

Potatoes: A Critical Review

by Pablo Villalobos and Jaime Sevilla

Executive summary

- Nunn and Qian study the effect of the introduction of potatoes in the Old World on population growth between 1700 and 1900.
- We think the paper credibly establishes that between one-sixth and one-quarter of the growth is a consequence of the introduction of potatoes.
- The main reason for doubt is the possibility of spurious correlation due to spatiotemporal autocorrelation and the fact that potatoes were mainly grown in Europe, which at the time was experiencing growth due to unrelated factors.
- After performing several tests to account for these concerns, we conclude they are not strong enough to reject the conclusion of the paper.

Introduction

During the 18th and 19th centuries we saw rapid urbanisation and a drastic increase of population throughout the world. Some attribute this to industrialization ; [Nunn and Qian](#) argue that we owe that to potatoes.

They are not coy about it either. “According to our most conservative estimates, the introduction of the potato accounts for approximately one-quarter of the growth in Old World population and urbanisation between 1700 and 1900”.

They back their hypothesis with data from several natural experiments, studying variability of city populations and adult heights between and within countries, and relate it to the timing when potatoes were introduced in different places.

In this article we will summarise, replicate and critically review Nunn and Qian’s paper. We follow the methods established in [\(Sevilla, 2021\)](#) to study whether their findings are robust and whether they successfully establish a causal relation.

Paper	Theory of causation (Historical exposure \Rightarrow mediator \Rightarrow long term outcome)	Correlation analysis Is there a non-spurious long term correlation?	Causality determination Is the correlation causal?
The Potato's Contribution to Population and Urbanization: Evidence from a Historical Experiment (Nunn & Qian, 2011)	Countries suitability for cultivating potatoes \Rightarrow higher yields from potatoes \Rightarrow increased growth in population and urbanisation	DiD Main issue: correlation of Europe with both growth and potato cultivation	Controls, within-country and within-continent comparisons, multiple independent datasets. The finding is robust to all of these.

Paper Summary

The study purports to evaluate the causal impact of the introduction of potatoes in the Old World in population growth and urbanisation rates.

The theory is that potatoes are more nutritious and provide 3x more calories per acre, so when they started being cultivated in the Old World they produced a positive shock in agricultural productivity, enabling higher populations and wealth per capita.

A previous study by [\(Mokyr, 1981\)](#) estimated the causal effect of potatoes on population growth in Irish counties in 1845¹, finding an effect size of 0.7. This implies counties with high potato cultivation grew an extra 0.15% that year, compared to counties with low cultivation². In comparison, Nunn and Qian's study uses data of the whole Old World, from 1300 to 1900, instead of just Ireland in a single time slice, and looks at urbanisation rates in addition to population.

To estimate the causal impact of potatoes, the authors exploit two sources of variation³ on each country's ability to grow potatoes: the time of introduction and the suitability of the land for potato cultivation. In their baseline analysis they find that around a quarter of the growth in population and urbanisation rates in that time period can be attributed to potatoes.

The main pitfall of this strategy is that it relies on there being no other shocks during that time which are correlated with suitability for cultivation. Unfortunately, Europe is much more suitable for potato cultivation than other Old World regions, and the potato was introduced at the same time when Europe was diverging from other countries due to several unrelated factors.

To check that the effect is indeed causal, the authors use several strategies:

- Adding several controls for alternative drivers of population and economic growth.
- Comparing only countries within the same continent
- Comparing cities within the same country
- Comparing the heights of soldiers in France, using their town of birth to estimate the importance of potatoes in their childhood diet.

¹ Using a 2SLS analysis.

² The regression coefficient is 0.7, dependent variable is yearly pop. growth rate and independent variable is percentage of total land devoted to potato cultivation. The counties with highest potato cultivation had 29% of the land devoted to potatoes, whereas for the lowest ones it was 7%. $0.007 * (0.29 - 0.07) = 0.00154 \approx 0.15\%$. Over 200 years, this corresponds to an extra growth of 36%, 0.3 in log scale.

³ They use a variant of the differences-in-differences (DD) method, where the difference in population growth between countries with varying potato suitability is compared before and after the introduction of potatoes. What makes this analysis slightly different from standard DD is that instead of collapsing all data into 'pre' and 'post' adoption periods, the study keeps the temporal structure of the data.

In all of those cases, they find a significant effect of potatoes on the outcome. For the within-continent and within-country analysis, the effect size is in the same order of magnitude as in the baseline regression⁴.

They also perform some analyses to determine which cutoff date to use for the introduction of potatoes:

- A regression interacting potato suitability with the time period, which shows an increasing effect for periods after 1750, consistent with that being the right cutoff.
- A series of regressions with a 'rolling window' of 400 years, taking the first 200 hundred years as the pre-adoption period and the last 200 years as the post-adoption period. Again the earlier windows show no effect of potatoes, while the latter ones (1600-1900) do show an effect.

I think the authors convincingly show that potatoes had a robust positive effect⁵. Moreover, the effect size is within what would be expected from theory and consistent with previous findings⁸.

Conclusion

We tentatively conclude the main claim in the paper is broadly correct. That is, around a quarter of Old World population growth from 1700 to 1900 was caused by the introduction of potatoes.

While the significance and effect size is diminished when taking into account multiple hypotheses and within-continent variation only, this is not enough to make the effect non-significant or much smaller. Spatial autocorrelation does not seem to be an issue due to the characteristics of the analysis, as found by [\(Kelly, 2020\)](#)⁶.

This analysis could be improved by a) finding other sources to estimate expected effect size, and b) verifying the correlated standard errors computed in [\(Kelly, 2020\)](#)⁷.

Acknowledgements

This review has been commissioned by the Forethought Foundation.

⁴ The effect size is 0.02 for within-continent comparisons and between 0.029 and 0.05 for within country comparisons.

⁵ Note: the empirical effects found in this study could be due to the introduction of the potato having a permanent effect on the growth rates of the countries, but it could also be due to it having a one-time effect on the long term population level. In this last case, since potato adoption was gradual, the level change would be spread over centuries. This means we can't distinguish these two hypotheses with the existing data.

⁶ See the appendix for the details on analysing multiple hypothesis testing, spatial autocorrelation, and the power of the analysis.

⁷ Note: Kelly expressed in correspondence that he no longer trusts the approach in his paper. For now we stand by our conclusion, and await his forthcoming publication with improved methodology.

Bibliography

Gelman, A. (2014). *Beyond Power Calculations: Assessing Type S (Sign) and Type M (Magnitude) Errors*. <https://journals.sagepub.com/doi/10.1177/1745691614551642>

Kelly, M. (2020). *Understanding Persistence*.
https://economics.yale.edu/sites/default/files/understanding_persistence_ada-ns.pdf

Mokyr, J. (1981). *Irish History with the Potato*
<https://journals.sagepub.com/doi/abs/10.1177/033248938100800102>

Nunn, N., Qian, N. (2011). *The Potato's Contribution to Population and Urbanization: Evidence from a Historical Experiment*
<https://scholar.harvard.edu/nunn/publications/potatos-contribution-population-and-urbanization-evidence-historical-experiment>

Sevilla, J. (2021). *Persistence: A Critical Review*
<https://docs.google.com/document/d/14ULAaTofWiQbTCP1ekuaenQJ6saXEzjgiKMznlBrXvQ>

Appendix: Table summary

Result of replication	Significant effect of potato introduction on population growth. Between $\frac{1}{2}$ and $\frac{1}{4}$ of world population growth from 1700 to 1900 seems to be attributable to potatoes.	
Statistical method of replication	Differences in differences regression	
Standardised $\beta \approx 0.051$ (0.018) [0.024] *	Adjusted p-value $\approx 1.1\%$	n=1552
Expected $\beta \approx 0.047^8$	# hypothesis = 2^9	Critical number of hypothesis = 10
Power $\approx 57\%^{10}$	Type S error rate $\approx 5e-9$	Exaggeration ratio ≈ 1.2
Moran's Z $\approx 4.04^{11}$	Moran's p $\approx 2e-5$	Persistence span = 1700 to 1900
Short commentary	<p>Based on the expected effect size, we find the study is adequately powered to detect an effect, and even estimate the effect size with some accuracy.</p> <p>While the degree of spatial autocorrelation is very high, (Kelly, 2020) finds the nature of the regression (panel data, fixed effects) prevents this correlation from exaggerating the results.</p> <p>The main regression likely overestimates the effect due to Europe having both high growth and high potato suitability. Adding continental controls reduces the effect size by $\frac{1}{3}$, still enough to explain $\frac{1}{2}$ of the total growth in that period.</p>	
Reproduction details	We replicate the regression in column (3) of table IV, as well as those in the first 3 columns of table VIII.	

Appendix: Reproduction Details and Calculations

⁸ Based on the ~ 0.3 effect over 200 years obtained by [\(Mokyr, 1981\)](#) and on theoretical calculations: if $\frac{1}{2}$ of the land is suitable for potatoes, their caloric yield is 3x that of Old World crops, and all the extra yield is absorbed by population growth, then we should expect a log increase of $\log(\frac{1}{2} + 3*\frac{1}{2}) \approx 0.3$, broadly consistent with the effect found by the paper.

⁹ They test for population growth and urbanisation rate.

¹⁰ Even if the expected effect size was 10x lower, the type S error rate would still be $\sim 13\%$, and the effect would still explain around 3% of the total growth over that period.

¹¹ Kelly (2020) finds a Z value of 6.44 for the regression in column (1) of table IV, which we were able to replicate. The same methodology yields this value for the baseline regression.

Our procedure is similar to the one in [\(Sevilla, 2021\)](#): In addition to reproducing the main results, we test for spatial autocorrelation¹², multiple hypotheses testing¹³, and the possibility that the analysis is underpowered¹⁴.

We reproduce the regressions on the first three columns of tables IV and VII. In the summary table we report the result of column (3) of table IV. The estimate of the effect size is 0.032, which when standardised is 0.051.

To give an intuitive understanding, this means that increasing the amount of suitable land for potato cultivation by 1 percent increases the population by 0.032 percent. To estimate the total impact of population in the whole 1700-1900 period, we compute the total population growth and the growth under the counterfactual where potatoes are not present (the potato variable is 0).

Total population growth in that period was 247%, which in log form is 0.9. The counterfactual growth is 195%, which in log form is 0.67. So the difference in logs is 0.23, or approximately 0.9/4, which is the total impact of potatoes.

The estimate of the effect size on column (3) of table VII, which takes into account only within-continent effects, is 0.020. This similarly translates to a total impact of 1/6 of the total growth.

We compute Moran's statistic following the procedure in [\(Kelly, 2020\)](#): taking the average of its value for each time period, calculated with a spatial kernel of 1 for the 5 closest points, and 0 otherwise¹⁵. We did not compute the correlated standard errors from [\(Kelly, 2020\)](#), instead we take his word that they are smaller than the clustered errors.¹⁶

We use design analysis as in Gelman (2014) to compute the power of the study and its type S and M error rates. To find an estimate of the true effect size, we look at the effect found by [\(Mokyr, 1981\)](#). He found an increase of 0.15% in the annual population growth rate, which over 200 years constitutes a total growth of 36%, for a logarithmic change of about 0.3. We also use a simple theoretical model which predicts a similar effect⁸.

¹² The phenomenon where spatially close places tend to be more similar to each other than would be expected if they were truly independent. The presence of spatial autocorrelation distorts estimates of standard errors and inflates t-statistics.

¹³ If the authors test the effect of potatoes on several different outcomes (multiple hypotheses), this increases the probability of spurious results.

¹⁴ The power is the probability of correctly rejecting the null hypotheses. If the true effect size is real but too small compared to the standard error, then the analysis will be underpowered and won't be able to detect an effect. This is mitigated by reducing the standard error, which usually means increasing the sample size.

¹⁵ Other kernels, like inverse distance and inverse exponential distance, found lower levels of spatial autocorrelation.

¹⁶ Kelly finds an effect size of 4.11 with a clustered standard error of 1.05, which is reduced to 0.44 after taking spatial correlation into account. These numbers are very different from the ones in the original paper and we are not sure how they are being computed. This is why we chose not to use them. In any case if Kelly is right, using his correction would only strengthen the conclusion.

To account for multiple hypotheses, we make the same corrections as [Sevilla \(2021\)](#)¹⁷. First we compute the critical number of hypotheses for the given significance threshold, which in this case is 10, and then we compare it with the actual number of hypotheses tested. In this case the authors test 2 hypotheses (effect of potatoes on population and urbanisation rate), which results in a small increase in the p-value of the study. Note that in this case the result of the analysis was positive for both hypotheses, which gives us further confidence that there are no statistical artefacts.

All the code used for the replication can be found [here](#).

¹⁷ That is, we adjust the p-value of the main regression using the [Šidák correction](#) $p_{\text{adj}} = 1 - (1 - p)^n$, where n is the number of hypotheses tested. In addition, we have to adjust the standard error σ so that it still represents a confidence interval of confidence level $\alpha \approx \text{erf}(1/\sqrt{2})$. To do this, we set $\sigma_{\text{adj}} = \sigma \sqrt{2} \text{erf}^{-1}(\alpha^{1/n})$.