The Mainstreaming of Marx: Measuring the Effect of the Russian Revolution on Karl Marx's Influence*

Phillip W. Magness† Michael Makovi‡

May 13, 2020

Abstract

Today, Karl Marx is considered one of the preeminent social scientists of the last two centuries, and ranks among the most frequently assigned authors in university syllabi. However in Marx's time, many competing sociological traditions and socialist political movements espoused similar ideas from different origin points. How did Marx emerge as preeminent? We hypothesize that the 1917 Russian Revolution is responsible for elevating Marx's fame and intellectual following above his contemporary competitors. Using the synthetic control method and Google Ngram data, we construct a synthetic counterfactual for Marx's citation patterns. This allows us to predict how often Marx would have been cited if the Russian Revolution had not happened. We find a significant treatment effect, meaning that Marx's intellectual influence may be partly due to political accidents.

Keywords: Marx, Russian Revolution, Russian Civil War, socialism, synthetic control

*We thank Daniel J. Wilson for discussing the harmonic mean p-value and its limitations, and Kevin Grier for detailed comments on an early draft. Drafts of this paper were presented at the 2018 Arts and Humanities Graduate Student Research Conference, 15 Nov. 2018 at Texas Tech University in Lubbock, Texas; at the 2019 conference of the Association of Private Enterprise Education (APEE), 5-8 April 2019, in Nassau, Bahamas; and at the internal brown bag seminar of the Free Market Institute at Texas Tech University, 28 Jan. 2020. We thank all the attendees at these conferences for their comments and suggestions.

†American Institute for Economic Research

‡Department of Agricultural and Applied Economics, Texas Tech University; Free Market Institute at Texas Tech University; soon to be Northwood University.

JEL Codes: B14, B24, B31, B51, Z10

1 Introduction

One hundred thirty-six years after his death, Karl Marx enjoys immense stature in multiple academic fields. He ranks among the primary formative figures in sociology, enjoys immense influence as a political theorist, and is a wellspring of several schools of thought in both history and literary criticism. Moving beyond more doctrinaire expressions of philosophical Marxism, his broader reach may also be seen in the analytical frameworks of critical theory, postcolonial theory, and several branches of textual deconstruction, hermeneutic analysis, and cultural studies from the mid-twentieth century to the present day.

With some irony, Marx's influence in economics remains an exception to this trend, outside of a few small heterodox schools or general study of his contributions to the history of economic thought. As measured through the other social sciences and many humanities though, Marx ranks among the most widely referenced thinkers in the academy. This reputation was succinctly summarized in a widely quoted passage by David McLellan from the *Blackwell Encyclopedia of Political Thought*: "Over the whole range of the social sciences, Marx has proved probably the most influential figure of the twentieth century" (Miller 1987, 322).

Empirical evidence of Marx's intellectual reach points strongly to his influence in American universities. At 3,856 classes, Marx's *Communist Manifesto* was one of the most frequently assigned texts on American college syllabi in 2015. Among other classics of the intellectual canon, only Plato's *Republic* (3,573) reached a comparable number. Marx's writings appeared at two or more times the frequency of principal works by thinkers including John Stuart Mill (1,969), Charles Darwin (1,701), Adam Smith (1,587), and Martin Luther King Jr. (1,985). While Marx's more sophisticated *Capital* fell below the comparatively accessible *Manifesto*, at 1,798 classes it was still assigned at greater frequency than not only Smith but also Jean-Jacques Rousseau's *Social Contract* (1,427), John Locke's *Second Treatise* (1,045), or John Rawls's *Theory of Justice* (1,248).[1] Marx's influence upon the scholarship of academics who assign his works in class is trickier to measure. Survey data also

[1] Figures calculated from the Open Syllabus Project, October 2015. As of February 2020, Marx's *Manifesto* (7,057) maintained near parity with Plato's *Republic* (7,088), with no other philosophical works in their vicinity. The few texts that surpassed Marx's count included a commonly used calculus textbook, a grammar manual by Diana Hacker, and the ubiquitous Strunk and White writing guide.

point to high levels of faculty association with Marxism in some disciplines such as sociology.[2]

Other attestations including an outpouring of academic commemorations in popular outlets on the two-hundredth anniversary of Marx's birth confirm the high stature he holds in academic writing today.[3] Intellectual discussions are populated, somewhat paradoxically, by what Thomas (1991, p. 23) calls "a galaxy of different Marxisms, within which the place of Marx's own thought is ambiguous." If anything, the socialist philosopher's current legacy may be characterized as having a broad general academic influence, but many competing claimants within.

Yet just over a century ago, this was not at all the case. The historical literature on Marx's reception in his own lifetime and in the first four decades that followed his death points to only modest notice of his contributions. Although Marx enjoyed a positive reputation in the socialist and communist political movements of the late nineteenth century, these commentaries often operated on the revolutionary periphery of labor organizing and had comparatively little penetration into the academy.

The exceedingly slow dissemination of Marxist thought into the English-speaking academy in particular has been documented by several studies of his British (Willis 1977) and American (Amini 2016) reception. Even in German-language sources, Marx's political reach experienced an acknowledged lag until well after his lifetime (Steinberg 1979). Within the economics profession, this pattern reflected a general scholarly relegation of Marxist doctrines to a state of obsolescence only a few years after he first proposed them. Marx constructed his system around the labor theory of value, closely tying his conceptualization of economic exploitation to the derivation of surplus value. This relationship purports to quantify the difference between the value of a good and the compensation to the laborer who produced it. While Marx perceived a fatal and inevitable flaw in capitalism linked to this measured difference, the economics profession almost entirely abandoned the premises of its

[2]Although the total number of Marxists in the academy was estimated at about 3% in a 2007 survey, this percentage is diluted by the inclusion of faculty in disciplines with less political or economic content, such as the hard sciences. The same survey revealed a much higher concentration of self-identified Marxists in the social sciences (17.6%) and particularly in sociology (25.5%). See Gross & Simmons (2007, table 12).

[3]See, e.g., Jason Barker, "Happy Birthday, Karl Marx. You were right!" *New York Times*, April 30, 2018; Andrew Hartman, "Marx at 200: Just Getting Started." *Dissent Magazine*, May 4, 2018; Adam Tooze, "Why Karl Marx is More Relevant than Ever." *Financial Times*, May 4, 2018.

underlying theory of value.

In 1871, only four years after the publication of *Capital*, economists William Stanley Jevons and Carl Menger almost simultaneously proposed an alternative solution to the theory of value. Value derives from subjective preferences as realized at the margin of an economic decision, rather than by aggregating the steps of labor in total. Although Marx does not appear to have engaged the unfolding "marginal revolution" in his lifetime, it quickly overtook his theory in the following decades. Philip Wicksteed (1885) published what is arguably the first marginalist critique of Marx just two years after his death. Eugen von Boehm-Bawerk extended Menger's theory of value into a book-length critique of Marx's system in 1896 (translated into English in 1898).

Marxist attacks on subjective valuation persisted beyond these early works, yet they never caught hold in the mainstream of an economics profession that increasingly embraced marginalism. The low esteem for Marx's *Capital* at the turn of the century was succinctly captured by C. Violet Butler's dismissive assessment, in the *Economic Journal* in 1907, of a new English translation: "Who should tilt at such a windmill?" By 1926, no less a source than John Maynard Keynes would describe Marx's *Capital* as "an obsolete economic textbook . . . without interest or application for the modern world." As Alan Ryan observed in his preface to Isaiah Berlin's (2013, p. xxvi) classic study of the thinker, "Marx's economics were not taken seriously other than on the Marxist Left" in the early twentieth century.

That Marx had, at best, only modest academic influence may also be seen in journal citation patterns of his work between his death and the First World War. Table 1 below examines Marx's citation patterns across multiple scholarly journals in the humanities and social sciences during this period. While not nonexistent, use of Marx's work lagged far behind not only earlier economic thinkers such as Smith but also his nineteenth-century contemporaries in overlapping areas of knowledge, John Stuart Mill and Herbert Spencer. Instead, Marx enjoyed near parity with another contemporary, Henry George, who was at times politically popular but also academically shunned.[4]

[4]The self-taught George experienced a notoriously tumultuous relationship with academic economists during the professionalization of the discipline in the late nineteenth century. Although his writings attained popular and political acclaim, he became a primary target of ire from the founding generation of the American Economics Association

In this study, we ask what elevated Marx from his turn-of-the-century position of only modest scholarly influence to a major figure in the philosophical canon only a few decades later. The answer appears to be intimately connected to a single event: the Bolshevik revolution and successful seizure of control of the Russian state by a Marxist movement in 1917.

This chain of events in 1917 occurred rapidly, although also with sudden and unpredictable turns as the political climate in Russia deteriorated. It took only seven months from the protest-instigated abdication of Tsar Nicholas II in March to the Bolshevik seizure of the successor provisional government. The October Revolution (technically November 7th on the western calendar) both overturned the competitor socialist government of Alexander Karensky and enabled the Bolsheviks to expel other radical parties and factions, including non-Marxist-Leninist socialists that diverged from their ideological vision of proletarian revolution.[5] The events plunged the country into almost five years of civil war, with Lenin's Bolsheviks eventually defeating the anti-communist White movement as the latter suffered a succession of military setbacks and political dysfunction. While Marxist ideology provided a pronounced influence upon the Bolshevik movement, military and political outcomes - including both the missteps of its many opponents and a fair amount of luck - determined its success in seizing control of the Russian state.

The Bolshevik political ascendance also drew widespread attention to Karl Marx's name and philosophical system - particularly as the western press sought to contextualize the unfolding revolution. Newspapers that seldom even noticed Marx prior to these events rushed to explain its obscure philosophical underpinnings to their readers. Other western publications that boasted closer ties to the political left or to labor movements revisited longstanding rifts in socialist politics, drawing scrutiny to Lenin's claims of being Marx's intellectual heir.[6] For many observers abroad, the

culminating in a prolonged public dispute with Francis Amasa Walker (Samuels 1983). George was aware of and generally antagonistic toward Marx's own doctrine, yet before the Soviet revolution George's economic philosophy exercised considerable influence on socialist politics. These influences included instances of Marx's comparative absence, as with the case of Sun Yat-sen's proclamation of a socialist republic in China in 1912 while citing explicit Georgist inspiration (Trescott 1994).

[5]The "February" Revolution was in March and the "October" Revolution was in November because of the disparity between "Old Style" (Julian) and "New Style" (Gregorian) dates. Russia did not adopt the Gregorian calendar until 1918.

[6]For examples see "The Ideas of the Russian Extremists," *Baltimore Sun* 4 November 1917; "Lenine by own word

once-unfamiliar Marx became a clue to understanding the "Bolshevik threat," particularly as both rumors and actual attempts to instigate similar upheavals in Europe spawned a Red Scare and associated political backlash.

In addition to effecting a revolutionary upheaval on explicitly Marxist grounds, V. I. Lenin's political rise provided a sizable boost to the intellectual study of Marx's doctrine and by extension the scholarly appreciation of Marxism. The Soviet state itself aggressively promoted Marxist economic theory through its educational and research institutions and its publishing houses. For example, in 1919, the Soviet state established the Marx-Engels Institute, directed by David Riazanov (Ryazanov, Rjazanov).[7] In conjunction with the Frankfurt School of Social Research, Riazanov and the Marx-Engels Institute published twelve volumes (of forty-two planned) of the *Marx-Engels-Gesamtausgabe* ("MEGA^1") in German (Levine & Rojahn 2002, p. 28; Bellofiore & Fineschi 2009, pp. 2f.; Fineschi 2010). The first eight volumes were published from 1927 to 1932 in Frankfurt and Berlin, and the last four volumes were published in the Soviet Union because of Hitler's rise to power (Levine & Rojahn 2002, p. 28).[8] This "MEGA^1" was the first to publish several of Marx's works, including the *Economic and Philosophical Manuscripts of 1844* in 1930, the *Criticism of Hegelian State Law* (*Kritik des Hegelschen Staatsrechts*) in 1927, and *The German Ideology* in 1932 (Bellofiore & Fineschi 2009, pp. 2f.; Fineschi 2010).

Among the institute's other publications were an eighteen volume Russian edition of the works

brands self a despot," *New York Tribune* 30 December 1918; "Blame for Russia's tragedy shared by the socialists" *Boston Globe* 2 June 1918; "The Philosophy of the Bolsheviks," *Manchester Guardian* 23 January 1918

[7]In February 1931, Riazanov was exiled following the Menshevik Trial - one of Stalin's purge trials (Barber 1981, p. 122; Fineschi 2010; Levine & Rojahn 2002). Riazanov was executed in 1938 (Levine & Rojahn 2002, p. 28).

[8]The contents of the twelve published volumes are summarized by Bellofiore & Fineschi (2009, pp. 10f., table 1.1). We have also personally examined the title pages and tables of contents of these volumes. The first several volumes - beginning in 1927 - indicate "im auftrage des [on behalf of the] Marx-Engel-Instituts Moskau" and "herausgegeben [edited by] von D. Rjazanov" in the middles of the title pages; as well as either "Marx-Engels-Archiv Verlagsgesellschaft M. B. H. [publishing company] Frankfurt A. M." or else "Marx-Engels-Verglag G. M. B. H. Berlin" at the bottoms of the title pages. In 1932, "im auftrage des Marx-Engels-Instituts" in the middle of the title page is replaced with "im auftrage des Marx-Engels-Lenin-Instituts" and "Rjazanov" is replaced by "V. Adoratskij" but the publisher at the bottom of the page remains in Berlin. In 1935, the publisher becomes "Marx Engels Verlag Moskau."

of Marx and Engels and two journals: *Arkhiv Karla Marksa i Friderikha Engel'sa* ("Archive of Karl Marx and Frederick Engels") and *Letopis' marksizma* ("Chronicle of Marxism") (Barber 1981, p. 16.). In November 1931, the Marx-Engels Institute was merged with the Lenin Institute to form the Marx–Engels–Lenin Institute (Barber 1981, p. 122). Thus, the early life of the Marx-Engels Institute - 1919 through 1931 - coincides very closely with our statistical treatment window, 1917-1932.

Later, the Soviet state became the primary translator and publisher of Marx's own works through the government-funded Progress Publishers, founded in 1931. Marx played a similarly prominent role in Soviet propaganda through artwork and statuary, dating to Lenin's personal direction (Brown & Taylor 1993). Indeed, Lenin initiated the practice of pilgrimage to Marx's grave in 1903 and made the first of several unsuccessful Soviet attempts to have the philosopher's remains relocated to Moscow as an object of veneration. While other factors certainly contributed to the further evolution of Marx's thought in the mid-twentieth century, including notably the flight of the German-speaking academic Left to Britain and North America in the face of Nazi persecution, the crucial factor in the elevation of Marx's intellectual stature appears to arise from the events of 1917.

To investigate this observation, we obtain yearly citation patterns of Marx's work in print materials as derived from the Google Ngram database of historical print publications. These data allow us to map long-term patterns of intellectual engagement with Marx's ideas between his lifetime and the present day. Our hypothesis holds that Marx was an occasionally acknowledged but relatively minor intellectual figure between his death and the events of 1917. After 1917, Marx's stature quickly rose to the preeminent level of influence. If, per our hypothesis, changing citation patterns can be conclusively linked to the Bolshevik upheaval, it follows that the subsequent scholarly "mainstreaming" of Marx cannot be attributed solely to growing intellectual validation of his ideas through scholarly inquiry, but it owes much to a sudden and chance outcome in geopolitical history.

To test our hypothesis, we use the synthetic control method (SCM) to construct a projected citation pattern for Marx's work relative to its pre-1917 trajectory. SCM allows us to project a post-1917 citation pattern for a synthetic Karl Marx that is constructed as a composite of other authors: contemporary writers of note in the late nineteenth and early twentieth centuries and canonical thinkers in general. These authors are algorithmically weighted so that collectively, they are cited at the same rate as Marx was before the Russian Revolution. This allows us to solve the chief problem of causal inference in social science, namely the identification of a counterfactual absent treatment.

For the period after 1917, we compare the historical Marx to his synthetic counterpart to assess the divergence between his projected citation pattern absent a Soviet revolution and the actual pattern. We find that the historical Karl Marx is cited at a noticeably higher frequency than his synthetic counterpart in the immediate aftermath of the revolution. This suggests that Marx's post-1917 influence is not due only to the growing intellectual influence or intrinsic quality of his philosophical system. Instead, historical accident played a central role in elevating Marx's intellectual prominence from a secondary or tertiary figure into a preeminent thinker across several academic disciplines. This finding is consistent with observations within the historical literature on Marx and suggests that his modern intellectual prominence must be reconciled with the historical role of the Soviet Union in elevating and disseminating Marxism.

2 Data and Method

Ideally, we would estimate the effect of the Russian Revolution on Marx by observing citations in two different universes: one in which the Russian Revolution occurred and another in which it did not. Because this is not feasible, we use the synthetic control method (SCM), which is ideal for causal inference in case studies with one treated unit (Abadie & Gardeazabal 2003; Abadie *et al.* 2010, 2014; Abadie 2019).[9] SCM is a data-driven method that combines aspects of the matching and difference-in-difference techniques to facilitate counterfactual comparisons. It has been used in a variety of fields, including political science (Abadie *et al.* 2010, 2014; Grier & Maynard 2016; Geloso & Grier 2019), economic policy and growth (Billmeier & Nannicini 2013; Cavallo *et al.* 2013; Lawson *et al.* 2019), health and drug policy (Abadie *et al.* 2010; Kreif *et al.* 2016; Furton 2018), criminology (Saunders *et al.* 2014), immigration (Powell *et al.* 2017; Nowrasteh *et al.* 2019), and urban economics (Gautier *et al.* 2009). As far as we know, we are the first to use SCM to analyze

[9]We implement SCM in Stata (Abadie *et al.* 2011; Galiani & Quistorff 2017). In appendix A, we provide additional technical details about our synthetic control methodology and parameters.

literature.[10]

SCM uses a data-driven process to approximate one unit's outcome using a weighted average of the outcomes of other units, called donors. These weights are chosen to minimize the outcome's RMSPE (root mean squared prediction error) during the pretreatment period. To illustrate, suppose that prior to 1917 Marx's citations could be estimated as being equal to 0.5 times Kropotkin's plus 0.3 times Henry George's plus 0.2 times Auguste Comte's. This set of weights forms a synthetic control because it is designed to predict how often Marx should have been cited — absent treatment — based on how often he had been cited in the past. Using these same donor weights, the synthetic control's outcome is predicted after 1917. The synthetic Marx is our best estimate of the counterfactual universe in which the Russian Revolution never occurred. If a weighted average of donors (authors) reliably predicts Marx's citations before 1917 but not afterwards, then the deviation between the real Marx and the synthetic Marx after 1917 is considered a treatment effect.

Instead of requiring the researcher to arbitrarily and subjectively choose weights, SCM chooses these weights by a data-driven process. Furthermore, the weights are constrained to be non-negative and sum to one to avoid extrapolation bias. No unit's synthetic prediction can lie outside the region of common support (the convex hull). In other words, our prediction of how often Marx should have been cited after 1917 is constrained to lie within the range of other authors' real-world citations.

SCM's identification principle is that no other treatment or idiosyncratic shock intervenes which affects some units differently than others (Abadie 2019, s.v. "Availability of a comparison group"). For example, SCM would not be biased by the invention of a new publication process in the middle of the pretreatment period as long as this new publication process equally affected every author. Only unit-specific treatments would introduce bias—for example, if this new publication process were available only to some authors and not others. Thus, like difference-in-difference, SCM is not biased by time-varying unobservables as long as these unobservables are not systematically

[10] Cf. Grundmann & Stehr (2001, p. 272), who use a method resembling difference-in-difference, comparing the number of citations for Werner Sombart and Martin Heidegger, before and after World War II, to reject the claim that Sombart lost his popularity in sociology due to his association with Nazism. Gentzkow *et al.* (2019), a review of "text as data," briefly discuss n-grams (p. 539), but do not mention the synthetic control method.

correlated with specific observational units.

SCM includes elements of matching as well because the weights are chosen not only to minimize the RMSPE (i.e., the outcome's prediction error), but also to achieve balance on indicator variables that are believed to affect the outcome. Thus, Marx's synthetic control will be composed of authors who not only predict Marx's citation count, but who also share other attributes with Marx. For example, we selected and investigated contemporary writers who broadly fall into categories of (a) contemporary economic, sociology, and philosophy writers from the time of Marx and (b) contemporary socialist and other revolutionary thinkers. Donor weights and indicator weights are then chosen to both minimize RMSPE and minimize indicator imbalance.[11]

Our donor list is compiled from three sources: first, we brainstormed a list of the most important and relevant political, social, or socialist thinkers we could think of. This list included 50 authors. Second, we consulted two primary source readers in political philosophy - viz. Rosen *et al.* (1999) and Cohen (2018) - and added any authors from Marx's era or earlier. Many of the authors in these books were already included in our original 50 authors, and so these two books contributed 19 authors to our list.[12] Finally, we added all authors from the first 20 volumes of the *Harvard Universal Classics* (Eliot 1909), also known as *Dr. Eliot's Five Foot Shelf*. This added 30 authors to our list.[13] Our final author list includes 98 authors, including Karl Marx himself. Therefore, there is a maximum of 97 potential donors to the synthetic Marx, and a maximum of 97 possible

[11]Concerning the iterative process of choosing two sets of weights to minimize two functions, see Grier & Maynard (2016, p. 3, n. 8).

[12]From these two readers, we excluded all authors who post-date Marx as well as the following authors: Pericles and Frederick the Great because it is not clear which books they authored, which means we cannot code years of publication and translation; Christine de Pizan because she receives outcomes of zero in almost all years; Gerrard Winstanley because he receives outcomes of zero in almost all years until 1905; Lenin and Engels because they are clearly correlated with Marx and subject to the same treatment effect; G. W. F. Hegel because it is difficult to distinguish his name from the historian Friedrich Wilhelm Karl, Ritter von Hegel (1813–1901); and Publius and Brutus because these are pseudonyms whose identification is difficult using Google Ngram.

[13]We restricted ourselves to the first 20 volumes because by this point, our sample was so large that our Stata script nearly exhausted the computational resources of any available desktop workstation, nearly requiring the use of a supercomputing cluster. From the first 20 volumes, we omitted the following authors: the *Thousand and One Nights* (i.e. *The Arabian Nights*) because the author's identity is unknown; Richard Brinsley Sheridan because he had a grandson and politician of the same name; and the Brothers Grimm because it is difficult to distinguish their

placebo (permutation) tests for hypothesis testing. These authors - and their associated indicator values - are listed in table 2.

Our outcome variable is citation counts from the Google Books Ngram Viewer, as will be explained shortly. Our list of indicator variables includes the following: (1) citation counts averaged over three years, taken every six years;[14] (2) year of publication of the author's primary or most notable work; (3) whether the author published in English, German, French, Greek, or Latin (five binary indicators); (4) year of translation to English (equal to the year of publication if originally published in English); (5) whether the author was a socialist; and (6) whether the author was "political." The publication and language indicators are designed to ensure that Marx is matched with authors who wrote at roughly the same time and who were roughly equal in their accessibility to English-speaking audiences. The socialist and political indicators are designed to ensure that Marx is matched with authors who wrote on similar genres and who would have been read and cited by similar people. Ideally, our synthetic Marx will be composed of socialists who wrote at the same time Marx did and whose works were available in English at approximately the same time.

Coding an author as "socialist" is self-explanatory. Examples of socialists in our dataset include (but are not limited to) Marx, Kropotkin, Proudhon, Bakunin, and a number of writers associated with pre-Marxian "utopian" socialism.[15] "Political" is meant to refer to any author whose works are often cited for their clear political implications. All socialists are political, but not all political authors are socialists. Examples of nonsocialist political authors include (but are not limited to) Adam Smith, David Ricardo, Friedrich List, Machiavelli, Aristotle, and Plato. Slightly more than half the authors in our dataset are political. Nonpolitical authors include (but are not limited to)

name(s) from the philologist Jacob Grimm.

[14]For example: citations averaged over 1914–16, citations averaged over 1908–10, etc. We average citations over several years because annual citations can be erratic and subject to sharp, temporary spikes. We seek to form a synthetic control that tracks the treated unit over the long term, ignoring temporary spikes and avoiding over-fitting. For a similar reason, we leave gaps of three years in between every set of moving three-year averages in order to avoid over-fitting. Kaul *et al.* (2018) show that including too many outcomes as indicators causes other indicators to receive too little weight.

[15]The one questionable author is Henry George. In a robustness test (not reported), we switch George's coding to nonsocialist, and the results do not meaningfully change.

philosophical, scientific, and sociological thinkers such as Charles Darwin, Auguste Comte, Emile Durkheim, and Immanuel Kant. These authors' works may have some political implications, but this is not what they are chiefly known for. For example, Kant is more often cited for his ethical philosophy than for his political liberalism. Coding authors as socialist or political is somewhat subjective, but this is unavoidable. It is important that all authors be matched with authors who are somewhat similar to themselves.[16] In table 2, we list every author and their indicator values.

Fortunately, these codings do not enter the SCM as explicit regressors. Instead, the only function of the indicators is to help match the authors under consideration with similar authors selected as potential candidates for comparison. Once the author weights are chosen, the indicators play no role in estimating the treatment effect. As long as the political/socialist codings are not so biased that an obvious match is excluded or an awful match is included, the bias should be minimal. As long as the set of authors receiving positive weights appears plausible and reasonable, the risk of bias is minimal. If the synthetic authors' citations successfully track the real authors' citations over the pretreatment period, and if the synthetic authors are plausibly similar to the real authors in other respects, then the treatment effect estimated by SCM should be reliable.

Our treatment is statistically defined as 1917, when the Russian Revolution began. In reality, the treatment developed over the course of the entire Russian Civil War, from 1917 until 1922.

[16] We considered alternative means of coding authors, but other methods were subject to bias as well. We could not use any classifications made after 1917 because this would threaten to introduce hindsight bias. For example, we considered coding whether an author's modern-day encyclopedia article mentions "Marx" or whether Marx's article mentions another author. However, this would threaten to introduce the biases of the encyclopedia article's author as a reflective exercise encompassing the effects of Marx's rise to his current state of intellectual influence. We also considered coding whether an author was a "sociologist" or "descriptivist" as opposed to a "prescriptivist." Marx and other social scientists might thereby be distinguished from the so-called "utopian socialists." However, this would, once again, introduce contemporary biases. In the early twentieth century, fields such as sociology, political philosophy, moral philosophy, and political economy had yet to be clearly distinguished. Similarly, contemporary distinctions between descriptivism and prescriptivism may not be identical to the way such distinctions would have been drawn by individuals living in the early twentieth century. Trying to code these authors according to contemporary academic classifications might introduce biases that are themselves products of political events such as the Russian Revolution. In the end, merely coding an author as "socialist" and/or "political" seemed less arbitrary than any other alternative.

Unlike panel linear regression methods, backdating a treatment date does not introduce bias in SCM (Abadie 2019, s.v. "No anticipation"). Therefore, we conservatively set the treatment date to its earliest possible date, in 1917, knowing that there is no biased estimation if the treatment actually began later.

P-values are obtained from in-space placebo tests, and this process is automated by the *synth_runner* module by Galiani & Quistorff (2017). After we conduct the SCM procedure on our treatment unit—Karl Marx—the same procedure is conducted for every one of our untreated units, as if they had been treated. The p-value is equal to the proportion of untreated units whose estimated treatment effect is larger than Marx's estimated treatment effect. This corresponds to the formal definition of a p-value, which is the probability that the treatment effect would be observed if the null hypothesis is true. Intuitively, the estimated treatment effect for a treated unit should be larger than the treatment effect for an untreated unit. If more than a certain percentage (e.g., 1%, 5%, or 10%) of the untreated units experience larger treatment effects than Marx, we suspect that Marx's treatment effect is due to random chance.

Typically, these p-values are then standardized. Each treatment effect is divided by that unit's pretreatment RMSPE. Intuitively, if a unit was subject to large pretreatment prediction error, meaning that the predicted values often diverged from the true values prior to treatment, then we would place less weight on any post-treatment deviations. By contrast, if a unit had a very low pretreatment prediction error, with its true and predicted values being almost identical prior to treatment, then we would place greater weight on any post-treatment deviations. Furthermore, dividing treatment effects by RMSPE normalizes the effect sizes. Suppose Plato is cited fifty times more often than Marx. We would expect any random pretreatment prediction errors to be about fifty times larger as well. Thus, dividing treatment effects by RMSPE both normalizes effect sizes and discounts post-treatment deviations by the size of pretreatment deviations. Therefore, we report standardized p-values, not ordinary p-values. In robustness tests, we explicitly normalize our effects by normalizing all authors' outcomes to exactly one in a set of arbitrary years.

However, even these standardized p-values can be difficult to interpret because each post-treatment year has its own treatment effect and p-value. Our post-treatment period is 1917–32, and we have sixteen different p-values. It is not obvious whether overall statistical significance requires that every single p-value be significant or whether it is sufficient for merely most of them to be. Furthermore, we would expect any treatment effect to take several years to manifest for two reasons. First, a book takes several years to write and publish, and our citation counts from Ngram generally exclude periodicals that fall outside of the Google scanned library. The second reason arises from the fact that the Russian Revolution itself unfolded over several years and was not truly concluded until the end of the Russian Civil War in 1922. In 1917, it was not yet necessarily clear to the world that the Bolsheviks, claiming inspiration from Marx, would prevail. If the Russian Revolution did cause more authors to cite Marx, we would expect the treatment effect to slowly increase over time, as more people became convinced that Marx was worth writing about or investigating for intellectual reasons and as enough time elapsed for an author to write a book and publish it.

In the early years, the treatment effect should not necessarily be statistically significant, and it is not clear in which specific year the p-values should finally begin to become significant. Therefore, we chiefly rely on a single "joint post standardized" ("joint post std") p-value, which is derived from the ratio of post-treatment RMSPE to pretreatment RMSPE. Intuitively, we calculate how much worse the post-treatment fit was than the pretreatment fit. A treatment effect should cause the post-treatment period to be estimated with more error, causing the ratio of post-treatment RMSPE to pretreatment RMSPE to become greater than one. We count how many authors had a larger ratio than Marx, meaning their post-treatment period was estimated with even more relative error than Marx's. A lower joint post standardized p-value is evidence of a more statistically significant treatment effect. The year-by-year p-values are still useful, however, because we can at least expect these p-values to become generally smaller over time, and we can expect most of these p-values to be significant toward the end of the post-treatment period. Therefore, examining the year-by-year p-values helps us impose a sanity check on the joint post standardized p-value.

Note that our procedure for obtaining p-values from placebo tests constrains the set of possible p-values and imposes a stepwise distribution. For example, if Marx has the largest treatment effect, his p-value is zero. But if one author out of ninety-seven donors has a larger treatment effect than Marx, then Marx's p-value is one-ninety-eighty, or 0.0102. If two authors out of ninety-eight have larger treatment effects, his p-value is two-ninety-eighths, or 0.0204. Therefore, significance at the

1% level will be very difficult to achieve, because if even a single author achieves a larger treatment effect by random chance, the p-value is already 0.0102, which is not statistically significant at the 1% level. Hence, the 5% and 10% levels are more reasonable.

Our outcome data - i.e. citations - come from the Google Books Ngram Viewer (Google n.d.; Google Ngram Viewer Team n.d.; Michel *et al.* 2011; Lin *et al.* 2012).[17] An n-gram is a string of n one-grams, where a one-gram is a string of characters without spaces (Michel *et al.* 2011, p. 176). This tool allows the user to measure and visualize the frequency that any arbitrary phrase occurs in the digitized Google Books corpus in arbitrary years. For example, a user may search for the one-gram "Marx" or the two-gram "Karl Marx." Thus, we proxy citations by measuring how often an author's name occurs as a phrase in Google's n-gram corpora.[18]

The original Google n-gram corpora were generated by OCR (optical character recognition) in 2009 and 2012, but the corpora are continually updated as Google scans new books, and Google continues to improve the accuracy of its OCR (Google n.d.). In early 2011, Google Ngram included

[17] https://books.google.com/ngrams

[18] Ideally, we would more directly measure citations using something like journal impact factors. Unfortunately, no systematic record of citation counts exists for our period of study, 1878–1932. We considered using checkout counts from a national library such as the British Library or the Library of Congress, but it is not clear that reliable checkout data exist for our period of study. (We thank Jake Syma, associate librarian at Texas Tech University, for helping us evaluate the feasibility of this option.) Even if citation or checkout data did exist for this period, it is possible that they would suffer from their own measurement biases. For example, some records may have been lost, and these losses may affect some authors more than others. We also considered using the number of copies of each book that was printed, with the assumption that books with larger print runs were also bought and read in larger numbers. However, many books were printed by multiple printers in multiple countries, and some of these printers are no longer in business. It is not clear whether data exist for every single print run for every single book. Although checkout rates and print runs are theoretically superior to n-grams as measures of an author's popularity, the potential measurement bias is even larger, and this bias may differ from one author to another. For example, if different kinds of authors were published by different kinds of publishers in different countries, and if these publishers kept records of differing levels of quality, and if wars and other events in each country had differing effects on maintenance of these records, then measurement error of printing runs may be systematically correlated with author or genre. Furthermore, if books were printed in different countries, counts of licensed prints may be biased by different levels of enforcement of copyright protection; an accurate count of prints would have to include illegal prints as well.

5,195,769 digitized books, or approximately 4% of all books ever published (Michel *et al.* 2011, p. 176). By 2012, this had been expanded to 8,116,746 books, or about 6% of all books ever published (Lin *et al.* 2012, p. 170). Google's n-gram corpora are only a subset of its Google Books corpus. In 2011, the Google Books corpus included more than 15 million books, or about 12% of all books published (Michel *et al.* 2011, p. 176). Google chose the 2011 *n*-gram subset of 5 million books on the basis of the quality of the metadata and OCR, and periodicals were excluded by Google (Michel *et al.* 2011, p. 176). The 2012 expansion followed the same procedures (Lin *et al.* 2012, p. 170). We used the English 2012 corpus, and our procedures for obtaining these data are detailed in appendix B.

The Google Books Ngram Viewer measures *n*-grams in percentage points, as the number of instances of a given *n*-gram in a given year divided by the total number of *n*-grams that year (Michel et al. 2011: 176). This normalizes by the number of books published each year to avoid skewing results (Google n.d.). In addition, a given *n*-gram is measured only if it occurs in at least forty books, in order to reduce the dataset to a manageable size (Google n.d.). Because raw citations counts are approximately one one-millionth (or one-ten-thousandth of a percent), we multiplied all citations by one million simply to remove excess decimal zeroes. For example, a citation count of 1.0e-06 simply becomes 1.0.[19]

The Google Books Ngram Viewer is not a perfect measure of citations for several reasons. First, as noted, the Google *n*-gram corpus includes only about 12% of all books published, and it excludes periodicals. Thus, our measures do not include citations in magazines, newspapers, and so on. Second, we do not measure citations per se, but only occurrences of phrases. For example, if a book includes the phrase "Das Kapital" or "the Communist Manifesto" without mentioning Marx's name, we do not include this. Third, if an author's name can be rendered or spelled in multiple ways, we count only one rendering or spelling.

This limitation on spelling is important because some authors have names that are difficult to identify, forcing us to rely on a form that is more easily identifiable but not always consistent

[19]Note that this merely makes the graphs easier to read by removing excess decimal places. Our results are based on comparing authors' levels to one another, so multiplying all levels by one million has no more substantive effect than if one converted all inches to miles in order to reduce the number of digits in a graph.

among authors. For example, Adam Smith's last name is too common by itself, so we counted "Adam Smith" instead. By contrast, it is uncertain whether Kropotkin should be called "Pyotr Alexeyevich Kropotkin" or just "Pyotr Kropotkin," so we simply counted "Kropotkin." In general, for every author, we chose the name we felt was most easily identifiable and least likely to be conflated with another famous individual with the same last name. Unfortunately, this means that some authors are identified by last name only, while others are identified by first and last names. This may introduce bias because if, in reality, authors were being cited by last name, then counting only "firstname lastname" will undercount their citations, especially compared to authors whom we count by "lastname" only.[20]

Some authors' names were so difficult to reliably isolate that we simply could not include them in the dataset at all. For example, "Hegel" may refer to the philosopher Georg Wilhelm Friedrich Hegel (1770–1831) or to the historian Friedrich Wilhelm Karl, Ritter von Hegel (1813–1901). These two authors have such similar names that we cannot use "Friedrich Hegel" or "Wilhelm Hegel" either. Nor is it clear whether more people would have referred to the philosopher as "Georg Hegel" or the full name "Georg Wilhelm Friedrich Hegel." Thus, we were forced to exclude Hegel from our dataset entirely. Similarly, "Claude Henri de Rouvroy, comte de Saint-Simon" is variously known as "Henri de Saint-Simon," "Comte de Saint-Simon," "St. Simon," and other monikers. And there are at least ten different geographical regions of France and Canada named "Saint-Simon." Therefore, we could not include the Comte de Saint-Simon in our dataset either.

In the face of these multiple sources of measurement bias, our identification strategy is the assumption that Google n-grams are reliable for identifying relatively rates of change over time. Suppose, for example, that in reality, Karl Marx was cited six times as often as the phrase "Karl Marx" occurred in Google's n-gram corpus. Nevertheless, if a weighted average of several authors' names reliably predicts outcomes for the name "Karl Marx" prior to 1917 but not after 1917, we consider this evidence of a treatment effect affecting the relative rate of change over time, even though we cannot identify absolute levels.

Because we are comparing authors to each other at given points in time, our procedure also controls for unobservable time-varying confounders that affect all authors equally. For example,

[20]E.g., "As Mr. Smith says in *The Wealth of Nations. . .*" would *not* be counted by us.

Pechenick *et al.* (2015) question the validity of using Google's n-grams because scientific and academic texts have constituted an increasing proportion of Google's corpus over time. However, as long as any author is as likely to be affected by this as any other, our results should not be biased. Like difference-in-difference, SCM is not biased by time-varying unobservables as long as these are uncorrelated with individual units.[21] Furthermore, SCM is even robust to some forms of author-specific time-varying unobservables. If a set of donor units reliably predicts outcomes for the treated unit for a long period of time prior to treatment, then it is likely that these donors share unobservable traits with the treated unit, even if we have little idea what these traits may be. Although SCM cannot guarantee control for all time-varying unobservables, SCM is nevertheless remarkable for its ability to control for time-varying unobservables to a greater degree than most statistical methods. Assuming perfect balance on observed indicator variables, the bias from unobserved shocks decreases as the unobserved shocks become smaller in magnitude and as the preintervention time period becomes longer (Abadie 2019, s.v. "Bias Bound").

The chief advantage of Google Ngram is that it is internally valid. Even though Google Books does not include every book ever published, it is likely that for any book that Google has scanned, the phrase "Karl Marx" is as likely to be correctly recognized by OCR as the phrase "Adam Smith." And we have no reason to believe that Google censored its choice of books in a way that systematically favors, say, John Locke over Karl Marx. Thus, any measurement error is likely to be random and uncorrelated with author. While other measures might be more externally valid, they are less internally valid or are subject to more systematic measurement bias. Using Google Ngram ensures that we maximize our ability to identify relative rates of change with internal validity. Our identification strategy will assume that Google Ngram accurately measures relative changes over time even if it cannot identify absolute levels. But in order to ensure that our results are not being biased by mismeasurement of levels, we conduct robustness tests in which we normalize all authors to have a citation count of precisely one in a set of arbitrary years. This will ensure that our results

[21] Some researchers have pointed out other forms of potential bias from using Google's n-grams. For example, Google counts only the frequency at which a phrase occurred in print, not how often that book was read, discussed, bought, or printed (Pechenick *et al.* 2015). We argue, however, that this bias is random - that is, not correlated with any specific authors.

are being driven by changes in relative rates of growth and not by absolute levels.

3 Results

3.1 Primary Results

In figure 1, we graphically depict the results of our primary SCM regression, in which the treated unit is "Karl Marx," all other authors are donors (controls), the pretreatment period is 1878–1916, and the (post-)treatment period is 1917–32. The solid line is the measure of outcomes (citations) for the real-world Karl Marx, while the dotted line is the synthetic Karl Marx, designed to track the real Karl Marx over the pretreatment period. Visually, we can see that the dotted and solid lines are very similar during the pretreatment period. However, after 1917, they begin to diverge. This indicates a treatment effect caused by the Russian Revolution.

During nearly the entire postreatment period, the real Marx's citations exceed those of the synthetic Marx. Furthermore, it is worth observing that the outcomes of the real Marx are somewhat erratic in a manner that is consistent with the history of the Russian Revolution. The real Marx's citations begin to increase in 1917, temporarily reach a peak in 1921, then decline until 1923, before increasing almost monotonically 1923-1932. We suggest that Marx's citations were erratic between 1917 and 1923 because this period reflects the Russian Revolution not as a single instantaneous event, but as an episode including the initial event's successful culmination and political entrenchment. The revolution itself was short—lasting from March 8–18 in 1917 (the February Revolution) to November 7–8 in 1917 (the October Revolution)—but it was followed immediately by the Russian Civil War. In the Civil War, Lenin's Red Army—the Bolsheviks—battled the loose alliance of anti-Soviet factions under the White Army. The civil war lasted from November 7, 1917, until October 25, 1922—nearly five years. During this time, there may have been uncertainty as to whether the Red Army or the White Army would prevail. Furthermore, there had been a previous, abortive socialist revolutionary movement in 1905, often referred to by Lenin as the "dress rehearsal" for the events of 1917. Observers may have wondered whether the 1917 revolution would similarly fail with a return to the *status quo ante*. This may explain why Marx's citations rise in 1917, then fall in

1921, and then rise again in 1923. The final rise beginning in 1923 matches closely the culmination of the civil war, when it became most clear that the new communist order would endure.

In table 3, we list annual outcomes for both the real Marx and the synthetic Marx. We also list the difference and the percent difference. In the percent difference, the synthetic is the base, so the percent difference is the percentage by which the actual Marx exceeds the synthetic Marx, treating the synthetic Marx as the measure of how often Marx should have been cited. We see that by 1932, the real Marx was being cited 192.146% more than the synthetic Marx, meaning that the real Marx's citations nearly tripled in the wake of the Soviets' seizure of power. At the bottom of the table, we average these outcomes over the entire period (1878–1932), over the pretreatment period (1878–1916), and over the treatment period (1917–32). We see that during the pretreatment period, the average percent difference was -1.226%, whereas in the post-treatment period, the average percent difference was 99.247%. This means that before 1917, the real and synthetic Marxes were similar, but that after 1917, the real Marx was cited almost twice as often as the synthetic. Thus, the treatment effect is somewhere between approximately 99.247% and 192.146%, depending on whether one uses the average difference during 1917-32 or whether one uses the difference in the last year, 1932.

Next, in table 4, we list the values of the indicators of the real (treated) Marx and the synthetic Marx to see whether the synthetic Marx achieved indicator balance. We wish to ensure that the pretreatment synthetic Marx resembles the real Marx in terms of causes of outcomes, not just outcomes themselves. We see that the two Marxes resemble each other in every way, including their years of publication, their years of translation, the languages in which they did and did not write, their genres (Political and Socialist), and their pretreatment outcomes (”Ngram averaged over . . .”).

Then, in table 5, we list the author composition of the synthetic Marx. We see that synthetic Marx is composed of 62.7% Ferdinand Lassalle, 20.4% Rodbertus, and small amounts of Oscar Wilde, Nietzsche, Henry George, Rousseau, Abraham Lincoln, and Plato. Lassalle and Rodbertus together comprise 83.1% of the synthetic Marx, which means that the synthetic Marx is being predominately composed by contemporary German socialists whose ideas were similar to Marx's own.

After that, in table 6, we list the p-values for hypothesis testing. The number of placebos indicates how many authors besides Marx were used to generate treatment effects. This number is slightly smaller than the full sample because some placebos fail to converge. The joint post std p-value is a normalized value, measured as "the proportion of placebos that have a ratio of post-treatment RMSPE over pretreatment RMSPE at least as large as the average ratio for the treated units."[22] If Marx's treatment effect is genuine, we should expect few placebos to have a proportion worse than his. Our joint post std p-value is exactly zero, meaning that Marx had the largest proportional prediction error in the posttreatment period.

We also list annual p-values, which are also standardized by the pretreatment RMSPE. We expected the treatment effect to develop slowly over time, so that the initial p-values would be *not* significant. Remarkably, our annual p-values are statistically significant almost from the outset, beginning in 1918. After 1918, the one subsequent year in which the effect is not statistically significant-1923-is likely due to the erratic, "spiky" nature of our data.[23]

In conclusion, we find that Marx's treatment effect is both practically large as well as statistically significant. Before 1917, the real and synthetic Marxes resemble each other on both indicators and outcomes, but after 1917, the real Marx is cited at a far greater rate than the synthetic Marx, whereas other authors' real citations do not diverge from their synthetic citations to such a great degree.

3.2 Robustness Tests

Thus, Marx appears to have a large and statistically significant treatment effect caused by the Russian Revolution and the ensuing Russian Civil War. However, we proceed to subject Marx to a series of robustness tests to ensure the result is genuine. Our robustness tests will include (1) dropping all nonsocialists from the sample, restricting the sample to socialists only; (2) similarly,

[22] Quoting the *synth_runner* help file.

[23] We have two reasons to expect the treatment effect to have developed slowly over time: first, publications take time. Google Ngram measures books but not periodicals, so any treatment effect will be delayed by the time required to write and publish a book. Second, as previously noted, the Russian Revolution was not an instantaneous event, and until at least the conclusion of the Russian Civil War in 1922, the outcome of the revolution would have been subject to uncertainty.

dropping all socialists from the sample (except Marx), restricting the sample to nonsocialists only; (3) using citations for "Marx" rather than "Karl Marx," (4) normalizing outcomes (citations) to one in a set of arbitrary years; (5) dividing the pretreatment period into separate training and validation periods; (6) an in-time placebo in which the treatment occurs in 1889; and (7) using German-language and French-language citations rather than English. We find significant results in items 1–5 but not in item 6, suggesting that our primary results show a true treatment effect and not the effect of mere random chance. Results in item 7 are significant in German but not in French, which indicates mixed support for our hypothesis, as well as a suggestion for future research.

3.2.1 Socialists Only

First, we conduct SCM using only socialists in our sample.[24] It is possible that Marx is outperforming nonsocialists but not socialists. But it is important to ensure that Marx's citations increased relative to fellow socialists, who are his most important counterfactuals. If Marx's citations increased only relative to nonsocialists but not relative to fellow socialists, we might doubt the importance or meaningfulness of the treatment effect.

Furthermore, it is important in SCM to restrict the donor pool to authors who are *prima facie* similar to the treated unit. SCM avoids extrapolation bias by restricting donor weights to be nonnegative and sum to one, but SCM cannot avoid interpolation bias if donors whose indicator values are much smaller than the treated unit's are averaged with donors whose indicator values are much larger, resulting in a synthetic control whose indicator values match the donor's, but whose donors individually are very different than the treated unit (Abadie 2019, "Availability of a comparison group). Therefore, we repeat our procedure using a sample of only socialists. This will ensure that Marx's treatment effect is being estimated with respect to those authors who are most relevant to him, and that Marx's donors are all authors who are very similar to himself.

In figure 2, we graphically depict the outcome, which is similar to what we have seen previously. In table 7, we list the p-values. Our joint-post-std p-value is exactly zero. This means that Marx

[24] Because every author in this subsample is a socialist, and because all socialists are political, we drop the Socialist and Political indicators from the SCM regression. In addition, because none of our socialist authors wrote in Greek or Latin, we drop those indicators as well.

achieved a larger treatment effect than literally every other socialist in our sample. Our annual p-values are significant in every year except 1917-19 and 1922-25, which is consistent with what we have seen previously, *viz.* that the effect becomes more significant over time, and especially after the conclusion of the Russian Civil War in 1922.

3.2.2 Nonsocialists Only

However, we are concerned about including socialists in our donor pool. The identifying principle of SCM is that no donor experienced a treatment or idiosyncratic shock (Abadie 2019, s.v. "No interference"). If the Russian Revolution affected any authors besides Marx, then the results of SCM might be invalid.

Furthermore, SCM has a second identifying principle: the Stable Unit Treatment Value Assumption (SUTVA). This assumption holds that the outcome for an untreated unit should not be affected by assignment of treatment to the treated unit. Or as Abadie (2019, s.v. "No interference") says, an intervention should not have spillover effects on untreated units. This assumption might be violated if the effect of the Russian Revolution was one of displacement—that is, Marx's citations grow relative to those of his fellow socialists because the latter's citations all decrease. Thus, it is possible that the effect of the Russian Revolution was to decrease citations for all socialists besides Marx. These would still be important treatment effects, but their interpretations would be different from what we have hypothesized thus far. If part of the total treatment effect is actually a decrease in other socialists' citations, this would mean that to some extent, the Russian Revolution did not increase Marx's absolute prominence in social science generally, but merely elevated him among socialists.

Therefore, to eliminate this potential violation of SUTVA, we repeat the SCM procedure but with only nonsocialists in our donor pool.[25] Thus, we evaluate whether Marx's citations increased relative to those of nonsocialists, who were less likely to have been directly affected by the Russian Revolution. If our treatment effect remains significant, then we can place more confidence in our hypothesis that the Russian Revolution elevated Marx in absolute importance among all social

[25] Because the Socialist indicator takes a value of zero for all nonsocialists, we drop this variable from the SCM indicator list.

scientists and not merely in relative importance among socialists.

As Abadie (2019, s.v. "No interference") notes, there is a tension between restricting the donor pool so that all donors are *prima facie* similar to the treated unit on the one hand, versus restricting the donor pool to eliminate spillovers on the other hand. Although we cannot fully satisfy both demands at once, we at least repeat SCM multiple times, each time addressing a different concern, to evaluate whether our results are sensitive or robust.

Note that by excluding all socialist authors from our donor pool, we implicitly perform another robustness test as well. In the SCM literature, it is common to test for robustness by excluding whichever donor received the greatest weight in the primary result, and/or by iteratively excluding each of the donors who received positive weight, one at a time (the so-called "leave one out" test). These tests are meant to ensure that the significant results are not being driven by only a single donor. In our case, by excluding *all* socialists from the donor pool, we implicitly perform the more common tests of excluding the donors who received positive weights.[26]

In figure 3, we graphically depict the results of the SCM procedure using only nonsocialist donors.[27] We see that the treatment effect continues to be visually apparent. In table 8, we list the p-values. Once again, our joint post std p-value is exactly zero. This means that Marx achieved a larger relative treatment effect than literally every nonsocialist author. The annual p-values are significant in every year except 1918 and 1923-24, which is similar to what we have seen previously. By using only nonsocialist donors, we mitigate the concern that the donors themselves might have experienced a treatment, violating the identifying assumptions of SCM. Furthermore, we corroborate our hypothesis that the Russian Revolution caused Marx to become more prominent among all social scientists in general, and not merely among fellow socialists in particular.

[26]In an unreported robustness test, we explicitly perform this test by dropping only the two largest donors-Ferdinand Lassalle and Rodbertus-from our sample. Results are very similar to those from the full sample. These results are available from the authors upon request.

[27]Because all our donors are nonsocialists, we drop the Socialist indicator variable from the SCM regression.

3.2.3 Citations of "Marx"

We are concerned that our results may be an artifact of our inability to precisely identify citation levels. Google Ngram Viewer only identifies phrases, and because of the idiosyncrasies of how names are spelled, we were forced to mix "lastname" with "firstname lastname" in our sample. This may advantage some authors over others because we would expect that *ceteris paribus*, "lastname" occurs more often than "firstname lastname." For example, an author cited as "Proudhon" or "Kropotkin" may be advantaged over "Karl Marx."

In theory, it is possible that the SCM procedure will not be affected by mis-measured levels as long as relative rates of change are correctly identified. If one set of donor authors predicts the treated unit's "lastname" just as well as another set of donor authors predicts the treated unit's "firstname lastname," and if the treated unit's outcomes for "firstname lastname" are proportional to the treated unit's outcomes for "lastname," then SCM may produce very similar results in both cases. In other words, if one graph is just a proportionally scaled version of the other graph, and if both graphs have equally good pretreatment prediction, then their posttreatment predictions may be very similar as well. Therefore, our identification strategy has been the assumption that as long as we can reliably identify within-author relative changes over time, we can identify a treatment effect, even if we cannot identify absolute levels. However, this assumption is untestable. Therefore, our inability to confidently identify absolute levels is concerning and calls for robustness testing.

In figure 4, we show the Google Ngram Viewer plot for "Marx" and for "Karl Marx * 6"—that is, for "Karl Marx" times six. We see that the two curves nearly overlap, meaning that "Marx" is mentioned almost exactly six times as often as "Karl Marx." This shows that while our procedure may be identifying relative rates of change, we cannot be confident that we are identifying levels. If our treatment effect is genuine, we should be able to find the same effect whether we use "Karl Marx" or "Marx." If the SCM procedure estimates a similar treatment effect even when essentially re-scale our treated unit's outcomes by a large multiple, this will help establish confidence in our assumption that identification of within-author changes over time is sufficient.

Therefore, we repeat the SCM procedure using "Marx" rather than "Karl Marx." In figure 5, we depict the SCM results graphically. We see that the real "Marx" is far more often than the synthetic "Marx" after 1917. An exception occurs in 1923, when the real and synthetic Marxes intersect. We suggest this is due to the uncertainty inherent in the Russian Civil War, which ended in 1922. In this graph, the treatment effect does not appear quite as clean and obvious as it did in previous graphs, but the real Marx nevertheless appears to generally outperform the synthetic Marx to a large degree during the posttreatment period.

In table 9, we list p-values for hypothesis testing. We see the joint post std p-value is 4.2%, which is statistically significant at the 5% level. We also see that in individual years, the results are almost always statistically significant, with notable exceptions in 1918 and 1922-24, which are similar to what we have seen previously. Two other years-1926 and 1928-fail to achieve significance as well. We suggest this is due to the erratic, "spiky" nature of our data.

In table 10, we list the synthetic author composition. We see that the majority of the weight is contributed by Nietzsche, at 53.5%. This is concerning because Nietzsche is not a socialist, and there is no clear reason why Nietzsche should be such an important counterfactual for Marx. The only socialist in the list is Rodbertus, and his weight is only 18.2%. It is self-evident that the synthetic "Marx" fails to achieve adequate indicator balance for the socialist indicator variable. Thus, we are concerned that when the treated unit is "Marx" rather than "Karl Marx," we fail to obtain a synthetic donor list which composed of what most would consider a reasonable set of counterfactuals for Marx. This casts some doubt on the validity of these results, because we wish for our synthetic Marx to resemble the real Marx not in terms of pretreatment outcomes, but also in terms of causes of those outcomes. This does not impugn the estimated treatment effect itself, but it does call into question the validity of using outcomes for "Marx" to test the robustness of the effect to different measurement of levels. In other words, these results for "Marx" do not contradict our previous results, but their value as supporting evidence is questionable. Thus, this suggests the need for additional robustness testing.

3.2.4 Normalization

To continue testing whether our primary results are robust to measuring absolute levels differently, we proceed to normalize our outcomes by setting each author's outcome (citations) to exactly one in a set of arbitrary years. Thus, each author is put on the same level as every other author, and our SCM identifies changes in relative rates of growth, not absolute levels. If Marx truly experienced a treatment effect, then this effect should persist even if we normalize his outcome. We normalize separately in every single year, run SCM, perform a placebo test, and then repeat the same procedure in the next year. For example, normalization in 1879 means that every author has all their outcomes for 1878–1932 divided by their outcome in 1879.[28] In addition, we use *synth_runner*'s "trends" parameter, which normalizes in the last pretreatment year.

Table 11 lists the joint-post-std p-value for each test. Because we have 39 pretreatment years, we have 39 normalization tests. Each year's test has its own number of placebos, its own joint post std p-value, and its own treatment effect (percent difference between treated and synthetic Marx in 1932). We see that four of thirty-nine tests are significant at the 1% level, thirty-four of thirty-nine are significant at the 5% level, and thirty-eight of thirty-nine are significant at the 10% level. Only one of thirty-nine tests is not significant, in 1897, with a p-value of 0.124. Thus, we see that our treatment effect is nearly always statistically significant when we normalize our outcomes to one in a set of arbitrary years.

In table 11, we also list the percent difference between the actual and synthetic outcomes in 1932 so that we can see whether we are getting a similar treatment effect no matter what year we normalize in. These values can be compared to the percent difference for 1932 in table 3, which was 192.146%. We see that the mean percent difference in 1932 is 122.615%, which is smaller than our primary result in table 3 — 192.146% — but still similar in meaning. A percent difference of 122.6159% means that Marx's citations more than doubled, while 192.146% implies his citations almost tripled. Although these estimates are quantitatively different, their meaningful interpretations are very similar.

We also perform meta-analysis of these thirty-nine p-values to determine their joint significance. Using *metap* by Tobias (1999), we combine the thirty-nine p-values using Fisher's (1932) method, Edgington's (1972a) additive method, and Edgington's (1972b) normal curve method.[29] These

[28] A few authors had outcomes of zero in a few years. For these authors, we replaced zero with the minimum nonzero value they had in any year as the denominator in the normalization.

[29] Fisher's test is undefined when some p-values are equal to zero, so for that test, we replace zeroes with 1/(number of placebos), which is the smallest possible nonzero p-value.

methods all assume independent tests, which is not true in our case. Nevertheless, because these tests are so common and well-known, it is worthwhile performing and reporting them.[30] We see in table 11 that the p-values of these three tests are respectively 2.635E-26, 5.924E-45, and 1.106E-24. For all practical purposes, the p-value of all thirty-seven tests together is zero.

However, the previous methods of meta-analysis assume our thirty-nine tests are independent, which is not true. Therefore, we also compute the harmonic mean p-value (HMP), a method for conducting meta-analysis of p-values that are not independent but whose covariance is unknown (Wilson 2019b,a).[31] This method is straightforward: the HMP is the harmonic mean of several ordinary p-values. The harmonic mean is undefined when some p-values are equal to zero, so we replace zeroes with 1/(number of placebos), which is the smallest possible nonzero p-value for any one SCM regression. Thus, we are conservatively reducing the joint significance. Note that this procedure makes it impossible to achieve joint significance at the 1% level, because our smallest possible nonzero p-value is $^1/_{97}$, or 0.0103. Therefore, we will test only at the 5% and 10% levels. We obtain both unweighted and weighted HMP values. Unweighted HMPs give equal weight to each joint post std p-value, while weighted HMPs use the number of placebos in each test to generate weights.

Our weighted and unweighted HMPs are both 0.020. The HMP cannot be directly interpreted, however. To interpret this value, we must either compare it to a table of critical values or convert it into an asymptotically exact p-value, which can be interpreted directly like an ordinary p-value (Wilson 2019a). According to Wilson (2019b, table 1), the critical HMP for $\alpha = 0.05$ ($L = 100$ tests) is 0.036.[32] Thus, our HMPs of 0.020 imply confidence at the 5% level. We can also compute the asymptotically exact p-value (AEP)—which, like a normal p-value, can be directly interpreted—by integrating the Landau distribution from HMP to infinity (Wilson 2019a).[33] Our unweighted and

[30] Since our tests all use the same data, they are not independent, but neither is their covariance known, because of the nonparametric method by which our p-values have been generated. Nevertheless, we can say that the effective number of tests is less than thirty-nine, because our tests are all highly correlated with one another. Therefore, these tests will over-estimate the joint significance and under-estimate the meta p-value.

[31] Wilson (2019b) builds on Good (1958), as noted by Held (2019) and Wilson (2019e).

[32] The lookup table gives critical values for L = 10 tests and for $L = 100$ tests, where a larger number of tests implies a smaller critical value. We have thirty-nine tests, so we conservatively choose $L = 100$.

[33] HMPs and AEPs are both obtained using the R package *harmonicmeanp* (Wilson 2019c). HMPs are computed

weighted AEPs are both 0.023, which is less than 5%.

However, depending on the form of covariance or dependence among tests, Wilson's HMP method may not always have the desired strength, i.e. it may have a higher type I (false-positive) rate than is desired (Goeman *et al.* 2019; Wilson 2019d). It may be advisable to combine the HMP with the Bonferroni correction and/or the Simes method (Wilson 2019d). Therefore, we also perform Simes's (1986) method, which is generally valid when tests are dependent (Samuel-Cahn 1996; Sarkar & Chang 1997; Rødland 2006).[34] Simes's method essentially penalizes small p-values by the number of tests. P-values are ordered from smallest to largest, and each p-value is multiplied by ${}^{N}/_{i}$. For example, with thirty-nine tests, the smallest p-value is multiplied by ${}^{39}/_{1}$, the second-smallest by ${}^{39}/_{2}$, etc., until the largest p-value is multiplied by ${}^{39}/_{39}$. The meta p-value is the minimum value of all these p-values. By construction, if even one p-value is equal to zero, then the meta p-value must be zero as well, since zero multiplied by ${}^{N}/_{i}$ will always be zero. Since our p-values were exactly zero in 1882, 1893, 1908, and 1913, the Simes method returns a p-value of exactly zero as well.

However, this result is somewhat paradoxical, and possibly misleading. Suppose, for example, that one had performed one million separate tests. And suppose that one p-value was zero while all the rest were 0.8. Intuitively we would expect the results to be insignificant overall. But the Simes p-value would be zero. Now, suppose that the one p-value which was zero were instead 0.01.

with the *hmp.stat* and *pharmonicmeanp* commands, while AEPs are computed with the *p.hmp* command. To run R commands within Stata, we used the *rcall* module by Haghish (2019b). To install *rcall* in Stata, we used the *github* module by Haghish (2019a). For details, see Haghish (2019c). We thank Alicia Plemmons for assistance with R.

[34]We thank Daniel J. Wilson for personally discussing the HMP with us, alerting us to Wilson (2019d), and recommending Simes's method. We eschew the Bonferroni correction because it is obviously too conservative. The Bonferroni correction multiplies each p-value by N, or equivalently, it sets the critical p-value at α divided by N. In our case, this means a critical p-value of ${}^{0.10}/_{39} = 0.0026$ at the 10% level. But our minimum possible nonzero p-value is ${}^{1}/_{97} = 0.0103$. This means that by construction, it is almost impossible for our p-values to be less than the Bonferroni-corrected critical p-value. Equivalently: when our smallest possible nonzero p = 0.0103 is multiplied by N=39, we obtain a Bonferroni-corrected p-value of 0.4017. Thus, the Bonferroni correction would virtually guarantee statistical insignificance no matter what our results are, purely by construction. This is unless we accept that p-values of zero in table 11 - in 1882, 1893, 1908, and 1913 - are literally zero, so that when multiplied by thirty-nine they are still zero. The Bonferroni correction's conservatism is well-known, as discussed by Wilson (2019b).

The first p-value of 0.01 would be multiplied by $N/i = {}^{1{,}000{,}000}/_{1}$, becoming 10,000. The smallest p-value would be the last one, viz. 0.8 multiplied by $N/i = {}^{1{,}000{,}000}/_{1{,}000{,}000}$. The overall test would be 0.8, which is not significant, as we expected. Thus, the Simes test is extremely sensitive to p-values of exactly zero, and it is not clear whether the test is valid in this case. In typical regression methods, p-values are obtained parametrically, and values of exactly zero are not possible. The non-parametric means of obtaining p-values with SCM allows p-values of exactly zero, which the Simes method was not necessarily designed to anticipate. If we replace all p-values of zero with the minimum nonzero p-value - i.e. 1/(number of placebos) - as we did for the HMP, then the Simes p-value is 0.037. Thus, the Simes method shows that our normalization tests are jointly significant at the 5% level whether or not we replace p-values of zero with 1/(number of placebos).

To summarize: the Fisher and Edgington methods all return meta p-values of almost zero, while Simes's method returns a p-value of exactly zero using the unmodified p-values. If we replace p-values of zero with the smallest possible nonzero p-values, then the Simes method returns a p-value of 0.037, which is significant at the 5% level. The harmonic mean p-value (HMP) method - which is defined only when p-values of zero are replaced with nonzero values - returns an asymptotically exact p-value of 0.023, which is significant at the 5% level as well.

Thus, our thirty-nine normalized tests are jointly significant at the 5% level according to every method of meta-analysis. This implies that our results are not being biased by our failure to identify absolute levels of citation. The practical significance of our results is nearly the same as well; our primary result - depicted in table 3 - found a percent difference of 192.146% in 1932 between the real Marx and his synthetic control, while our normalization tests - depicted in in table 11 - found a mean percent difference of 122.615%. Whether Marx's citations doubled or tripled by 1932, the meaningful interpretation is the same; the Russian Revolution caused a substantial increase in Marx's citations.

Our identification strategy has been the assumption that Google Ngram Viewer can identify changes in relative rates of growth even if we cannot accurately identify levels. By using "Marx" instead of "Karl Marx" and then by normalizing our outcomes, we guarantee that our results are determined by rates, not by levels. Our ability to maintain significance - both practical and statistical - helps us maintain our confidence in our identification strategy.

3.2.5 Training and validation

Next, we wish to make sure that we are not overfitting the pretreatment period and thereby biasing our results with idiosyncratic variations in the donor units' outcomes. To do so, we use the "cross-validation" technique from Abadie *et al.* (2014, pp. 501f.). We divide our pretreatment period - 1878-1916 - into two periods: training in 1878–98 and validation in 1899–1916. Using only indicator variables from the training period, we select variable weights which minimize the RMSPE in the validation period. These same weights are then used to predict outcomes in the posttreatment period. As Abadie *et al.* (2014, p. 502) say, "Intuitively, the cross-validation technique selects the weight v_m that minimize out-of-sample prediction errors."[35]

In figure 6, we graphically depict this procedure. We see that the synthetic Marx does not fit the pretreatment period quite as well as in previous SCM regressions, where the RMSPE is minimized for the entire pretreatment period. Nevertheless, the synthetic Marx appears visually to fit reasonably well, despite using indicators from one half of the pretreatment period to minimize RMSPE for the second half.

We also see that in the posttreatment period, the divergence between the real and synthetic Marxes becomes much larger than in the pretreatment period. In table 12, we list the p-values. Our joint-post-std p-value is 0.064, indicating significance at the 10% level. Thus, our results remain qualitatively similar and statistically significant when dividing the pretreatment period into training and validation periods for cross-validation.

However, in figure 6, there does seem to be a concerning lack of fit from 1905 onward. In 1917, when the treatment begins, the synthetic Marx is already under-performing compared to the real

[35] Klößner *et al.* (2018) show that the cross-validation technique is not well-defined. In a typical SCM regression, there is a unique set of weights which satisfies the optimization criteria. However, with the cross-validation technique, there is not necessarily a unique set of weights because of the two-stage procedure. They show that the cross-validation technique may produce different results merely by using a different software package - e.g. R instead of Stata - or by merely sorting the donors differently. Klößner *et al.* (2018) argue that this does not weaken SCM but it does cast doubt on the cross-validation technique as a robustness test. Abadie (2019, note 8) replies that the concern is simply that "researchers should aim to demonstrate that their results are not overly sensitive to particular choices of V," i.e. to different methods of estimating the variable weights matrix. Thus, lack of a unique optimum for this robustness test is not a concern.

Marx, which may suggest that the treatment began in 1905, not in 1917. We suggest the divergence is due to the abortive 1905 Russian Revolution, which we believe had a treatment effect of its own.[36]

On the other hand, table 12 shows that the p-value in 1917 is 0.681, meaning that the difference in 1917 between the synthetic and actual Marxes is not statistically significant at the moment when the treatment begins. In other words, although the synthetic and real Marx *visually* appear to have already diverged from one another by 1917 in figure 6, this difference is not statistically significant. In addition, the "avg pre rmspe p" value is 0.5, meaning that 50% of authors had a worse pretreatment fit. This means that our training-and-validation procedure has still produced a better fit for "Karl Marx" than it has for half our authors. For comparison, in our primary regression, the "avg pre rmspe p" value was 0.615.[38] Thus, our pretreatment fit may not be too bad to draw inferences from, and statistically, it may not be very much worse than the pretreatment fit in our original SCM regression, which visually looked much better.

Thus, our cross-validation procedure-in which indicator variables from the first half of the pre-treatment period are used to minimize the RMSPE for the second half of the pretreatment period-has preserved the statistical significance of our treatment effect. This shows that the effect of the 1917 revolution is robust to a different set of indicator weights. But the divergence between the real and synthetic Marxes which begins in 1905 in figure 6 is concerning. Nevertheless, this divergence is consistent with Lenin's statement that the 1905 revolution was the "dress rehearsal" of the 1917 revolution. Moreover, the effect of the 1905 revolution is statistically less significant than the effect of the 1917 revolution, which is consistent with the fact that the 1905 revolution was ulti-

[36]Lenin's pamphlet *What Is to Be Done?* was published in 1902—shortly before the 1905 revolution—and it helped precipitate the divide between the Bolsheviks and Mensheviks in 1903. In 1920, Lenin referred to the 1905 revolution as the "dress rehearsal" of the 1917 revolution, saying that without the 1905 revolution, "victory of the October Revolution in 1917 would have been impossible" (Ascher 1988, p. 1; 2014, p. 51).[37] According to Ascher (1988, p. 2; 2014, p. 51), "Bolshevism emerged as a distinct political movement during the first revolution." Furthermore, it was after 1904 that "Lenin and his followers began to adopt the strategies and tactics that became the key features of Bolshevism, distinguishing it fully from other strands of Marxism. It was also during the ferment of 1905 that the soviet, which played a central role in 1917 and for many years thereafter, was founded" (Ascher 2014, p. 51).

[38]Our tables do not report the avg pre rmspe p because it is merely a diagnostic test of pretreatment fit.

mately abortive. Overall, these results are consistent with our hypothesis that Marx experienced a treatment effect in 1917.[39]

3.2.6 In-time placebo test: 1889 as treatment

Next, we perform an in-time placebo test. If the 1917 revolution had a genuine effect, then we should fail to find a treatment effect if we set the treatment time spuriously, to a year in which we know there should not have been any effect. We set the treatment year to 1889, so that the end of the treatment period is 1904, just before the 1905 Russian Revolution. We expect the synthetic Marx to track the real Marx over both the pre- and post-treatment periods because nothing actually happened in 1889. If we do find a treatment effect in 1889, this will be evidence against the genuineness of our 1917 treatment effect.

In figure 7, we graphically depict SCM with a treatment year of 1889. We see the dotted, synthetic Marx tracks the solid, real Marx over the entire period from 1850 to 1904. This means that 1889 had no effect. In table 13, we list the p-values. The joint-post-std p-value is 0.293, which

[39]Nevertheless, the visual lack of pretreatment fit from 1905 through 1917 remains concerning. Therefore, we also ran SCM with 1905 set as the treatment year, in order to determine whether 1905 indeed constitutes a turning point for Marx. Because the 1905 revolution was later considered the precursor of the 1917 revolution, it is reasonable to suppose that 1905 has an effect too, although the effect should be smaller in 1905 than in 1917 if our hypothesis is correct. These results are not reported, but are available upon request. Using 1905 as the treatment date, we find a statistically significant treatment effect, with a joint-post-std p-value of 0.094. This suggests the 1905 Russian Revolution had a treatment effect, but it is only just barely statistically significant — meaning the evidence in support of an effect in 1905 is real but weaker than the evidence in support of an effect in 1917. Moreover, the last two annual p-values in 1915 and 1916 are 0.135 and 0.146 respectively). This suggests that the effects of the 1905 revolution were fleeting; by the time the 1917 revolution had begun, the effect of the 1905 revolution was no longer statistically significant. This may be evidence in favor of our treatment effect in 1917. The fact that the 1905 "dress rehearsal" of the 1917 revolution has an effect similar to that of the 1917 revolution itself suggests that Marxist revolutions in Russia in general all have similar effects on citations of Marx. At the same time, the evidence in support of an effect due to the 1905 revolution is weaker than the evidence for 1917. The p-values in 1915 and 1916 are greater than 0.1, suggesting that by 1917, the treatment effect had begun to dissipate. This may be because the 1905 revolution was ultimately abortive, and its effects did not last. By the time the 1917 revolution began, the real and synthetic Marx had statistically begun to resemble one another, as evidenced by a p-value of 0.146 in 1916 (in the unreported results for 1905) and a p-value of 0.681 in 1917 (in table 12).

fails to achieve any statistical significance. Thus, we find statistical significance in years in which a Marxist revolution did occur but we fail to find statistical significance in a year in which no Marxist revolution occurred. This helps us place greater confidence in our primary result for 1917.

3.2.7 German- and French-language Citations

Our primary concern has been with Marx's reception in English-speaking countries, and so we have used Ngram data from the English (2012) corpus. However, it is interesting to investigate whether Marx's reception in other countries was affected as well by the 1917 Russian Revolution. Therefore, we repeat our SCM procedure using Ngram data from the German (2012) and French (2012) corpuses.[40] We omit the "year of translation to English" indicator variable.

In figure 8, we graphically depict German citations in the period 1878-1940. We see that until 1932, the synthetic Marx greatly outperforms the real Marx, much as we saw in prior English results. We extend the visual results period to 1940 to show that after 1932-33, citations for the real Marx plummet until they equal the synthetic Marx. This neatly coincides with Adolf Hitler's rise to power as Chancellor of Germany in 1933. Nevertheless, until the rise of Hitler, the Russian Revolution appears to have greatly increased Marx's citations in German as well. In table 14, we list p-values, obtained for the period 1917-1932, not until 1940. That is, the graphical depiction in figure 8 is extended to 1940, but the hypothesis testing in table 14 is performed on the same results period as every other test (*viz.* 1917-1932). We see the joint post std p-value is 0.043, indicating that this effect is statistically significant.

In figure 9, we we graphically depict French citations in the period 1878-1932, and in table 15, we list the p-values. Graphically, we can see that there is no treatment effect, because the synthetic Marx closely tracks the real Marx not only during the pretreatment period but during the posttreatment period as well. Indeed, the joint post std p-value is 0.681.

[40]Some authors' names were changed for the purpose of obtaining Ngram data for German and French. These changes are: "Aristotle" is "Aristoteles" in German; "Plato" is "Platon" in German and in French; "Cicero" is "Cicéron" in French; "Dostoyevsky" is "Dostojewski" in German and "Dostoïevski" in French; "Thucydides" is "Thukydides" in German and "Thucydide" in French; "Augustine" is "Augustinus" in German and "Augustin" in French; "Aquinas" is "Aquin" in German and "d'Aquin" in French; "Aesop" is "Äsop" in German and "Ésope" in French; and"John Calvin" is "Johannes Calvin" in German and "Jean Calvin" in French.

Thus, the Russian Revolution appears to have greatly boosted Marx's citations in German just as it did in English - at least until 1932, when Hitler came to power. By contrast, however, the Russian Revolution appears to have had no effect in French. Future research may investigate why the Russian Revolution had so much less of an effect in French-speaking countries than in English- and German-speaking countries.

4 Conclusion and Interpretation

Our findings provide clear empirical evidence that the scholarly mainstreaming of Karl Marx is intimately connected to the events of the Russian Revolution of 1917 and subsequent political consolidation by the Soviets in its immediate aftermath. Our findings using the SCM to estimate a synthetic Marx for comparison to the real Marx's Ngram citation pattern are robust to a variety of tests and alternative specifications. Prior to treatment, the real Marx was cited approximately as often as his synthetic counterfactual. But after treatment, the real Marx was cited approximately 2 to 3 times as often. This evidence helps to explain how Marx, a relatively obscure figure in his own lifetime and an outsider to mainstream scholarly discussion for the first three decades after his death, acquired a preeminent position of scholarly influence in the twentieth century. In addition to their geopolitical implications, the events of 1917 marked a clear reversal in fortune for Marx's intellectual reputation—from an obsolete economic theorist operating on the scholarly periphery to a figure of central scholarly influence. As intellectual historians have long hinted, Marx's stature was secured by the political successes of Lenin and the Soviets.

Our findings present two important implications for the interpretation of Marx's work in the present day. First, as our empirical evidence illustrates, Marx's intellectual reputation received an important boost not only from scholarly assessment of the merits of his theories, but also from the chance shock of an external political event wherein revolutionary figures, acting in his name, seized and consolidated their control over the government of a major world power. By implication, Marx's high academic stature today might have followed a dramatically different course were it not for that chance event—perhaps comparable to the stature of George or another of his lesser-known contemporaries. Socialist political thought might have followed a trajectory arising from another figure such as Bakunin or Rodbertus, operating on the political periphery. It further raises the possibility that Marx's strong academic influence—particularly in sociology, political theory, literary criticism, and a number of derivative philosophical offshoots within the critical-theory tradition—might have never taken hold.

Second, our findings renew the challenging questions of how social scientists should grapple with Marx's stature in light of its inextricable connections to the Soviet Union's dubious historical legacy. While much of the discussion emanating from the bicentennial of Marx's birth sought to separate the discussion of his modern relevance from the brutal totalitarian track record of twentieth-century communism, the elevation of Marx's stature provided by the most significant communist political event of the same period illustrates that the two cannot be easily separated. It is insufficient to portray Soviet communism as an aberration or a deviation from true Marxist doctrine, as Marxist doctrine owes its intellectual prominence to that very same Soviet episode. In assessing whether and to what degree this historical link shapes current interpretations of Marxist doctrine, one must grapple with the historical implications of Marxism's early twentieth-century intellectual ascendance as, primarily, a Soviet project.

We offer these questions as avenues of future research into the curious position of Marx's intellectual stature today. Our findings place this line of inquiry on an empirical foundation, with potential applications to other aspects of the intellectual dissemination of Marx and contemporary figures who both paralleled and diverged from that pattern. What emerges from the data is a robust means of testing the origins of the academic mainstreaming of Marx's theories and thus grappling with the historical implications of those origins.

References

Abadie, Alberto. 2019. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*. `https://economics.mit.edu/files/17847`.

Abadie, Alberto, & Gardeazabal, Javier. 2003. The economic costs of conflict: A case study of the Basque Country. *American economic review*, **93**(1), 113–132.

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

Abadie, Alberto, Diamond, Alexis, & Hainmueller, Jens. 2011. *Synth: Stata module to implement synthetic control methods for comparative case studies.*

Abadie, Alberto, Diamond, Alexis, & Hainmueller, Jens. 2014. Comparative politics and the synthetic control method. *American Journal of Political Science*, **59**(2), 495–510.

Amini, Babak. 2016. A brief history of the dissemination and reception of Karl Marx's Capital in the United States and Britain. *World Review of Political Economy*, **7**(3), 334–349.

Ascher, Abraham. 1988. *The revolution of 1905: Russia in disarray.* Vol. 1. Stanford University Press.

Ascher, Abraham. 2014. *The Russian Revolution: A Beginner's Guide.* Oneworld Publications.

Barber, John. 1981. *Soviet historians in crisis, 1928-1932.* Springer.

Bellofiore, Riccardo, & Fineschi, Roberto. 2009. *Re-reading Marx: New Perspectives after the Critical Edition.* Palgrave Macmillan.

Berlin, Isaiah. 2013. *Karl Marx.* Princeton University Press.

Billmeier, Andreas, & Nannicini, Tommaso. 2013. Assessing economic liberalization episodes: A synthetic control approach. *Review of Economics and Statistics*, **95**(3), 983–1001.

Brown, Matthew Cullerne, & Taylor, Brandon. 1993. *Art of the Soviets: painting, sculpture, and architecture in a one-party state, 1917–1992.* Manchester University Press.

Butler, C. Violet. 1907. Karl Marx. Capital; Frederick Engels, Ernest Untermann. Capital. Vol. II. The Process of Circulation of Capital. *The Economic Journal*, **17**(68), 558–561.

Cavallo, Eduardo, Galiani, Sebastian, Noy, Ilan, & Pantano, Juan. 2013. Catastrophic natural disasters and economic growth. *Review of Economics and Statistics*, **95**(5), 1549–1561.

Cohen, Mitchell. 2018. *Princeton Readings in Political Thought: Essential Texts since Plato-Revised and Expanded Edition*. Princeton University Press.

Edgington, Eugene S. 1972a. An additive method for combining probability values from independent experiments. *The Journal of Psychology*, **80**(2), 351–363.

Edgington, Eugene S. 1972b. A normal curve method for combining probability values from independent experiments. *The Journal of Psychology*, **82**(1), 85–89.

Eliot, Charles William. 1909. *The Harvard Classics*. Vol. 50. PF Collier & son.

Fineschi, Roberto. 2010 (9). *The Marx-Engels-Gesamtausgabe (MEGA) Project (MEGA-website)*. `https://marxdialecticalstudies.blogspot.com/2010/09/marx-engels-gesamtausgabe-mega-project.html`. Accessed: 5 Feb. 2020.

Fisher, Ronald Aylmer. 1932. *Statistical methods for research workers*. London: Oliver & Boyd.

Furton, Glenn. 2018. *High Crimes? The Effect of Marijuana Legalization on Crime in Colorado*. `https://ssrn.com/abstract=3271993`.

Galiani, Sebastian, & Quistorff, Brian. 2017. The synth_runner package: Utilities to automate synthetic control estimation using synth. *The Stata Journal*, **17**(4), 834–849.

Gautier, Pieter A, Siegmann, Arjen, & Van Vuuren, Aico. 2009. Terrorism and attitudes towards minorities: The effect of the Theo van Gogh murder on house prices in Amsterdam. *Journal of Urban Economics*, **65**(2), 113–126.

Geloso, Vincent, & Grier, Kevin. 2019. *Did Separatism Hurt Quebec's Economy? The Causal Effects of the Election of the Parti Québécois in 1976*. `https://ssrn.com/abstract=3303210`.

Gentzkow, Matthew, Kelly, Bryan, & Taddy, Matt. 2019. Text as data. *Journal of Economic Literature*, **57**(3), 535–74.

Goeman, Jelle J., Rosenblatt, Jonathan D., & Nichols, Thomas E. 2019. The harmonic mean p-value: Strong versus weak control, and the assumption of independence. *Proceedings of the National Academy of Sciences*.

Good, I John. 1958. Significance tests in parallel and in series. *Journal of the American Statistical Association*, **53**(284), 799–813.

Google. n.d.. *Google Ngram Viewer*. https://books.google.com/ngrams.

Google Ngram Viewer Team. n.d.. *About Ngram Viewer*. https://books.google.com/ngrams/info.

Grier, Kevin, & Maynard, Norman. 2016. The economic consequences of Hugo Chavez: A synthetic control analysis. *Journal of Economic Behavior & Organization*, **125**, 1–21.

Gross, Neil, & Simmons, Solon. 2007. The social and political views of American professors. *In: Working Paper presented at a Harvard University Symposium on Professors and Their Politics*.

Grundmann, Reiner, & Stehr, Nico. 2001. Why is Werner Sombart not part of the core of classical sociology? From fame to (near) oblivion. *Journal of Classical Sociology*, **1**(2), 257–287.

Haghish, E. F. 2019a. *github*. Stata module. Version 2.0.0. net install github, from("https://haghish.github.io/github/").

Haghish, E. F. 2019b. *rcall*. Stata module. Version 2.4.1. github install haghish/rcall, stable.

Haghish, E.F. 2019c. Seamless interactive language interfacing between R and Stata. *The Stata Journal*, **19**(1), 61–82.

Held, Leonhard. 2019. On the Bayesian interpretation of the harmonic mean p-value. *Proceedings of the National Academy of Sciences*, **116**(13), 5855–5856.

Kaul, Ashok, Klößner, Stefan, Pfeifer, Gregor, & Schieler, Manuel. 2018. Synthetic control methods: Never use all pre-intervention outcomes together with covariates. http://www.oekonometrie.uni-saarland.de/papers/SCM_Predictors.pdf.

Keynes, John Maynard. 1931. The end of laissez-faire and A short view of Russia. *In: Essays in persuasion*.

Klößner, Stefan, Kaul, Ashok, Pfeifer, Gregor, & Schieler, Manuel. 2018. Comparative politics and the synthetic control method revisited: A note on Abadie et al.(2015). *Swiss journal of economics and statistics*, **154**(1), 11.

Kreif, Noémi, Grieve, Richard, Hangartner, Dominik, Turner, Alex James, Nikolova, Silviya, & Sutton, Matt. 2016. Examination of the synthetic control method for evaluating health policies with multiple treated units. *Health economics*, **25**(12), 1514–1528.

Lawson, Robert, Grier, Kevin, & Absher, Samuel. 2019. You say you want a (Rose) Revolution? The effects of Georgia's 2004 market reforms. *Economics of Transition and Institutional Change*, **27**(1), 301–323.

Levine, Norman, & Rojahn, Jürgen. 2002. Introduction. *Rethinking Marxism*, **14**(4), 28–29.

Lin, Yuri, Michel, Jean-Baptiste, Aiden, Erez Lieberman, Orwant, Jon, Brockman, Will, & Petrov, Slav. 2012. Syntactic annotations for the google books ngram corpus. *Pages 169–174 of: Proceedings of the ACL 2012 System Demonstrations.* Association for Computational Linguistics.

Michel, Jean-Baptiste, Shen, Yuan Kui, Aiden, Aviva Presser, Veres, Adrian, Gray, Matthew K, Pickett, Joseph P, Hoiberg, Dale, Clancy, Dan, Norvig, Peter, Orwant, Jon, *et al.* 2011. Quantitative analysis of culture using millions of digitized books. *Science*, **331**(6014), 176–182.

Nowrasteh, Alex, Forrester, Andrew C, & Blondin, Cole. 2019. How mass immigration affects countries with weak economic institutions: A natural experiment in Jordan. *The World Bank Economic Review*. in press.

Pechenick, Eitan Adam, Danforth, Christopher M, & Dodds, Peter Sheridan. 2015. Characterizing the Google Books corpus: Strong limits to inferences of socio-cultural and linguistic evolution. *PloS one*, **10**(10), e0137041.

Pipes, Richard. 1991. *The Russian Revolution*. NY: Alfred A. Knopf.

Powell, Benjamin, Clark, Jeff R, & Nowrasteh, Alex. 2017. Does mass immigration destroy institutions? 1990s Israel as a natural experiment. *Journal of Economic Behavior & Organization*, **141**, 83–95.

Rødland, Einar Andreas. 2006. Simes' procedure is 'valid on average'. *Biometrika*, **93**(3), 742–746.

Rosen, Michael, Wolff, Jonathan, & McKinnon, Catriona. 1999. *Political thought*. Oxford University Press.

Samuel-Cahn, Ester. 1996. Is the Simes improved Bonferroni procedure conservative? *Biometrika*, **83**(4), 928–933.

Sarkar, Sanat K, & Chang, Chung-Kuei. 1997. The Simes method for multiple hypothesis testing with positively dependent test statistics. *Journal of the American Statistical Association*, **92**(440), 1601–1608.

Saunders, Jessica, Lundberg, Russell, Braga, Anthony A, Ridgeway, Greg, & Miles, Jeremy. 2014. A synthetic control approach to evaluating place-based crime interventions. *Journal of Quantitative Criminology*, **31**(3), 413–434.

Simes, R John. 1986. An improved Bonferroni procedure for multiple tests of significance. *Biometrika*, **73**(3), 751–754.

Steinberg, Hans-Josef. 1979. *Sozialismus und deutsche Sozialdemokratie: zur Ideologie der Partei vor dem 1. Weltkrieg*. Dietz.

Thomas, Paul. 1991. Critical reception: Marx then and now. *The Cambridge Companion to Marx*, **1**, 23.

Tobias, Aurelio. 1999. *metap*. Stata module.

Trescott, Paul B. 1994. Henry George, Sun Yat-sen and China: More Than Land Policy Was Involved. *American Journal of Economics and Sociology*, **53**(3), 363–375.

Vega Yon, George G, & Quistorff, Brian. 2012. *Parallel: Stata module for Parallel Computing.*

Vega Yon, George G, & Quistorff, Brian. 2019. parallel: A command for parallel computing. *The Stata Journal*, **19**(3), 667–684.

Wicksteed, Philip H. 1885. The Jevonian Criticism of Marx. *To-day: monthly magazine of scientific socialism*, **3**(16), 177–179.

Willis, Kirk. 1977. The Introduction and Critical Reception of Marxist Thought in Britain, 1850–1900. *The Historical Journal*, **20**(2), 417–459.

Wilson, Daniel J. 2019a. Correction for Wilson, The harmonic mean p-value for combining dependent tests. *Proceedings of the National Academy of Sciences*, **116**(43), 21948.

Wilson, Daniel J. 2019b. The harmonic mean p-value for combining dependent tests. *Proceedings of the National Academy of Sciences*, **116**(4), 1195–1200.

Wilson, Daniel J. 2019c. *harmonicmeanp: Harmonic Mean p-Values and Model Averaging by Mean Maximum Likelihood*. `https://CRAN.R-project.org/package=harmonicmeanp`. R package version 3.0.

Wilson, Daniel J. 2019d. Reply to Goeman et al.: Trade-offs in model averaging using multilevel tests. *Proceedings of the National Academy of Sciences*.

Wilson, Daniel J. 2019e. Reply to Held: When is a harmonic mean p-value a Bayes factor? *Proceedings of the National Academy of Sciences*, **116**(13), 5857–5858.

Tables

Table 1: Journal Mentions before 1917

	Founded	Karl Marx	Herbert Spencer	Adam Smith	Henry George	John Stuart Mill
Journal of Education	1875	4	138	24	18	39
Publications of the Modern Language Association (PMLA)	1884	1	2	0	1	1
Political Science Quarterly	1886	53	47	165	70	63
Quarterly Journal of Economics	1886	18	11	95	25	34
American Journal of Psychology	1887	0	35	3	0	10
Annals of the American Academy of Political and Social Science	1890	16	34	70	15	31
The Economic Journal	1891	64	32	299	33	44
Yale Law Journal	1891	1	5	7	3	4
Journal of Political Economy	1892	10	5	59	12	25
American Historical Review	1895	5	8	24	3	4
American Journal of Sociology	1895	26	90	59	15	29
TOTAL		198	407	805	195	284

Table 2: List of authors

Name	Political	Socialist	Original Language	Year of Publication	Year of Trans. to English	Magnum Opus
Abraham Lincoln	1	0	English	1861	1861	First inaugural address
Adam Smith	1	0	English	1776	1776	Wealth of Nations
Aeschylus	0	0	Greek	-456	1663	many plays
Aesop	0	0	Greek	-546	1484	Fables
Alexander Hamilton	1	0	English	1787	1787	Federalist Papers
Aquinas	0	0	Latin	1265	1911	Summa Theologica
Aristophanes	0	0	Greek	-386	1655	many plays
Aristotle	1	0	Greek	-322	1600	Politics
Auberon Herbert	1	0	English	1885	1885	The Right and Wrong of Compulsion by the State
Auguste Comte	0	0	French	1844	1865	A General View of Positivism
Augustine	0	0	Latin	426	1871	City of God
Bakunin	1	1	French	1882	1883	God and the State
Bastiat	1	0	French	1850	1867	The Law, What is Seen and What is Unseen, Economic Sophisms
Benjamin Constant	1	0	French	1816	1819	The Liberty of Ancients Compared with that of Moderns
Benjamin Franklin	1	0	English	1793	1793	Autobiography
Bentham	1	0	English	1780	1780	The Principles of Morals and Legislation
Blanqui	1	1	French	1866	1971	Instructions for an Armed Uprising
Cervantes	0	0	Spanish	1605	1620	Don Quixote
Charles Darwin	0	0	English	1859	1859	Origin of Species
Charles Fourier	1	1	French	1808	1901	Théorie des quatre mouvements et des destinées générales
Cicero	1	0	Latin	-51	1854	Many; De Re Publica
Dante Alighieri	0	0	Italian	1320	1785	Divine Comedy
David Ricardo	1	0	English	1817	1817	On the Principles of Political Economy and Taxation
de Gouges	1	0	French	1791	1979	Declaration of the Rights of Women and the Female Citizen
Dostoyevsky	0	0	Russian	1866	1885	Crime and Punishment
Dryden	0	0	English	1677	1677	All for Love
Durkheim	0	0	French	1893	1997	De la division du travail social
Edmund Burke	1	0	English	1790	1790	Reflections on the Revolution in France
Edward Bellamy	1	1	English	1888	1888	Looking Backward
Epictetus	0	0	Greek	135	1750	Enchiridion
Euripides	0	0	Greek	-406	1781	many plays
Ferdinand Lassalle	1	1	German	1863	1884	On the Labor Problem: Lassalle's Speech on the 16th of April [1863] at a Leipzig Workers' Meeting; The Working Man's Programme: An Address
Fichte	0	0	German	1794	1868	The Science of Knowledge
Francis Bacon	0	0	English	1597	1597	Essays
Frederick Douglass	1	0	English	1845	1845	A Narrative of the Life of Frederick Douglass, an American Slave

Table 2 continued from previous page

Name	Political	Socialist	Original Language	Year of Publication	Year of Trans. to English	Magnum Opus
Friedrich List	1	0	German	1841	1909	Das Nationale System der Politischen Ökonomie
Friedrich Schiller	0	0	German	1799	1800	Wallenstein
Goethe	0	0	German	1808	1821	Faust
Hans Christian Andersen	0	0	English	1848	1848	Tales
Hegel	0	0	German	1807	1967	The Phenomenology of Spirit
Henry George	1	1	English	1879	1879	Progress and Poverty
Herbert Spencer	1	0	English	1851	1851	Social Statics
Hobbes	1	0	English	1651	1651	Leviathan
Hume	1	0	English	1763	1763	A Treatise of Human Nature
Immanuel Kant	0	0	German	1785	1895	Groundwork of the Metaphysics of Morals
Izaak Walton	0	0	English	1653	1653	The Compleat Angler
James Fitzjames Stephen	1	0	English	1862	1862	Many articles
James Madison	1	0	English	1787	1787	Federalist Papers
John Bunyan	0	0	English	1678	1678	Pilgrim's Progress
John C. Calhoun	1	0	English	1851	1851	A Disquisition on Government
John Calvin	0	0	Latin	1536	1559	Institutes of the Christian Religion
John Locke	1	0	English	1689	1689	Two Treatises
John Milton	1	0	English	1644	1644	Areopagitica
John Ruskin	1	1	English	1860	1860	Unto This Last
John Stuart Mill	1	0	English	1848	1848	Principles of Political Economy
John Woolman	1	0	English	1774	1774	The Journal of John Woolman
Jonathan Swift	0	0	English	1729	1729	A Modest Proposal
Karl Marx	1	1	German	1867	1887	Capital
Kempis	0	0	Latin	1418	1470	The Imitation of Christ
Kropotkin	1	1	English	1902	1902	Mutual Aid
Lord Acton	1	0	English	1877	1877	The History of Freedom in Antiquity/Christianity
Lord Byron	0	0	English	1819	1819	Don Juan
Ludwig Feuerbach	0	0	German	1841	1854	Essence of Christianity
Machiavelli	1	0	Italian	1532	1640	Prince
Malthus	1	0	English	1798	1798	An Essay on the Principle of Population
Marcus Aurelius	0	0	Greek	170	1634	Meditations
Marlowe	0	0	English	1589	1589	Dr. Faustus
Martin Luther	0	0	German	1517	1830	Many
Montesquieu	1	0	French	1748	1750	The Spirit of the Laws
Nietzsche	0	0	German	1883	1896	Thus Spoke Zarathustra
Oliver Goldsmith	0	0	English	1773	1773	She Stoops to Conquer
Oscar Wilde	1	1	English	1891	1891	The Soul of Man Under Socialism
Percy Bysshe Shelley	0	0	English	1819	1819	The Cenci
Plato	1	0	Greek	-400	1763	Republic

Table 2 continued from previous page

Name	Political	Socialist	Original Language	Year of Publication	Year of Trans. to English	Magnum Opus
Pliny the Younger	0	0	Latin	-113	1900	letters
Plutarch	0	0	Greek	-120	1579	Lives
Proudhon	1	1	French	1840	1890	What is Property?
Ralph Waldo Emerson	0	0	English	1841	1841	Essays
Richard Cobden	1	0	English	1835	1835	England, Ireland and America, by a Manchester Manufacturer
Robert Browning	0	0	English	1843	1843	A Blot in the 'Scutcheon
Robert Burns	0	0	English	1793	1793	various poems and songs, Scots Wha Hae, Auld Lang Syne
Robert Owen	1	1	English	1813	1813	A New View of Society
Rodbertus	1	1	German	1850	1898	Overproduction and Crises
Rousseau	1	0	French	1762	1913	Social Contract
Sismondi	1	1	French	1815	1847	Political Economy
Sophocles	0	0	Greek	-405	1649	many plays
Spinoza	0	0	Latin	1677	1856	Ethics
Thomas Browne	0	0	English	1643	1643	Religio Medici
Thomas Carlyle	1	1	English	1837	1837	The French Revolution: A History
Thomas Jefferson	1	0	English	1776	1776	Declaration of Independence
Thoreau	0	0	English	1854	1854	Walden
Thucydides	1	0	Greek	-400	1628	The Peloponnesian Wars
Tocqueville	1	0	French	1835	1835	Democracy in America
Virgil	0	0	Latin	-19	1697	Aenid
Voltaire	1	0	French	1731	1950	History of Charles XII (1731), The Age of Louis XIV (1751), and Essay on the Customs and the Spirit of the Nations (1756)
William Godwin	1	0	English	1793	1793	An Enquiry Concerning Political Justice
William Graham Sumner	1	0	English	1883	1883	What Social Classes Owe to Each Other
William Penn	1	0	English	1682	1682	Fruits of Solitude
Wollstonecraft	1	0	English	1792	1792	A Vindication of the Rights of Woman

Table 3: SCM, 1878–1932: outcomes

Year	Y-Actual	Y-Synthetic	Difference	% Difference (Base = Synthetic)
1878	0.158	0.196	-0.039	-19.738
1879	0.206	0.198	0.008	4.024
1880	0.119	0.203	-0.084	-41.399
1881	0.152	0.216	-0.064	-29.610
1882	0.122	0.204	-0.082	-40.319
1883	0.240	0.272	-0.032	-11.893
1884	0.217	0.215	0.003	1.259
1885	0.235	0.228	0.006	2.779
1886	0.359	0.398	-0.039	-9.815
1887	0.225	0.225	0.000	-0.015
1888	0.218	0.274	-0.056	-20.530
1889	0.174	0.225	-0.051	-22.752
1890	0.309	0.369	-0.060	-16.255
1891	0.612	0.326	0.286	87.598
1892	0.323	0.426	-0.103	-24.154
1893	0.364	0.447	-0.084	-18.724
1894	0.441	0.391	0.050	12.720
1895	0.343	0.358	-0.015	-4.224
1896	0.254	0.352	-0.098	-27.901
1897	0.183	0.279	-0.097	-34.544
1898	0.394	0.295	0.099	33.406
1899	0.289	0.450	-0.160	-35.680
1900	0.285	0.401	-0.115	-28.804
1901	0.362	0.361	0.001	0.152
1902	0.315	0.313	0.001	0.479
1903	0.316	0.299	0.017	5.789
1904	0.220	0.307	-0.088	-28.491
1905	0.255	0.419	-0.164	-39.084
1906	0.579	0.693	-0.114	-16.490
1907	0.667	0.494	0.174	35.188
1908	0.688	0.415	0.273	65.757
1909	0.556	0.480	0.076	15.779
1910	0.856	0.452	0.404	89.268
1911	0.679	0.601	0.078	13.032
1912	0.801	0.516	0.286	55.352
1913	0.480	0.528	-0.047	-8.949
1914	0.664	0.520	0.144	27.661
1915	0.449	0.551	-0.102	-18.484
1916	0.540	0.541	-0.001	-0.187
1917	0.449	0.637	-0.188	-29.511
1918	0.763	0.492	0.270	54.925
1919	0.967	0.411	0.556	135.325
1920	1.181	0.516	0.665	129.042
1921	1.392	0.561	0.831	148.252
1922	0.887	0.467	0.420	89.832
1923	0.675	0.481	0.194	40.255
1924	0.870	0.538	0.333	61.883
1925	0.825	0.490	0.335	68.425
1926	0.943	0.488	0.456	93.530
1927	1.215	0.489	0.725	148.220
1928	1.010	0.550	0.460	83.723
1929	1.065	0.509	0.556	109.368
1930	1.509	0.617	0.892	144.606
1931	1.580	0.725	0.855	117.934
1932	1.656	0.567	1.089	192.146
Averages				
1878-1932	0.575	0.418	0.157	28.003
1878-1916	0.376	0.370	0.005	-1.226
1917-1932	1.062	0.534	0.528	99.247

Table 4: SCM, 1878–1932: indicator balance

Indicator	Treated	Synthetic
Year of publication	1867	1851.002
Year of translation to English	1887	1887.340
Wrote in English	0	0.126
Wrote in German	1	0.854
Wrote in French	0	0.015
Wrote in Greek	0	0.005
Wrote in Latin	0	0.000
Socialist	1	0.949
Political	1	0.977
Ngram averaged over 1914-1916	0.551	0.537
Ngram averaged over 1908-1910	0.700	0.449
Ngram averaged over 1902-1904	0.284	0.306
Ngram averaged over 1896-1898	0.277	0.309
Ngram averaged over 1890-1892	0.415	0.374
Ngram averaged over 1884-1886	0.270	0.280
Ngram averaged over 1878-1880	0.161	0.199

Table 5: SCM, 1878–1932: synthetic author composition

Author	Weight
Ferdinand Lassalle	0.627
Rodbertus	0.204
Oscar Wilde	0.096
Nietzsche	0.023
Henry George	0.022
Rousseau	0.015
Abraham Lincoln	0.008
Plato	0.005

Table 6: SCM, 1878–1932: p-values

No. placebos:	96
Joint post std p:	0
Year	Std p
1917	0.219
1918	0.073
1919	0.021
1920	0
1921	0
1922	0
1923	0.188
1924	0.073
1925	0.052
1926	0.021
1927	0.021
1928	0.052
1929	0
1930	0
1931	0
1932	0

Table 7: SCM, 1878–1932, socialists only: p-values

No. placebos:	12
Joint post std p:	0
Year	Std p
1917	1
1918	0.250
1919	0.333
1920	0
1921	0
1922	0.167
1923	0.500
1924	0.333
1925	0.333
1926	0.083
1927	0
1928	0
1929	0
1930	0
1931	0
1932	0

Table 8: SCM, 1878–1932, NON-socialists only: p-values

No. placebos:	82
Joint post std p:	0
Year	Std p
1917	0.012
1918	0.122
1919	0
1920	0.012
1921	0
1922	0.024
1923	0.207
1924	0.110
1925	0.024
1926	0.073
1927	0.037
1928	0.061
1929	0.024
1930	0
1931	0
1932	0

Table 9: SCM, 1878–1932, Author = "Marx": p-values

No. placebos:	96
Joint post std p:	0.042
Year	Std p
1917	0.010
1918	0.177
1919	0.031
1920	0.042
1921	0
1922	0.427
1923	0.875
1924	0.948
1925	0.052
1926	0.229
1927	0.083
1928	0.188
1929	0.031
1930	0.010
1931	0.010
1932	0.031

Table 10: SCM, 1878–1932, Author = ”Marx”: synthetic author composition

Author	Weight
Nietzsche	0.535
Rodbertus	0.182
Rousseau	0.112
Abraham Lincoln	0.1
Machiavelli	0.039
Plato	0.022
Augustine	0.008

Table 11: Normalization

Normalization year	No. placebos	Joint post std p	Pct diff btw treated & synth - 1932
1878	96	0.052	130.440
1879	96	0.031	188.284
1880	97	0.041	53.446
1881	95	0.021	98.603
1882	96	0	112.192
1883	96	0.042	128.352
1884	96	0.021	197.873
1885	96	0.031	147.653
1886	96	0.021	104.291
1887	96	0.042	113.178
1888	96	0.010	198.293
1889	96	0.021	172.372
1890	97	0.010	130.564
1891	96	0.010	128.124
1892	95	0.021	206.685
1893	97	0	112.676
1894	97	0.021	86.986
1895	97	0.021	92.271
1896	97	0.031	153.849
1897	97	0.124	39.633
1898	97	0.031	90.999
1899	96	0.021	170.696
1900	97	0.062	83.589
1901	97	0.021	106.962
1902	97	0.031	146.722
1903	96	0.031	115.931
1904	97	0.021	145.989
1905	97	0.031	67.012
1906	97	0.010	117.450
1907	96	0.083	54.614
1908	97	0	102.693
1909	96	0.031	96.433
1910	96	0.021	154.315
1911	97	0.041	152.146
1912	97	0.021	126.698
1913	97	0	97.372
1914	96	0.010	176.205
1915	96	0.083	42.837
1916 "Trends"	97	0.010	137.544
Mean pct diff:			122.615
No. tests:			39
No. tests signif @ α <0.01:			4
No. tests signif @ α <0.05:			34
No. tests signif @ α <0.10:			38
No. tests not significant:			1
Fisher's meta-p:			2.635E-26
Edgington's additive meta-p:			5.924E-45
Edgington's normal curve meta-p:			1.106E-24
Harmonic mean p-value (HMP), unweighted:			0.020
Harmonic mean p-value (HMP), weighted:			0.020
Critical HMP for $\alpha = 0.05$, L = 100 tests:			0.036
Asymtotically exact p-value, unweighted:			0.023
Asymtotically exact p-value, weighted:			0.023
Simes's test, unmodified p-values:			0
Simes's test, zeroes replaced w/ 1/No. placebos:			0.037

Table 12: SCM, 1878–1932, cross-validation (pretreatment period divided into training and validation periods): p-values

No. placebos:	94
Joint post std p:	0.064
Year	Std p
1917	0.681
1918	0.149
1919	0.043
1920	0.149
1921	0.021
1922	0.043
1923	0.138
1924	0.106
1925	0.223
1926	0.074
1927	0.074
1928	0.213
1929	0.021
1930	0.021
1931	0.011
1932	0.085

Table 13: SCM, 1889 in-time placebo: p-values

No. placebos:	92
Joint post std p:	0.293
Year	Std p
1889	0.826
1890	0.370
1891	0.065
1892	0.315
1893	0.054
1894	0.435
1895	0.272
1896	0.500
1897	0.913
1898	0.283
1899	0.098
1900	0.120
1901	0.315
1902	0.500
1903	0.391
1904	0.467

Table 14: SCM, 1878–1932, German-language citations: p-values

No. placebos:	94
Joint post std p:	0.043
Year	Std p
1917	0.883
1918	0.032
1919	0.021
1920	0.021
1921	0.032
1922	0.011
1923	0.032
1924	0.053
1925	0.074
1926	0.181
1927	0.096
1928	0.043
1929	0.106
1930	0.074
1931	0.096
1932	0.032

Table 15: SCM, 1878–1932, French-language citations: p-values

No. placebos:	94
Joint post std p:	0.681
Year	Std p
1917	0.585
1918	0.021
1919	0.255
1920	0.628
1921	0.234
1922	0.894
1923	0.660
1924	0.766
1925	0.798
1926	0.372
1927	0.500
1928	0.234
1929	0.926
1930	0.660
1931	0.691
1932	0.979

Figures

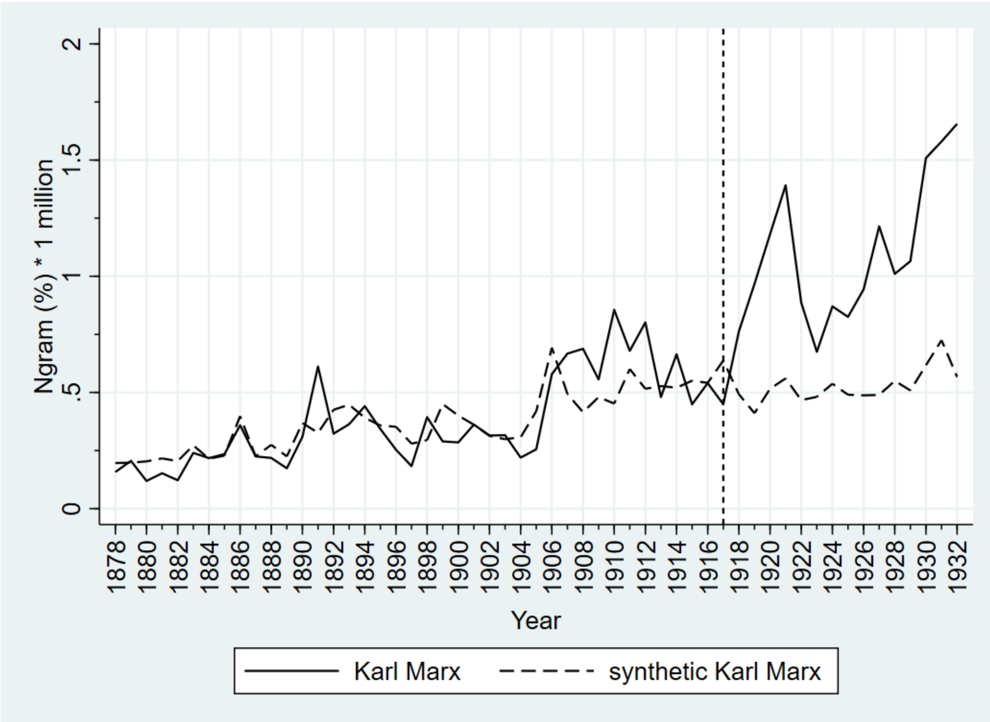

Figure 1: SCM, 1878–1932: graphical representation

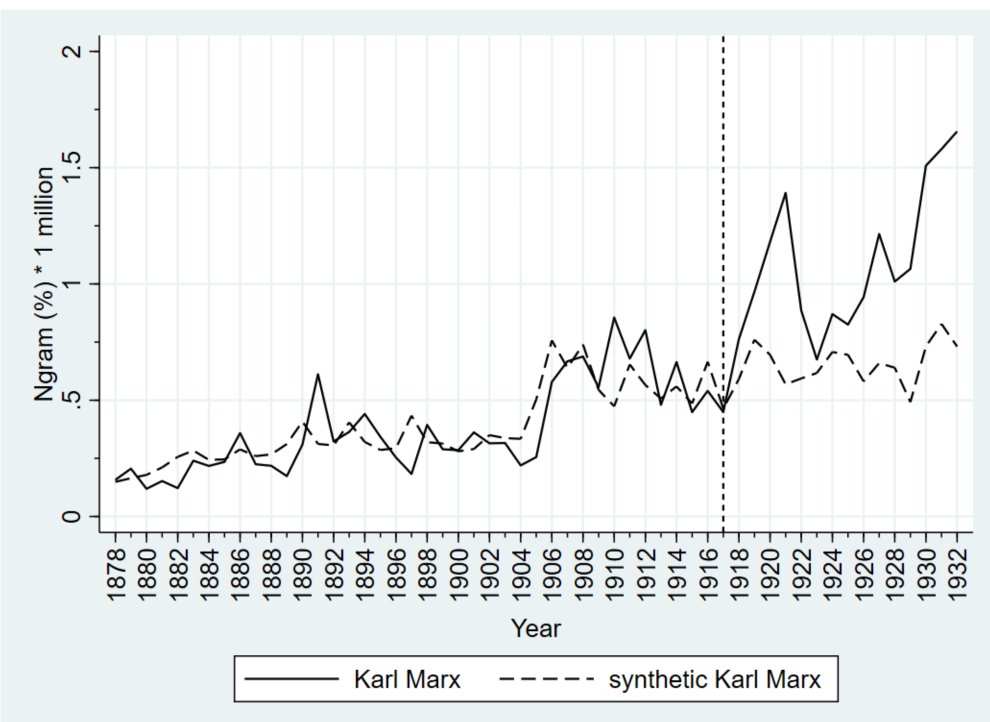

Figure 2: SCM, 1878–1932, socialists only: graphical representation

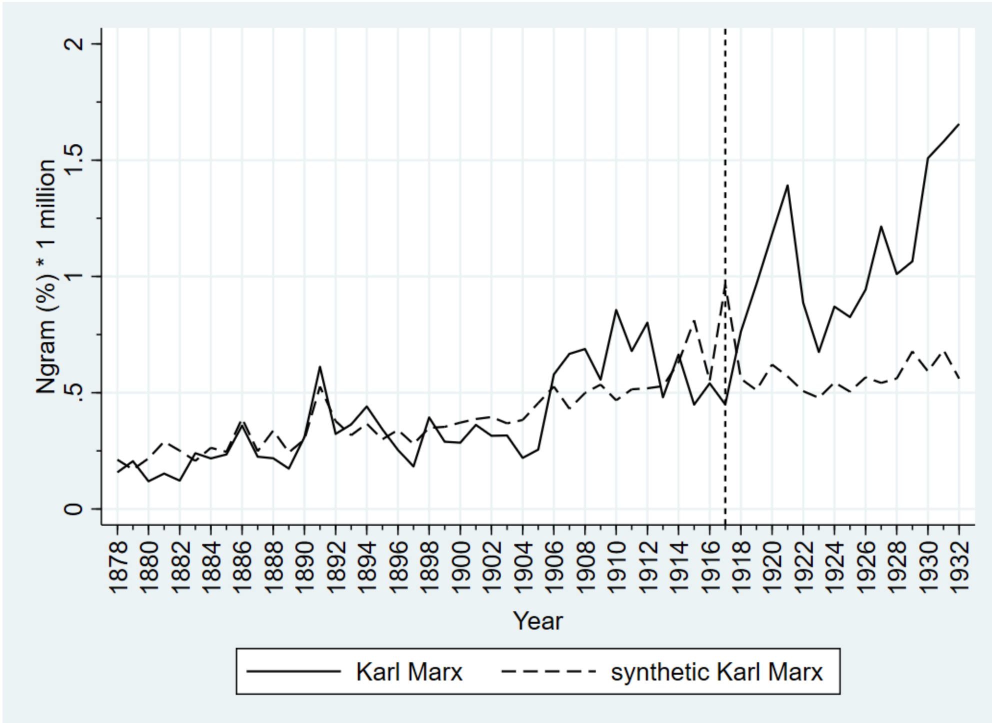

Figure 3: SCM, 1878–1932, NON-socialists only: graphical representation

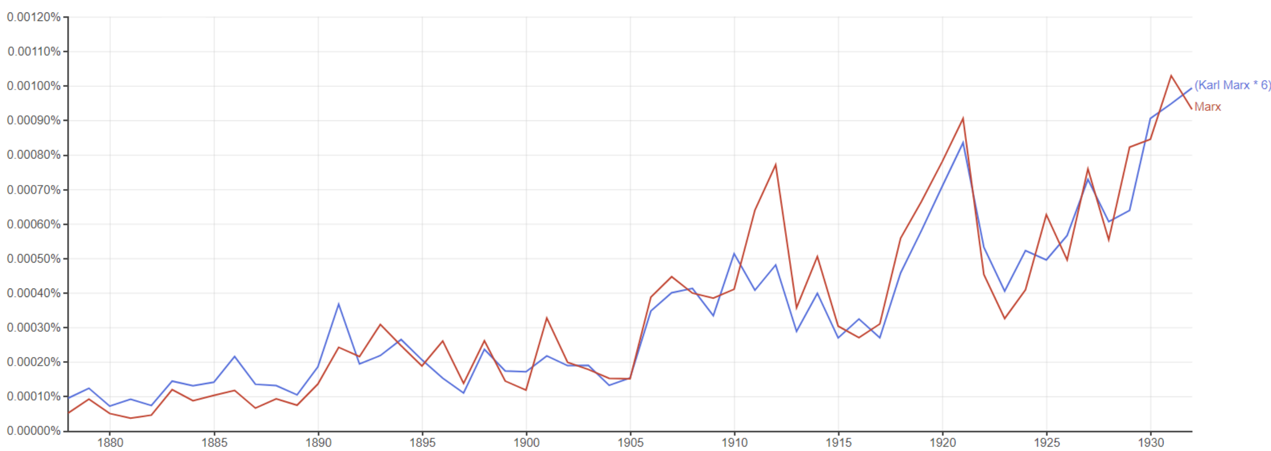

Figure 4: Google Ngrams Viewer, plotting "Marx" against "Karl Marx * 6," showing that "Marx" is mentioned almost exactly six times as often as "Karl Marx." This also shows that even if we are successfully identifying relative rates of change, we cannot identify levels.

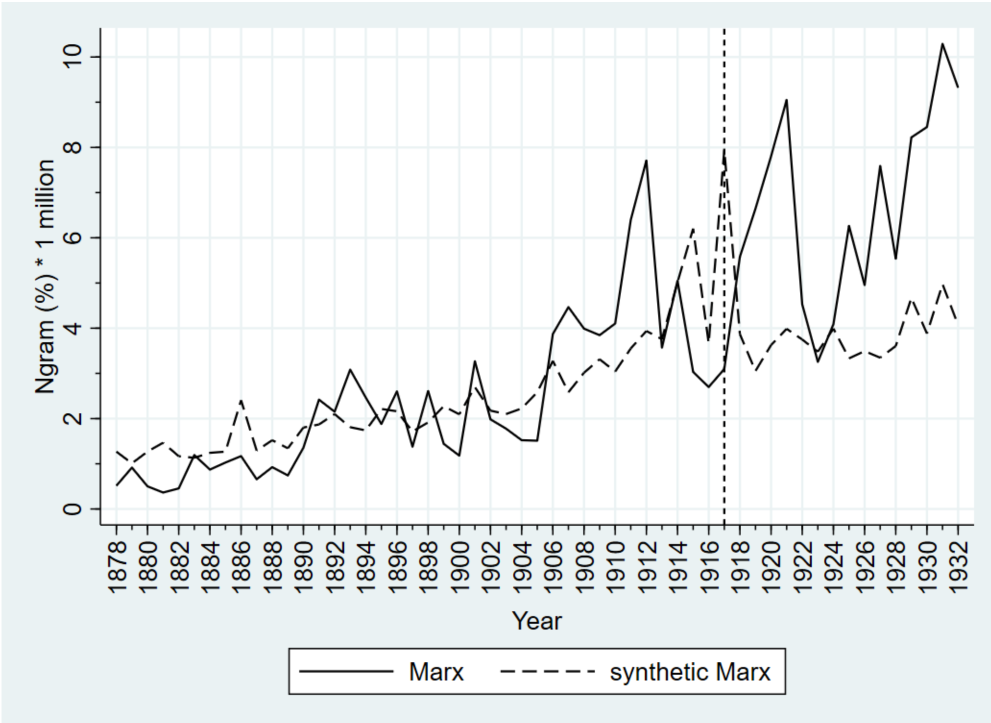

Figure 5: SCM, 1878–1932, author = "Marx": graphical representation

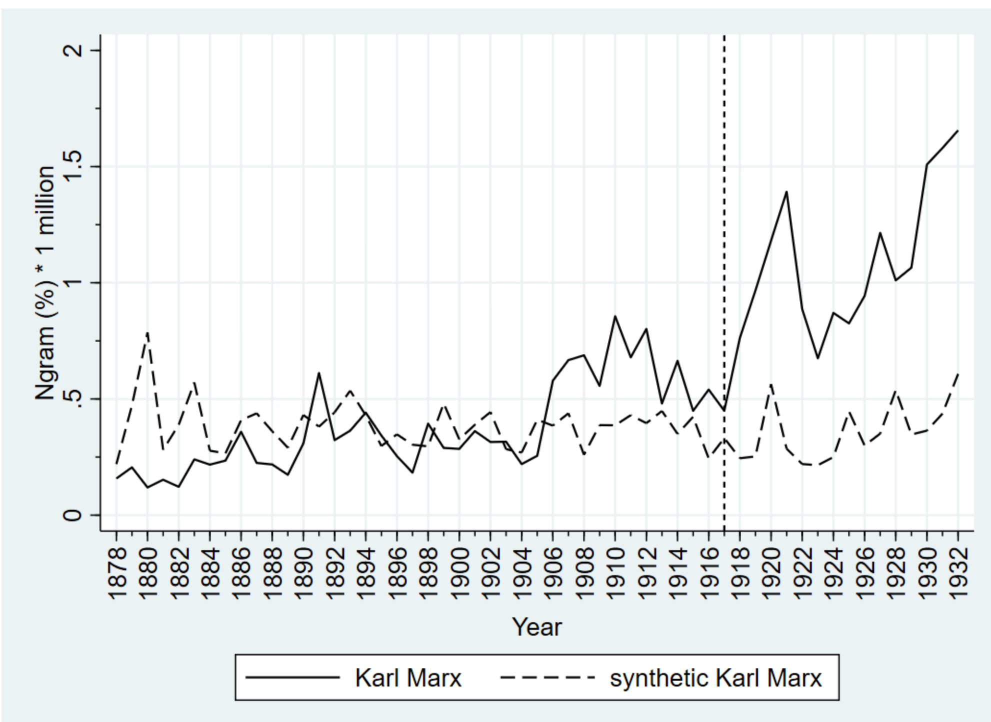

Figure 6: SCM, 1878–1932, cross-validation (pretreatment period divided into training and validation periods): graphical representation

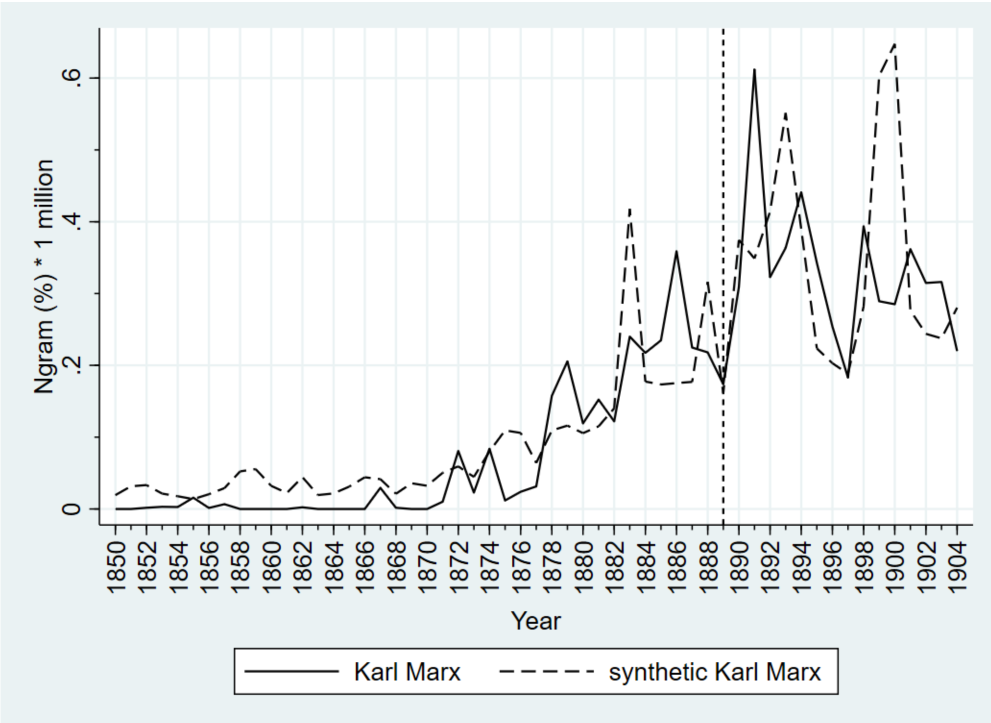

Figure 7: SCM, 1889 in-time placebo: graphical representation

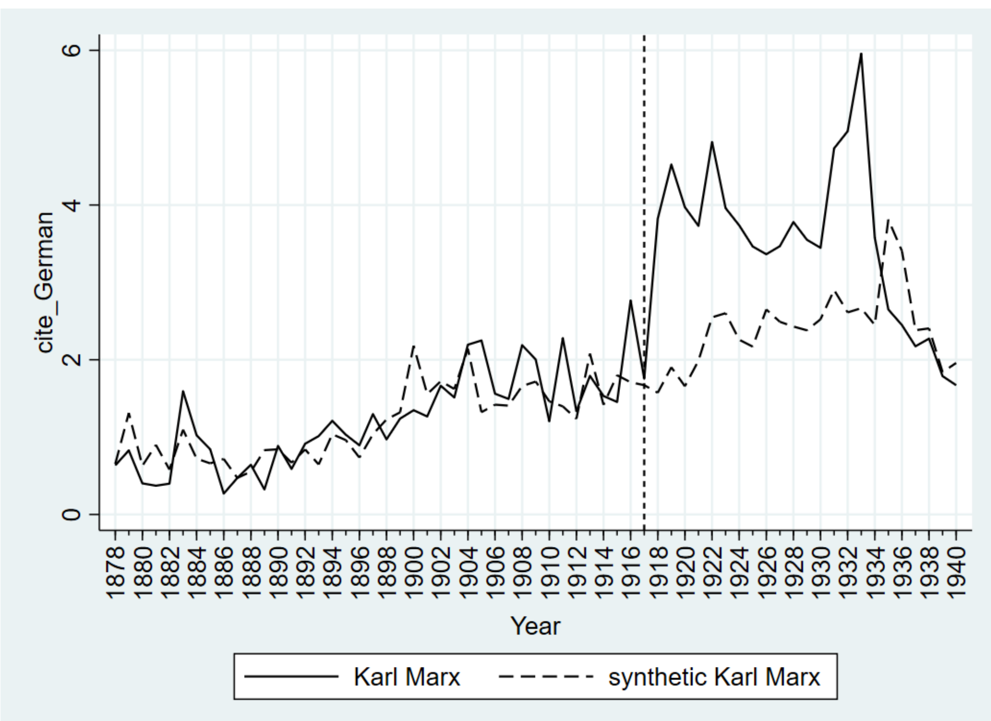

Figure 8: SCM, 1878–1940, German-language citations: graphical representation

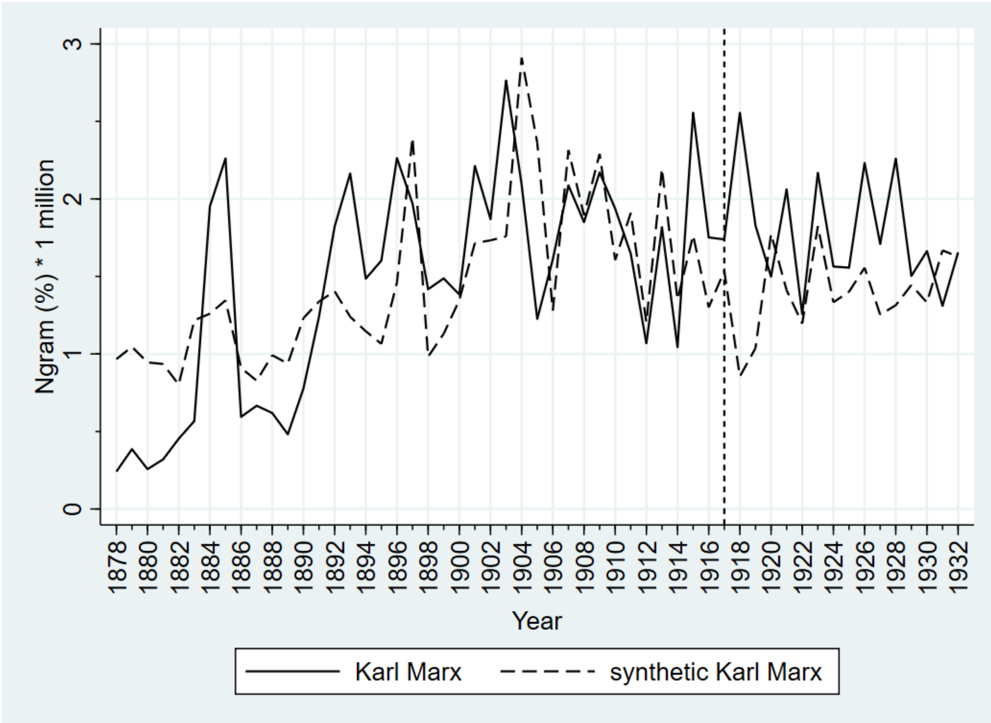

Figure 9: SCM, 1878–1932, French-language citations: graphical representation

A Technical Details on Synthetic Control and Hypothesis Testing

In this appendix, we provide technical details about our synthetic control regressions. We used two Stata modules to perform synthetic control: *synth* by Abadie *et al.* (2011))[41] and *synth_runner* by Galiani & Quistorff (2017).[42] *Synth_runner* is a wrapper for *synth* which automates the process of obtaining p-values from in-space placebos. In addition, we use the *parallel* module by Vega Yon & Quistorff (2012, 2019) to accelerate *synth_runner* by executing each in-space placebo on a different CPU core, reducing execution time at the cost of increasing RAM usage.[43]

All of our regressions were based on these templates, which estimate the treatment effect in 1917:

```
synth cite_English ///
    YearofPublication wrote_English wrote_German wrote_French wrote_Greek wrote_Latin ///
    YearofTranslationtoEnglish Socialist Political ///
    cite_English(1914(1)1916) cite_English(1908(1)1910) ///
    cite_English(1902(1)1904) cite_English(1896(1)1898) ///
    cite_English(1890(1)1892) cite_English(1884(1)1886), ///
    cite_English(1878(1)1880), ///
    nested allopt technique(dfp) ///
    trunit(#) trperiod(1917) resultsperiod(1878(1)1932) mspeperiod(1878(1)1916) ///
    fig keep("synth_results") replace
```

[41] Code to install in Stata:

```
ssc install synth, all
```

[42] Code to install in Stata:

```
net install synth_runner, from(https://raw.github.com/bquistorff/synth_runner/master/) replace
```

[43] Code to install in Stata:

```
net install parallel, from(https://raw.github.com/gvegayon/parallel/master/) replace
mata mata mlib index
```

```
synth_runner cite_English ///
    YearofPublication wrote_English wrote_German wrote_French wrote_Greek wrote_Latin ///
    YearofTranslationtoEnglish Socialist Political ///
    cite_English(1914(1)1916) cite_English(1908(1)1910) ///
    cite_English(1902(1)1904) cite_English(1896(1)1898) ///
    cite_English(1890(1)1892) cite_English(1884(1)1886), ///
    cite_English(1878(1)1880), ///
    nested allopt technique(dfp) ///
    trunit(#) trperiod(1917) mspeperiod(1878(1)1916) ///
    gen_vars keep("synth_runner_results") replace parallel
```

These two regressions are nearly identical, since *synth_runner* is just a wrapper for *synth*. Our outcome variable is the number of citations from Google Ngrams, without smoothing, named cite_English. Author descriptors are used as indicator variables: YearofPublication, wrote_English, wrote_German, wrote_French, wrote_Greek, wrote_Latin (a series of binary indicators of whether the author wrote in the specified language or not), YearofTranslationtoEnglish, Socialist, and Political.

In "trunit(#)," the "#" represents the integer ID number of the treated unit, where the ID is assigned according to Stata's *tsset* panel data setup. This number will depend on the sample chosen, e.g. all authors, socialists only, etc.

In synthetic control regressions, pretreatment outcomes are included as indicator variables as well. The rest of our indicators are therefore cite_English(X(1)Y), meaning the value of cite_English averaged over the years X through Y. These values are averaged over three years because citations are so erratic. By averaging citations over three years, we ensure that the synthetic author tracks the real author's long-term trend, not their short-term erratic deviations. We use three years as an indicator, then we leave a three year gap, and then we use another three years as an indicator. Thus, we first use cite_English(1914(1)1916), because 1916 is the last pretreatment year. Then, we use three years prior, cite_English(1908(1)1910), etc., until cite_English(1878(1)1880). Thus, we use indicators for the entire pretreatment period, but with averaging and gaps to avoid over-fitting. Cf. Kaul *et al.* (2018), who show that including too many outcomes as indicators causes over-fitting, resulting in other indicators receiving too little weight. When running regressions for different periods of time—e.g. in-time placebos—we maintained the same structure of this regression, simply changing the treatment period and the set of indicators appropriately. For example, when the treatment year is 1889, we include cite_English(1850(1)1852) through cite_English(1886(1)1888).

Synth allows three different levels of precision: default, *nested*, and *nested allopt*. As explained in the *synth* help file, the synth module will, by default, use a regression-based method to obtain weights placed on each indicator variable (the V-matrix). The donor weights (the W-matrix) are then chosen to minimize the RMSPE given these variable weights (V-matrix). Specifying *nested* will cause *synth* to engage in an interative optimization procedure that searches among all diagonal positive semidefinite V-matrices, using the regression-based V-matrix as a starting point. It is possible, however, that *nested* may only find a local optimum and not a global optimum. To mitigate this concern, *allopt* performs the same procedure with three different starting V-matrices: the regression-based V-matrix, equal weights, and using Stata's *ml* search procedure. *Nested allopt* returns the best of these three. We use *nested allopt* in all our SCM regressions in order to mitigate the concern that our synthetic controls are merely local optima.

In addition, *synth* allows the inclusion of parameters to be passed to Stata's *ml* optimizer. We specify "technique(dfp)" (Davidon–Fletcher–Powell) because it requires only one-fifth the execution time and one-seventieth the RAM required by the default *nr* (modified Newton–Raphson) technique - at least for our data, econometric technique (SCM), software (Stata 16), and computer equipment (AMD Ryzen 7 2700X). The *bfgs* (Broyden–Fletcher–Goldfarb–Shanno) technique is nearly as fast and reduces RAM requirements nearly as much as *dfp*; *bfgs* requires approximately one-third more time to execute and double the RAM of *dfp*, which is still but a small fraction of what *nr* requires. All three techniques - *nr*, *dfp*, and *bfgs* - produce very similar results, although they differ very slightly because they are using different techniques to determine when to terminate the iterative convergence procedure. Performing SCM with *nested allopt* requires a prohibitively large execution time and RAM usage unless either *dfp* or *bfgs* is specified, or unless one uses a supercomputing cluster with dozens of CPU cores and hundreds of GB of RAM. Thus, we use *dfp*.

In our permutation tests, the number of placebos is often less than the full sample because some placebos do not converge. This can be mitigated by specifying "margin(0.01)," which indicates a 1% margin of constraint violation tolerance. We do not include this parameter in any of our reported results. But our results do not materially change when this parameter is included, which increases the effective sample size at the cost of precision.

Meta-analysis of multiple p-values is performed using the harmonic mean p-value method by Wilson (2019b). This method is implemented in R with *harmonicmeanp* by Wilson (2019c). Three R commands are used: *hmp.stat* to obtain harmonic mean p-values (HMPs), and *p.hmp* and *p.pharmonicmeanp* to obtain asymptotically exact p-values (AEPs). R commands are executed within Stata using *rcall* by Haghish (2019b,c), and *rcall* is installed in Stata using *github* by Haghish (2019a).[44]

B Obtaining Data From Google Ngram

As noted by the Google Ngram Viewer Team (n.d.), Google Ngram data can be downloaded from `https://books.google.com/ngrams/datasets`. However, these data are somewhat difficult to use. First, each wave of Ngrams is listed separately. Second, Ngram data are dispersed among different .GZ files, by n length and by first two letters. For example, "Karl Marx" is a 2-gram starting with the letters "ka," while "Kropotkin" is a 1-gram starting with the letters "kr." Each .GZ file is at least several hundred megabytes in size. Combining these data would require unzipping

[44] The sequence of events in Stata to install all the necessary packages is this:

net install github, from("https://haghish.github.io/github/")

github install haghish/rcall, stable

rcall: install.packages("harmonicmeanp", repos="http://cran.uk.r-project.org")

rcall: library(harmonicmeanp)

To illustrate the use of *rcall*, we use Stata to replicate one of the critical AEP values from Wilson (2019, table 1). An HMP of 0.040 with 10 tests and an HMP of 0.0233 with 1 billion tests both imply $\alpha < 0.05$. (Note that the following commands are executed in Stata, not in R; the "rcall:" command in Stata is followed by a R command using R syntax, and saves values in Stata r-class matrices.)

rcall: pharmonicmeanp(0.040, L=10, lower.tail=TRUE)

rcall: pharmonicmeanp(0.023, L=1000000000, lower.tail=TRUE)

rcall: x <- c(0.040, 0.040, 0.040, 0.040, 0.040, 0.040, 0.040, 0.040, 0.040, 0.040)

rcall: p.hmp(x)

The results of all these commands should be approximately 0.05.

and combining approximately as many files as we have authors. This is feasible, but tedious.

A simpler method is to simply download the data from the Google Ngram Viewer (Google n.d.) at `https://books.google.com/ngrams`. Type in the phrase one wishes to search for, e.g. "Karl Marx," but without quotation marks. We suggest typing in only one phrase at a time to avoid accidentally mixing up data for one author with another. Select the years, corpus, and smoothing. Available corpora include different languages in 2009 and 2012. We used 1800–2000, "English (2012),"[45] and smoothing 0 (zero).[46]

A graph will then appear, plotting the n-grams for that author. From there, view the HTML page source.[47] In the page source, there is a section containing comma-separated values of all the n-grams. To keep this example short enough to print, we give an example for the years 1850–1860. For "Karl Marx," 1850–1860, English, smoothing 0, the relevant section of the HTML source looks like this:

```
var data = [{"ngram": "Karl Marx", "type": "NGRAM", "timeseries":
[0.0, 0.0, 1.7127106399783543e-09, 3.0619566881995297e-09,
2.7978830363650786e-09,  1.5773920836181787e-08, 1.4656383884315005e-09,
6.72457067807386e-09, 0.0, 0.0, 0.0], "parent": ""}];
```

Simply copy the floating point numbers within brackets, which in this case is:

```
0.0, 0.0, 1.7127106399783543e-09, 3.0619566881995297e-09,
2.7978830363650786e-09,  1.5773920836181787e-08, 1.4656383884315005e-09,
6.72457067807386e-09, 0.0, 0.0, 0.0
```

[45]A detailed list of all corpora is found in the "About" page, i.e. Google Ngram Viewer Team (n.d.). In the viewer, available corpora include "American English (2012), "American English (2009)," "British English (2012)," "British English (2009), "English (2012)," "English (2009)," etc.

[46]Smoothing X means that data for a given year is averaged with data from X years before and after. Eg., smoothing 1 (one) for year K means that data for year K are averaged with data for years (K-1) and (K+1), forming a three year average. Similarly, smoothing 3 (three) is a seven year average. In all our regressions, we used smoothing 0. If the researcher desires a different level of smoothing, it is easy to manually smooth the data using smoothing 0. Simply *tsset* the data in Stata and use lags and leads to form a (weighted or unweighted) average of several years' data.

[47]In Google Chrome, right-click and choose "View page source," or press Ctrl+U.

At this point, there are two ways to format the numbers for analysis:

1. Paste these numbers into a text file, save as a .CSV (comma-separated values), and import into the software of one's choice.

2. In LibreOffice Calc (a free spreadsheet software), paste these numbers in the first cell. Select that cell—which contains the entire list of numbers—and click the "Data" menu. Select "Text to Columns" and make sure "Separator options → Comma" is selected. Click OK. Now all the numbers are arranged in a row, with each number in its own column. Now, select the entire row and cut (Ctrl+V). Next, click the first cell and select "Paste Special → Paste Special" (Ctrl+Shift+V) and make sure that "Options → Transpose" is checked. Now, each number is in its own row. Now add a column of years, and a row of variable names. Save, and the result is ready to be imported into the software of one's choice.

Although this process is somewhat cumbersome, it is still faster and simpler than unzipping large GZ files—at least for our purposes. If one wishes to automate this process using web-scraping, the URL for each phrase is straightforward. For example, the URL for "Karl Marx," 1800–2000, English, smoothing 0 is `https://books.google.com/ngrams/graph?content=Karl+Marx&year_start=1800&year_end=2000&corpus=15&smoothing=0`.

Because raw citations counts are within a few orders of magnitude of $^1/_{10}$ of one million (or $^1/_{100{,}000}$ of a percent), we multiplied all citations by one million simply to remove excess decimal zeroes. For example, in the example above, the third number in the list is "1.7127106399783543e-09." When this is multiplied by one million, the e-09 becomes e-03. Thus, the vertical axis on all our SCM graphs can span approximately 0 to 2 instead of 0 to $^1/_{500{,}000}$. This does not affect our results any more than if one were to convert inches to miles to reduce the number of digits printed on a graph.