ONLINE APPENDIX: TAXATION AND INNOVATION IN THE 20TH CENTURY

Ufuk Akcigit
John Grigsby
Tom Nicholas
Stefanie Stantcheva

May 24, 2021

Contents

A Data Definitions and Construction

A.1 Variable Definitions

In this section, we detail the construction of relevant variables for our analysis. Variables related to corporate tax base rules are described in detail in Section A.6.2.

- *Top Corporate Marginal Tax Rate (Corp. MTR)* - The additional tax burden accruing to a firm in the top tax bracket in state s for an additional one dollar of revenue if all of its operations were in s.
- *90^{th} Percentile Income Marginal Tax Rate (MTR90)* - The additional tax burden accruing to an individual at the 90^{th} percentile of the national income distribution for an additional one dollar of earnings. Calculated using the tax calculator by Bakija (2006).
- *90^{th} Percentile Income Average Tax Rate (ATR90)* - The total tax burden for an individual at the 90^{th} percentile of the national income distribution divided by that individual's total income. Calculated using the tax calculator by Bakija (2006).
- *Median Income Marginal Tax Rate (MTR50)* - The additional tax burden accruing to an individual at the 50^{th} percentile of the national income distribution for an additional one dollar of earnings. Calculated using the tax calculator by Bakija (2006). Data on median incomes come from the Census. Table P-53 reports the median income for men back to 1947. These data are available from `https://www.census.gov/data/tables/time-series/demo/income-poverty/historical-income-people.html`. We also use income data from the 1940 decennial Census, and log-linearly interpolate median incomes from 1940 through 1947.
- *Median Income Average Tax Rate (ATR50)* - The total tax burden for an individual at the 50^{th} percentile of the national income distribution divided by that individual's total income. Calculated using the tax calculator by Bakija (2006).
- *Inventor productivity* - An inventor's productivity in year t is defined to be the number of eventually-granted patents that the inventor has applied for as of year $t-1$. In robustness table C.26, inventor i's productivity in year t is defined to be the total number of citations ever received by patents applied for by i through year t. An inventor is said to be "high productivity" in year t if they are in the top 10% of the national inventor productivity distribution in year t. In robustness table C.23, an inventor is said to be high productivity if they are in the top 5% of the national productivity distribution in year t. Finally, robustness table C.25 allows an inventor to be high productivity if they are in the top 10% of the productivity distribution, of middle productivity if they are between the 75^{th} and 90^{th} percentile of the productivity distribution, and low productivity otherwise.
- *Personal MTR* - An inventor's effective marginal (average) tax rate is defined to be the marginal (average) tax rate faced by the 90^{th} percentile earner in the national income distribution if the inventor is high productivity, and the marginal (average) tax rate faced by a median earner if the inventor is low productivity. In appendix table C.25, middle productivity inventors have an effective tax rate equal to the tax rate faced by an individual earning at the 75^{th} percentile of the national income distribution. In all regressions, we use lagged effective tax rates as independent variables. Thus an inventor living in state s will face an effective tax rate for innovation output in year t which is the effective tax rate the inventor would have faced in year $t-1$ given their $t-1$ productivity level and the tax laws in place in year $t-1$.

- *Log Patents* - The natural logarithm of the number of eventually-granted patents applied for in state s in year t. In inventor-level regressions, this variable corresponds to the log of the number of eventually-granted patents applied for by inventor i in years t through $t+2$.
- *Log Citations* - The natural logarithm of the number of citations ever received by eventually-granted patents which were applied for in state s in year t. In inventor-level regressions, this variable corresponds to the log of the number of citations ever received by eventually-granted patents which were applied for by inventor i in years t through $t+2$. Citation counts adjusted according to the algorithm of Hall, Jaffe and Trajtenberg (2001), detailed for our data in Appendix A.5.
- *Log Inventors* - The natural logarithm of number of inventors in state s in year t as implied by the Lai et al. (2014) algorithm applied to our dataset. A detailed description of this algorithm is provided in Appendix A.2.
- *Corporate Patent* - A corporate patent is one which is assigned to a corporation after being granted.
- *Share Assigned* - The share of patents in state s in year t which are assigned to a corporation.
- *Has Patent* - An indicator variable, equal to 1 if the inventor has at least one successful patent application between years t and $t+2$. Inventors are included in the regression sample for the period between their first ever successful patent application, and their last ever successful patent application.
- *Has 10+ Cites* - An indicator variable, equal to 1 if the inventor's patents, applied for between years t and t+2, ever receive at least 10 citations in total between them. Inventors are included in the regression sample for the period between their first ever successful patent application, and their last ever successful patent application. Patent citation counts adjusted according to the algorithm of Hall, Jaffe and Trajtenberg (2001), detailed for our data in Appendix A.5.
- *Corporate Inventor* - An inventor is said to be a corporate inventor in year t if they are granted at least one corporate patent in the next three years.
- *Agglomeration* - The number of patents, in thousands, applied for by inventors $j \neq i$ who share inventor i's modal patent class in year t in state s.
- *Home State Flag* - An indicator variable equal to one if the state is the first state in which the inventor applies for a patent.
- *Assignee Has Patent* - An indicator variable equal to one if an inventor i's firm has at least one patent applied for in year t by an inventor $j \neq i$ in destination state s.
- *Inventor Tenure/Experience* - an inventor's tenure is the number of years that have passed since the inventor's first successful patent application.
- *Personal Income per Capita* - state level personal income per capita, included in regressions in units of tens of thousands of 1982 dollars. Downloaded from `https://www.bea.gov/data/economic-accounts/regional`, accessed on March 15, 2021.
- *R&D Tax Credit:* statutory credit rate adjusted for recapture and type of credit for a given state-year;

- *Population Density* - thousands of people per square kilometer.
- *Tax Base Index* - corporate tax base index. Higher values generally correspond to broader tax bases. Construction of the index is detailed in Appendix A.6.2.
- *State Governor: Democrat* - an indicator variable equal to 1 if the state's governor is a Democrat.
- *% State Upper House Democrat* - the percent of a state's Upper House of the legislative branch who is a Democrat.
- *% State Lower House Democrat* - the percent of a state's Lower House of the legislative branch who is a Democrat.
- *State Pat./State Patents* - the natural log of the total number of patents applied for in the inventor's residence state.
- *Kogan et al. (2017) patent value* - Patent value constructed by Kogan et al. (2017) by considering jumps in the total stock market value of an assignee in a short window around the successful patent grant, deflated to millions of 1982 dollars. This is only defined for patents granted to publicly-traded corporations.
- *Has High-Value Patent (Pat.)* - An indicator variable equal to one if an inventor has, over the subsequent three years, Kogan et al. (2017) stock market value of patents applied for among all corporate inventors active in year t.

A.2 Disambiguation Algorithm

This section describes the algorithm used to disambiguate inventor names and presents measures of its performance.

A.2.1 Algorithm Description

We employ the algorithm of Lai et al. (2014) to disambiguate inventors in our historical patent data.[1] The goal of disambiguation is to determine if two patent-inventor level records were produced by the same inventor. A problem of this sort may be distilled into a clustering problem well-suited to standard machine learning algorithms: given a training dataset and a set of features – such as inventor name, location, technology class, assignee, and coauthor networks – we wish to group records together into profiles which indicate that the two records were produced by the same inventor. The goal is to assign probabilities of an inventor match based on the characteristics of every pair of observations. The central idea is that two records coming from two very similar names (not necessarily identical: "John A Smith" vs "John Adam Smith" for instance) working in similar subject areas, working for the same company in roughly the same geographic location, are likely to be the same person.

Such a machine learning approach has three central benefits relative to other more rudimentary approaches, such as treating each individual name as a separate inventor, or hand-matching innovators' records to one another. First, the Lai et al. approach permits minor name typos or data entry errors, without incorrectly decoupling these inventors. Second, it provides probabilistic

[1]The code and associated files for the original disambiguation may be downloaded from `https://github.com/funginstitute/downloads`; accessed October 13, 2016.

matches based on more information than name and location, which helps disambiguate between common names – a John Smith working in software is likely different to a John Smith with patents in bootmaking. Finally, the algorithm does not impose any functional forms on the relationship between a pair's set of attributes and the probability that those pairs belong to the same inventor.

Of course, this machine learning approach is imperfect and will struggle to correctly match inventors who drastically change their names or have exceptional careers. For instance, if an inventor named Jane Smith changes her name after marrying a man with surname Robertson, the algorithm will struggle to adapt, as names are the most distinguishing piece of information amongst records. Similarly, if a software engineer living in California and working for Apple decides to change his career and move to Montana to open a new shoe factory, the algorithm is likely to suggest that these are two separate inventors, rather than one inventor with such an uncommon career trajectory.

The clustering exercise is subject to two principal challenges. First, one must produce a suitable training dataset from which to glean the probability that two patent records with a similarity profile of x belong to the same inventor. Here, one may follow two approaches. One could submit a hand-curated dataset of known matches to the disambiguation algorithm to determine the likelihood of a match. However, the construction of these datasets are often subject to bias if, for example, researchers are more likely to include better-known inventors. An alternative approach, and the one followed by Lai et al., is to allow the algorithm to produce its own training dataset based on features in the data. For example, a training dataset of known matches could be constructed by examining individuals with matching rare names.

Our baseline approach lies somewhere in between these two strategies. We use the matches of Lai et al. to form the basis of our training dataset. We draw twenty million pairs of records belonging to different inventors according to Lai et al. to complete our training dataset. Using this as a training dataset relies on two principal assumptions: first, we assume that the Lai et al. disambiguation correctly identifies inventors, and second we assume that the sets of features that were predictive of inventor clustering are stable over time, so that the same rules for determining matches in the modern sample of Lai et al. will apply to our historical sample. We choose this approach in order to best match the state-of-the-art disambiguation of inventors in the modern data.[2]

The second major challenge to the disambiguation exercise is computational. Ideally, one would compare every pair of records in our data, and build a similarity profile for each. However, with over 12 million unique patent-inventor records in our dataset, one would have to compare over 144 trillion record pairs in order to compare each record to each other, which is computationally infeasible. To circumvent this challenge, we follow Lai et al. in disambiguating successively larger blocks. We first group records into blocks of possible matches, based on the first characters of an inventor's name. Then we compare all records within a block to one another, but never compare across blocks. After disambiguating a set of narrow blocks, we expand the size of the block, for example by considering all record pairs that match the first three letters of an inventor's name, rather than the first five letters. By iteratively allowing progressively larger blocks, and assuming clusters within prior blocking rounds were successfully disambiguated, we greatly reduce the computational burden of the disambiguation.

Our starting point is the historical inventor data digitized by Akcigit, Grigsby and Nicholas (2017), combined with the patent data of Lai et al. (2014) available on the Harvard Dataverse Network (HDN).[3] We first manually clean inventor names and location to correct for obvious

[2]In early versions of the paper, we experimented with allowing the algorithm to find its own training sets, and found qualitatively similar headline results.

[3]Accessed from https://dataverse.harvard.edu/dataset.xhtml?persistentId=hdl:1902.1/15705 on Febru-

typos. The most common correction is to remove prefixes and suffixes, such as “DR,” “JR,” and “SR.” In addition, we standardize names to be all capital letters, and consider a person’s first name to be the first word of their name. Finally, we consider only the first patent class listed on a patent document to be that patent’s primary classification. If our OCR process generates missing information, such as a missing name, class or location, we supplement the data with patent data from patentsview and the NBER patent data to fill in information where possible. These auxiliary data sources only contain patents granted since 1975, and so cannot improve the dataset’s early years. Dropping the 1970s from our analysis does not meaningfully affect our results.

To compare records, we construct a similarity profile for every pair of records to be compared. A similarity profile x is a vector of similarity scores for the active attributes in the disambiguation. Specifically, a similarity profile is encoded as follows:

- First and Last names
  1. If one of the two records is missing the name
  2. If there is no clear misspelling or abbreviation employed, and the strings do not exactly match
  3. If there is a misspelling (defined as either missing 1 or 2 characters somewhere, or switching the place of a few characters)
  4. If exact match or, in the case of first names, if one string appears to be an abbreviation of the other in that it has the first 3 characters the same (e.g. “ROB” and “ROBERT”)
- Middle Names
  0. If have different middle names
  1. If one of the two records have missing middle name
  2. If both records have missing middle name
  3. If one record has a full middle name (e.g. “WILLIAM”) and the other has just the middle initial which matches the full middle name (e.g. “W”).
  4. If exactly the same name
- Location
  1. If over 50 miles apart
  2. If under 50 miles apart
  3. If under 25 miles apart
  4. If under 10 miles apart
  5. If under 1 mile apart
- Patent Classes
  0. If different strings
  1. If exactly the same string
- Assignees

ary 13, 2017.

5. If the Jaro-Winkler string distance between assignee names is at least 0.9
4. If JW distance > 0.8
3. If JW distance > 0.7
2. If one of the two records has a missing assignee
1. Otherwise

- Coauthors

  1. If coauthors exactly the same (coauthors entered as <First Initial> . <Last Name> and separated by comma in the variable)
  0. Otherwise

- Country

  0. If different country
  1. If the same non-US country
  2. If the same US country

Next, one may construct, for every observed similarity profile, the probability that this profile belongs to the same inventor or not, by comparing the frequency with which it occurs in the training dataset. Specifically, defining $\mathcal{M}$ to be the set of matched inventor pairs in the training dataset, and $\mathcal{N}$ to be the set of non-matched inventor pairs in the training dataset, one may use Bayes' rule to write the probability of a match as

$$P(\mathcal{M}|x) = \frac{P(x|\mathcal{M})P(\mathcal{M})}{P(x|\mathcal{M})P(\mathcal{M}) + P(x|\mathcal{N})\,(1 - P(\mathcal{M}))}$$

where $P(\mathcal{M})$ is the prior probability of a match, which we follow Lai et al. in setting as proportional to the ratio of the number of within-cluster pairs (i.e. disambiguated inventors from prior blocking rounds) in a block to the total number of pairs in that block.[4] For numerical reasons, it is more convenient to work with the posterior *odds* of a match, defined as

$$\frac{P(\mathcal{M}|x)}{1 - P(\mathcal{M}|x)} = \frac{P(x|\mathcal{M})}{P(x|\mathcal{N})} \cdot \frac{P(\mathcal{M})}{1 - P(\mathcal{M})}$$

In particular, we calculate the likelihood ratio, $r(x)$, for every observed similarity profile x. This likelihood ratio is defined as

$$r(x) = \frac{P(x|\mathcal{M})}{P(x|\mathcal{N})} \tag{1}$$

This can be determined directly from the training dataset by comparing the number of records with similarity profile x that belong in the matched training dataset (i.e. come from the same inventor), to the number of records with similarity profile x that belong in the unmatched training dataset

[4] The discrete nature of the similarity profile space described above makes the computation of this match probability much simpler.

(i.e. come from different inventors).[5] Once we have the likelihood ratios calculated, we invert them to calculate the probability that two records originated from the same inventor:

$$P(\mathcal{M}|x) = \frac{1}{1 + \frac{1-P(\mathcal{M})}{P(\mathcal{M})}\frac{1}{r(x)}} \tag{2}$$

We say that two records originated from the same inventor if this posterior probability of a match is at least 0.99.[6]

Our blocking routine proceeds as follows:[7]

Round 1. Block based on exact first and last name. Compare records based on middle name and patent location.

Round 2. Block based on exact first and last name. Compare records based on middle name, coauthor network, patent class, and assignee name.

Round 3. Block based on first five characters of first name, and exact last name. Compare records based on middle name, coauthor network, patent class, and assignee name.

Round 4. Block based on first three characters of first name, and exact last name. Compare records based on middle name, coauthor network, patent class, and assignee name.

In addition, we have tested the sensitivity of our results to various alternative disambiguation routines, such as ignoring middle name or assignee (for which information is more sparse in early periods), using automatically generated training datasets or taking the Lai et al. (2014) disambiguation as a training dataset, shifting the threshold for two records to be considered as belonging to the same inventor, and introducing additional blocking rounds into the disambiguation. The core results of our paper are robust to these alternative disambiguation approaches. Our baseline disambiguation is the one that best balanced the false positive and false negative rates according to our tests described below.

Finally, we subset our data to only consider US inventors. As was indeed the case in our time period, the most productive inventors are Kia Silverbrook, Shunpei Yamazaki, George Lyon, Donald Weder, and Melvin De Groote. We refer the reader to Lai et al. (2014) for additional statistics on the performance of the algorithm on modern data.

A.2.2 Disambiguation: Performance

This disambiguation procedure is subject to two types of error. The first, which we dub a Type I error, occurs when two records are linked together as if they were from the same inventor, but in reality was created by two distinct individuals. Such errors are most likely to occur from the "common name" problem, where multiple distinct individuals have the same name. This is particularly

[5] To account for small sample bias in rare similarity profiles, we follow Lai et al. in applying a Laplace correction to these likelihood ratio values.

[6] In the early stages of our analysis, we experimented with match thresholds of 0.98 and 0.95 to determine whether records originated from the same inventor. After examining the data by hand, we determined that this was too low, as common names such as Robert Smith were often spuriously considered the most prolific inventors in the data. This problem largely vanished with the threshold of 0.99.

[7] We experimented with additional rounds of blocking, as well as with allowing for inexact surname matches in the blocking routine. Manual checks of the data revealed that this routine minimized errors with common names, and correctly matched the most productive inventors as listed by outside data sources.

difficult if, for instance, parents gave eponymous names to their offspring. Another source of this problem arises from an inability to distinguish whether two records with the same name appear in multiple states in two subsequent years – it can be difficult to tell whether such records reflect one inventor who moves or two distinct inventors.

The second, a Type II error, occurs when two records fail to be linked together into one inventor ID despite being produced by the same inventor. These errors occur when the algorithm is too strict in its criteria for matching workers together. Increasing the threshold of similarity at which two records are considered the same inventor reduces the likelihood of Type I errors while increasing the probability of Type II errors.

We check the performance of the algorithm along these two dimensions as follows. To check Type I errors, we construct a dataset by considering records that share same first three letters of their last name. We consider the 8 most common first three letters: "SCH," "WIL," "HAR," "CHA," "SMI," "STE," "SIL," "JOH," "BRO." We then manually checked this dataset to determine the Type I error rate by comparing records' similarity profiles, as well as their application years (to remove the influence of eponymous parent-child pairs). The goal was to determine the share of patent $\times$ inventor records that were assigned the same inventor ID but appeared to belong to different people by human judgment. This process revealed that just 1.5% of records were incorrectly grouped into the same inventor ID, suggesting that the Type I error rate is low.

To check the prevalence of Type II errors, we consider the share of records with the same set of progressively comprehensive characteristics that are matched to the same inventor. Table A.1 reports these shares. Each row of the table reports the share of records with the same set of observable characteristics that have multiple inventor IDs. For instance, the first row reports the share of records which have the same name but multiple inventor IDs, while the second row shows the share of records that have the same name and same state, but multiple inventor IDs. As we progress down the table, the set of matching characteristics get more and more stringent - the final row reports the share of records which have the same name, class, assignee, year and city, but multiple record IDs. For the purposes of the discussion below, let us define two inventor IDs to be "part of an incorrect split" for a given set of characteristics if they share those same characteristics, but have different inventor IDs.

The four columns calculate this share in different ways. The first two columns consider inventor-weighted shares: that is, what share of inventor IDs are part of an incorrect split? The second two columns consider patent-weighted shares: what share of patents belong to inventor IDs that are part of an incorrect split? Columns 1 and 3 count all inventor IDs that are part of an incorrect split in the numerator of this share. That is, if there are three inventors IDs with the same set of characteristics, all three will count as being part of an incorrect split. This calculation somewhat overweights the set of incorrect splits. For instance, suppose that there are two inventors, labeled A and B, that have the same name and location. Inventor A has 1,000 patents, while inventor B only has one. One might reasonably claim that only inventor B's one patent record was incorrectly disambiguated. To account for this, columns 2 and 4 remove the inventor ID with the most patents within a set of criterion from the set of incorrectly split IDs.

The table shows that 34.7% of inventor IDs have another inventor ID with the same name. These inventors account for 22.4% of patents. However, only 7% of inventors have the same name, class and state as another inventor ID (accounting for 4.4% of patents). Removing the most prolific inventor from the set of incorrectly split inventors reveals that only 3.2% of inventor IDs (accounting for 1.5% of patents) were split from another inventor with the same state, name and patent class.

These numbers are small and are unlikely to meaningfully bias our core results. By way of comparison, Figure A.1 compares the numbers of Table A.1 with analogous numbers calculated

Table A.1: Disambiguation Performance: Type II errors

		Inventor-Weighted		Patent-Weighted	
		With Duplicates	No Duplicates	With Duplicates	No Duplicates
Criteria	**Location Match**	(1)	(2)	(3)	(4)
Name Only	No Location	34.7	11.4	22.4	9.6
	State	22.9	9.2	15.0	5.4
	County	12.5	5.4	8.2	2.8
	City	9.2	4.1	6.1	2.0
Name + Class	No Location	9.6	4.2	5.9	2.0
	State	7.0	3.2	4.4	1.5
	County	4.2	1.9	2.6	0.84
	City	3.0	1.4	1.9	0.60
Name + Assignee	No Location	11.4	4.9	6.9	2.4
	State	6.6	3.0	4.2	1.3
	County	2.4	1.2	1.6	0.50
	City	0.07	0.03	0.05	0.01
Name + Class + Assignee	No Location	3.3	1.6	2.0	0.64
	State	2.2	1.0	1.4	0.42
	County	0.80	0.39	0.52	0.16
	City	0.02	0.01	0.01	0.00
Name + Class + Assignee + Year	No Location	0.65	0.32	0.35	0.11
	State	0.45	0.22	0.25	0.08
	County	0.15	0.08	0.08	0.03
	City	0.00	0.00	0.00	0.00

Notes: Table reports the share of potential type II errors by a variety of criteria. The table shows the share of observations that have the same sets of a given set of criteria, but multiple different inventor IDs in our disambiguated patent data. Each row corresponds to a different set of criteria; for instance, the top row considers the share of records with the same name that have different inventor IDs. Each column represents a different method of counting identical records. Columns (1) and (2) consider the share of inventor IDs which are part of a duplicate pair, while columns (3) and (4) consider the share of patents that are part of a duplicate pair. Columns (2) and (4) only consider the inventor ID with the fewest patents to be incorrectly matched, while columns (1) and (3) consider both inventor IDs with the same values of the criteria to be matched.

Figure A.1: Type II Error Comparison – Harvard Dataverse NetworkLai et al. (2014) comparison with our disambiguation

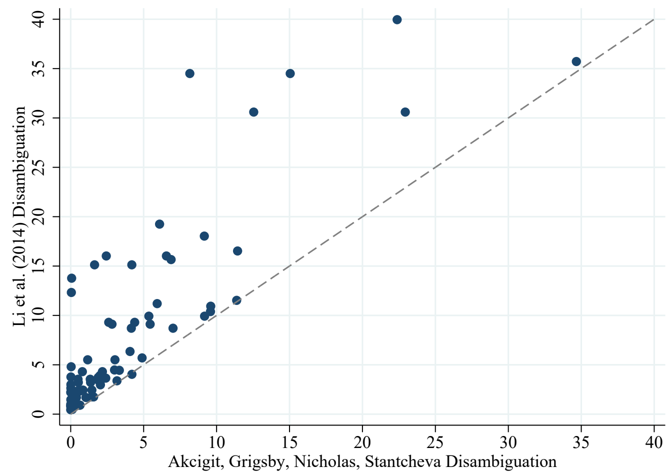

Notes: Figure plots the rate of Type II errors in our disambiguation and in the Lai et al. (2014) data. Each dot is a cell from Table A.1; see notes from that Table for details.

using the Lai et al. (2014) dataset. The horizontal axis plots the rate of incorrect splitting in our data, while the vertical axis plots it in the Lai et al. data. The dashed line is a 45-degree line, while each dot is a different set of characteristics and method of calculating the share of split inventors; that is, each dot represents a cell of Table A.1. The fact that every dot is above the 45-degree line indicates that our disambiguation has fewer Type II errors than the disambiguation of Lai et al. (2014) by every method used to calculate it. We consider this to be evidence that our disambiguation performs admirable, particularly coupled with the low Type I error rate of 1.5%.[8]

Comparison with External List of Prolific Inventors. In this section, we assess our disambiguation algorithm's performance on the most prolific US inventors. We compare the list of prolific inventors available from Wikipedia[9] to the number of patents of these inventors according to our disambiguation algorithm.

The first step of this exercise is to identify the inventors from our disambiguated dataset that correspond to the ones listed in Wikipedia. To do so, we match each of Wikipedia's inventors to a set of inventors in our dataset with the same first and last name. We then keep only inventors from our dataset that have either a missing middle name or a middle name starting with the same letter of the Wikipedia's counterpart. Subsequently, we choose the unique inventor ID in each set with total number of patents closest to Wikipedia's. Results are robust to choosing simply the most patented inventor per set. Then, we compute the percentage deviation of total patents of such inventors relative to Wikipedia's and we classify the results in different intervals.

[8]We also considered allowing inventors to have the same name if one inventor's name was a common nickname of another. A list of common English nicknames was obtained from `https://www.familiesunearthed.com/nicknames.htm`. Accounting for nicknames in this way had minimal impact on these tests, increasing the error rate by no more than 1% by any measure.

[9]Accessed from `https://en.wikipedia.org/wiki/List_of_prolific_inventors` on January 15, 2021. The source of the data from the Wikipedia table is mainly `http://patft.uspto.gov/netahtml/PTO/index.html`, or Google patents

Table A.2: Comparison of Wikipedia Patent Counts and Disambiguated Patent Counts for Prolific Inventors

Patent Count: Ratio of Our Disambiguation to Wikipedia	Threshold of Career Overlap					
	50% (1)	60% (2)	70% (3)	80% (4)	90% (5)	100% (6)
1.05+	11 (12%)	11 (17%)	8 (21%)	6 (29%)	5 (52%)	2 (25%)
1.01 – 1.05	8 (9%)	8 (13%)	8 (21%)	5 (24%)	1 (8%)	1 (13%)
0.99 – 1.01	15 (16%)	11 (17%)	7 (18%)	5 (24%)	3 (25%)	3 (38%)
0.95 – 0.99	17 (19%)	12 (19%)	6 (16%)	3 (14%)	3 (25%)	2 (25%)
0.90 – 0.95	14 (15%)	4 (6%)	1 (3%)	0 (0%)	0 (0%)	0 (0%)
0.75 – 0.90	16 (18%)	11 (17%)	5 (13%)	1 (5%)	0 (0%)	0 (0%)
0.5 – 0.75	6 (7%)	5 (8%)	1 (3%)	1 (5%)	0 (0%)	0 (0%)
0 – 0.5	4 (4%)	2 (3%)	2 (5%)	0 (0%)	0 (0%)	0 (0%)
Total	**91 (100%)**	**64 (100%)**	**38 (100%)**	**21 (100%)**	**12 (100%)**	**8 (100%)**

Notes: Table compares the patent counts of inventors in our disambiguation with those taken from an external Wikipedia data source. Each row reports the number of inventors with a particular number of patents in our disambiguation relative to the Wikipedia data. For instance, the row 1.05+ reports the number of inventors who have at least 5% more patents in our data than reported in the Wikipedia table. Each column segments the sample to inventors for whom we observe at least X% of their career in our data, which spans 1920-2010; for instance, column (1) selects on inventors of for whom we observe at least 50% of their career in our sample. The counts of inventors in each cell are reported in each row, with the percentage of inventors reported in parentheses beneath the counts.

It is important to note that our dataset includes only patents that were granted between 1920 and 2010, whereas most Wikipedia's inventors were still active in the decade 2010-2020. To account for this, we first drop the few inventors who were active before 1920 for the bulk of their careers, such as Thomas Edison (1847-1931). Then, we consider different baskets of inventors depending on what percentage of their career is covered in the timeframe of our dataset (1920-2010). For instance, if we enforce that we observe at least 50% of an inventor's career, then we would include an inventor whose period of activity was 2000-2020, but not if it was 2001-2020.

The results are presented in Table A.2. The total number of patents of the most prolific inventors are generally underestimated in our dataset. However, this is mainly due to the mismatch in the coverage period. When we condition on observing at least 70% of an inventor's career span, our deviation from Wikipedia's patent counts is no larger than 5% for 55% of our inventors. Enforcing that we observe the entirety of an inventor's career increases this share to 75%.

Of course, there is no guarantee that the crowd-sourced Wikipedia list is perfectly accurate. It may additionally include patents for inventors who live abroad, while our data only covers inventors while they live in the US. Nevertheless, the fact that our disambiguated data produces patent counts which are close to those in Wikipedia is heartening. This, coupled with the relatively low rates of Type I and II errors uncovered above, gives us confidence that our disambiguated data are of a high quality.

A.3 Assigning Inventors to States

Our patent data provides information on the residence address of the patent's inventors. However, we do not observe the residence of all inventors on a patent in the historical period. Specifically, we observe an inventor's state if either 1) they are the first inventor on the patent, or 2) the patent is contained in the Harvard Dataverse Network (HDN) data produced by Lai et al. (2014). In order to run our inventor-level regressions, we must assign each inventor to a particular home state. In this section, we detail our approach to doing so.

For all non-primary authors on historical patents, we impute a location using the following algorithm:

1. We assign all HDN and first author inventors to the state listed in the data

2. If an inventor is an HDN or first author inventor on one patent in a given year, but not on another patent, we assign that inventor to his first-author state. If he is first author in multiple states in that year, we assign him to the state listed on the patent if that state matches one of his first author states; otherwise we proceed to step 3 below (using alternative years)

3. We replace the inventor's state with the preceding year's state if state information is still missing.

4. We replace the inventor's state with the following year's state if state information is still missing.

5. If the inventor-patent record is still missing state information, but the inventor has multiple first-author states listed in that year, then we pick a random first-author state for that inventor-patent.

6. If all else fails, we assign the state of the first-author on the patent.

An additional challenge arises from the fact that a number of inventors have patents granted in multiple states in the same year. There may be many causes for multiple unique states within a given year for an inventor. The most common causes of these multi-state inventors are:

- An inventor may live in state A until midway through a particular year, and then move to state B. They file a patent application both in state A before moving and in state B after moving. They never file a patent in state B before moving, and never file a patent in state A after moving.

- Inventors may have multiple home addresses. As a result, they consistently file in both state A and state B in multiple years. For example, inventors may spend half of the year in Chicago, IL, and half of the year in Milwaukee, WI, and thus frequently have patents in both of these states in a given year.

- Inventors have multiple coauthors, who live in different states and who alternate in terms of who is the first listed author. For instance, Harvey Clayton Rentschler lives in Pittsburgh, PA, but frequently coauthors with J. Marden, who lives in Orange, NJ. Every time they coauthor a patent, the location is listed as Orange, NJ, but every time Harvey Rentschler sole authors a patent, his location appears to be Pittsburgh. These situations are particularly common among assigned patents, and seem to account for all individuals living in an exceptionally high number of states. Indeed, everyone who shows up in 7 or more states has a coauthor on their patents, while the share of those with a coauthor is 92.8% for those with multiple states, compared with just 66.3% for those in one state[10]

- Possible disambiguation errors: two individuals may have very similar names, work in similar classes, and live just across a state border from one another (so are close in latitude-longitude). As a result these two separate inventors may be classified as the same person by the disambiguator. This would inflate the number of states an individual lives in.

To address this concern, we assign multi-state inventors a home state using the following algorithm:

1. Each year, assign an inventor to the modal state in which we observe him/her operating as a sole author.

2. If the inventor does not have any sole authorships in that particular year, check if they have sole authorships in the preceding or subsequent year. If the preceding and subsequent year both have sole authorships in the same modal location, then assign the inventor to that state. This smoothes over off years for inventors and removes spurious migration.

3. If we still do not have a location for the inventor, then we assign them to the modal location we observe them in in the given year, regardless of whether the patent was sole authored or coauthored.

4. If the inventor has two modal states (e.g. has 2 patents in both Illinois and Wisconsin in the given year), then choose a random choice of those states and assign the inventor to that state.

A.4 Tenure versus quality

One might be concerned that our quality measures simply reflect career length. Those with longer careers are likely to have accumulated more patents, regardless of their underlying productivity. For our purposes, it suffices that those with more patents are more likely to be in the top ten percent of the income distribution, so a correlation between career length and quality measures would not invalidate our empirical approach. In addition, we control for a quadratic in inventor tenure, as well as an indicator for being high quality, in our regressions. Nevertheless, it is useful to explore the relationship between inventor quality and tenure in the dataset.

Figure A.2 plots the relationship between inventor quality and tenure. The horizontal axis measures an inventor's tenure in the dataset, defined as the number of years since the inventor first applied for a patent. The vertical axis plots the share of inventors with that tenure level who are counted as high type under our baseline measure of quality; that is, the share of inventors who are in the top ten percent of the distribution of patents accumulated in ones career up through that

[10]This is partially mechanical as these inventors are also more productive so have more chances to appear in multiple states.

Figure A.2: The Relationship Between Inventor Quality and Tenure

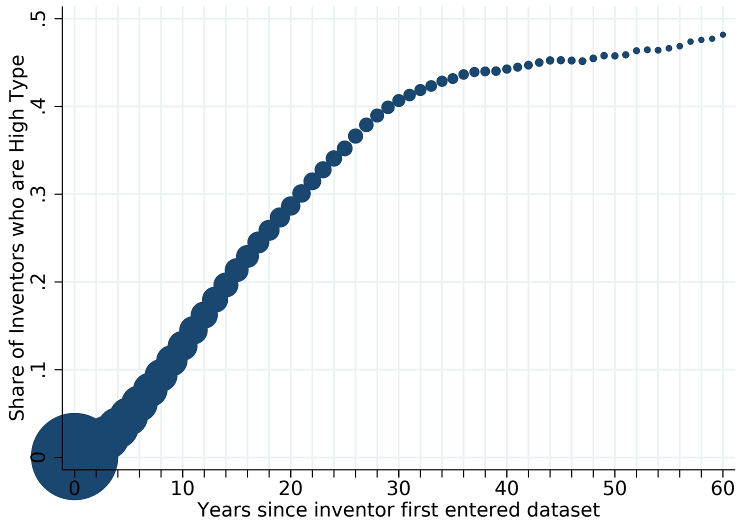

Notes: This figure plots the share of inventors who are counted as high quality in a particular year against their tenure in the dataset. Inventors are defined as high quality if they are in the top 10% of career patent counts up through the year in question. The size of each marker reflects the share of inventors in that tenure bin.

year. The size of the markers is proportional to the share of observations in the dataset who have a given tenure level.

Unsurprisingly, the figure shows a strong positive correlation between quality and tenure. However, it is not the case that high quality inventors are only those with high tenure. Approximately ten percent of inventors with 10 years of tenure are high quality. This number rises to 28% when considering 20-year inventors. Even at very high tenure levels (of which there are very few), no more than 45% of inventors are counted as high quality. Thus the relationship between tenure and quality, while significant, is far from one-for-one.

A.5 Citation Adjustment

Our data includes the full network of citations from patents granted from September 1947, when the USPTO began to note citation data in a systematic way, to 2015. Citations start in 1947 because a USPTO Notice was issued on December 19th, 1946, instructing examiners to add citations in the published format of the patent, a practice that was incorporated into the *Manual of Patenting Examining Procedure* (paragraph 1302.12).

For patents granted before 1947, the noted citation count is left censored: a patent granted in 1940 will only have citations from patents granted after 1947, but will not have citations from patents between 1941 and 1946. This artificially deflates the number of citations received by patents before 1947, confounding attempts to use citations as an objective measure of a patent's quality. Furthermore, aggregate citation trends may weaken the link between raw citation counts and patent quality. For instance, if patents granted in 1960 cite an average of 5 prior patents, but those granted in 1990 cite 20 patents, one might expect the average citation received from a 1960 patent to be more indicative of a high quality innovation than a citation received in 1990. We therefore adjust the number of citations received by each patent following the quasi-structural approach laid out in Hall, Jaffe and Trajtenberg (2001).

This approach relies on two critical assumptions. First, we assume that the citation process is *stationary*. That is, we assume that the evolution of citation shares does not change over time: a patent will on average receive a share $\pi_{k\tau}$ of its citations τ years after it is granted, regardless of the grant year. This allows us to project back our adjustment factors to patents filed before the citation data began in 1947. Second, we assume *proportionality*. That is, we assume that the shape of the citation evolution does not depend on the total number of citations received so that highly cited patents are more highly cited at all lags. This allows the application of the same adjustment factor to every patent in our data granted in a given period and belonging to a given patent class.

The adjustment proceeds as follows. We start with the full patent citation network data, keeping only those patents granted in the United States. Let C_{kst} be the total number of citations to patents in year s and technology category k coming from patents in year t.[11] Further, define P_{ks} to be the total number of citations received by patents granted in year s in technological category k. One can then define $\pi_{kst} = C_{kst}/P_{ks}$ to be the average share of citations received by patents in class k in year s from patents granted in year t We assume that π_{kst} is some multiplicatively separable function of grant year, patent category, and a citation lag. That is, we can write

$$\log[\pi_{kst}] = \alpha_0 + \alpha_s + \alpha_t + \alpha_k + f_k(L) \tag{3}$$

for $L = t - s$ the lag between cited and citing patent grant years, and $f_k(\cdot)$ some category-specific function of these lags. For our purposes, we define $f_k(L) = \tilde{\gamma}_{k,L}$. We may then estimate equation 3 using OLS to recover estimates of $\alpha_0, \alpha_s, \alpha_t, \alpha_k$, and $\tilde{\gamma}_{k,L}$ for each value of s, t, k and L in our data.[12] Taking exponentials of equation 3 yields

$$C_{kst}/P_{ks} = e^{\alpha_0} e^{\alpha_s} e^{\alpha_t} e^{\alpha_k} e^{\tilde{\gamma}_{k,(t-s)}} \tag{4}$$

This formulation allows us to standardize citation counts over time and across categories. Specifically, in order to adjust for patent class, cited year, and citing year effects, we weight each citation from a patent in year t to a patent in class k in year s by $\exp(-\hat{\alpha}_k - \hat{\alpha}_s - \hat{\alpha}_t)$. Each patent's citation counts are therefore reflective of the patent's quality relative to the average patent in some base year and category.[13]

While this procedure accounts for aggregate differences across patent classes and grant years, it does not yet correct for bias arising from the left truncation of citation records. To build intuition for the truncation correction, consider an example in which each of the estimated α coefficients were 0: the only bias in our citation data arises from the lag. In that case, the assumptions of proportionality and stationarity suggest a natural adjustment factor for a patent granted L years before the 1947 cutoff. Define $G_k(L)$ to be the CDF of the lag distribution: the share of an average patent's citations received within the first L years after its grant. The adjustment factor is then given by

$$\sigma_{k,L} = \frac{1}{1 - G_k(L)}$$

[11] For the purposes of the adjustment, we use technological categories as defined by the NBER patent data. For a detailed description of these data, see Hall, Jaffe and Trajtenberg (2001).

[12] It is rare for a patent to receive citations more than 30 years after its initial grant date, and thus we top-code the citation lag L to have a maximum value of 30. That is, we define $L = \min\{t - s, 30\}$.

[13] For our purposes, we choose each patent citation to be relative to a patent in the "Other" category granted in 1975, receiving citations from patents also granted in 1975. Mechanically, this corresponds to setting the omitted categories in estimation of equation 3 to be $k =$ "Other", $s = t = 1975$.

We would then predict that a patent in category k granted in year $1947-L$ and receiving c citations from patents granted after 1947 would have received $\sigma_{k,L}c$ citations had the USPTO kept track of citations before 1947.[14]

In order to incorporate the year and category fixed effects into this truncation adjustment framework, one must establish a notion of the CDF of the lag distribution conditional on year and category effects. To do so, we interpret the $\exp(\tilde{\gamma}_{k,L})$'s as weights for each patent in the citation data. For instance, if the estimated $\exp(\tilde{\gamma}_{k,L=2})$ is 2, then an average patent is twice as likely to receive a citation after 1 year than in the year of patent grant, conditional on year and category effects. To construct the CDF of citations by lag conditional on year and class effects, we can sum our estimates of $\exp(\tilde{\gamma}_{k,L})$, normalizing the estimated coefficients so that they sum to 1. This gives us our estimate of $G_k(L)$:

$$\hat{G}_k(L) = \frac{\sum_{l=1}^{L} \exp(\tilde{\gamma}_{k,l})}{\sum_{l=1}^{30} \exp(\tilde{\gamma}_{k,l})} \tag{5}$$

We can then calculate our truncation adjustment factor as before[15]

$$\hat{\sigma}_{k,L} = \frac{1}{1-\hat{G}_k(L)}. \tag{6}$$

To summarize, the citation adjustment proceeds in four steps:

1. Estimate equation (3) using OLS to recover $\alpha_0, \alpha_k, \alpha_t, \alpha_s$ and $\gamma_{k,L}$.

2. For each citation made from a patent p' granted in year t to a patent p in class k granted in year s is weighted by

$$\omega_{k,s,t} = e^{-\alpha_k-\alpha_t-\alpha_s}$$

   Define, for each cited patent p, the year- and category-adjusted citation count c to be the sum of the $\omega_{k,s,t}$ it received.

3. Calculate $\hat{G}_k(L)$ according to equation (5)

4. Using $\hat{G}_k(L)$, calculate the truncation adjustment factor $\hat{\sigma}_{k,L}$ according to (6). Finally, define a patent p's adjusted citation count to be $\tilde{c} = c \cdot \sigma_{k,L}$ if p is in class k and was granted L years before 1947.

A.6 Historical Corporate Tax Data

We collected the corporate tax rates from a large variety of sources. We also collect surtaxes or surcharges, as well as additional temporary taxes imposed on top of the main rates. They are sometimes imposed as a percentage of regular tax liabilities and sometimes as a rate to add to the main rate. We record them as rates to add to the main rate with applicable thresholds. We have not collected minimum taxes (they are very low and probably not applicable to the companies in our sample) and alternative minimum taxes.

[14] Ignoring year and category effects and adjusting citations in this way does not significantly change the results presented in the main body of the paper.

[15] Note that we only calculate the truncation adjustment up to $L = 20$, despite estimating $\gamma_{k,L}$ for L as large as 30. This is to bound $\hat{G}_k(L)$ away from 1, so that we do not divide by 0 in the adjustment. For L larger than 20, we apply the adjustment factor for $L = 20$.

A.6.1 Tax Rates

Historical data on the state corporate tax rates were collected from a number of sources. In particular, our tax rate data come from:

- HeinOnline Session Laws. This is an archive of state legislation enacted since U.S. territories were established and provides great historical coverage.
- ProQuest Congressional. It supplements to HeinOnline's database for the District of Columbia.
- HeinOnline State Statutes. This is an archive of historical state statutes. In the early 20th century, state statutes were periodically recodified into clean statutes. Between these recodifications, publishers released annotations that recorded updates. HeinOnline contains statutes and annotations, mostly prior to 1940.
- State Tax Review and State Tax Handbooks by the Commerce Clearinghouse (CCH). The Commerce Clearing House is a company that publishes tax guides as a resource for businesses and tax professionals. We located their products through University of Michigan and the Library of Congress.
- Council of State Governments Book of States. These are a biannual publication aiming to review state-level tax changes. The first usable volume was published in 1948. In particular, the state finances chapter summarizes income tax developments and includes a chart for corporate income tax rates.
- Tax Foundation Publications.
- Report of the Subcommittee on State Taxation of Interstate Commerce (1964), also referred as *"The Willis Commission Report "*.
- State Income Tax Administration by Penniman and Heller. It was published in 1959 and provides a good overview of the development of corporate income taxes in US states.
- National Tax Association Proceedings.
- The Progress of State Income Taxation Since 1911, Lutz (1920).

A.6.2 Tax Base Rules

Following Suárez Serrato and Zidar (2018), we use fifteen variables to control for changes in how the state corporate tax base is computed. One may contact the authors and cite this paper if one wished to utilize these raw corporate tax base data. Specifically, the variables collected are:

- *Investment Tax Credit:* investment tax credit rate for a given state-year;
- *R&D Tax Credit:* statutory credit rate adjusted for recapture and type of credit for a given state-year;
- *R&D Tax Credit Base - Incremental Moving Average:* an indicator variable equal to one if the R&D tax credit applies to an incremental base that is a moving average of past expenditures in a given state-year;

- *R&D Tax Credit Base - Incremental Fixed Period* indicator variable equal to one if the R&D tax credit applies to an incremental base that is fixed on a level of past expenditures in a given state-year;
- *Loss carryback rules:* the number of years a corporation may carry back any excess loss following the loss year;
- *Loss carryforward rules:* the number of years a corporation may carry forward any excess loss following the loss year;
- *Franchise tax:* indicator variable equal to one if a Franchise tax is levied on corporations in a given state-year;
- *Throwback rules:* indicator variable equal to one if the state eliminates "nowhere income" that would be untaxed by either the state with the corporation's nexus or the state in which the relevant sales were being made;
- *Combined Reporting Rules:* an indicator variable equal to one if a state requires a unitary business to submit combined reporting;
- *Federal Income Tax Deductible:* indicator variable equal to one if the federal income tax is deductible in a given state-year;
- *Federal Income as State Tax Base:* indicator variable equal to one if the federal income tax base is used as the state tax base in a given state-year;
- *Federal Accelerated Depreciation:* variable equal to one if the federal accelerated depreciation is allowed in a given state-year;
- *Accelerated Cost Recovery System (ACRS) Depreciation:* indicator variable equal to one if the accelerated cost recovery system is allowed in a given state-year;
- *Federal Bonus Depreciation:* indicator variable equal to one if the federal bonus depreciation is allowed in a given state-year;
- *Sales Apportionment Weight:* the share of national profits of multi-state firms that are allocated to sales (for tax purposes) in a given state.

The loss carryback rule, the loss carry forward rule, the federal income tax deducibility, the federal income as state tax base, the federal accelerated depreciation, the federal bonus depreciation from 1958 through 1980 are collected from the State Tax Handbooks published by the Commerce Clearing House (CCH). From 1980 through 2010, these same variables are supplemented by Suárez Serrato and Zidar (2018). The investment tax credit data from 1964 onward are supplemented by Chirinko and Wilson (2008). The Federal R&D Tax Credit was introduced for the first time in 1981; one year later, in 1982, Minnesota was the first state to follow suit and enact a state R&D tax credit. Hence, data for state-level R&D Tax Credits (both the rate and whether the tax credit base is incremental with a moving average or incremental with fixed base) from 1982 onward come from Suárez Serrato and Zidar (2018).

The apportionment weights are compiled from a variety of sources. These sources include:

- The Progress of State Income Taxation Since 1911, Lutz (1920).
- Report of the Subcommittee on State Taxation of Interstate Commerce (1964), also referred as *"The Willis Commission Report "*.

- Hearings Before the Committee on the Judiciary, United States Senate, Ninety-fifth Congress, First and Second Sessions on Interstate Taxation (1977-1978).

- Report of the Committee of independent experts on company taxation, also referred as the Ruding Report.

- "Supplement to the Appendix to the Journal of the Senate Legislature of the State of California" (1951).

- CSG Book of the States.

- CCH State Tax Reviews and CCH State Tax Handbooks.

Again, data from 1980 onward are supplemented from Suárez Serrato and Zidar (2018). Since we lack exact data for each state-year, especially for the earliest period, we have to interpolate some years of the data. In particular, we assume that an apportionment rule is in place until we have explicit confirmation that the rule has changed.

Before the Uniform Division of Income for Tax Purposes Act (UDITPA) in 1957, different states had different ways of dealing with the taxation of multi-state companies. Although not all states adopted it, the UDITPA made these apportionment and allocation rules of the business income of multi-state companies more uniform, with a three-factor formula based on equal weights to the shares of a corporation's payroll, property, and sales in the state. In the past twenty years, the weight on sales has started to increase, which should arguably decrease the importance for a company of corporate income tax in states in which it has property and employment (but a low share of its sales).

While most states apportion based solely on the location of payroll, property, or sales, some states occasionally used concepts such as production costs. When states offer different apportionment rules to manufacturing versus service sectors, we use the manufacturing rules since most innovation occurred within the manufacturing sector in the latter half of the 20th century.

We additionally collect information on franchise taxes. In several states, statutes make direct corporate taxes unconstitutional and franchise taxes get around this problem. Some states have one or the other, sometimes both, but companies only pay one. Types of franchise taxes include taxes on net income (which are extremely similar to corporate income taxes and which we consider as such), Business enterprise tax (in New Hampshire), Gross receipts tax or commercial activity tax (which is the gross receipts tax in Ohio), Business and occupation tax (West Virginia, Washington, or Ohio, sometimes different for different industries), net worth/capital stock/asset value/shareholder equity combination taxes, or a value-added tax (Michigan's single business tax which is a franchise tax, not a sales tax).

To construct our index of "corporate tax base breadth," we follow Suárez Serrato and Zidar (2018) in constructing an index of the "corporate tax base breadth." We regress state corporate tax revenues as a share of GDP on all tax base and apportionment variables, as well as state and year fixed effects. That is, we estimate regressions of the form

$$\left(\frac{Revenue}{GDP}\right)_{st} = \alpha_s + \phi_t + \gamma \text{Corp. MTR}_{st} + \mathbf{X}'_{st}\mathbf{\Psi}^{BASE}_{st} + u_{st}$$

where $\left(\frac{Revenue}{GDP}\right)_{st}$ is the ratio of state corporate tax revenues to state GDP, α_s and γ_t, Corp. MTR_{st} is the state's top corporate tax rate, u_{st} is an error term and $\mathbf{X}_{st}$ is a vector of corporate tax base rules. Specifically, $\mathbf{X}$ contains R&D tax credits, the sales apportionment weight, years of allowable loss carryback and carryforward, and indicators for whether R&D tax credits are applied on an incremental base that is a moving average or fixed on a level of past expenditures, whether the state tax is a franchise tax, whether federal income taxes are deductible from income, whether the federal income tax base is used for state taxation, whether the state allows federal accelerated depreciation, whether the state has throwback rules, whether the state uses combined reporting rules, and state investment tax credits. We have state revenue data from 1980-2010, and we weight the regressions by each state's average GDP over the sample.

The index is the predicted value from this regression (excluding state and year fixed effects), standardized to have zero mean and unit standard deviation. Formally, we begin by defining a raw base index given by

$$\tilde{b}_{st} = \gamma \text{Corp. MTR}_{st} + \mathbf{X}'_{st} \mathbf{\Psi}^{BASE}_{st}$$

Then we standardize this raw index to have zero mean and unit standard deviation over the full sample. Define $\sigma_{\tilde{b}}$ to be the standard deviation of $\tilde{b}$ over the period 1958-2000 when we observe the state base rules. Letting S be the number of states in our sample, and T the number of years between 1958 and 2000, our final base index is defined as

$$b_{st} = \frac{\frac{1}{ST} \sum\limits_{s,t} \tilde{b}_{st}}{\sigma_{\tilde{b}}}$$

This index may be interpreted as the number of standard deviations higher revenue a state might expect to receive from its tax base rules, relative to an average state-year. Since the state documents containing tax base information are missing for 1979, we linearly interpolate our tax base index over this year.

B Calculating Multinomial Logit Elasticities

The elasticity of an inventor i residing in state s to the personal tax rate may be expressed as[16]

$$\eta^i_{st} = \frac{d \ln P^i_{st}}{d \ln(1 - ATR^{pj[i]}_{st})} = \alpha \cdot (1 - P^i_{st})$$

We may then define the elasticity of location choices to tax rates by taking the weighted average of these η^i_{st}. We may do this separately inventors' home state and non-home states. Letting I^d_s and I^o_s denote home and non-home inventors in state s, the elasticity of locating in state s to personal taxes may be expressed as

$$\eta^{s,d}_p \equiv \frac{d \log \sum_{i \in I^d_s} P^i_{st}}{d \ln(1 - ATR_{st})} = \frac{\alpha \sum_{i \in I^d} P^i_{st}(1 - P^i_{st})}{\sum_{i \in I^d_s} P^i_{st}}. \tag{7}$$

Likewise, for inventors in non-home states,

$$\eta^{s,d}_p \equiv \frac{d \log \sum_{i \in I^o_s} P^i_{st}}{d \ln(1 - ATR_{st})} = \frac{\alpha \sum_{i \in I^o} P^i_{st}(1 - P^i_{st})}{\sum_{i \in I^d_s = o} P^i_{st}}. \tag{8}$$

Average home and non-home elasticities are then defined as the weighted average of these elasticities across all countries:

$$\eta^d_p \equiv \sum_s \left(\frac{\sum_{i \in I^d_t} P^i_{st}}{\sum_s \sum_{i \in I^d_t} P^i_{st}} \right) \eta^{s,d}_p, \tag{9}$$

and

$$\eta^o_p \equiv \sum_s \left(\frac{\sum_{i \in I^o_t} P^i_{st}}{\sum_s \sum_{i \in I^o_t} P^i_{st}} \right) \eta^{s,o}_p, \tag{10}$$

where I^d_t and I^o_t represent the set of inventors locating in their home state or out-of-state, as in Section 2. Finally, we may aggregate these elasticities to an overall mobility elasticity as

$$\eta_p \equiv \gamma^d \eta^d + (1 - \gamma^d)\eta^o. \tag{11}$$

[16]Elasticities to the corporate tax may be defined analogously.

C Additional Tables and Figures

C.1 Summary Statistics

Table C.1: Disambiguation Output: Unique Inventor Counts

Sample	# Inventors	# Patents
1920-2004, US only	2,953,471	5,336,672
1940-2000, US only	1,744,224	2,775,100
Lai et al. Patents, our disambiguation	1,374,891	2,179,599
Lai et al. Disambig (US)	1,462,626	2,179,599

Notes: Table shows performance of the Lai et al. disambiguation algorithm as applied to our historical patent data. Each row contains performance information for a different subsample. The category "Lai et al. Patents, our disambiguation" reports the performance of our algorithm on the patent records included in the original Lai et al. sample. Likewise, "Lai et al. Disambig (US)" reports the number of unique inventors that Lai et al. find when applying their algorithm to U.S. patents. The first column shows the number of unique inventors found by the disambiguation algorithm, while the second shows the unique number of patents in each subsample.

Table C.2: Additional Summary Statistics

	Mean (1)	S.D. (2)	1940-59 (3)	1960-79 (4)	1980-99 (5)
Inventor-Level Data: Controls/Other Outcomes					
Pr{Has corporate patent in 3 years}	0.564	0.496	0.437	0.566	0.619
Pr{Has non-corporate patent in 3 years}	0.165	0.371	0.244	0.154	0.136
Agglomeration forces	0.068	0.215	0.017	0.029	0.114
Tenure	9.726	13.245	9.380	9.942	9.745
Pr{Assignee has patents in multiple states}	0.790	0.407	0.797	0.794	0.784
State-Level Data: Other Outcomes					
Log Patents	5.96	1.51	5.54	5.98	6.35
Log Inventors	5.97	1.53	5.44	6.00	6.49
Log Citations	8.52	1.77	7.69	8.30	9.56
Mean Kogan et al. (2017) Patent Value	23.71	16.71	23.94	26.99	20.25
Real Manufacturing Value Added ($ billions)	5.08	6.40	2.67	5.56	7.02
Real Manufacturing Total Payrolls ($ billions)	2.30	2.93	1.43	2.73	2.74
Real Average Weekly Earnings	102.47	22.61	80.59	113.34	113.47
Employees per Establishment	51.10	20.21	50.21	53.97	49.13
Real Personal Income Per Capita ($ 0000s)	3.31	1.27	2.04	3.27	4.62
Share of Workforce in Manufacturing	0.22	0.12	0.29	0.23	0.15
State-Level Data: Corporate Tax Rules					
Share with Sales Apportionment Weight	0.95	0.22	0.83	0.97	1.00
Sales Apportionment Weight	42.27	23.50	39.44	37.90	47.33
Property Apportionment Weight	30.13	15.54	35.14	31.34	26.32
Payroll Apportionment Weight	25.62	13.39	19.95	28.69	26.32
Years Losses Allowed to be Carried Back	1.19	1.45	0.10	1.16	1.29
Years Losses Allowed to be Carried Forward	6.69	5.87	1.41	3.19	9.76

Notes: Table reports additional summary statistics for our estimation sample. This includes all mainland US states, excluding Louisiana, from 1940-2000. Columns (1) and (2) report the mean and standard deviation, respectively, for the full sample period, while columns (3)-(5) report the averages in each 20-year period from 1940 to 2000. "Tenure" corresponds to the number of years the inventor has been in the sample since her first patent. Inventors are included between the years of their first successful patent application and their last successful patent application. Summary statistics for all inventor controls are averaged over inventor-year observations to reflect summary statistics of our estimation sample; thus inventors with long careers will appear more than once in the average. Kogan et al. (2017) Patent Value is expressed in millions of 1982 dollars. Corporate tax base rules are defined in detail in Appendix A.6, and are defined conditional on having non-zero state corporate tax rates.

Table C.3: Summary Statistics on Inventor Careers

	Mean	Median	SD	90^{th}	95^{th}	99^{th}
Years Active	3.33	1.00	6.24	8.00	15.00	33.00
Number of States	1.07	1.00	0.38	1.00	2.00	3.00
Number of Patents	2.60	1.00	5.79	5.00	9.00	25.00
Patents Per Year	1.02	1.00	0.53	1.00	2.00	3.00
Total Citations Received	65.34	12.71	409.31	110.73	229.11	890.12
Citations Per Year	21.88	8.26	70.62	41.69	74.23	244.71
Number of Classes	1.62	1.00	1.75	3.00	4.00	9.00

Notes: Table reports summary statistics of our sample of disambiguated inventors. The categories "Number of States," "Number of Patents," "Total Citations Received," and "Number of Classes" refer to statistics over an inventor's entire career, while "Patents Per Year" and "Citations Per Year" refer to average numbers per year of an inventor's career. Each column reports a different moment of the distribution for the variable considered in the row. For instance, the 90^{th} percentile of the distribution of total career length ("Years Active") is 11 years, the 95^{th} percentile of that distribution is 19 years and the mean is 4.14 years. Inventors are considered active between the first year in which they have a successful patent application and the last year in which they have a successful patent application. All statistics are inventor-weighted, in contrast to Tables I and C.2 which present statistics averaged over inventor $\times$ year observations. The table considers inventors active in our primary estimation sample: U.S. inventors between 1940 and 2000. The table therefore represents 1,744,224 inventors.

Figure C.1: Patents per 10,000 State Residents Over Time

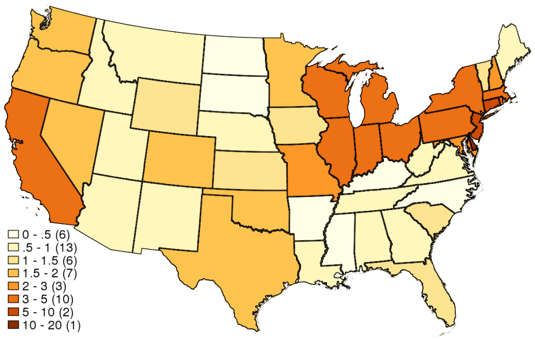

Panel A: 1940

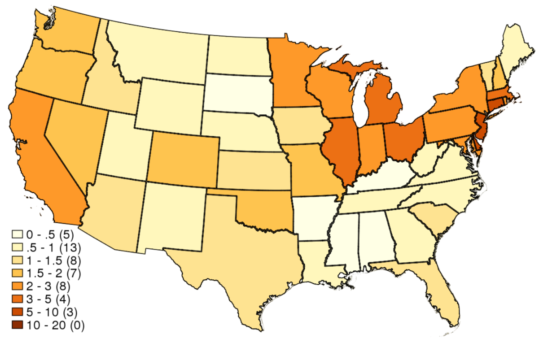

Panel B: 1950

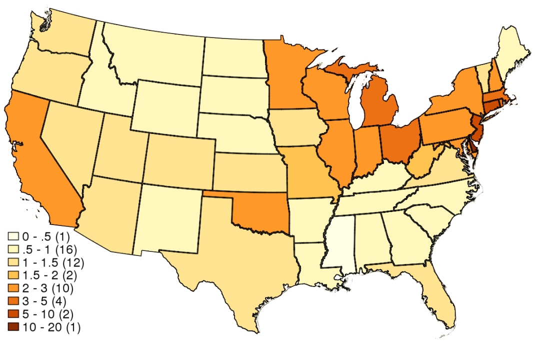

Panel C: 1960

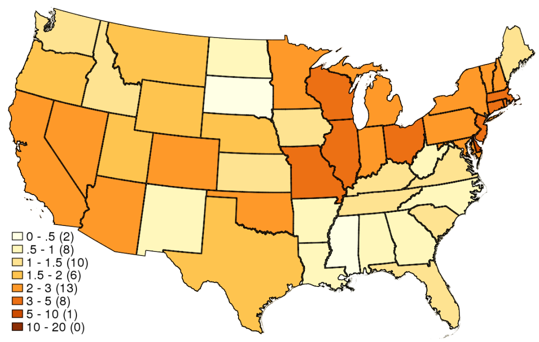

Panel D: 1970

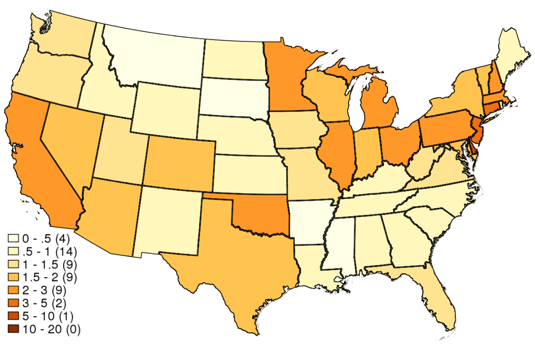

Panel E: 1980

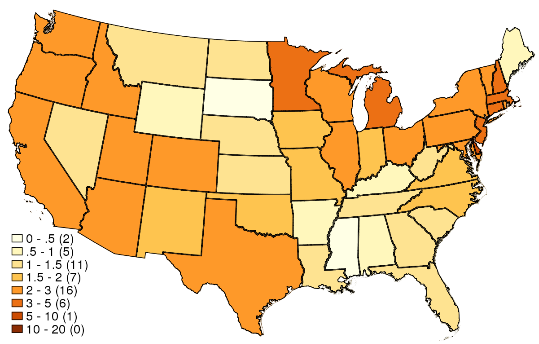

Panel F: 1990

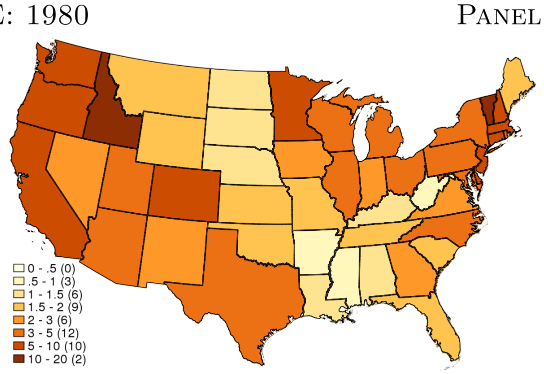

Panel G: 2000

Notes: Figure plots the number of patents per 10,000 state residences in each state every ten years from 1940 through 2000. Darker colors indicate more innovation.

Figure C.2: Inventors per 10,000 State Residents Over Time

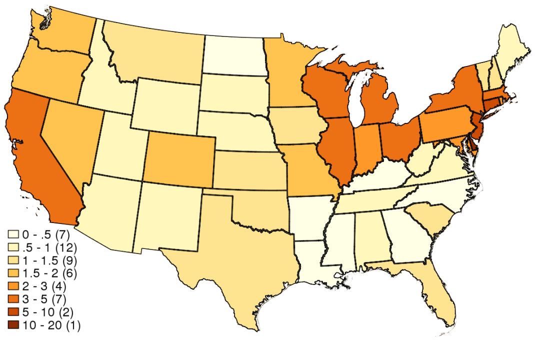

Panel A: 1940

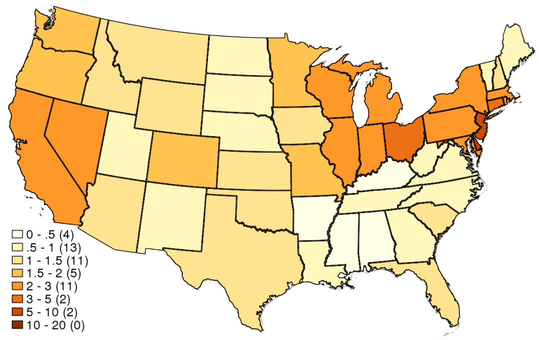

Panel B: 1950

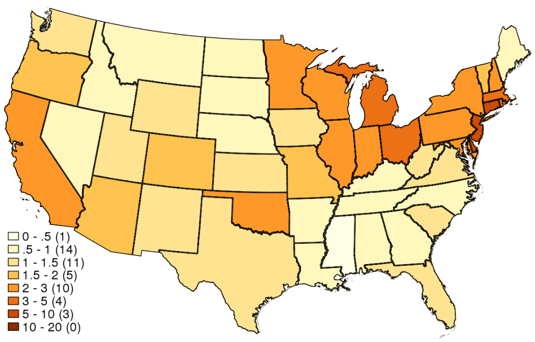

Panel C: 1960

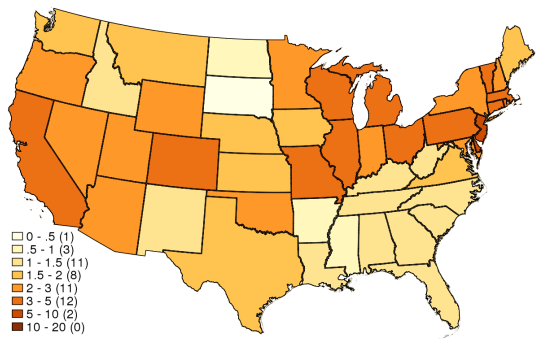

Panel D: 1970

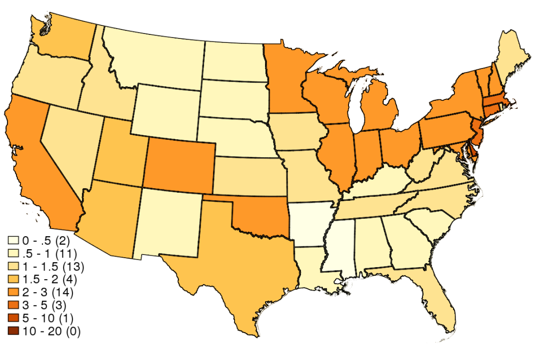

Panel E: 1980

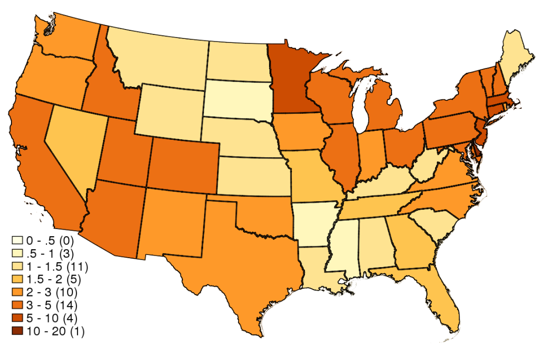

Panel F: 1990

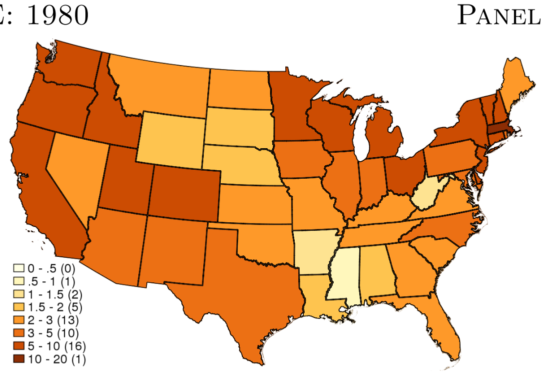

Panel G: 2000

Notes: Figure plots the number of unique inventors per 10,000 state residences in each state every ten years from 1940 through 2000. Darker colors indicate more innovation.

Figure C.3: Share of Corporate Patents and Corporate Inventors

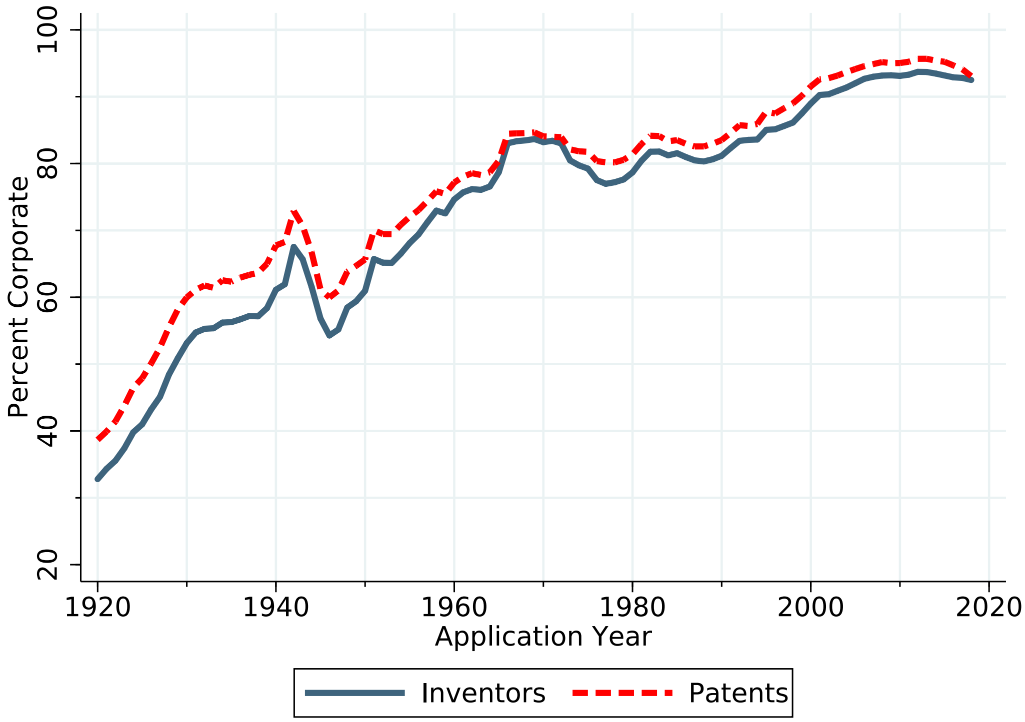

Notes: The graph shows the share of patents assigned to corporations (dashed line) and the share of inventors who patent for corporations (solid line).

C.2 Summary of Tax Variation

Figure C.4: Introduction Year of State Personal Income Taxes

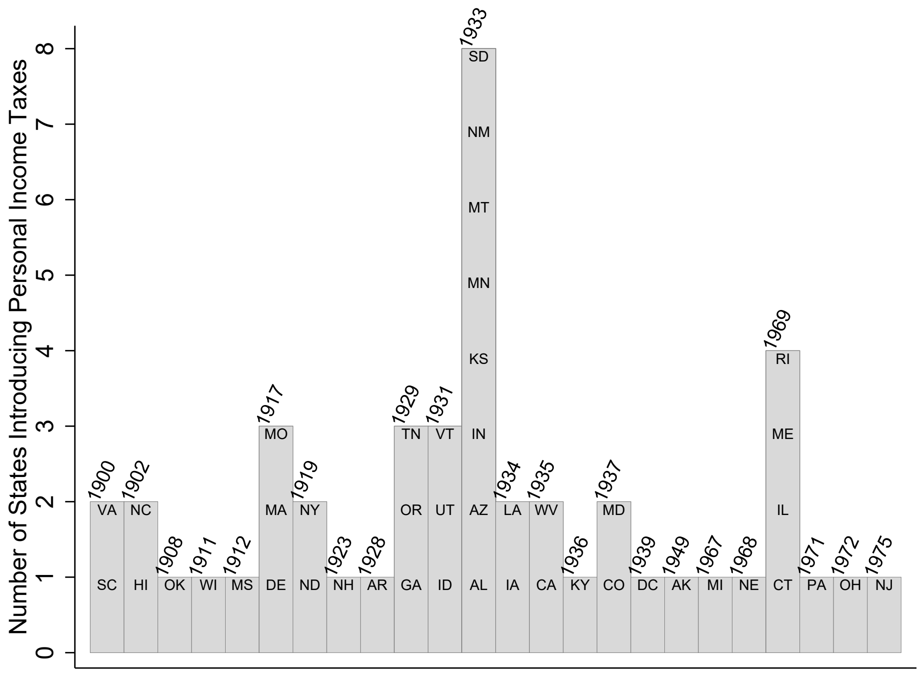

Notes: Figure plots the first year in which each state has a statutory personal income tax rate.

Figure C.5: The Evolution of Personal Income Taxes

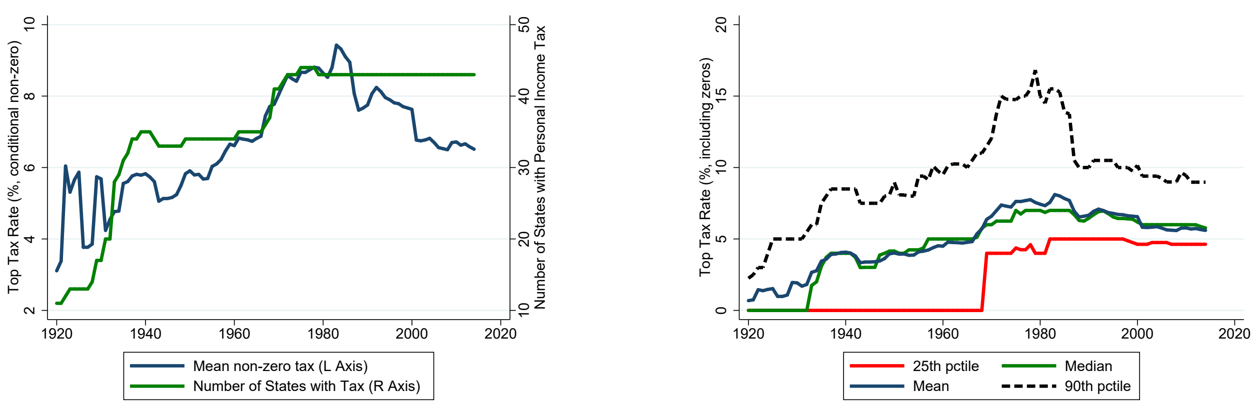

Panel A: Intensive and Extensive Margin Panel B: Distribution of Statutory Tax Rates

Notes: Figure plots the share of states with a personal income tax, as well as the distribution of those taxes over time.

Figure C.6: State Personal Marginal Tax Rates at the Median Income

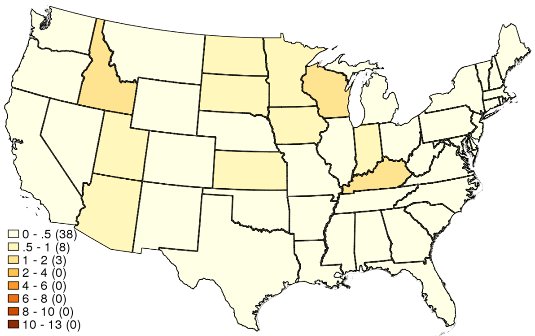

Panel A: 1940

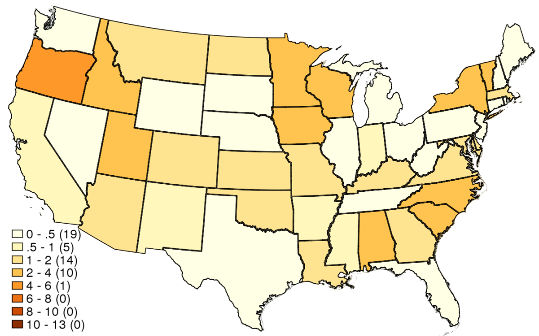

Panel B: 1950

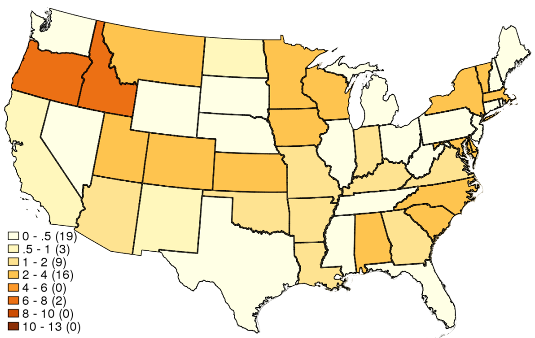

Panel C: 1960

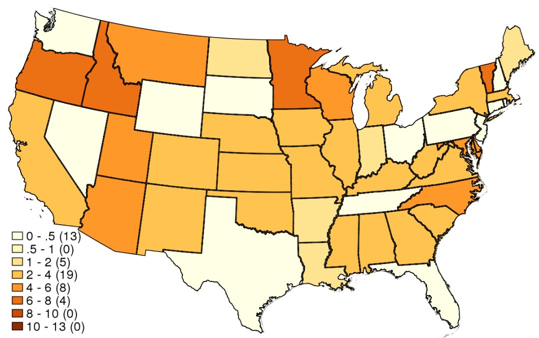

Panel D: 1970

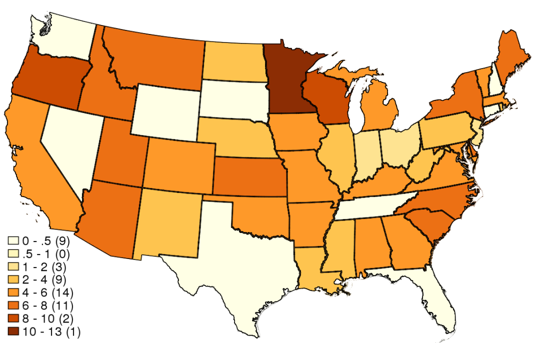

Panel E: 1980

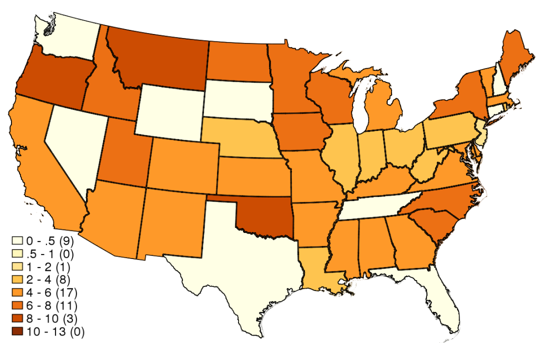

Panel F: 1990

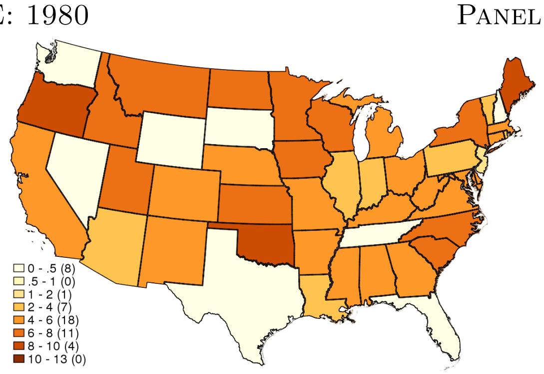

Panel G: 2000

Notes: Figure plots the state statutory marginal personal income tax rates faced by individuals whose adjusted gross income is equal to the median of the nation income distribution for men every ten years from 1940 through 2000. Darker colors indicate higher tax rates. Tax rates measured in percentage points.

Figure C.7: State Personal Marginal Tax Rates at 90^{th} Income Percentile

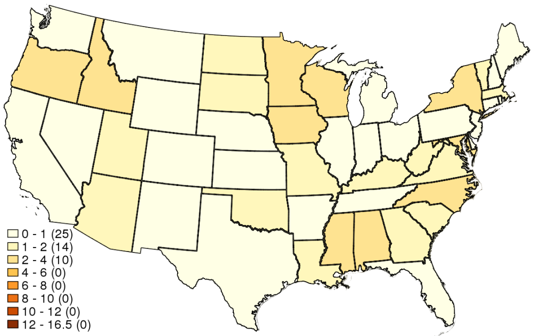

Panel A: 1940

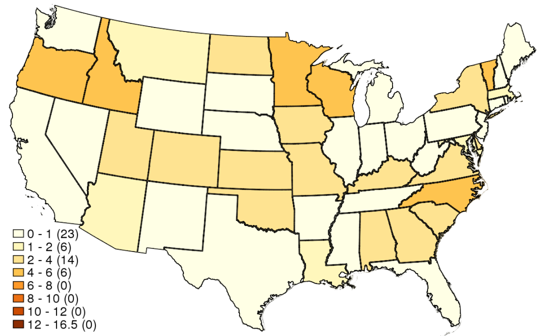

Panel B: 1950

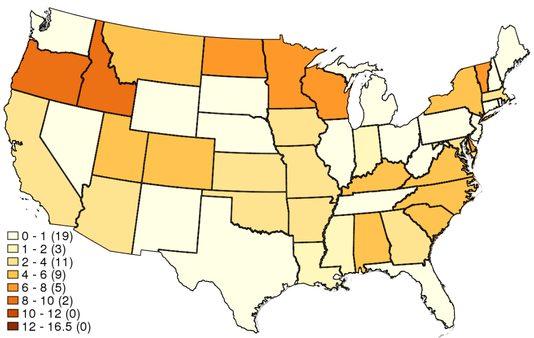

Panel C: 1960

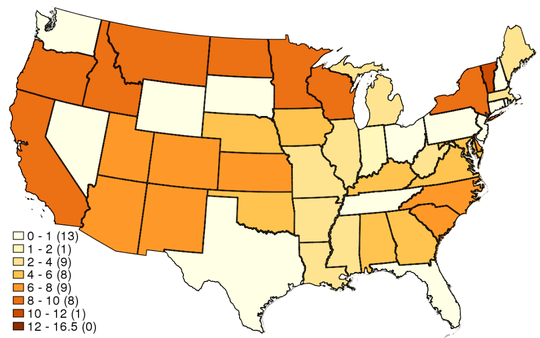

Panel D: 1970

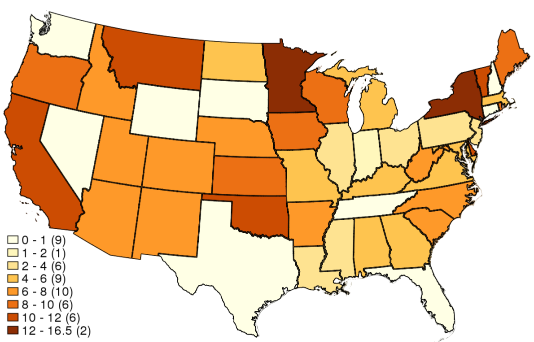

Panel E: 1980

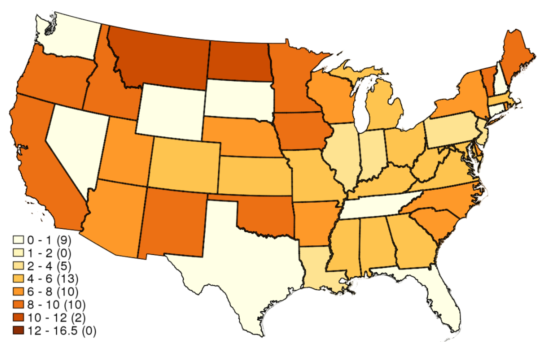

Panel F: 1990

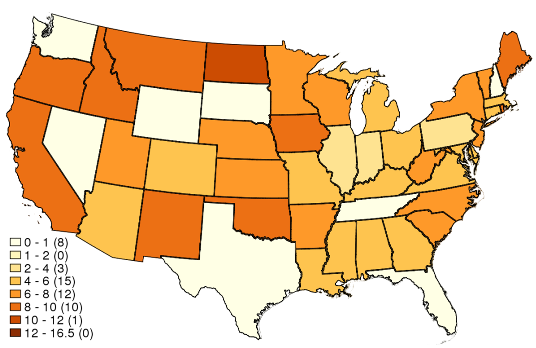

Panel G: 2000

Notes: Figure plots the state statutory marginal personal income tax rates faced by individuals whose adjusted gross income is equal to the 90^{th} percentile of the nation income distribution for men every ten years from 1940 through 2000. Darker colors indicate higher tax rates. Tax rates measured in percentage points.

Figure C.8: The Evolution of Personal Income Taxes in Select States

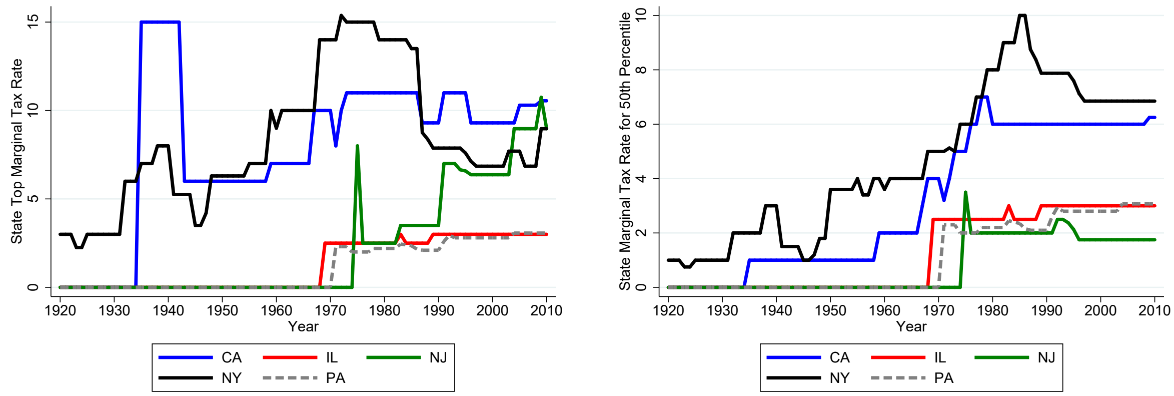

Panel A: Time Series of Key States' Top Statutory MTR

Panel B: Time Series of Key States' MTR for Median Earner

Notes: Figure plots the time series of marginal personal income tax rates for the five most innovative states in our sample. Tax rates are measured in percentage points. Panel A shows the top statutory personal income tax rate, while Panel B plots it for the median earner.

Figure C.9: Trends in State Statutory Tax Policy Changes

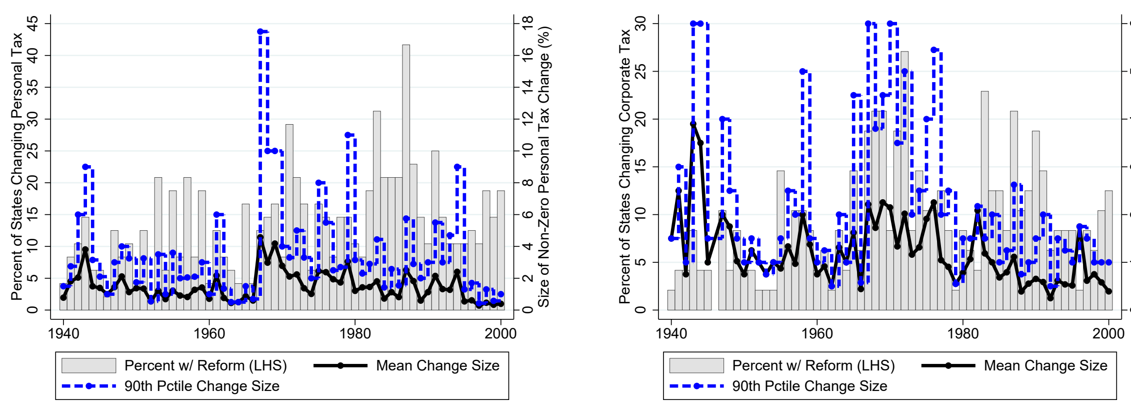

Panel A: Personal Income Tax

Panel B: Corporate Income Tax

Notes: Figure plots the time series of the share of states experiencing a statutory personal income (panel A) and corporate income (panel B) tax rate change. The gray bars, plotted against the left axis, show the share of all states that experience a statutory top tax rate change. The black solid line plots the mean size (positive or negative) of non-zero tax changes, while the blue dashed line represents the size of a 90^{th} percentile non-zero tax rate change in that year. Tax rate changes are measured in percentage points. Black and blue lines are plotted against the right axis.

Figure C.10: Size of marginal tax rate changes for median and 90th percentile workers

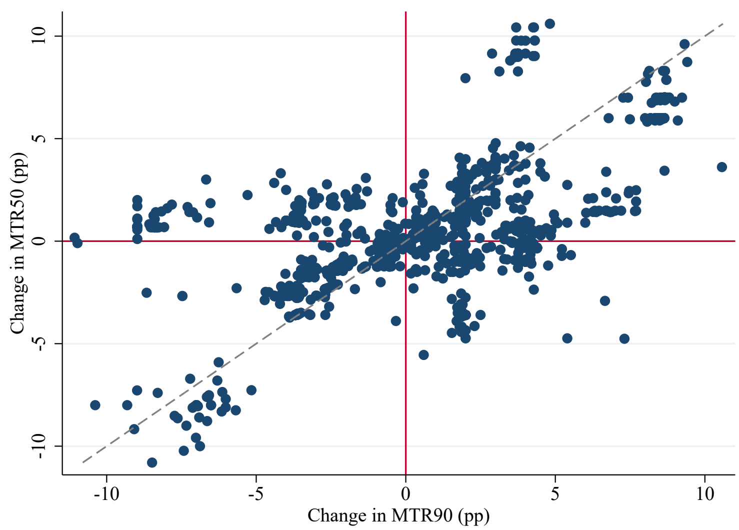

Notes: Figure is a scatter plot showing the relative size of tax rate changes for median versus 90th percentile workers. Each dot is a different state-year cell and the dashed line is a 45 degree line. Of the state-years that experience a change in the tax rate for either the 90th percentile or median worker, 44% had a larger change for top earners while 56% had a larger change for median earner.

Table C.4: Share of Tax Changes Generated by Each Source

	Federal tax rate (1)	State tax rate (2)	Income distribution (3)
Personal tax rate for median earner	0.763	0.237	0.040
Personal tax rate for 90th percentile	0.783	0.217	0.046
Corporate top tax rate	0.766	0.288	–

Notes: Table reports the share of state-level changes in combined (i.e. federal plus state) tax liabilities that are caused by various forces. Column (1) reports the share of tax rate changes that are accompanied by a federal tax rate change. These changes may also be accompanied by state tax changes; indeed, since federal taxes are deductible for personal taxes in many states, nearly every change in federal tax change is associated with a change in personal state rates. Column (2) reports the share of tax rate changes that arise from changes in the statutory state tax rate. Since nearly all changes in personal federal tax rates are accompanied by state tax changes, column (2) only reports the share of changes that have no federal component for the personal tax rates. Thus column (1) and (2) add up to 1 in the first two rows. For corporate taxes, we report the share of changes that have a state-level component, regardless of whether federal taxes change as well. Finally, column (3) shows the share of personal income tax changes for percentile p that result from that percentile crossing across a tax bracket. We have no such tax bracket information for corporate taxes, thus this column is blank on the bottom row.

Figure C.11: The Evolution of Corporate Taxes

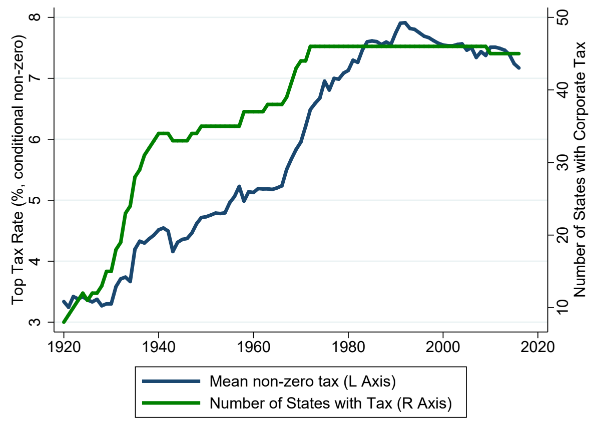

Panel A: Intensive and Extensive Margin

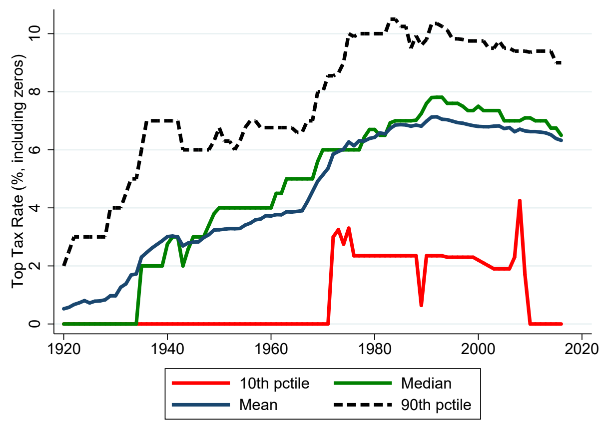

Panel B: Distribution of Statutory Tax Rates

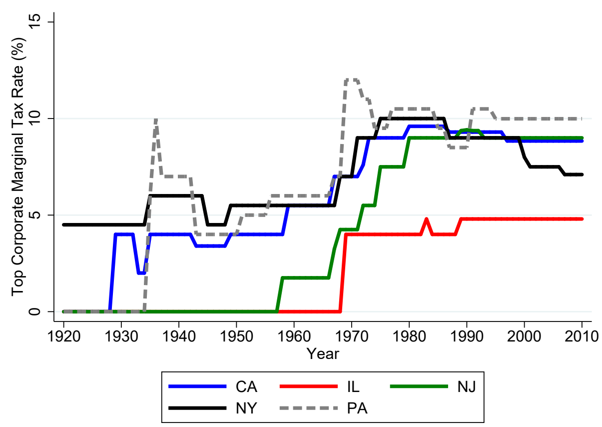

Panel C: Time Series of Select States

Notes: Figure plots the time series of the distribution and proliferation of state corporate tax rates. Panel A shows the number of states with a corporate income tax and the mean non-zero tax rate. Panel B plots the distribution of top state corporate tax rates over time. Panel C shows the evolution of top state corporate tax rates for the five most innovative states in our sample. All tax rates are measured in percentage points.

Figure C.12: Introduction Year of State Corporate Taxes

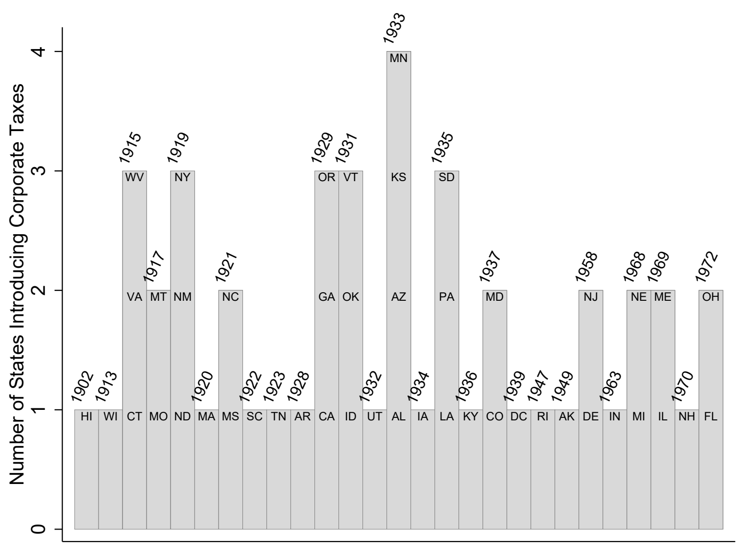

Notes: Figure plots the first year in which each state has a statutory corporate income tax rate.

Figure C.13: Top State Corporate Marginal Tax Rates over Time (Percentage Points)

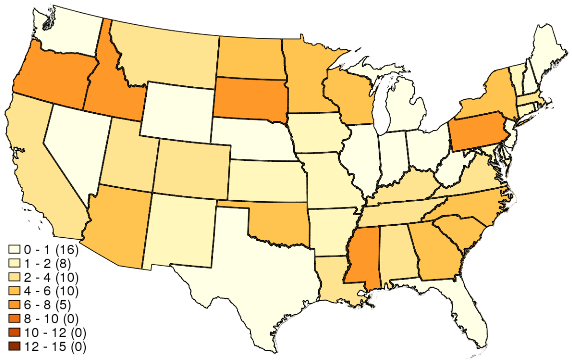

Panel A: 1940

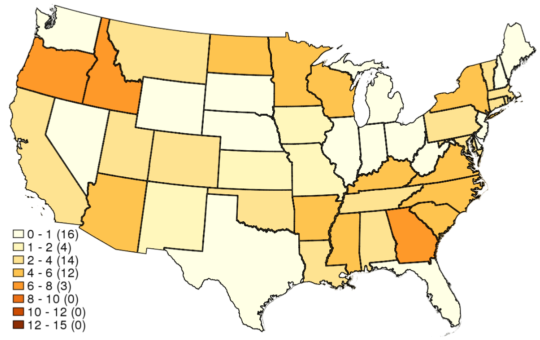

Panel B: 1950

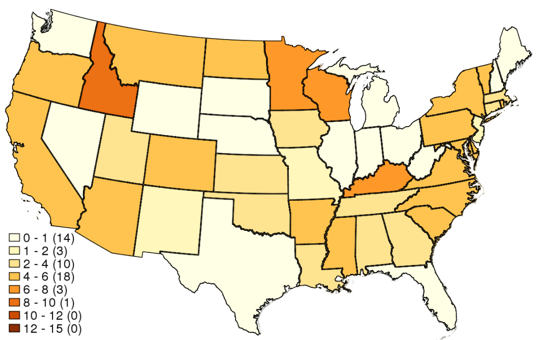

Panel C: 1960

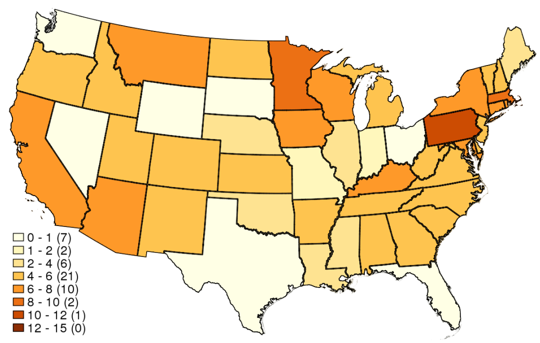

Panel D: 1970

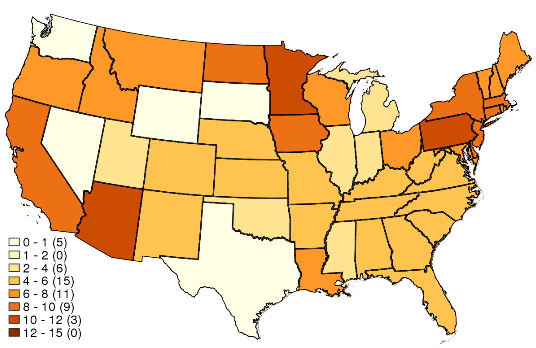

Panel E: 1980

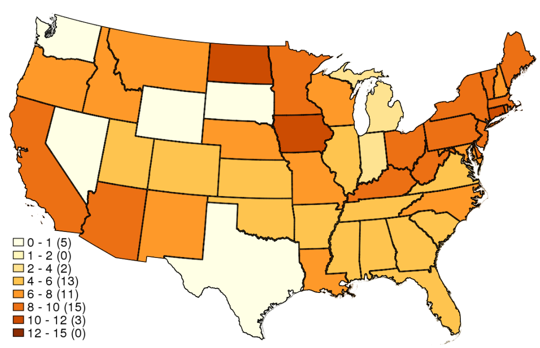

Panel F: 1990

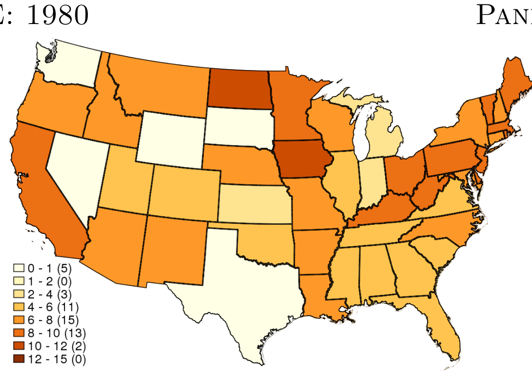

Panel G: 2000

Notes: Figure plots the state top statutory marginal corporate income tax rates every ten years from 1940 through 2000. Darker colors indicate higher tax rates. Tax rates measured in percentage points.

Figure C.14: Evolution of State Corporate Income Apportionment Rules Through Time

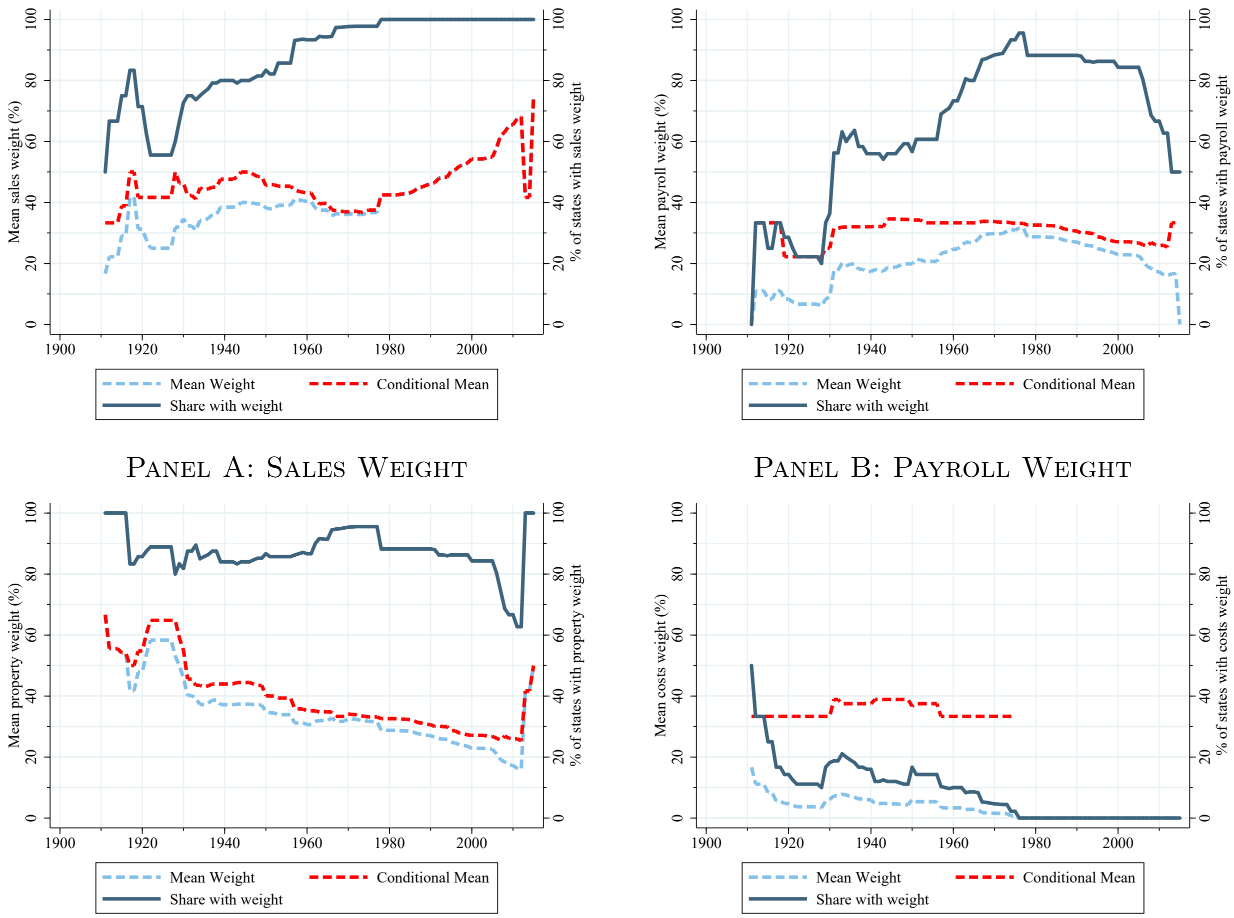

Panel A: Sales Weight

Panel B: Payroll Weight

Panel C: Property Weight

Panel D: Costs Weight

Notes: Figure plots the evolution of state corporate income tax apportionment rules. The solid blue lines plot the share of states that have at least some weight on a particular factor, conditional on having a corporate tax at all. The light blue dashed lines plot the average weight placed on the factor across all states with a corporate tax, while the red dashed lines plot the average weight placed on the factor by states which have at least some weight on that apportionment factor. Each panel considers a different apportionment factor. Data collection and definitions described in detail in Appendix A.6.2.

Figure C.15: Distribution of Time Between Large Tax Changes

Panel A: Distribution of Time Between Large Top Statutory Personal Income Tax Changes

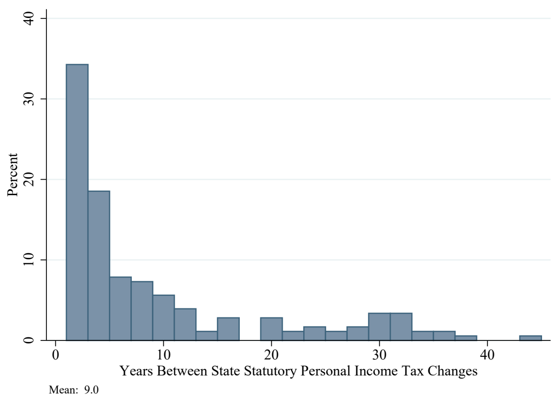

Panel B: Mean Time Between Large Top Statutory Personal Income Tax Changes

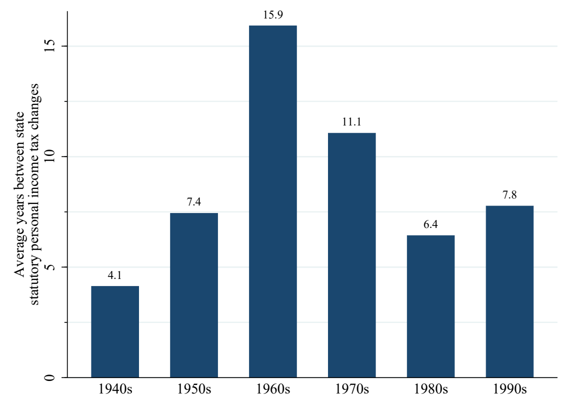

Panel C: Distribution of Time Between Large Top Statutory Corporate Income Tax Changes

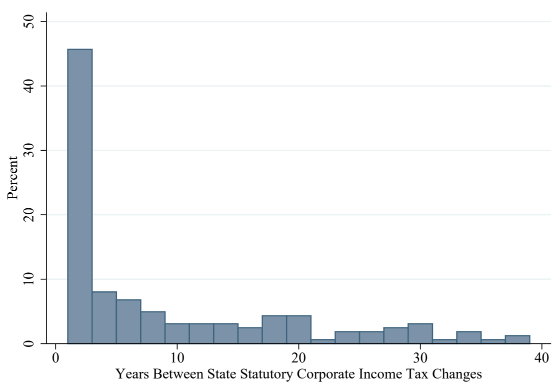

Panel D: Mean Time Between Large Top Statutory Corporate Income Tax Changes

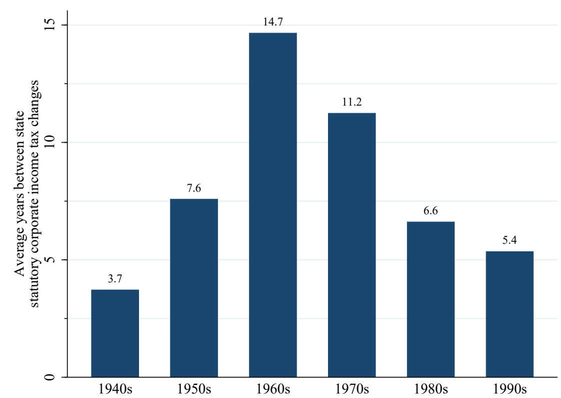

Notes: Figure summarized the number of years between large (i.e. top 10%) state tax reforms, the likes of which are used to define our event studies. Panels A and B consider state reforms to the top statutory personal tax rate, while Panels C and D consider state reforms to the top statutory corporate tax rate. Panels A and C show the distribution of times between large reforms, while Panels B and D show the average time between large reforms by decade. Decades are assigned to the second reform in a pair; for instance, if there is a reform in 1957 and 1962, a time gap of 5 would be coded in 1962.

C.3 Robustness and Extensions: State-Level Regressions

Table C.5: State-level OLS regressions with all control coefficients reported

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1-MTR90)$	1.803***	1.516***	1.784***	0.056
	(0.450)	(0.507)	(0.427)	(0.071)
ln(1− Corp. MTR)	2.759***	2.382***	2.308***	0.573***
	(0.701)	(0.770)	(0.640)	(0.141)
Real Personal Income per Capita	0.069	0.308*	0.074	0.025
	(0.131)	(0.154)	(0.125)	(0.020)
Population Density	-0.009	0.002	-0.005	-0.010***
	(0.011)	(0.013)	(0.010)	(0.002)
R&D Tax Credit	0.001	0.010	0.002	0.002
	(0.010)	(0.013)	(0.009)	(0.002)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	7.07	9.65	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical to that table, but reports the coefficients additional control variables as well.

Table C.6: State-level OLS regressions on additional outcomes

	Log Unadj. Citations (1)	Log Corp. Patents (2)	Log Non-Corp. Patents (3)	Log Mean Pat. Value (4)	Log Av. Estab. Emp. (5)
$\ln(1-MTR90)$	1.527***	1.810***	1.080**	-0.859*	-0.045
	(0.512)	(0.552)	(0.419)	(0.437)	(0.236)
ln(1− Corp. MTR)	2.923***	4.000***	1.979***	1.791***	1.399***
	(0.832)	(0.915)	(0.547)	(0.556)	(0.367)
Observations	2867	2867	2867	2831	2867
R^2	0.96	0.95	0.96	0.75	0.87
Mean of Dep. Var.	9.55	6.71	5.68	3.02	3.98
S.D. of Dep. Var.	1.57	1.49	1.13	0.58	0.32

	Log Value Added (6)	Log Total Payroll (7)	Log Av. Weekly Earn. (8)	Log Income Per Capita (9)	Share in Manufact. (10)
$\ln(1-MTR90)$	0.862	1.181**	0.129	0.296**	12.001***
	(0.556)	(0.548)	(0.118)	(0.113)	(4.473)
ln(1− Corp. MTR)	3.875***	3.438***	0.338***	0.747***	22.468***
	(0.681)	(0.689)	(0.125)	(0.165)	(7.621)
Observations	2867	2867	2867	2867	2844
R^2	0.95	0.95	0.96	0.98	0.93
Mean of Dep. Var.	8.74	7.96	4.64	8.08	25.15
S.D. of Dep. Var.	1.19	1.21	0.24	0.42	10.94

Notes: Table reports estimates from a state-level regression following equation (3), using alternative outcome variables. The outcome variables are as follows. Column (1) log total unadjusted citation counts; (2) Log number of patents assigned to corporations; (3) Log number of patents that are not assigned to corporations; (4) log average Kogan et al. (2017) patent value among patents applied for in that state; (5) Log average establishment size; (6) Log total state manufacturing value added; (7) Log total state manufacturing payrolls; (8) Log average weekly earnings of manufacturing employees in the state; (9) Log real per capita income; (10) Percent of workers in the manufacturing sector. Robust standard errors two-way clustered at state × five-year and year level in parentheses. All regressions control for lagged population density, real personal income per capita (excluded in column 9), R&D tax credits, state and year fixed effects and are weighted by state population in 1940. Tax rates are lagged by 3-years and measured as log net-of-tax rates. "ATRs" correspond to average tax rates, while "MTRs" correspond to marginal tax rates. Mainland states, excluding Louisiana, included for the period 1940-2000. $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table C.7: State-level OLS regressions using alternative tax rates

	Log Patents	Log Citations	Log Inventors	Share Assigned
$\ln(1 - MTR50)$	2.749**	2.738**	2.649**	-0.053
	(1.092)	(1.201)	(0.996)	(0.169)
ln(1− Corp. MTR)	2.783***	2.323***	2.345***	0.600***
	(0.745)	(0.813)	(0.678)	(0.146)
$\ln(1 - ATR90)$	3.396***	2.982**	3.323***	-0.113
	(1.084)	(1.202)	(1.001)	(0.173)
ln(1− Corp. MTR)	2.614***	2.235***	2.173***	0.612***
	(0.744)	(0.816)	(0.678)	(0.149)
$\ln(1 - ATR50)$	5.529***	5.457***	5.020***	0.139
	(1.803)	(2.015)	(1.632)	(0.288)
ln(1− Corp. MTR)	2.781***	2.326***	2.371***	0.577***
	(0.740)	(0.799)	(0.673)	(0.142)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	7.07	9.65	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except it includes the different personal income tax rates as independent variables. Each panel, separated by a horizontal line, reports separate sets of regressions. faced by both median ($MTR50$) and 90^{th} percentile earners as independent variables. "ATRs" correspond to average tax rates, while "MTRs" correspond to marginal tax rates.

Table C.8: State-level OLS regressions including the personal marginal income tax rate for both median and 90^{th} percentile earners as independent variables

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	1.384***	1.031*	1.391***	0.086
	(0.447)	(0.517)	(0.438)	(0.069)
$\ln(1 - MTR50)$	1.668	1.932	1.561	-0.120
	(1.203)	(1.344)	(1.118)	(0.178)
ln(1− Corp. MTR)	2.574***	2.168***	2.135***	0.587***
	(0.719)	(0.798)	(0.653)	(0.147)
Observations	2867	2867	2867	2867
R^2 - Overall	0.96	0.96	0.97	0.87

Notes: See notes to Table II. This table is identical except it includes the marginal personal income tax rate faced by both median ($MTR50$) and 90^{th} percentile earners as independent variables.

Table C.9: State-level OLS regressions using tax rates for married individuals with two dependents

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1-MTR90)$	1.680***	1.599***	1.667***	0.053
	(0.493)	(0.549)	(0.476)	(0.074)
ln(1− Corp. MTR)	2.768***	2.293***	2.351***	0.546***
	(0.692)	(0.754)	(0.637)	(0.140)
Observations	2773	2773	2773	2773
Mean of Dep. Var.	7.08	9.67	7.10	0.72
S.D. of Dep. Var.	1.32	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except it uses the personal income tax rate faced by a married individual with two dependents, rather than a single individual with no dependents.

Table C.10: State-level OLS regressions using tax rates which force itemized deductions

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1-MTR90)$	2.061***	1.735**	2.091***	-0.002
	(0.578)	(0.669)	(0.558)	(0.094)
ln(1− Corp. MTR)	2.784***	2.357***	2.358***	0.561***
	(0.712)	(0.773)	(0.652)	(0.140)
Observations	2773	2773	2773	2773
Mean of Dep. Var.	7.08	9.67	7.10	0.72
S.D. of Dep. Var.	1.32	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except it calculates the personal income tax rate assuming that all individuals itemize their deductions rather than taking the standard deduction if it is optimal to do so.

Table C.11: State-level OLS regressions using Auerbach-Poterba Effective Federal Corporate Tax

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	1.702***	1.474***	1.636***	0.079
	(0.430)	(0.498)	(0.408)	(0.076)
ln(1− Corp. MTR)	3.048***	2.588***	2.673***	0.540***
	(0.788)	(0.851)	(0.724)	(0.159)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	7.07	9.65	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except, where possible, federal corporate tax rates are calculated using the effective Federal tax rate calculated by Auerbach and Poterba (1987) and updated by Auerbach (2007).

Table C.12: State-level OLS regressions using alternative lags on tax rates

PANEL A: 1-YEAR LAGGED TAXES				
	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	1.911***	1.695***	1.887***	0.082
	(0.499)	(0.577)	(0.475)	(0.072)
ln(1− Corp. MTR)	2.958***	2.495***	2.476***	0.604***
	(0.673)	(0.747)	(0.615)	(0.144)
PANEL B: 2-YEAR LAGGED TAXES				
$\ln(1 - MTR90)$	1.845***	1.603***	1.819***	0.064
	(0.465)	(0.528)	(0.441)	(0.072)
ln(1− Corp. MTR)	2.847***	2.419***	2.380***	0.597***
	(0.688)	(0.761)	(0.627)	(0.148)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	7.07	9.65	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except it lags tax rates by different amounts. Tax rates are lagged by 1 year in Panel A and by 2 years in Panel B, compared with 3 years in Table II.

Table C.13: State-level OLS regressions with 10-year Newey-West standard errors

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	1.803***	1.516**	1.784***	0.056
	(0.642)	(0.717)	(0.613)	(0.089)
ln(1− Corp. MTR)	2.759***	2.382**	2.308***	0.573***
	(0.831)	(0.962)	(0.759)	(0.162)
R&D Tax Credit	0.076	0.994	0.219	0.168
	(1.415)	(1.821)	(1.309)	(0.214)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	7.07	9.65	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.34	0.14

Notes: See notes to Table II. This table is identical except that Newey-West standard errors allowing for autocorrelation in errors for up to 10 years are reported in parentheses, rather than the baseline of standard errors two-way clustered at the state × five-year and year levels.

Table C.14: State-level OLS regressions excluding New York and California

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	2.866***	2.861***	2.842***	0.050
	(0.546)	(0.603)	(0.493)	(0.112)
ln(1− Corp. MTR)	2.453***	1.891**	1.998***	0.567***
	(0.665)	(0.723)	(0.600)	(0.145)
Observations	2745	2745	2745	2745
Mean of Dep. Var.	6.79	9.35	6.81	0.72
S.D. of Dep. Var.	1.26	1.49	1.27	0.15

Notes: See notes to Table II. This table is identical except that New York and California are dropped from the estimation sample.

Table C.15: State-level OLS regressions excluding the 1970s

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	2.108***	2.036***	2.121***	-0.007
	(0.578)	(0.662)	(0.546)	(0.104)
ln(1− Corp. MTR)	3.792***	3.259***	3.162***	0.804***
	(0.873)	(0.987)	(0.810)	(0.175)
Observations	2350	2350	2350	2350
Mean of Dep. Var.	7.07	9.68	7.08	0.71
S.D. of Dep. Var.	1.35	1.62	1.37	0.15

Notes: See notes to Table II. This table is identical except that the 1970s are dropped from the estimation sample.

Table C.16: State-level OLS including controls for state political leaning

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1 - MTR90)$	1.087**	0.743	1.104***	-0.026
	(0.421)	(0.449)	(0.395)	(0.079)
ln(1− Corp. MTR)	2.100***	2.185***	1.790***	0.589***
	(0.592)	(0.651)	(0.537)	(0.148)
State Governor: Democrat	0.062***	0.061**	0.059***	0.004
	(0.021)	(0.024)	(0.020)	(0.004)
% State Upper House Democrat	0.001	0.001	0.001	0.000
	(0.002)	(0.002)	(0.001)	(0.000)
% State Lower House Democrat	-1.217***	-1.212***	-1.110***	-0.147***
	(0.179)	(0.208)	(0.165)	(0.025)
Observations	2764	2764	2764	2764
Mean of Dep. Var.	7.10	9.69	7.12	0.72
S.D. of Dep. Var.	1.31	1.55	1.33	0.14

Notes: See notes to Table II. This table is identical except that it includes additional political controls, specifically the party of the state governor, and the percent of the state legislature which is Democrat.

Table C.17: State-level OLS regressions including fixed effects but excluding all controls

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1-MTR90)$	1.770***	1.616***	1.771***	0.031
	(0.441)	(0.482)	(0.417)	(0.075)
ln(1− Corp. MTR)	2.910***	2.136***	2.359***	0.804***
	(0.649)	(0.694)	(0.582)	(0.151)
Observations	2914	2914	2914	2914
Mean of Dep. Var.	7.06	9.64	7.08	0.72
S.D. of Dep. Var.	1.33	1.56	1.35	0.15

Notes: See notes to Table II. This table is identical except that it drops all controls from the regression, but continues to include state and year fixed effects.

Table C.18: Unweighted state-level OLS regressions

	Log Patents (1)	Log Citations (2)	Log Inventors (3)	Share Assigned (4)
$\ln(1-MTR90)$	2.083***	1.850***	2.004***	0.025
	(0.399)	(0.470)	(0.385)	(0.081)
ln(1− Corp. MTR)	1.452**	1.683**	1.079*	0.519***
	(0.660)	(0.749)	(0.605)	(0.148)
Observations	2867	2867	2867	2867
Mean of Dep. Var.	6.04	8.62	6.06	0.66
S.D. of Dep. Var.	1.44	1.70	1.46	0.18

Notes: See notes to Table II. This table is identical except that it does not weight states by their 1940 population.

Table C.19: First Stage of Core Macro Instrumental Variables regression: 5-year Lagged Income Distribution

Dep. Var.:	MTR - 90^{th}	Top Corp. MTR	MTR - 50^{th}	Top Corp. MTR
	(1)	(2)	(3)	(4)
90th Pctile Income MTR Instr	0.650***	-0.014		
	(0.081)	(0.025)		
Median Income MTR Instr			0.779***	-0.046
			(0.043)	(0.035)
Top Corporate MTR Instr	0.060	0.752***	0.063**	0.756***
	(0.055)	(0.060)	(0.025)	(0.060)
	(0.001)	(0.000)	(0.000)	(0.000)
Observations	2867	2867	2867	2867
R^2	0.366	0.605	0.592	0.606

Notes: Table reports the first stage regression of the state-level instrumental variables regression reported in Panel C of Table II. Personal tax rates and corporate tax rates are instrumented for by the predicted tax rates given by equations (6) and (8) respectively. Standard errors two-way clustered at state × five-year and year level are reported in parentheses below point estimate. IV convolves five-year lags of both the state tax law and income distribution with current federal tax rates. All regressions include controls for lagged personal income per capita, population density, and R&D Tax Credits, as well as state and year fixed effects. Each column corresponds to a different dependent variable.

Table C.20: Time Series Regressions on Federal Tax Rates

	Log Patents	Log Citations	Log Inventors	Share Assigned
	(1)	(2)	(3)	(4)
$\ln(1 - MTR90)$	1.896***	2.494***	1.875***	0.071
	(0.226)	(0.228)	(0.240)	(0.045)
$\ln(1-$ Corp $\text{MTR}^{AP})$	0.111	0.463**	0.512	0.073
	(0.334)	(0.190)	(0.307)	(0.093)
Trend	0.067***	0.163***	0.077***	0.003
	(0.010)	(0.016)	(0.012)	(0.004)
$\ln(Income/Cap)$	-2.108***	-4.274***	-2.342***	-0.066
	(0.402)	(0.654)	(0.515)	(0.168)
Observations	39	39	39	39

Notes: Table reports estimates from time series regressions of aggregate innovation measures on federal tax rates. "$MTR90$" corresponds to the federal tax rate that would apply to someone whose adjusted gross income is equal to the 90^{th} percentile of income in the US in that year, while "Corp. MTR^{AP}" represents to the effective federal corporate tax rate as calculated by Auerbach and Poterba (1987). Tax rates are lagged by three years. "Trend" is a linear trend, while $Income/Cap$ is aggregate personal income per capita, lagged by one year. Newey-West standard errors allowing for up to 10-years of serial correlation are reported in parentheses. Auerbach-Poterba corporate tax rates are not available before 1959; therefore, the sample consists of all years from 1962-2000.

C.4 Robustness and Extensions: Inventor-Level Regressions

Table C.21: Effects of Taxes at the Individual Inventor Level (IV)

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.616***	0.490***	0.855***	1.340***	0.246**
	(0.063)	(0.060)	(0.140)	(0.157)	(0.092)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.353	0.404	0.504	0.622	0.737
ln(1− Personal MTR)	0.554***	0.427***	0.791***	1.075***	0.367***
	(0.051)	(0.056)	(0.134)	(0.146)	(0.069)
ln(1− Corp. MTR)	-0.257	-0.227	-0.781*	-1.161*	-0.090
	(0.163)	(0.166)	(0.406)	(0.638)	(0.191)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	5634789	5634789	3576456	3459553	1150717
Mean of Dep. Var.	0.638	0.414	0.602	2.958	0.609
S.D. of Dep. Var.	0.481	0.492	0.709	1.470	0.488
R^2	0.358	0.405	0.493	0.624	0.729

Notes: See the notes to Table IV. Personal tax rates and corporate tax rates are instrumented for by the predicted tax rates given by equations (6) and (8) respectively. That is, the instruments use current federal tax law and five-year lagged state tax laws, using a five-year lagged income distribution. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

Table C.22: First Stage of Core Inventor-level Instrumental Variables Regression

LHS Tax:	Personal (1)	Personal (2)	Corporate (3)
Personal MTR Instrument	0.918***	0.988***	-0.017
	(0.076)	(0.086)	(0.015)
Corporate MTR Instrument	0.029		0.396***
	(0.051)		(0.097)
State × Year FE	N	Y	N
Inventor FE	Y	Y	Y
Observations	6899376	8305203	8287859
Mean of Dep. Var.	-0.246	-0.245	-0.571
S.D. of Dep. Var.	0.115	0.109	0.157
R^2	0.881	0.924	0.221

Notes: Table reports the first stage regression of the inventor-level instrumental variables. Personal tax rates and corporate tax rates are instrumented for by the predicted tax rates given by equations (6) and (8) respectively. All tax rates are input as log retention rates, so if the tax rate is τ, the dependent variable and instrument are included in the regression as $\ln(1-\tau)$. Standard errors two-way clustered at state and year level are reported in parentheses below point estimate. IV convolves five-year lags of both the income distribution and the inventor's residence state tax law with current federal tax rates. See notes to Table IV for list of control variables. Columns 1 and 2 have the personal marginal tax rate as the dependent variable, while column 3 uses the top corporate marginal tax rate as a dependent variable.

Table C.23: Inventor-level OLS regressions defining counting inventors as being high quality if they are in the top 5% of patent counts

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.425***	0.367***	0.672***	0.728***	0.482***
	(0.058)	(0.056)	(0.176)	(0.186)	(0.114)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.420	0.516	0.578	0.715	0.806
ln(1− Personal MTR)	0.448***	0.417***	0.849***	1.109***	0.288***
	(0.051)	(0.055)	(0.162)	(0.170)	(0.068)
ln(1− Corp. MTR)	0.080	0.074	0.093	0.120	0.004
	(0.065)	(0.064)	(0.118)	(0.219)	(0.109)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6460412	6460412	4566398	4415498	1485493
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.420	0.517	0.578	0.716	0.808

Notes: See notes to Table IV. This table is identical except it defines inventors to be high quality if they belong to the top 5% of the distribution of cumulative patent counts, rather than the top 10%.

Table C.24: Inventor-level OLS regressions defining counting inventors as being high quality if they are in the top 25% of patent counts

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.391***	0.352***	0.687***	1.075***	0.186**
	(0.062)	(0.055)	(0.121)	(0.136)	(0.086)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.422	0.519	0.581	0.718	0.809
ln(1− Personal MTR)	0.369***	0.330***	0.661***	0.982***	0.198**
	(0.056)	(0.051)	(0.106)	(0.125)	(0.079)
ln(1− Corp. MTR)	0.077	0.075	0.103	0.109	0.010
	(0.061)	(0.059)	(0.104)	(0.205)	(0.110)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6460412	6460412	4566398	4415498	1485493
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.421	0.518	0.579	0.717	0.807

Notes: See notes to Table IV. This table is identical except it defines inventors to be high quality if they belong to the top 25% of the distribution of cumulative patent counts, rather than the top 10%.

Table C.25: Inventor-level OLS regressions using three inventor quality levels

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.426***	0.348***	0.694***	1.066***	0.224***
	(0.067)	(0.051)	(0.095)	(0.112)	(0.048)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.382	0.427	0.538	0.659	0.771
ln(1− Personal MTR)	0.406***	0.330***	0.666***	0.970***	0.220***
	(0.061)	(0.047)	(0.082)	(0.102)	(0.045)
ln(1− Corp. MTR)	0.069	0.072	0.098	0.121	-0.043
	(0.057)	(0.052)	(0.090)	(0.186)	(0.088)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	4669149	4669149	2795112	2714189	908602
Mean of Dep. Var.	0.599	0.386	0.598	2.961	0.661
S.D. of Dep. Var.	0.490	0.487	0.726	1.494	0.473
R^2	0.381	0.426	0.536	0.658	0.769

Notes: See notes to Table IV. This table is identical except it assigns inventors to three quality levels. Personal MTR in this table is defined as the marginal tax rate faced by the 90^{th} percentile earner in state s in year t for high productivity inventors (top decile of lagged patent counts), by the 75^{th} percentile earner for middle productivity inventors (between the 75th and 90th percentiles of lagged patent counts) and the marginal tax rate rate faced by the median earner for low productivity inventors.

Table C.26: Inventor-level OLS regressions ranking inventors according to their cumulative citation counts

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.343***	0.322***	0.545***	0.866***	0.111
	(0.067)	(0.075)	(0.159)	(0.204)	(0.077)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.420	0.517	0.577	0.718	0.808
ln(1− Personal MTR)	0.317***	0.294***	0.530***	0.783***	0.154**
	(0.057)	(0.066)	(0.142)	(0.174)	(0.070)
ln(1− Corp. MTR)	0.103	0.097	0.149	0.178	0.028
	(0.068)	(0.068)	(0.122)	(0.231)	(0.112)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6460412	6460412	4566398	4415498	1485493
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.419	0.516	0.576	0.717	0.807

Notes: See notes to Table IV. This table is identical except it defines inventors to be high quality if they belong to the top 10% of the distribution of cumulative citation counts, rather than the top 10% of cumulative patent counts.

TABLE C.27: INVENTOR-LEVEL OLS REGRESSIONS INTERACTING MTR90 WITH PERCENTILE BINS OF INVENTOR QUALITY

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Cites (4)	Has High-Value Pat. (5)
ln(1− MTR90) × 50th-75th	0.149**	0.080**	0.097***	0.210***	-0.055**
	(0.056)	(0.035)	(0.033)	(0.043)	(0.025)
ln(1− MTR90) × 75th-90th	0.300***	0.194***	0.313***	0.580***	0.005
	(0.091)	(0.062)	(0.077)	(0.103)	(0.038)
ln(1− MTR90) × 90th-95th	0.420***	0.288***	0.527***	0.837***	0.099*
	(0.104)	(0.072)	(0.112)	(0.139)	(0.053)
ln(1− MTR90) × 95th-99th	0.565***	0.448***	0.890***	1.325***	0.181***
	(0.117)	(0.083)	(0.151)	(0.176)	(0.059)
ln(1− MTR90) × 99th+	0.735***	0.679***	1.696***	2.270***	0.269***
	(0.133)	(0.112)	(0.246)	(0.293)	(0.070)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.393	0.439	0.549	0.667	0.782
ln(1− MTR90) × 50th-75th	0.142***	0.077**	0.091***	0.175***	-0.050**
	(0.051)	(0.031)	(0.027)	(0.043)	(0.020)
ln(1− MTR90) × 75th-90th	0.291***	0.189***	0.296***	0.526***	0.013
	(0.084)	(0.057)	(0.068)	(0.095)	(0.032)
ln(1− MTR90) × 90th-95th	0.410***	0.281***	0.500***	0.768***	0.107**
	(0.097)	(0.067)	(0.105)	(0.129)	(0.047)
ln(1− MTR90) × 95th-99th	0.553***	0.440***	0.858***	1.254***	0.191***
	(0.110)	(0.079)	(0.144)	(0.167)	(0.056)
ln(1− MTR90) × 99th+	0.722***	0.671***	1.665***	2.204***	0.281***
	(0.128)	(0.110)	(0.239)	(0.290)	(0.065)
ln(1− Corp. MTR)	0.090	0.098*	0.168	0.227	-0.005
	(0.060)	(0.056)	(0.107)	(0.191)	(0.088)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	4777949	4777949	2883935	2798747	940657
Mean of Dep. Var.	0.604	0.387	0.588	2.943	0.652
S.D. of Dep. Var.	0.489	0.487	0.723	1.490	0.476
R^2	0.392	0.438	0.546	0.666	0.780

Notes: Table reports estimated coefficients from regression which interacts the marginal tax rate faced by an individual at the 90^{th} percentile of the income distribution with percentile bins of the inventor cumulative patents distribution. Regressions also include indicators for inventor quality bins, but are otherwise identical to that in Table IV. Estimated coefficients represent (semi-)elasticities of innovation to the marginal tax rate faced by a 90^{th} percentile earner in the inventor's residence state, relative inventors in the bottom 50% of the inventor quality distribution, who serve as a control. The upper bound elasticity of innovation to taxes is the coefficient on $\ln(1-MTR90)\times$ 99th+. A lower bound on the elasticity is this coefficient minus the coefficient on $\ln(1-MTR90)\times$ 95th-99th.

Table C.28: Inventor-level OLS regressions excluding New York and California

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.495***	0.447***	0.902***	1.233***	0.323***
	(0.070)	(0.069)	(0.171)	(0.186)	(0.086)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.428	0.523	0.581	0.717	0.813
ln(1− Personal MTR)	0.454***	0.410***	0.849***	1.133***	0.338***
	(0.060)	(0.061)	(0.153)	(0.165)	(0.078)
ln(1− Corp. MTR)	0.052	0.044	-0.022	0.021	0.006
	(0.060)	(0.059)	(0.085)	(0.212)	(0.124)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	4910129	4910129	3467757	3350408	1126399
Mean of Dep. Var.	0.706	0.416	0.448	2.735	0.497
S.D. of Dep. Var.	0.455	0.493	0.654	1.428	0.500
R^2	0.426	0.522	0.580	0.716	0.811

Notes: See notes to Table IV. This table is identical except it drops inventors living in New York and California from the estimation sample.

Table C.29: Inventor-level OLS regressions excluding the 1970s

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.395***	0.319***	0.543***	0.915***	0.124
	(0.089)	(0.085)	(0.195)	(0.222)	(0.109)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.434	0.533	0.600	0.735	0.812
ln(1− Personal MTR)	0.359***	0.286***	0.523***	0.827*	0.157
	(0.074)	(0.075)	(0.172)	(0.460)	(0.100)
ln(1− Corp. MTR)	0.159	0.196**	0.400**	0.413	0.147
	(0.100)	(0.097)	(0.166)	(1.269)	(0.181)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	5516741	5516741	3920471	3786699	1244162
Mean of Dep. Var.	0.711	0.431	0.472	2.833	0.500
S.D. of Dep. Var.	0.453	0.495	0.674	1.485	0.500
R^2	0.433	0.533	0.599	0.735	0.811

Notes: See notes to Table IV. This table is identical except it drops the 1970s from the estimation sample.

Table C.30: Inventor-level OLS regressions excluding all control variables

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.397***	0.321***	0.611***	0.710***	0.445***
	(0.053)	(0.051)	(0.138)	(0.133)	(0.127)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.420	0.517	0.578	0.716	0.805
ln(1− Personal MTR)	0.353***	0.286***	0.574***	0.650***	0.444***
	(0.045)	(0.045)	(0.121)	(0.113)	(0.112)
ln(1− Corp. MTR)	0.179*	0.161*	0.297	0.225	-0.025
	(0.101)	(0.093)	(0.218)	(0.247)	(0.109)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6465423	6465423	4569456	4418422	1485861
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.419	0.516	0.576	0.715	0.804

Notes: Table reports estimates from an inventor-level OLS regression, similar to Table IV only excluding all control variables. See notes to that table for details. Regressions only include controls for inventor quality and the listed fixed effects.

Table C.31: Inventor-level OLS regressions using tax rates for a married couple with two dependents

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.324***	0.310***	0.585***	0.888***	0.156
	(0.067)	(0.064)	(0.158)	(0.167)	(0.101)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.421	0.518	0.580	0.719	0.809
ln(1− Personal MTR)	0.294***	0.273***	0.539***	0.759***	0.173*
	(0.053)	(0.052)	(0.138)	(0.146)	(0.095)
ln(1− Corp. MTR)	0.084	0.075	0.108	0.137	0.016
	(0.063)	(0.063)	(0.118)	(0.214)	(0.116)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6344207	6344207	4491944	4345516	1472687
Mean of Dep. Var.	0.708	0.424	0.458	2.785	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.457	0.500
R^2	0.420	0.517	0.579	0.718	0.808

Notes: See notes to Table IV. This table is identical except it uses the tax rate faced by a married couple with two dependents, rather than the tax rate faced by a single individual. To calculate the married tax rate, we assume that an earner at the j^{th} percentile has a spouse who earns 50% as much as the j^{th} percentile.

Table C.32: Inventor-level OLS regressions excluding inventor fixed effects

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.135***	0.263***	0.093	0.112	-0.212***
	(0.029)	(0.022)	(0.105)	(0.167)	(0.070)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	N	N	N	N	N
R^2	0.154	0.094	0.114	0.173	0.163
ln(1− Personal MTR)	0.141***	0.247***	0.116	0.137	-0.150**
	(0.026)	(0.025)	(0.096)	(0.172)	(0.059)
ln(1− Corp. MTR)	0.078	0.108	0.132	0.186	-0.099
	(0.060)	(0.115)	(0.195)	(0.614)	(0.279)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	N	N	N	N	N
Observations	6460412	6460412	4566398	4415498	1485493
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.153	0.092	0.111	0.169	0.156

Notes: See notes to Table IV. This table is identical except it excludes inventor fixed effects.

TABLE C.33: ALTERNATIVE INVENTOR-LEVEL IV REGRESSIONS, USING STATE TAX LAWS FROM INVENTORS' HOME STATE

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.636***	0.505***	0.880***	1.406***	0.291***
	(0.075)	(0.067)	(0.152)	(0.170)	(0.087)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.353	0.404	0.504	0.622	0.737
ln(1− Personal MTR)	0.590***	0.456***	0.841***	1.191***	0.381***
	(0.066)	(0.062)	(0.149)	(0.155)	(0.072)
ln(1− Corp. MTR)	-0.274	-0.241	-0.805*	-1.218*	-0.097
	(0.165)	(0.168)	(0.413)	(0.647)	(0.192)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	5634718	5634718	3576424	3459521	1150707
Mean of Dep. Var.	0.638	0.414	0.602	2.958	0.609
S.D. of Dep. Var.	0.481	0.492	0.709	1.470	0.488
R^2	0.358	0.405	0.493	0.624	0.729

Notes: See notes to Table C.21. This table is identical except that the instruments use current federal tax law and five-year lagged home state tax laws, using a five-year lagged income distribution, rather than the inventor's current residence state.

Table C.34: First Stage of Alternative Inventor-level IV using state tax laws from inventors' home state

LHS Tax:	Personal (1)	Personal (2)	Corporate (3)
Personal MTR Instrument	0.847***	0.843***	-0.015
	(0.068)	(0.063)	(0.011)
Corporate MTR Instrument	0.098		0.395***
	(0.065)		(0.097)
State × Year FE	N	Y	N
Inventor FE	Y	Y	Y
Observations	6899298	8305081	8287738
Mean of Dep. Var.	-0.246	-0.245	-0.571
S.D. of Dep. Var.	0.115	0.109	0.157
R-Squared	0.853	0.897	0.220

Notes: Table presents first stage estimates for the inventor-level instrumental variable using the state tax laws of the inventor's home state. See notes to Tables C.22 and C.33.

Table C.35: Inventor-level OLS Regressions Allowing Coefficients to Vary pre versus post 1970

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
$\ln(1-MTR90)\times$ 1940-1969	0.629***	0.612***	1.023***	1.360***	0.059
	(0.123)	(0.135)	(0.330)	(0.375)	(0.117)
$\ln(1-MTR90)\times$ 1970-2000	0.122*	0.148**	0.349**	0.422**	0.052
	(0.066)	(0.056)	(0.149)	(0.182)	(0.088)
ln(1−Corp. MTR) × 1940-1969	0.087	0.192**	0.299	0.774**	-0.032
	(0.071)	(0.091)	(0.180)	(0.323)	(0.211)
ln(1−Corp. MTR) × 1970-2000	0.122*	0.106*	0.168	0.211	0.027
	(0.065)	(0.056)	(0.112)	(0.149)	(0.115)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.42	0.52	0.58	0.72	0.81
$\ln(1-MTR90)\times$ 1940-1969	0.669***	0.569***	1.028***	1.063***	0.712***
	(0.144)	(0.137)	(0.345)	(0.331)	(0.184)
$\ln(1-MTR90)\times$ 1970-2000	0.059	0.142**	0.318**	0.688***	-0.226
	(0.063)	(0.059)	(0.152)	(0.200)	(0.180)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6465345	6465345	4569398	4418390	1485854
Mean of Dep. Var.	0.71	0.42	0.46	2.78	0.50
S.D. of Dep. Var.	0.46	0.49	0.66	1.45	0.50
R^2	0.42	0.52	0.58	0.72	0.81

Notes: Table reports inventor-level OLS regressions which allow the effect of taxes to be different before and after 1970. We augment regression (12) by interacting tax rates with indicators for whether the time period t is before or after 1970. Otherwise, the specification is identical to that of Table IV. See footnotes to that table for more information.

Table C.36: Inventor-level OLS regressions showing the effect of agglomeration on tax elasticities, controlling for state patenting in any class

Dependent Variable:	Has Patent (1)	Has 10+ Cites (2)	Log Patents (3)	Log Citations (4)	Has High-Value Pat. (5)
ln(1− Personal MTR)	0.247	0.375**	0.828**	1.297***	-0.025
	(0.210)	(0.172)	(0.375)	(0.448)	(0.213)
ln(1− Personal MTR) × Agglom.	-0.247***	-0.313***	-0.420***	-0.698***	-0.034
	(0.086)	(0.078)	(0.152)	(0.199)	(0.040)
State Patents × ln(1− Personal MTR)	0.030	0.008	0.002	-0.012	0.035*
	(0.022)	(0.018)	(0.039)	(0.050)	(0.021)
State × Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
R^2	0.421	0.518	0.580	0.717	0.809
ln(1− Personal MTR)	0.309***	0.328***	0.845***	1.200***	0.107
	(0.097)	(0.103)	(0.249)	(0.319)	(0.194)
ln(1− Personal MTR) × Agglom.	-0.194**	-0.286***	-0.423**	-0.667***	-0.030
	(0.086)	(0.086)	(0.171)	(0.238)	(0.054)
State Patents × ln(1− Personal MTR)	0.016*	0.007	-0.011	-0.021	0.021
	(0.009)	(0.010)	(0.022)	(0.034)	(0.019)
State Patents	0.031***	0.033***	0.090***	0.076**	0.027***
	(0.007)	(0.008)	(0.026)	(0.031)	(0.010)
ln(1− Corp. MTR)	0.043	0.040	-0.017	0.035	-0.037
	(0.057)	(0.053)	(0.101)	(0.193)	(0.124)
State FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Inventor FE	Y	Y	Y	Y	Y
Observations	6460412	6460412	4566398	4415498	1485493
Mean of Dep. Var.	0.707	0.422	0.458	2.779	0.500
S.D. of Dep. Var.	0.455	0.494	0.664	1.454	0.500
R^2	0.420	0.517	0.579	0.717	0.807

Notes: Table plots estimates of the interaction of agglomeration forces with tax rates. Regression is identical to IV, except that it includes a control for the total number of patents in the state, and interactions of the tax rate with both our agglomeration measure and the total number of patents in the state.

C.5 Case Studies

We present here three special episodes of tax reform in New York, Delaware, and Michigan to provide some sharp visual evidence of the effects of taxes on innovation. Figures C.16-C.17 show the results from each of these episodes. In each case, the black solid line represents the time series in the state under consideration, while the dashed line represents a control state. The control state is constructed according to the synthetic control method of Abadie, Diamond and Hainmueller (2010). That is, it is a weighted average of other states in the sample, where the weights are chosen to best match the average innovation outcome of interest (patents, inventors, or citations), as well as real personal income per capita and population density for the period before the tax change in the state of interest. It is difficult to compute implied elasticities from these reforms as they changed both the corporate and personal income tax rates.

For the case of New York, the control state turns out to be California. For Michigan and Delaware, it is a combination of other states. For the post-tax change period, the synthetic state represents a plausible counterfactual of what may have happened in the state of interest absent the tax change. The first panel shows log patents, the second shows log inventors and the third shows log citations. The dashed vertical lines (or, for Michigan the gray area) represents the timing of the tax change.

New York 1968 vs. California

The first case study is shown in Figure C.16 and concerns New York's 1968 tax reform bill, in which the top marginal personal income tax rate increased from 10% to 14% and its state top corporate tax rate increased from 5.5% to 7%. The control state here is California, where the top tax rate increased as well from 7% to 10%, but remained lower while the corporate tax rate remained the same at 5.5%. All variables are normalized at their 1965 levels. Before the tax bill, New York and California follow remarkably similar trends for all three innovation outcomes. However, after the reform, they diverge and New York performs much worse in terms of innovation relative to the synthetic control.

Michigan 1967-1968

Figure C.17 shows the case study of Michigan. Michigan introduced its personal state tax rate in 1967 at 2.6%. One year later, in 1968, it introduced its corporate state tax at 5.6%. The synthetic control for Michigan is composed of several variations on California, Ohio/Pennsylvania, and, for some of the outcome variables, a bit of Texas. While the control state and Michigan evolve very similarly before 1967, Michigan starts performing significantly worse for the innovation outcome measures after the introduction of its tax regime.

Delaware 1969-1970

The third case study concerns Delaware. In July 1969, the corporate tax rate increased from 5% to 6%, and in August 1971 a temporary surcharge of 20% was added on top of the 6% corporate tax rate. In 1970, the personal tax rate increased from 11% to 18%. In this case, the best-fitting synthetic control is comprised of Nevada, California, and Connecticut. Figure C.18 shows that the effects on patents, citations, and inventors were noticeably large with the negative trend setting in at the time of the tax reform.

These case studies provide particularly clear visual evidence of a strong negative relationship between taxes and innovation. When combined with the macro state-level regressions, the instrumental variable approach and the border county analysis, the results overall bolster the conclusion that taxes were significantly negatively related to innovation outcomes at the state level.

Figure C.16: Synthetic Control Analysis: New York 1968

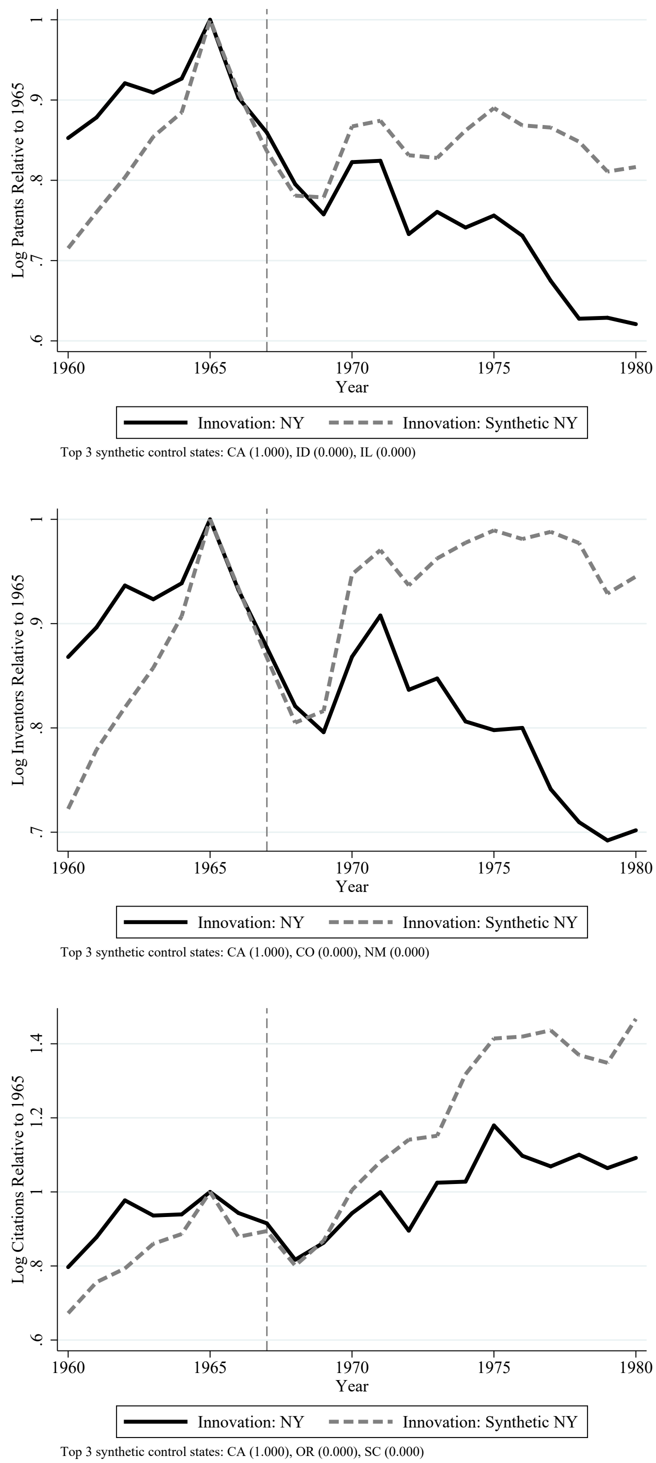

Notes: Figure plots synthetic control analyses for New York's 1968 tax reform bill, in which the top marginal personal income tax rate increased from 10% to 14%, and its state corporate tax rate increased from 5.5% to 7%. The first row shows the patterns for log patents, the middle row for log inventors, and the bottom row for log citations. We normalize the patent counts for synthetic and actual New York to be one in 1965.

Figure C.17: Synthetic Control Analysis: Michigan 1967-68

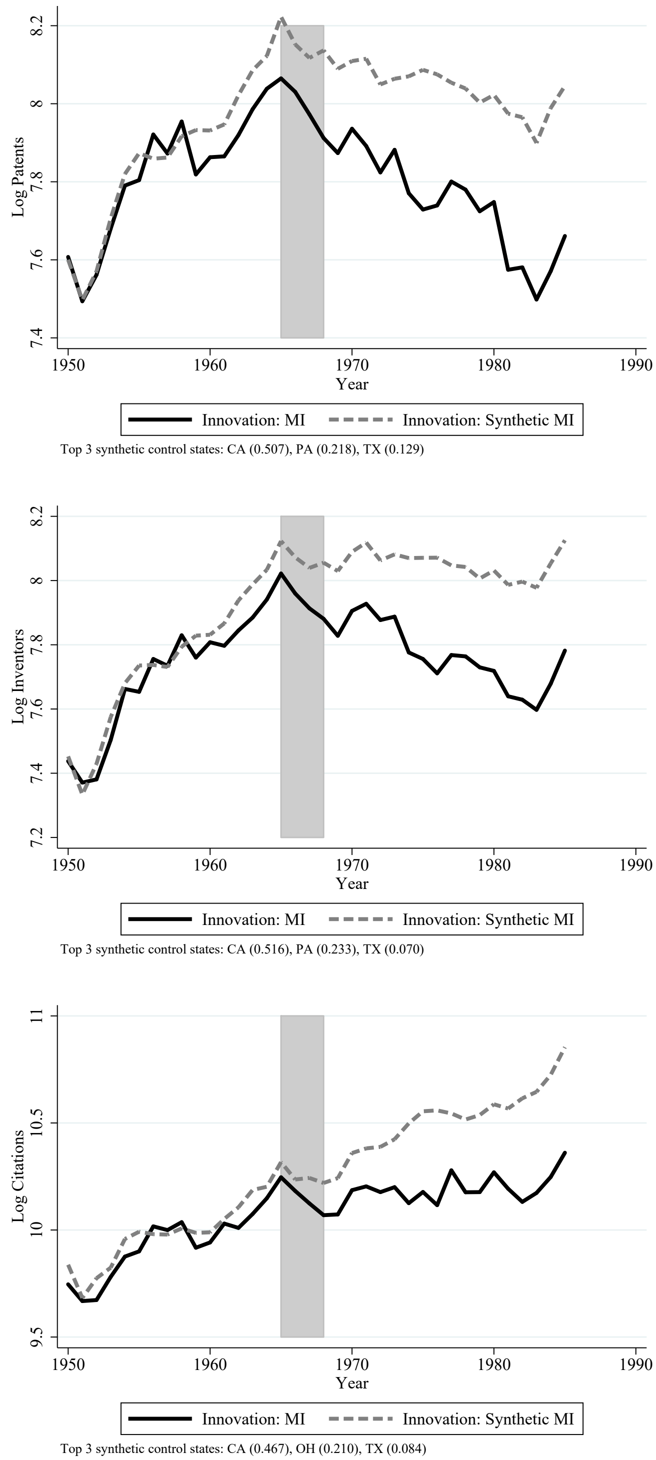

Notes: Figure plots synthetic control analyses Michigan around its major reforms in 1967 and 1968. In 1967, Michigan introduced its personal income tax, at a rate of 2.6%. In 1968, it then introduced its corporate income tax, at a rate of 5.6%. The first row shows the patterns for log patents, the middle row for log inventors, and the bottom row for log citations.

Figure C.18: Synthetic Control Analysis: Delaware 1969-71

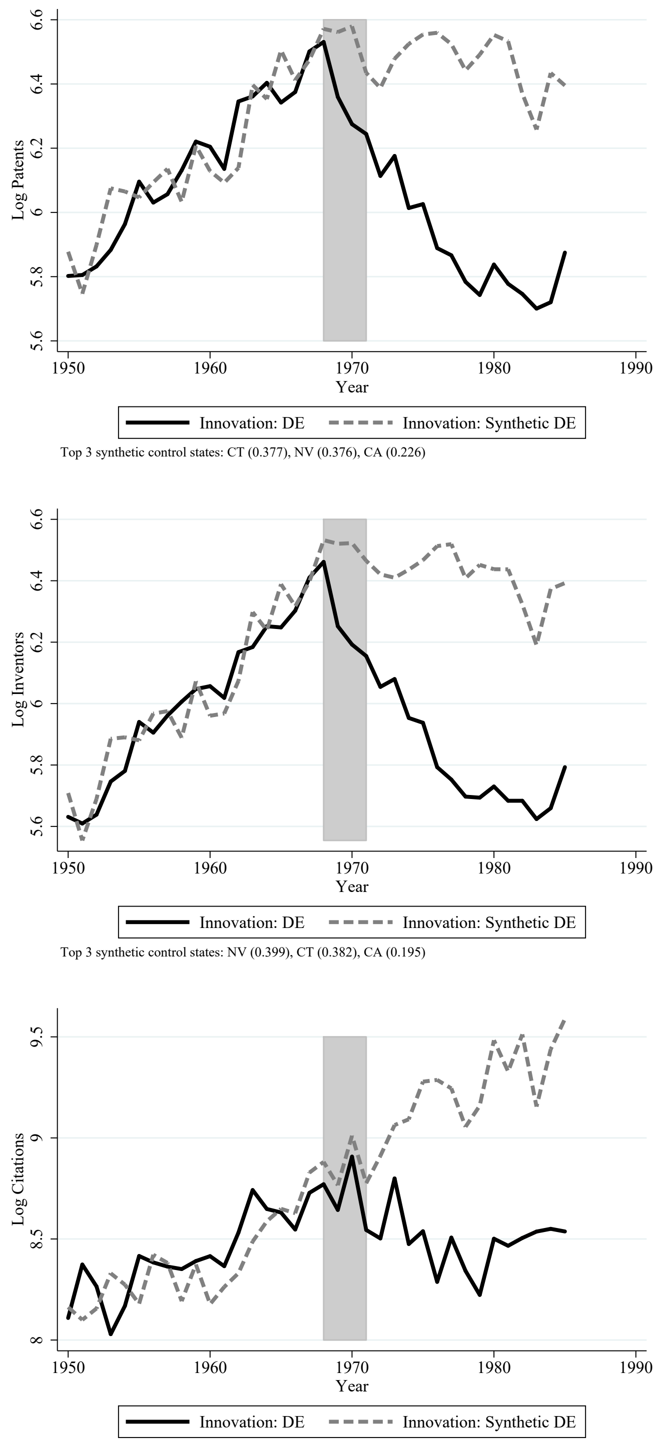

Notes: Figure plots synthetic control analyses around Delaware's tax reforms. In July 1969, the corporate tax rate increased from 5% to 6%, and in August 1971 a temporary surcharge of 20% was added on top of the 6% corporate tax rate. In 1970, the personal tax rate increased from 11% to 18%. The first row shows the patterns for log patents, the middle row for log inventors, and the bottom row for log citations.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic control methods for comparative case studies: Estimating the effect of california's tobacco control program", *Journal of the American Statistical Association*, *105*(490) (2010), 493–505.

Akcigit, Ufuk, John Grigsby, and Tom Nicholas, "The rise of american ingenuity: Innovation and inventors of the golden age". NBER Working Paper No. 23047, 2017.

Auerbach, Alan J., "Why have corporate tax revenues declined? Another look", *CESifo Economic Studies*, *53*(2) (2007), 153–171.

Auerbach, Alan J. and James Poterba. "Tax loss carryforwards and corporate tax incentives", In Martin Feldstein, ed., *Tax Policy and the Economy, Vol. 3*, 305–342. (National Bureau of Economic Research, 1987).

Bakija, Jon, "Documentation for a comprehensive historical U.S federal and state income tax calculator program". Williams College Working Paper, 2006.

Chirinko, Robert S and Daniel J Wilson, "State investment tax incentives: A zero-sum game?", *Journal of Public Economics*, *92*(12) (2008), 2362–2384.

Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg, "The NBER patent citations data file: Lessons, insights and methodological tools". NBER Working Paper No. 8498, 2001.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological innovation, resource allocation, and growth", *Quarterly Journal of Economics*, *132*(2) (2017), 665–712.

Lai, Ronald, Alexander D'Amour, David M. Doolin, Guan-Cheng Li, Ye Sun, Vetle Torvik, Amy Yu, and Lee Fleming, "Disambiguation and co-authorship networks of the U.S. patent inventor database", *Research Policy*, *43*(6) (2014), 941–955.

Lutz, Harley L, "The progress of state income taxation since 1911", *The American Economic Review*, *10*(1) (1920), 66–91.

Suárez Serrato, Juan Carlos and Owen Zidar, "The structure of state corporate taxation and its impact on state tax revenues and economic activity", *Journal of Public Economics*, *167* (2018), 158–176.