impact of the Harrying can still be felt in 2020. I find no measurable impact of the Harrying today, either in GDP by area or in GDP per capita. This provides evidence that Harried locations (eventually) recovered and returned to their position in the spatial distribution. Table 8 in the appendix displays results from the analogous rank-rank regression.

Table 5 The Impact of the Harrying of the North over time

	Value 1066 by area (SD)	Value 1086 by area (SD)	Log value 1086 per capita	GDP 2020 by area (SD)	Log GDP 2020 per capita
Harried	0.107	-1.537***	-0.679***	0.00215	-0.029
	(0.300)	(0.172)	(0.127)	(0.0238)	(0.062)
Constant	0.123	0.576***	-1.22***	-0.241***	10.14***
	(0.115)	(0.164)	(0.076)	(0.0207)	(0.044)
Observations	54	54	54	54	54
R2	0.00255	0.590	0.357	0.000151	0.004

Notes: This table shows the estimated impact of the Harrying of the North on the Harried areas relative to other rebellious areas. The geographic unit of analysis is modern local authorities. Standard errors are robust. In column one, I show results for 1066 value by area. In Column two, for 1086 value by area, and in Column three, for 1086 value per capita. In columns four and five, I show results for GDP in 2020 by area and per capita, respectively. Observations and R-squared are given below. There are 25 treated (Harried) local authorities and 29 control (rebellious but not Harried) local authorities.

We can also consider results from the more demanding two-way fixed effects specification (using value by area), controlling for local authority-level time-invariant characteristics. Table 7 in the appendix reports the results of running a 2WFE specification. Relative to 1066, this specification implies that Harried areas in 1086 are 1.78 (SE = 0.289) standard deviations poorer (in terms of value per area) and Harried areas in 2020 are not statistically significantly different (coefficient of -0.25 (SE = 0.33)).

This difference in differences analysis relies on the usual assumption that, in the absence of treatment (the Harrying), treated locations would have, on average, evolved in a manner similar to that of the control locations. Over 1,000 years, this is obviously a strong assumption, although similar such assumptions (or RD equivalents) are made in various leading papers in this literature, see for example: Dell [2010], and Nunn and Qian [2011]. These papers, and others in this literature, provide as evidence for the parallel trends assumption, the randomness of treatment assignment. This is the same argument I take, arguing that William could have "Harried" any of the areas that revolted, but only did so in the North for reasons orthogonal to the North's counterfactual outcome. The DiD analysis also requires the assumption of no contemporaneous shocks, although this was certainly a tumultuous period; no other event in the North during this period stands out in the historical record as being on anywhere near the same scale as Williams Harrying.

Although including intermediate periods does not provide evidence for the identifying assumption, it may allow us to better understand the dynamics of the adjustment process. Using data on lay subsidies in 1332 and 1525 and agricultural shares in 1831 from Heldring et al. [2022] I can conduct an event-study type analysis. To do this, I first follow Heldring et al. [2022] and rescale variables by taking away the minimum value, dividing by the range and then standardising.[18] I then perform a two-way fixed effects regression on the sample of Harried and control, not harried but rebellious, local authorities, allowing the treatment coefficient to vary with year. Figure 13 in the appendix shows the results. Perhaps surprisingly, but consistent with previous work, we see a bounce back in 1332 as the North boomed in a rebuilding period [Darby et al., 1979]. However, it appears this boom was short-lived as the north remains negatively impacted in 1525 and, although to a lesser extent, in 1831. As previously found, this negative impact is no longer visible in 2020. This exercise shows that although path-dependence does not appear to link 1086 with 2020, the shadow of the harrying is arguably still visible as late as 1831. Although in this paper we are interested in connecting the modern day to 1086, this analysis raises questions for future work regarding the length of the shadow that such large shocks may cause, and the intermediate mechanisms through which it operates.

The results in this section show that a large shock to the population of a region did not cause it to permanently change trajectory, although it may have caused a long-term shadow [Heath Milsom, 2024]. This is indicative evidence, in this case, that local fundamentals are more likely to determine the long-run distribution of economic activity, at least over the almost 1,000-year period that we are considering. However, this exercise is also not without caveats. First, the time frame involved is extremely long, straining the parallel trends assumption considerably. Second, the randomness of choosing the put down the Northern rebellion and not others is not empirically verifiable, and we rely on appeals to the historical literature and contemporaneous sources.

5 Conclusion

This paper has leveraged data from the Doomsday book to show that the spatial distribution of income we see today in England was partly in place almost 1,000 years ago. I show that areas that were 10% richer in 1086 are, on average, between 1% and 2% richer today and that this relationship is robust to a barrage of corrections for potentially spurious spatial correlation. This persistence is not mediated by geographic characteristics of locations, nor

[18] Doomsday values, lay subsidy taxes, agricultural employment shares, and modern incomes are obviously somewhat difficult to compare. Details on each data source and how they are constructed are given in the appendix section A.1.

the existence of pre-existing infrastructure, but rather by time-invariant measures of market access. I then show some suggestive evidence, using a dynamic quantitative spatial economics model and the natural experiment of the Harrying of the North, that this long-run correlation is not driven by path dependence, but rather by local fundamentals.

Recently (and historically), politicians in the UK have championed spatially redistributive policies such as the Northern Powerhouse and the Levelling Up initiative. This work gives evidence to suggest that such initiatives will not rebalance the long-run distribution of economic activity unless they alter local fundamentals — and it is here that governments aiming for spatial equity should focus.

A Appendix

A.1 Intermediate periods data details

In this sub-section, I describe in more detail the data on intermediate periods, and how such data is combined in the analysis in sub-section 4.2. These data are downloaded from the replication files kindly provided in Heldring et al. [2022], and this description closely follows that of Heldring et al. [2022]. For more information, see Heldring et al. [2022] and the citations there within.

Lay subsidies of 1332 and 1525. I use data from the 1332 and 1525 lay subsidy tax returns. These proxy local income at the ancient parish level and are given in per-taxpayer terms. The lay subsidies taxed movable wealth. The 1525 lay subsidy taxes were the main source of income of the head of each household. Income was taxed at a flat rate of four pence per pound, goods at six pence per pound for value under 20 pounds and twelve pence for over 20 pounds, and landed incomes at twelve pence per pound (note that at the time there was 240 pence to the pound). If the household did not earn at least 1 pound in wages a year, have 1 pound in landed income a year, or possess two pounds in goods, it is not recorded in the survey. The lay subsidy for 1332 was similar. It taxed one-tenth of all movable wealth above a threshold but excepted personal effects like household goods. Surviving records from both subsidies have various missing values; once 0's are removed, we do not have data for 59% and 53% of parishes for the 1332 and 1525 returns, respectively. In particular, much of the far north is missing.

Agricultural employment in 1831. I use data on the agricultural employment share from the digitised 1831 census at the parish level, again accessed from Heldring et al. [2022]. This variable captures the proportion of the over-20 male population employed in agriculture.

Data processing. To allow comparisons across different series at different geographies over time, I perform various processing steps described here. First, I re-scale the lay subsidy data following Heldring et al. [2022] by removing the minimum from each series and dividing by the range. In addition, I remove 0 observations and winsorise at the 1 and 99 percent levels. Second, I transform variables from the parish level to modern local authorities using a GIS file from Satchell et al. [2023]. The matching process is fairly successful, but it is not possible to match all geographies consistently, and 20% of observations are lost. Third, I average parish-level variables over matched modern local authorities. In the analysis in sub-section 4.2, I will weight regressions by the number of parishes used in each local authority level average to account for heterogeneous measurement error introduced in this process. Fourth, I standardise each variable (lay subsidies, doomsday values, agricultural employment shares, modern incomes) to allow comparisons across series. Finally, using this data, I construct a panel dataset at the local-authority level over time for use in analysis.

A.2 Additional figures and tables

Figure 3 Spatial persistence over the very long run

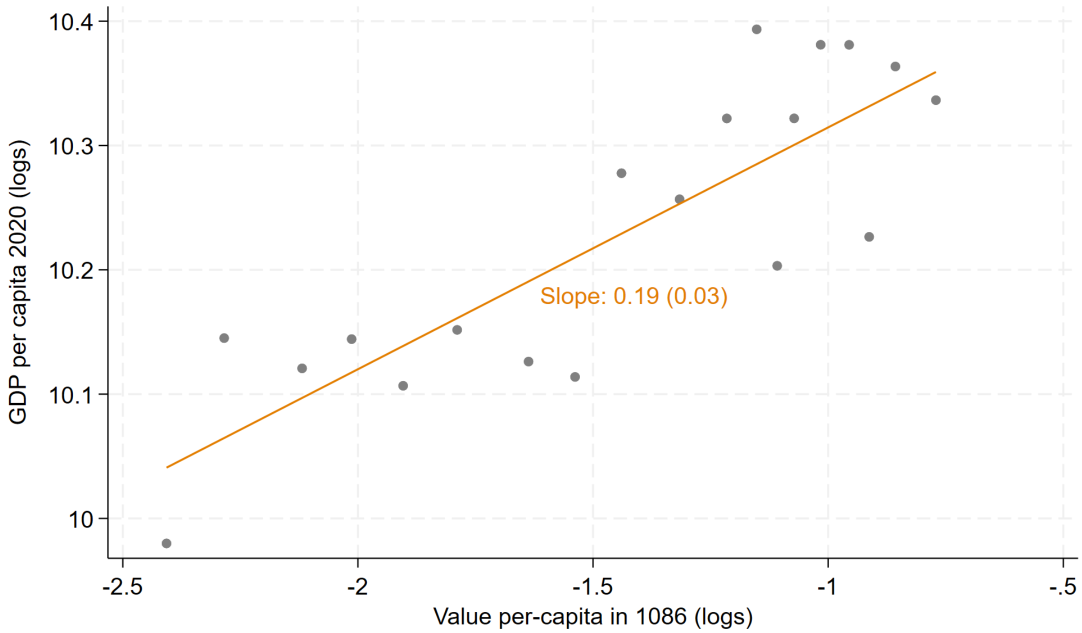

Notes: This figure plots log value per capita in 1086 against log GDP per capita in 2020 in a binscatter figure. A linear line of best fit is shown in orange, the slope and robust standard errors of which are given. The R-squared of the relationship is 0.1. The geographical units are modern local authorities. Standard errors are robust.

Figure 4 Conley T-statistic over specified distance cutoff

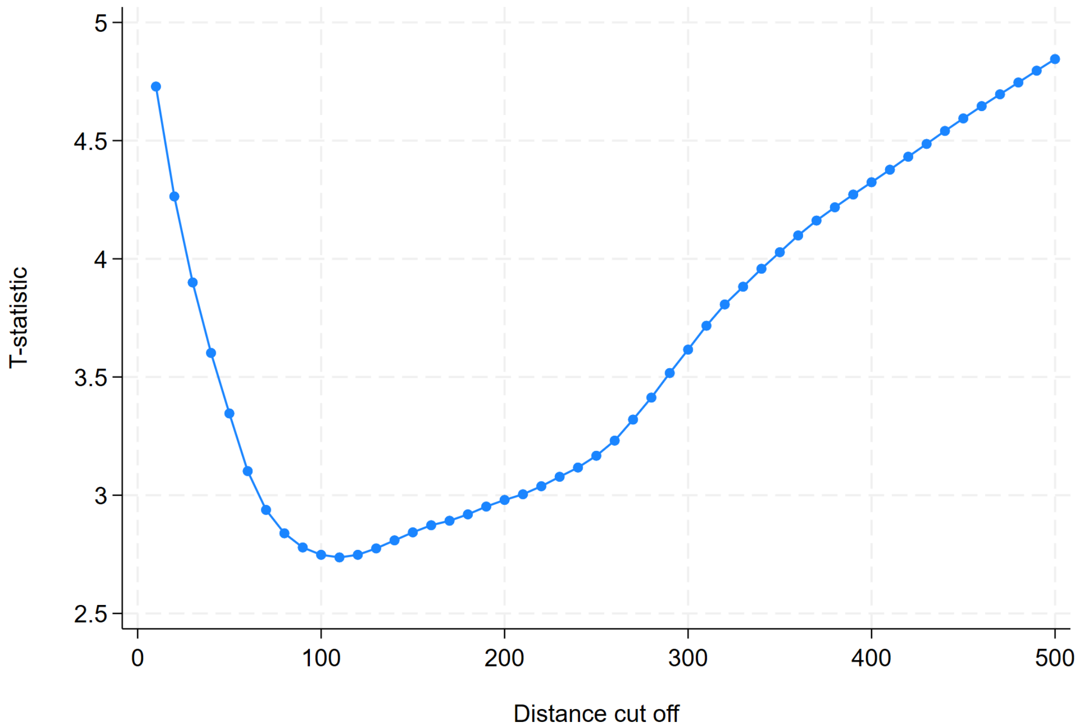

Notes: This figure shows the estimated t-statistics on the main long-run persistence coefficient of interest, β in the following regression: $\ln(2020GDPpercapita)_i = \beta \ln(1086Valuepercapita)_i + \varepsilon_i$. For each regression, I adjust the standard errors following the procedure due to Conley [1999] and report the recovered test statistic over various cutoffs indicated on the x-axis.

Figure 5 Harried and other rebellious control local authorities

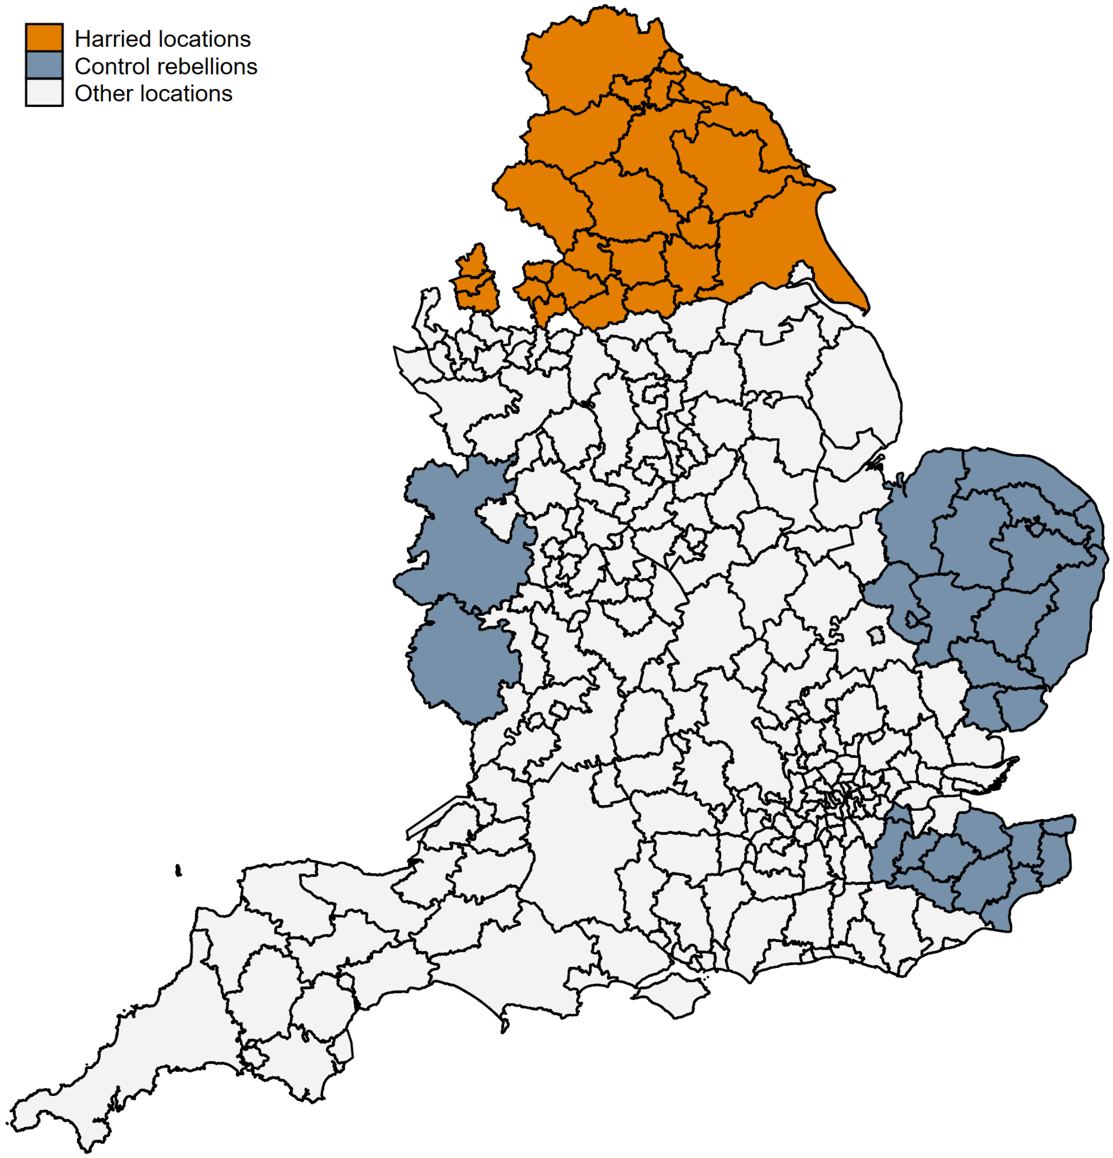

Notes: This figure shows the Harried and other rebellious control local authorities in orange and blue, respectively.

Figure 6 Short(er) run impact of the Harrying of the North

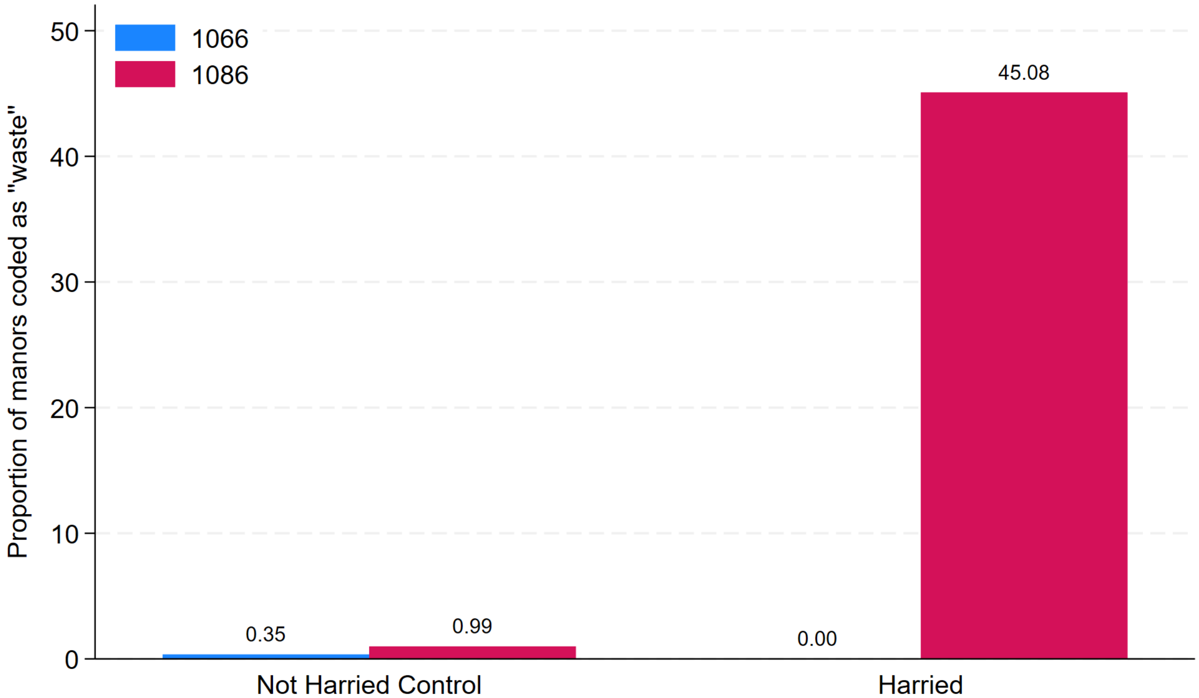

Notes: This figure shows the proportion of manors noted as “waste” in the Doomsday Book in 1066 and 1086 among Harried local authorities and non-Harried control local authorities.

Figure 7 Model results robustness

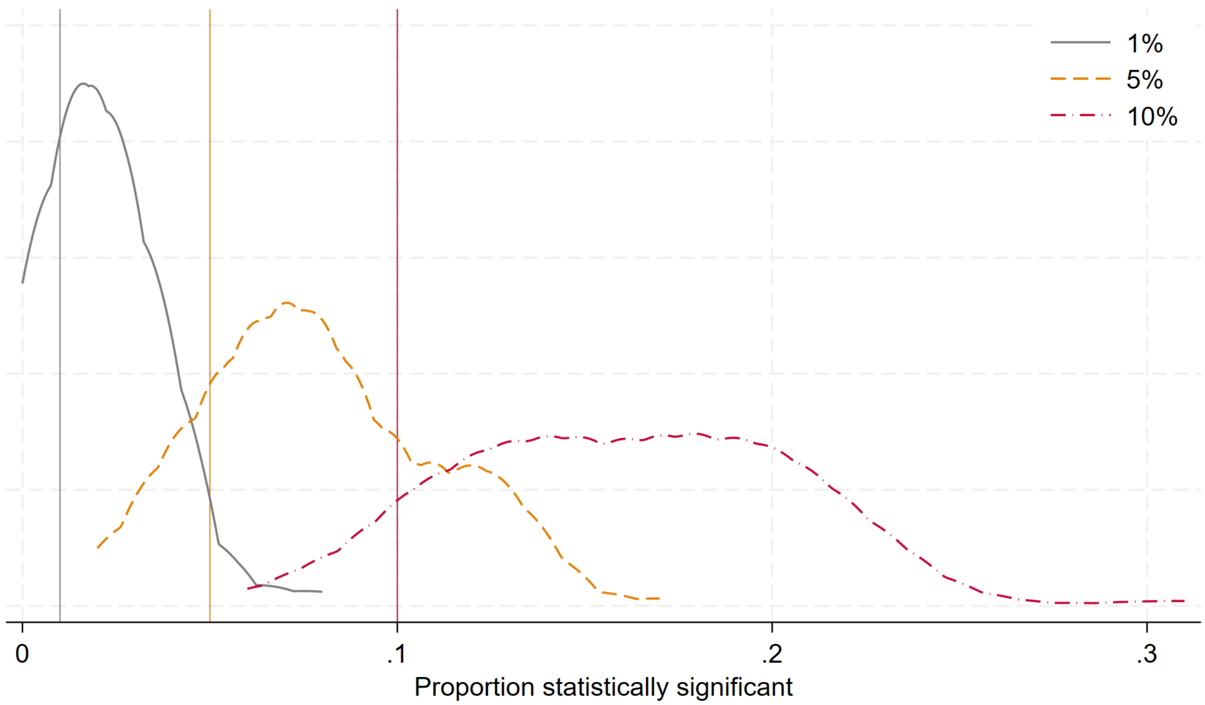

Notes: This figure shows the distribution over model results, where I use a different random permutation of parameter estimates in each iteration. For each random parameter vector, I calculate 100 random vectors of location fundamentals and iterate the model forward to this vector. For each 100 random vectors of fundamentals, I then estimate the elasticity between the counterfactual modern income distribution and that from 1086, controlling for spatial correlation via the thin-plate-spline procedure. Over these 100 iterations, I calculate the proportion of elasticities significant at the 1%, 5%, and 10% levels. In the figure, I plot the distribution of these statistics for each of the 100 outer iterations over the random parameter vectors. Vertical lines at 0.01, 0.05, and 0.1 have been added for visual comparison.

Figure 8 Model results robustness parameter values

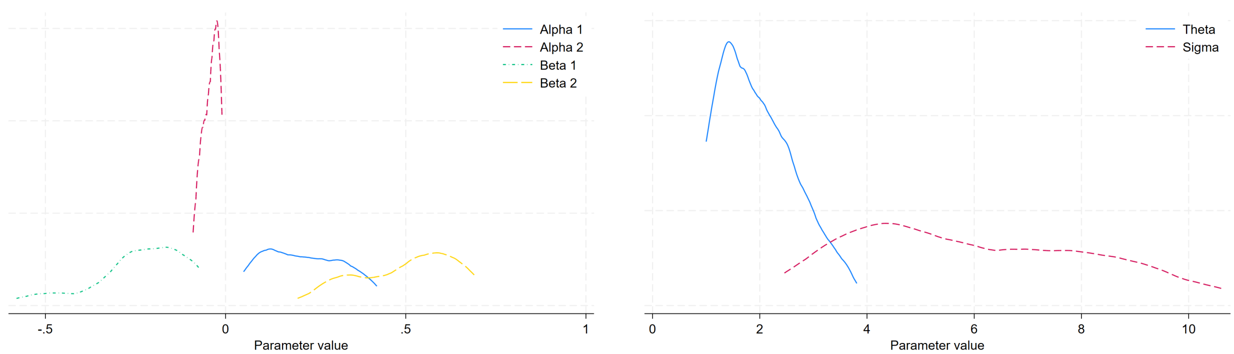

Notes: These figures show the distribution of random parameter vectors used to estimate the distributions shown in figure 7.

Figure 9 Correlation between amenities and local characteristics

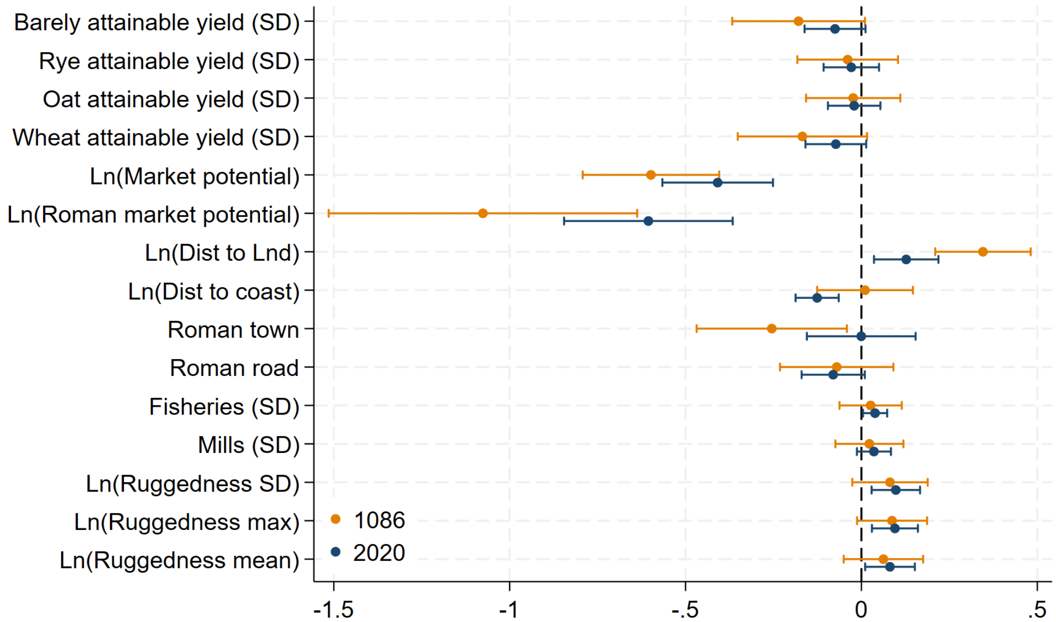

Notes: This figure shows the correlation between the estimated local amenities and time-invariant local characteristics. 95% confidence intervals constructed with Conley standard errors with a 50km cut-off are shown.

Figure 10 Correlation between 1086 productivities and local characteristics

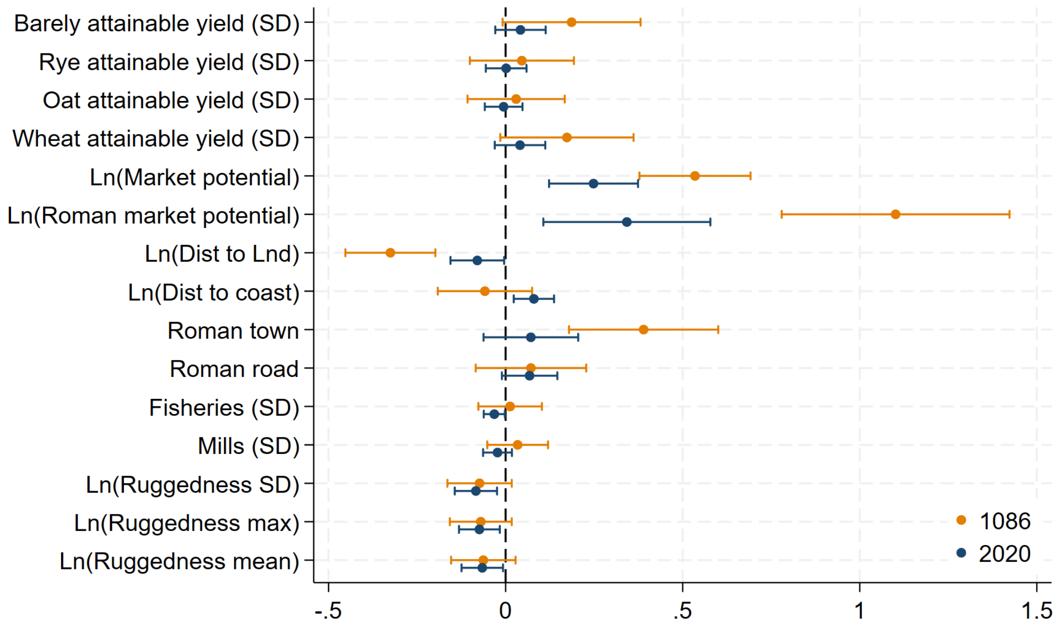

Notes: This figure shows the correlation between the estimated local productivities and time-invariant local characteristics. 95% confidence intervals constructed with Conley standard errors with a 50km cut-off are shown.

Table 6 Spatial persistence over the very long run robustness

	(1) Without winsorizing	(2) Excluding Harried LA's	(3) Excluding Modern London	(4) Median 2020 Wages	(5) Average 2020 Wages	(6) In Levels
Value per-capita 1086 (logs)	0.159***	0.204***	0.189***	0.117***	0.153***	
	(0.0261)	(0.0353)	(0.0325)	(0.0133)	(0.0161)	
Value per-capita 1086						25154.1***
						(4271.0)
Constant	10.46***	10.53***	10.49***	6.352***	6.570***	22471.0***
	(0.0454)	(0.0556)	(0.0552)	(0.0225)	(0.0274)	(1102.9)
Observations	283	257	251	283	282	283
R^2	0.094	0.093	0.103	0.185	0.214	0.090

Notes: This table shows the robustness of the main persistence result to various alternative specifications. In column one, I show the results without first winsorising either variable. In column two, I exclude local authorities in the North that are Harried. In column three, I exclude local authorities that consist of modern-day London. In column four, I use modern median wages in a local authority from the ASHE data as the dependent variable. In column five, I use modern average wages in a local authority from the ASHE data as the dependent variable. In column six, I show the results using each variable in levels as opposed to logged values.

Table 7 Harried two way fixed effects table

	Baseline	Weight by 1086 pop	Weight by 2020 pop
Harried × 1086	-1.644***	-2.323***	-1.574***
	(0.269)	(0.370)	(0.271)
Harried × 2020	-0.104	-0.745*	-0.197
	(0.301)	(0.400)	(0.323)
Observations	162	162	162
R2	0.644	0.727	0.633

Notes: This table shows the results from estimating a two-way fixed effects model with local authority and period fixed effects. It considers value per density over three time periods: 1066 (omitted category), 1086, and 2020. Standard errors are clustered at the local authority level. In column one, the regression is unweighted. In column two, I weight by 1086 population. In column three, I weight by the 2020 population.

Table 8 The Impact of the Harrying of the North over time: rank regressions

	Value 1066 by area (rank)	Value 1086 by area (rank)	GDP 2020 by area (rank)
Harried	-29.22	-159.4***	18.83
	(22.66)	(13.14)	(18.09)
Constant	170.6***	200.2***	86.82***
	(9.663)	(10.65)	(12.31)
Observations	54	54	54
R2	0.0324	0.734	0.0205

Notes: This table shows the estimated impact of the Harrying of the North on the Harried areas relative to other rebellious areas. Variables have been measured in year-specific ranks over all 283 modern local authorities; larger ranks imply higher values. The geographic unit of analysis is modern local authorities. Standard errors are robust. In column one, I show results for the 1066 value by area. In Column two, for 1086 value by area, and in Column three, for 2020 GDP by area. Observations and R-squared are given below. There are 25 treated (Harried) local authorities and 29 control (rebellious but not Harried) local authorities.

Figure 11 Estimated local productivities in 1086 and 2020

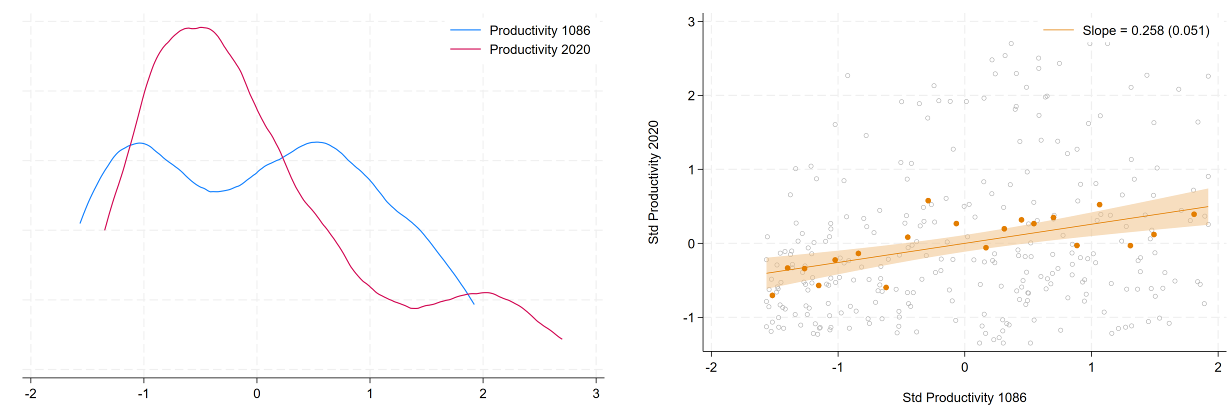

Notes: This figure displays the estimated productivity fundamentals in both 1086 and 2020. In the left-hand panel, I plot the distribution of the standardised data series in blue in 1086 and in red in 2020. In the right-hand panel, I plot the relationship between these two series in a scatter plot with an overlayed binscatter. The x-axis plots 1086 values, and the y-axis plots 2020 values. In orange, I show ventile averages (binscatter plot) and overlay a corresponding linear line of best fit and confidence interval. The top right-hand corner gives the slope of this line and the associated standard errors.

Figure 12 Estimated local amenities in 1086 and 2020

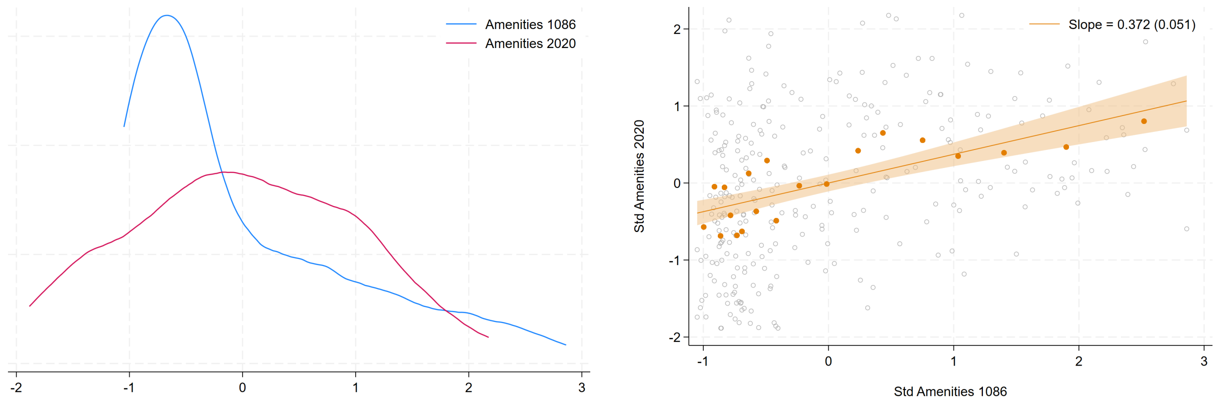

Notes: This figure displays the estimated amenity fundamentals in both 1086 and 2020. In the left-hand panel, I plot the distribution of the standardised data series in blue in 1086 and in red in 2020. In the right-hand panel, I plot the relationship between these two series in a scatter plot with an overlayed binscatter. The x-axis plots 1086 values, and the y-axis plots 2020 values. In orange, I show ventile averages (binscatter plot) and overlay a corresponding linear line of best fit and confidence interval. The top right-hand corner gives the slope of this line and the associated standard errors of this line.

Figure 13 Impact of the Harrying with intermediate periods

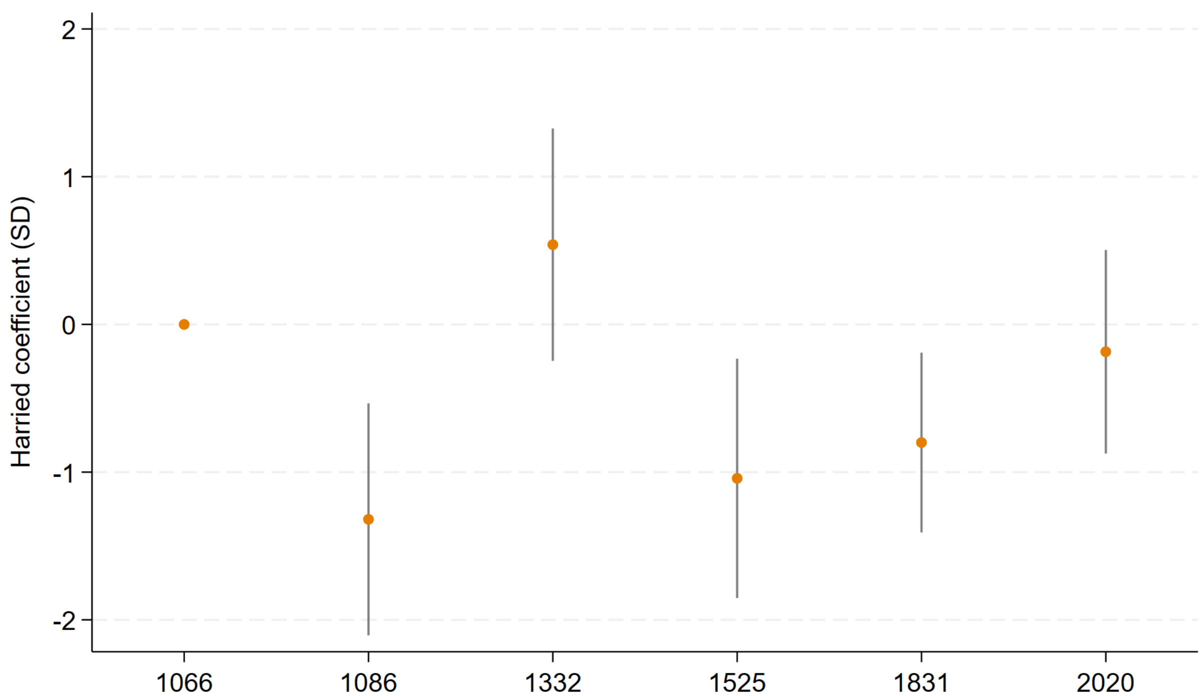

Notes: This figure shows the results from estimating a two-way fixed effects model on a constructed panel data set at the local-authority level on a sample of Harried and rebellious but not Harried local authorities. Each coefficient plotted corresponds to the impact of the Harrying in that year, measured in terms of standard deviations. 95% confidence intervals are depicted in grey. Data for 1066 and 1086 are taken from the Domesday Book as described in the text. Data for 1332 and 1525 are from the lay subsidy taxes and taken from Heldring et al. [2022]. Data from 1831 are from the census and proxy income with agricultural employment. Data from 2020 are modern-day values described in the text. The appendix section A.1 gives more details on the data sources and construction.

A.3 Details on rebellions against Williams rule

This section uses information from Barlow [2014], Rex [2014].

- North: In 1068, Edwin and Morcar, the earls of Mercia and Northumbria, started a rebellion supported by Edgar, who was crowned king by the English after the battle of Hastings and subsequently gave up his crown to William. The rebellion started because William had made the earl's lands smaller and imposed a heavy tax. William crushed the rebellion but forgave Edwin and Morcar, who were kept as guests in his court.

- North: In 1069, Edgar joined forces with Sweyn of Denmark and attacked York; afterwards, their forces scattered, causing small rebellions all over the country. William paid the Danes to leave and then pursued a campaign of destruction in the North- the Harrying of the North.

- East Anglia: In 1070-71, the Danes arrived in the fens in the east of England and were helped by a local ruler, Hereward the Wake. The rebellion ended when the Normans captured their center of power in Ely.

- East Anglia and Northumbria: The revolt of the earls 1075. Ralph de Guader, Earl of East Anglia, Roger de Breteuil, Earl of Hereford, and Waltheof, Earl of Northumberland, revolted. Waltheof almost immediately gave up and confessed to the rebellion. Rodger was stopped by the English bishop Wulfstan's fyrd. Ralph was similarly stopped near Cambridge and eventually fled to Denmark.

- Exeter: After the battle of Hastings, Gytha, Harold's mother, fled to Exeter and from there fomented rebellion against William. After Exeter refused to swear fealty or pay taxes to William in 1068, he laid siege to Exeter. After 18 days, William and his forces entered the city. By all accounts, William gave very generous terms, allowing the city to not pay taxes in return for pledging fealty, and the Williams' soldiers were denied their traditional right of looting the surrendered city.

- Welsh borders: William never conquered Wales, and indeed never appeared to intend to. However, Welsh forays and attacks in the border regions caused William to establish a series of earldoms in the borderlands at Chester, Shrewsbury, and Hereford.

- Kent: In Kent, a rebellion quickly surfaced against the local ruler Odo in 1067, who was extremely unpopular with his new subjects. This was put down by Odo without the need for William to personally intervene.

B Dynamic quantitative spatial economics model equilibrium

As described in Allen and Donaldson [2022] the equilibrium of the DQSE can be set out as follows. For any initial population vector $\{L_{i0}\}$ and vectors of geographic fundementals $\{\bar{A}_{it}, \bar{u}_{it}, \tau_{ijt}, \mu_{ijt}\}$ such that for all i, t the following holds.

1. A locations income equals the value of purchases from it: $w_{it}L_{it} = \sum_j X_{ijt}$. Which implies that $w_{it}^{\sigma} L_{it}^{1-\alpha_1(\sigma-1)} = \sum_j K_{ijt} L_{jt}^{\beta_1(\sigma-1)} W_{jt}^{1-\sigma} w_{jt}^{\sigma} L_{jt}$. Where all exogenous and predetermined variables have been bundled together into the Kernel:
$K_{ijt} = (\tau_{ijt}(\bar{A}_{it} L_{it-1}^{\alpha_2} \bar{u}_{jt} L_{jt-1}^{\beta_2})^{-1})^{1-\sigma}$.

2. Trade is balanced. Income is fully spent $w_{it}L_{it} = \sum_j X_{jit}$. Which implies that: $w_{it}^{1-\sigma} L_{it}^{\beta_1(1-\sigma)} w_{it}^{\sigma-1} = \sum_j K_{ijt} L_{jt}^{\alpha_1(\sigma-1)} w_{jt}^{1-\sigma}$.

3. Total population equals the sum of those arriving. $L_{it} = \sum_j L_{ijt}$. This implies that: $L_{it} V_{it}^{-\theta} = \sum_j \mu_{ijt}^{-\theta} \Pi_{jt}^{-\theta} L_{jt-1}$.

4. Total population in the previous period equals the sum of those leaving. $L_{it-1} = \sum_j Lijt$. Which implies that: $\Pi_{it}^{\theta} = \sum_j \mu_{ijt}^{-\theta} W_{jt}^{\theta}$.

We can simplify this system by imposing symmetry in trade costs, which implies that 1. and 2. can be combined. Thus we are left with a 3-equation model in each i, t with three unknowns $\{L_{it}, W_{it}, \Pi_{it}\}$, and unknown parameters $\{\alpha_1, \alpha_2, \beta_1, \beta_2, \sigma, \theta\}$.

References

Daron Acemoglu, Simon Johnson, and James A Robinson. The colonial origins of comparative development: An empirical investigation. *American economic review*, 91(5):1369–1401, 2001.

Alberto Alesina, Paola Giuliano, and Nathan Nunn. On the origins of gender roles: Women and the plough. *The quarterly journal of economics*, 128(2):469–530, 2013.

Treb Allen and Dave Donaldson. Persistence and path dependence in the spatial economy. Technical report, National Bureau of Economic Research, 2022.

Treb Allen, Costas Arkolakis, and Yuta Takahashi. Universal gravity. *Journal of Political Economy*, 128(2):393–433, 2020.

Treb Allen, Costas Arkolakis, and Xiangliang Li. On the equilibrium properties of spatial models. *American Economic Review: Insights*, 6(4):472–489, 2024.

Charles Angelucci, Simone Meraglia, and Nico Voigtländer. How merchant towns shaped parliaments: From the norman conquest of england to the great reform act. *American Economic Review*, 112(10):3441–3487, 2022.

Sam Asher and Paul Novosad. Rural roads and local economic development. *American economic review*, 110(3):797–823, 2020.

Clare Bambra. *Health divides: where you live can kill you.* Policy Press, 2016.

Frank Barlow. *The feudal kingdom of England: 1042-1216.* Routledge, 2014.

Hoyt Bleakley and Jeffrey Lin. Portage and path dependence. *The quarterly journal of economics*, 127(2):587–644, 2012.

Richard H Britnell and Bruce MS Campbell. *A commercialising economy: England 1086 to c. 1300.* Manchester University Press, 1995.

Jack Britton, Laura van der Erve, Ben Waltmann, and Xiaowei Xu. London calling? higher education, geographical mobility and early-career earnings. *Institute for Fiscal Studies Report, https://ifs. org. uk/publications/london-calling-higher-education-geographical-mobility-and-early career-earnings*, 2021.

Stephen Broadberry, Bruce MS Campbell, Alexander Klein, Mark Overton, and Bas Van Leeuwen. *British economic growth, 1270–1870.* Cambridge University Press, 2015.

Gregory Casey and Marc Klemp. Historical instruments and contemporary endogenous regressors. *Journal of Development Economics*, 149:102586, 2021.

Ting Chen, James Kai-sing Kung, and Chicheng Ma. Long live keju! the persistent effects of china's civil examination system. *The economic journal*, 130(631):2030–2064, 2020.

Martina Cioni, Giovanni Federico, and Michelangelo Vasta. Persistence studies: a new kind of economic history? *Review of Regional Research*, 42(3):227–248, 2022.

Gregory Clark. The long march of history: Farm wages, population, and economic growth, england 1209–1869 1. *The Economic History Review*, 60(1):97–135, 2007.

Diego Comin, William Easterly, and Erick Gong. Was the wealth of nations determined in 1000 bc? *American Economic Journal: Macroeconomics*, 2(3):65–97, 2010.

Timothy G Conley. Gmm estimation with cross sectional dependence. *Journal of econometrics*, 92(1):1–45, 1999.

Timothy G Conley and Morgan Kelly. The standard errors of persistence. *Journal of International Economics*, 153:104027, 2025.

Nicholas Crafts. *Forging ahead, falling behind and fighting back: British economic growth from the industrial revolution to the financial crisis*. Cambridge University Press, 2018.

Nicholas FR Crafts. British economic growth during the industrial revolution. 1985.

Carl-Johan Dalgaard, Nicolai Kaarsen, Ola Olsson, and Pablo Selaya. Roman roads to prosperity: Persistence and non-persistence of public infrastructure. *Journal of Comparative Economics*, 50(4):896–916, 2022.

Paul Dalton. *Conquest, anarchy and lordship: Yorkshire, 1066-1154*. Number 27. Cambridge University Press, 2002.

Henry Clifford Darby and Henry Clifford Darby. *Domesday England*. Cambridge University Press, 1986.

Henry Clifford Darby, Robin E Glasscock, John Sheail, and George R Versey. The changing geographical distribution of wealth in england: 1086–1334–1525. *Journal of Historical Geography*, 5(3):247–262, 1979.

Martin J Daunton. *The Cambridge urban history of Britain*. Cambridge University Press, 2001.

Donald R Davis and David E Weinstein. Bones, bombs, and break points: the geography of economic activity. *American economic review*, 92(5):1269–1289, 2002.

Vincent Delabastita and Sebastiaan Maes. The feudal origins of manorial prosperity: Social interactions in eleventh-century england. *The Journal of Economic History*, 83(2):464–500, 2023.

Melissa Dell. The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903, 2010.

Dave Donaldson and Richard Hornbeck. Railroads and american economic growth: A "market access" approach. *The Quarterly Journal of Economics*, 131(2):799–858, 2016.

Sebastian Ellingsen. Long-distance trade and long-term persistence. *University of Pompeu Fabra Working Paper*, 2025.

FAO and IIASA. Global agro ecological zones version 4 (gaez v4). *Accessed January 2024. http://www.fao.org/gaez/*, 2021.

Raphaël Franck and Oded Galor. Flowers of evil? industrialization and long run development. *Journal of Monetary Economics*, 117:108–128, 2021.

Vivian H Galbraith. An episcopal land-grant of 1085. *The English Historical Review*, 44 (175):353–372, 1929.

Oded Galor and Ömer Özak. The agricultural origins of time preference. *American Economic Review*, 106(10):3064–3103, 2016.

Stephen Gibbons, Stephan Heblich, and Edward W Pinchbeck. The spatial impacts of a massive rail disinvestment program: The beeching axe. *Journal of Urban Economics*, 143: 103691, 2024.

Luigi Guiso, Paola Sapienza, and Luigi Zingales. Long-term persistence. *Journal of the European Economic Association*, 14(6):1401–1436, 2016.

Mark Hagger. *William: King and Conqueror*. Bloomsbury Publishing, 2021.

Hanson. An urban geography of the roman world, 100 b.c. to a.d. 300. *University of Oxford D.Phil.*, 2016a.

Hanson. Cities database (oxrep databases). version 1.0. access 05/03/2024. *DOI: ¡https://doi.org/10.5287/bodleian:eqapevAn8*, 2016b.

Luke Heath Milsom. Persistent spatial equilibria. evidence from a sudden river. Technical report, The Institute for Fiscal Studies, 2024.

Stephan Heblich, Alex Trew, and Yanos Zylberberg. East-side story: Historical pollution and persistent neighborhood sorting. *Journal of Political Economy*, 129(5):1508–1552, 2021.

Stephan Heblich, Stephen J Redding, and Hans-Joachim Voth. Slavery and the british industrial revolution. 2023.

Michael Hechter. Regional inequality and national integration: The case of the british isles. *Journal of Social History*, 5(1):96–117, 1971.

Leander Heldring, James A Robinson, and Sebastian Vollmer. The economic effects of the english parliamentary enclosures. Technical report, National Bureau of Economic Research, 2022.

J Vernon Henderson, Tim Squires, Adam Storeygard, and David Weil. The global distribution of economic activity: nature, history, and the role of trade. *The Quarterly Journal of Economics*, 133(1):357–406, 2018.

David Horspool. *The English Rebel: One Thousand Years of Trouble-making from the Normans to the Nineties*. Penguin UK, 2009.

E.H. Hunt. Industrialization and regional inequality: wages in britain, 1760–1914. *The Journal of Economic History*, 46(4):935–966, 1986.

Remi Jedwab, Edward Kerby, and Alexander Moradi. History, path dependence and development: Evidence from colonial railways, settlers and cities in kenya. *The Economic Journal*, 127(603):1467–1494, 2017.

Remi Jedwab, Noel D Johnson, and Mark Koyama. Pandemics and cities: Evidence from the black death and the long-run. *Journal of Urban Economics*, 139:103628, 2024.

Morgan Kelly, Joel Mokyr, and Cormac Ó Gráda. The mechanics of the industrial revolution. *Journal of Political Economy*, 131(1):59–94, 2023.

Sanghoon Lee and Jeffrey Lin. Natural amenities, neighbourhood dynamics, and persistence in the spatial distribution of income. *The Review of Economic Studies*, 85(1):663–694, 2018.

Jeffrey Lin and Ferdinand Rauch. What future for history dependence in spatial economics? *Regional Science and Urban Economics*, 94:103628, 2022.

Philip McCann. Perceptions of regional inequality and the geography of discontent: Insights from the uk. *Regional Studies*, 54(2):256–267, 2020.

Peter McClure. Patterns of migration in the late middle ages: the evidence of english place-name surnames. *The Economic History Review*, 32(2):167–182, 1979.

M. et al. McCormick. Roman road network (version 2008). *DARMC Scholarly Data Series, Data Contribution Series 2013-5. DARMC, Center for Geographic Analysis, Harvard University, Cambridge MA 02138.*, 2013.

John McDonald. Efficiency in the domesday economy, 1086: evidence from wiltshire estates. *Applied Economics*, 42(25):3231–3240, 2010.

John McDonald and Graeme D Snooks. The determinants of manorial income in domesday england: evidence from essex. *The Journal of Economic History*, 45(3):541–556, 1985.

Guy Michaels and Ferdinand Rauch. Resetting the urban network: 117–2012. *The Economic Journal*, 128(608):378–412, 2018.

Edward Miguel and Gerard Roland. The long-run impact of bombing vietnam. *Journal of development Economics*, 96(1):1–15, 2011.

Ulrich K Müller and Mark W Watson. Spatial correlation robust inference. *Econometrica*, 90(6):2901–2935, 2022.

Pamela Nightingale. *Trade, money, and power in medieval England*. Taylor & Francis, 2023.

Nathan Nunn and Diego Puga. Ruggedness: The blessing of bad geography in africa. *Review of Economics and Statistics*, 94(1):20–36, 2012.

Nathan Nunn and Nancy Qian. The potato's contribution to population and urbanization: evidence from a historical experiment. *The quarterly journal of economics*, 126(2):593–650, 2011.

Nathan Nunn and Leonard Wantchekon. The slave trade and the origins of mistrust in africa. *American economic review*, 101(7):3221–3252, 2011.

David M Palliser. Domesday book and the'harrying of the north'. *Northern History*, 29(1): 1–23, 1993.

J. Palmer. Electronic edition of domesday book: Translation, databases and scholarly commentary, 1086; second edition [computer file]. *http://dx.doi.org/10.5255/UKDA-SN-5694-1*, 2016.

Stephen J Redding and Daniel M Sturm. The costs of remoteness: Evidence from german division and reunification. *American Economic Review*, 98(5):1766–1797, 2008.

Peter Rex. *The English Resistance: The Underground War Againt the Normans*. Amberley Publishing Limited, 2014.

David Roffe. *Decoding Domesday*. Boydell & Brewer Ltd, 2015.

Michele Rosenberg and Federico Curci. Factory location: resistance to technology adoption and local institutions. *Available at SSRN 4428398*, 2024.

Josiah Cox Russell. British medieval population. 1948.

A.E.M Satchell, P.K Kitson, G.H Newton, L. Shaw-Taylor, and E.A Wrigley. 1851 england and wales census parishes, townships and places. [data collection]. *UK Data Archive. 10.5255/UKDA-SN-852232*, 2023.

Michael E Smith. Peasant mobility, local migration and premodern urbanization. *World Archaeology*, 46(4):516–533, 2014.

Pauline Stafford. Women in domesday. *Reading Medieval Studies*, 15:75–94, 1989.

Anna Stansbury, Daniel Turner, and Edward Balls. Tackling the uk's regional economic inequality: Binding constraints and avenues for policy intervention. *M-RCBG Associate Working Paper Series*, 2023.

MJ Strickland. *Armies, chivalry and warfare in medieval Britain and France*, volume 7. Paul Watkins, 1998.

Marcel Timmer, Gaaitzen J de Vries, and Klaas De Vries. Patterns of structural change in developing countries. In *Routledge handbook of industry and development*, pages 65–83. Routledge, 2015.

Felipe Valencia Caicedo. The mission: Human capital transmission, economic persistence, and culture in south america. *The Quarterly Journal of Economics*, 134(1):507–556, 2019.

Ordericus Vitalis. *The ecclesiastical history of England and Normandy*, volume 3. HG Bohn, 1854.

Nico Voigtländer and Hans-Joachim Voth. Persecution perpetuated: the medieval origins of anti-semitic violence in nazi germany. *The Quarterly Journal of Economics*, 127(3): 1339–1392, 2012.

Hans-Joachim Voth. Persistence–myth and mystery. In *The handbook of historical economics*, pages 243–267. Elsevier, 2021.

James T Walker. National income in domesday england. In *Money, prices and wages: Essays in honour of Professor Nicholas Mayhew*, pages 24–50. Springer, 2015.

Jeffrey G Williamson. Regional inequality and the process of national development: a description of the patterns. *Economic development and cultural change*, 13(4, Part 2):1–84, 1965.