

Essays in Causal Inference



INAUGURALDISSERTATION
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von
Yann Müller
aus Seeheim-Jugenheim

Mannheim
2025

Abteilungssprecher	Prof. Thomas Tröger
Vorsitzender	Prof. Carsten Trenkler
Referent	Prof. Antonio Ciccone
Koreferent	Prof. Miren Azkarate-Askasua
Verteidigung	Mannheim, den 31. Oktober 2025

Abbildungserklärung

Alle Abbildungen in den Hauptkapiteln dieser Dissertation wurden von den Autoren erstellt, sofern nicht ausdrücklich anders angegeben. Bilder auf der Einleitungsseite der Dissertation sowie auf den Kapiteleinleitungsseiten wurden mit Hilfe von OpenAIs ChatGPT (DALL·E-Bildgenerierungstool) erstellt.

Danksagung

Und ob ich schon wanderte im finstern Tal, fürchte ich kein Unglück [...]

Psalm 23, Lutherbibel

Die vorliegende Dissertation ist das Ergebnis der zahlreichen Unterstützung, die ich im Laufe der letzten Jahre erfahren habe und für die ich mich im Folgenden ausdrücklich bedanken möchte.

Auf akademischer Seite möchte ich mich zuerst bei meinem Betreuer Prof. Antonio Ciccone bedanken. Antonios bescheidene und nahbare Art haben mir in der Promotion sehr geholfen. Immer wieder war ich beeindruckt und inspiriert von der Tiefgründigkeit, mit der Antonio akademische Fragestellungen betrachtet und zu den verschiedensten Themenbereichen konstruktive Unterstützung geben kann. Ich hoffe in der Zukunft viel von dieser Tiefgründigkeit mitnehmen zu können.

Ebenso bedanke ich mich für die Unterstützung von Prof. Ana Moreno-Maldonado, insbesondere für die vielen Ratschläge zur Navigation durch die Promotion. Auch die vielen Kommentare der Mannheimer Macro-Gruppe waren hilfreich im Prozess dieser Arbeit. Des Weiteren bedanke ich mich bei den Mitgliedern meiner Kommission, Prof. Miren Azkarate-Askasua und Prof. Carsten Trenkler.

Die teilweise sehr harten Phasen der Promotion waren durch meine akademischen Weggefährten Herr Bjarne Horst, Dr. Max Riegel und Herr Henrik Sauer weitaus erträglicher. Die akademischen und politischen Diskussionen waren eine große Bereicherung und ich werde immer positiv auf den regelmäßigen Austausch zurückblicken.

Zudem möchte ich mich bei den Mitarbeitern der Universität Mannheim bedanken, insbesondere Frau Marion Lehnert, Frau Kristina Kadel, Herr Ulrich Kehl, Dr. Golareh Khalilpour, Frau Stephanie Lilwall und Frau Sylvia Rosenkranz.

Auf privater Seite möchte ich die Freundschaft und Gefährtschaft der Mumble Calls Gruppe hervorheben, die seit der Kindheit eine enge Begleitung in meinem Leben ist und mit welcher ich prägende Lebenserfahrungen bis heute teilen kann. Auch möchte ich mich bei Herrn Philipp Ollendorff für die jahrelange Freundschaft und gegenseitige intellektuelle Stimulation in allen Bereichen des Lebens und gegen den Wahnsinn der Zeit bedanken.

Abschließend möchte ich den wichtigsten Teil dieser Reise hervorheben: die allgegenwärtige Unterstützung meiner Familie, die meinen Lebensweg überhaupt erst ermöglicht hat und somit auch der Ursprung dieser Dissertation ist. Diese Arbeit ist somit eine Widmung an meine verstorbene Oma Wally sowie insbesondere an meine Eltern *Dagmar* und *Horst*, welche mich über mein Leben hinweg in einer selbstlosen Weise unterstützt haben.

Contents

Danksagung	IV
Contents	VI
List of Tables	X
List of Figures	XI
Introduction	1
References	4
Chapter I: Quantifying Waterway Supply Chain Shocks	5
1.1 Introduction	7
1.2 Data	9
1.2.1 Trade Flow	9
1.2.2 Water Levels	10
1.3 Shock Construction	12
1.3.1 Capacity Restrictions	12
1.3.2 Shock Construction	13
1.3.3 Shock Diagnostics	15
1.4 Empirical Framework	17
1.4.1 Outcome Variables	17
1.4.2 Empirical Strategy	18
1.5 Empirical Results	20
1.5.1 Industrial Production	20

1.5.2	Truck Index	22
1.5.3	Train Transportation	24
1.5.4	Energy Imports	24
1.5.5	Energy Prices	26
1.6	Conclusion	28
	References	31
	Appendix	33
1.A	Robustness Checks	33
1.A.1	Shock Regression	33
1.A.2	Baseline Specification	41
1.A.3	Additional Results	43
1.B	Data Details	44
1.B.1	Trade Flow Data	44
1.B.2	Water Level and Related Data	47
1.B.3	Outcome Variables	47
1.B.4	Capacity and Shock Series	50
	Chapter II: The Amateur's Curse of Positive Income Shocks	53
2.1	Introduction	55
2.2	Background: Verbandspokal	58
2.3	Data	59
2.3.1	Consecutive Ranking	59
2.3.2	Board of Directors	61
2.3.3	Newspaper Coverage	61
2.3.4	Team Managers and Players	61
2.3.5	Elections	61
2.4	Empirical Strategy	61
2.5	Empirical Results	64
2.5.1	Effect on League Ranking	64
2.5.2	Effect on Financial Stability	65
2.5.3	Effect on Team Manager Tenure and Team Composition	68
2.5.4	Effect on Board of Directors	70
2.5.5	Effect on Political Elections	72

2.6	Conclusion	73
	References	75
	Appendix	77
2.A	Data	77
2.A.1	<i>Verbandspokal</i> Finals	77
2.A.2	Board of Directors	80
2.B	Additional Results	81
2.B.1	Ranking Distributions	81
2.B.2	Event Study Robustness Specifications	83
2.B.3	Event Study Relegation/Promotion	85
2.B.4	Event Study Board of Directors	86
2.B.5	Event Study Elections Results	88
	Chapter III: The Causal Effect of Domestic Market Potential	89
3.1	Introduction	91
3.2	Background	94
3.2.1	Oases	94
3.2.2	Introduction of Motorized Mechanical Wells	96
3.2.3	Agricultural Adaptation after Technological Shock	98
3.2.4	Change in Population Structure	98
3.2.5	Moroccan Economy	100
3.3	Data	102
3.3.1	Micro-Climates	102
3.3.2	Market Potential	105
3.4	Empirical Strategy	106
3.4.1	Exposure to Optimal Geographic Conditions	106
3.4.2	Empirical Framework	109
3.5	Empirical Results	110
3.5.1	Effects on Employment Growth	110
3.5.2	Predictions by Market Potential and Employment Density	112
3.6	Robustness Checks	113
3.6.1	Interaction with Oasis Locations	113
3.6.2	Limiting Potential Markets	114
3.6.3	Changes in Elasticities	116

3.6.4	Alternative Instrument Construction	116
3.7	Conclusion	118
	References	120
	Appendix	124
3.A	Data Appendix	124
3.A.1	Optimal Geographic Conditions for Oases Emergence	124
3.A.2	Geographic Control Variables	127
3.A.3	Census and Annual Reports Data	128
3.B	Additional Results	134
3.C	Technical Appendix	140
3.C.1	Transportation Network: Least-Cost Route	140
3.C.2	Predicting Oasis Location with Geodata using Gradient Boosting	140
3.D	Theory Appendix	144
3.D.1	Optimal Allocation	144
3.D.2	Comparative Statics	145
3.D.3	Technological Shock	146
	Resources	150
	Declaration	152
	Curriculum Vitae	153

List of Tables

2.1	Team Characteristics of Treatment and Control Clubs	63
2.A.1	Regional <i>Verbandspokal</i> Finals Decided by Penalty Shootout	77
2.B.1	Board of Directors Characteristics of Treatment and Control Clubs .	86
3.1	Regression of Population on Oases Exposure	99
3.2	Overview of Population Development in Morocco, 2004–2024.	102
3.3	Summary of Optimal Geographic Conditions for an Oasis Ecosystem	103
3.4	Regression Estimates Distance-Measured Market Potential	111
3.1	Overview Census and Annual Reports Data	128
3.B.1	Regression Estimates Distance-Measured Market Potential - First Stage	134
3.B.2	Regression Estimates Least-Cost Market Potential - First Stage . . .	135
3.B.3	Regression Estimates Least-Cost Market Potential	136
3.B.4	Regression Estimates Distance-Measured Market Potential and Oasis Locations (1/2)	137
3.B.5	Regression Estimates Distance-Measured Market Potential and Oasis Locations (2/2)	138
3.B.6	Regression Estimates Machine Learning Computed Least-Cost Market Potential	139

List of Figures

1.1	Overview of Goods Transported for each Rhine Section.	9
1.2	Means of Transportation for certain Goods Categories.	10
1.3	Administratively relevant measuring stations along the Rhine.	11
1.4	Nautical Definitions.	12
1.5	Number of Days with Restrictions at Selected Stations.	14
1.6	Aggregate Monthly Capacity	15
1.7	IRF of Goods Transportation on the Rhine (Cross-Border Receiving)	16
1.8	IRFs of German Industrial Production	20
1.9	IRF of Industrial Production for Different States and Sectors	21
1.10	IRF of Truck Index	23
1.11	IRF of Train Transportation (Cross-Border Receiving)	24
1.12	IRF of Petroleum and Crude Oil Imports	25
1.13	IRF of Refined Petroleum and Crude Oil Imports by Country	26
1.14	IRF of PPI components	27
1.15	PPI Heating Oil by Marketplace	28
1.16	IRF of PPI Heating Oil by Marketplace	29
1.A.1	IRF of Goods Transportation on the Rhine with Adjusted Shock Series (Cross-Border Receiving)	33
1.A.3	IRF of Industrial Production for States and Sectors with Adjusted Shock Series	34
1.A.2	IRF of German Industrial Production with Adjusted Shock Series . . .	34
1.A.4	IRF of Truck Index with Adjusted Shock Series	35
1.A.5	IRF of Train Transportation with Adjusted Shock Series (Cross-Border Receiving)	35
1.A.6	IRF of Petroleum and Crude Oil Imports with Adjusted Shock Series	36
1.A.7	IRF of Refined Petroleum and Crude Oil Imports by Country with Adjusted Shock Series	36
1.A.8	IRF of German PPI components with Adjusted Shock Series	37

1.A.9	IRF of Goods Transportation on the Rhine with Adjusted Shock Series (Cross-Border Receiving)	37
1.A.10	IRF of German Industrial Production with Adjusted Shock Series . .	38
1.A.11	IRF of Industrial Production for States and Sectors with Adjusted Shock Series	38
1.A.12	IRF of Truck Index with Adjusted Shock Series	39
1.A.13	IRF of Train Transportation with Adjusted Shock Series (Cross-Border Receiving)	39
1.A.14	IRF of Petroleum and Crude Oil Imports with Adjusted Shock Series	40
1.A.15	IRF of Refined Petroleum and Crude Oil Imports by Country with Adjusted Shock Series	40
1.A.16	IRF of German PPI components with Adjusted Shock Series	41
1.A.17	IRF of German Industrial Production	42
1.A.18	IRF of German Industrial Production (Energy-Intensive Sectors Only)	42
1.A.19	IRF of PPI Heating Oil by Marketplace	43
1.B.1	Trade Flows by Category and Rhine section in 2021 (1/3)	44
1.B.2	Trade Flows by Category and Rhine section in 2021 (2/3)	45
1.B.3	Trade Flows by Category and Rhine section in 2021 (3/3)	46
1.B.4	Sector Shares of Value Added	49
1.B.5	Number of days with restrictions at selected stations	50
1.B.6	Average Number of Days with Restrictions	51
1.B.7	Average Number of Days with Restrictions (High vs. Low Water Levels)	51
1.B.8	Final Rhine Capacity Shock Series	52
2.1	League Ranking Before/After Verbandspokal Win	58
2.2	Illustration Consecutive Ranking in the Year 2000	60
2.3	League Ranking (Event Study)	65
2.4	League Tier (Event Study)	66
2.5	Insolvency Proceedings (Event Study)	67
2.6	Debt-Related Newspaper Coverage (Event Study)	68
2.7	Team Manager Tenure Duration (Event Study)	68
2.8	New Players per Season (Event Study)	70
2.9	Average Player Tenure (Event Study)	71
2.10	Age of Board Members (Event Study)	71
2.11	Board Members with Doctoral Degrees (Event Study)	72
2.12	Seat Share of Majority Party (Event Study)	73
2.A.1	Illustrative German Commercial Register Document	80
2.B.1	Ranking Distribution in (Pre-)Treatment Years	81
2.B.2	Ranking Distribution in Post-Treatment Years	82
2.B.3	Robustness Specification - Pre-Treatment Binning	83

2.B.4	Robustness Specification - Sample Restriction Pre-2015 Finals	84
2.B.5	Robustness Specification - Not Yet Treated Control Group	84
2.B.6	Likelihood of Relegation Before/After <i>Verbandspokal</i> Win (Event Study)	85
2.B.7	Likelihood of Promotion Before/After <i>Verbandspokal</i> Win (Event Study)	86
2.B.8	Board of Directors Size Before/After <i>Verbandspokal</i> Win (Event Study)	87
2.B.9	Tenure Duration of Board Members Before/After <i>Verbandspokal</i> Win (Event Study)	87
2.B.10	Change in Seat Share of Majority Party After Controlling for League Ranking (Event Study)	88
3.1	Overview of Oases in Morocco and Illustration of a Wadi Oasis.	95
3.2	Population Development in Morocco.	96
3.3	Agricultural Development: Livestock	97
3.4	Overview of Economic Development in Morocco.	101
3.5	Instrument Construction: Exposure to Optimal Geographic Conditions.	107
3.6	Employment Growth Predictions (2004-2014)	112
3.7	Employment Growth Predictions (2014-2024)	113
3.8	Regression Coefficients of Distance-Restricted Market Potential Com- putations.	115
3.9	Regression Coefficients of Market Potential Computations with Varying Elasticities.	117
3.10	Machine Learning Computed Exposure to Optimal Geographic Condi- tions.	118
3.A.1	Geographic Conditions in Morocco	124
3.A.2	Instrument Construction: Fulfilled Geographic Requirements	125
3.A.3	Geographic Control Variables	127
3.A.4	Census Data Illustration	129
3.A.5	Share of Males in Morocco	130
3.A.6	Share of Children Age 0-14	131
3.A.7	Age First Marriage	131
3.A.8	Share Primary/Secondary Education	132
3.A.9	Fertility Rate	132
3.A.10	Tourism Development: Hotel Rooms	133
3.C.1	Least-Cost Path from Azilal to Casablanca	141
3.C.2	Province Density Maps of XGBoost-Predicted Oasis Probabilities . . .	143



Der Student über dem Nebelmeer

Introduction

Starting a PhD was accompanied by clear expectations of what was coming next. These expectations were soon to be updated with new information that made me realize how big of a universe the field of Economics actually is. There are countless avenues one can pursue, but at times these opportunities can feel like a maze. Starting out as someone interested in the typical themes of Macroeconomics, I gradually realized that what really fascinates me is the challenge at the root of research questions in basically every field: The quest of constructing a robust identification strategy to answer some $X \text{ causes } Y$ hypothesis. In Economics, and other social sciences, this can be quite difficult at times since many relevant questions of our field cannot be answered in an experimental setup. The solution lies in *quasi-experimental* research designs that attempt to exploit an exogenous variation to answer a research question. The process of deriving the causal effect of X on Y in a complex environment is called *Causal Inference*. Since my chapters happen to be on categorically very different topics, and given the fact that each chapter emphasizes the construction of a robust identification strategy, I feel the most adequate title for this dissertation summarizing the chapters and my passion within the field of Economics is *Essays in Causal Inference*.

Books like *Mostly Harmless Econometrics* by [Angrist and Pischke \(2009\)](#) and *Causal Inference* by [Cunningham \(2021\)](#) summarize the beauty and details in this field, so I will refrain from going into further detail here. Rather, I want to highlight the impact this field had on me beyond the academic scope. In our day and age, which is arguably becoming increasingly polarized, we more often than not encounter some of these $X \text{ causes } Y$ statements. At first, they seem reasonable and are supported by aggregate correlations or even some observational studies. These statements are internalized over time and henceforth considered a premise. Simple examples I have often thought about include the correct way for humans to walk or the optimal food to eat. It becomes more complicated once an $X \text{ causes } Y$ statement is made about some political issue or real social dynamics. I abstain from further specification of such examples and leave it to the

reader's imagination. Nevertheless, it has become like an obsession to profoundly think of basic assumptions I had about life, collect information, and adapt my beliefs accordingly to what actually seems to be true. Admittedly, way too often I came across some *scientific evidence* that was communicated as the source for some *X causes Y* statement that barely fits any standards I have come to know throughout the PhD. I do not insinuate bad will. Rather I have accepted that many statements cannot be proved rigorously (yet).¹ This realization made me aware about the necessary scrutiny and humility in the quest for truth and I intend to carry that spirit long after the PhD.

But these elaborations shall suffice for now. In the following, I will briefly present the chapters of the dissertation, in which I have carefully addressed the quest of constructing a robust identification strategy across a variety of topics. I hope that each of these chapters can help to make a – small – step towards justifying or rejecting some of these *X causes Y* statements, respectively.

Chapter I, co-authored with Bjarne Horst, presents a novel method to identify the causal effect of supply chain disruptions on waterways. Over the past years, the issue of global supply chain disruptions has been analyzed intensively in the economics literature (compare [Bai et al., 2024](#); [Baqae and Farhi, 2022](#); [Acemoglu and Tahbaz-Salehi, 2025](#)), specifically following the COVID-19 pandemic and the Russian invasion of Ukraine. However, less attention was paid to regional supply chain disruptions that might not affect global container volumes but still affect national economies. We exploit variations in shipping capacity on waterways induced by high and low water level events, which are exogenous to the current state of the economy, and adjust for anticipation effects. Our method provides a measure that allows for a consistent comparison of supply chain disruptions at basically any waterway system with available water level and trade flow data. We employ our method to the river Rhine and analyze the effects on the German economy. Although the Rhine accounts for only 5–8% of goods transported in Germany, transportation restrictions on the Rhine have a significant negative effect on the industrial production throughout Germany. This effect is more pronounced in energy-intensive sectors. We find a link to regionally heterogeneous effects on energy prices. Also, we find that these supply chain disruptions were followed by a relatively persistent increase in petroleum imports from Russia, likely to substitute for interrupted imports via the Rhine. This finding raises questions regarding future welfare effects of waterway transportation disruptions on the Rhine, since such a substitution pattern seems politically less feasible in the near future.

¹Ultimately, and especially in social sciences, many statements can be traced to some underlying assumptions about the very nature of what constitutes a human being.

Chapter II dives into the niche of sport economics and presents an empirical design to study the causal effect of positive income shocks on amateur soccer clubs. There is already a broad literature that analyzes how different kinds of economic agents react to income shocks (compare [Hall, 1978](#); [Lindqvist et al., 2020](#); [Brollo et al., 2013](#)). However, these theoretical predictions do not always match the data (e.g., [Deaton, 1991](#); [Malmendier and Nagel, 2011](#)). I present a novel example, in which a positive income shock does not produce the predicted outcomes in terms of sporting success of amateur soccer clubs. To this end, I exploit the institutional settings in German amateur soccer and utilize penalty shootouts in regional tournaments, the *Verbandspokal*, as a quasi-random decision process. I present evidence of longer team manager tenure and propose a link to the Peter principle, a decision-making bias where past performance is overvalued compared to other observable characteristics in the evaluation process (compare [Peter and Hull, 1969](#)).

Chapter III explores the role of domestic market potential, which is a measure for the access to nearby markets. Historically, market potential was a major determinant in the spatial allocation of economic activity, and regions with a high market potential, like the Rust Belt in the US or Rhein-Ruhr in Germany, experienced a substantial increase in economic activity in the past. However, more recent research based on observational data suggests that the relevance of market potential is declining ([Brühlhart et al., 2020](#)). The remaining question is whether this observation might be driven by institutional factors and locational fundamentals or actually captures the causal effect, which highlights the common identification issues in literature on market potential ([Redding, 2022](#)). In the past, one approach to identify an exogenous variation was to exploit innovations in transportation technologies that affected market potential (e.g., [Jacks and Novy, 2018](#)). I present a novel approach that attempts to capture an exogenous variation in the location of modern cities in the emerging economy of Morocco. By exploiting information about the geographical features and human cultivation of oases in North Africa, I construct an instrumental variable for domestic market potential. Employing this instrument in a 2SLS framework reveals that market potential had no significant effect in the past decade on employment growth as suggested by naïve OLS estimates.


Although Chapters II and III are single-authored, the first-person plural (“we”) is used throughout for clarity and stylistic consistency. This usage should not be read as implying joint authorship.

For readers of the print version who want to access the links mentioned in the footnotes, an additional resources section is provided at the very end of the document.

References

- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2025): “The Macroeconomics of Supply Chain Disruptions,” *The Review of Economic Studies*, 92, 656–695.
- ANGRIST, J. D. AND J.-S. PISCHKE (2009): *Mostly harmless econometrics: An empiricist’s companion*, Princeton university press.
- BAI, X., J. FERNÁNDEZ-VILLAYERDE, Y. LI, AND F. ZANETTI (2024): “The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory,” *NBER Working Paper 32098*.
- BAQAEE, D. AND E. FARHI (2022): “Supply and demand in disaggregated Keynesian economies with an application to the Covid-19 crisis,” *American Economic Review*, 112, 1397–1436.
- BROLLO, F., T. NANNICINI, R. PEROTTI, AND G. TABELLINI (2013): “The political resource curse,” *American Economic Review*, 103, 1759–1796.
- BRÜLHART, M., K. DESMET, AND G.-P. KLINKE (2020): “The shrinking advantage of market potential,” *Journal of Development Economics*, 147, 102529.
- CUNNINGHAM, S. (2021): *Causal inference: The mixtape*, Yale university press.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–48.
- HALL, R. E. (1978): “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence,” *Journal of Political Economy*, 86, 971–987.
- JACKS, D. S. AND D. NOVY (2018): “Market potential and global growth over the long twentieth century,” *Journal of International Economics*, 114, 221–237.
- LINDQVIST, E., R. ÖSTLING, AND D. CESARINI (2020): “Long-run effects of lottery wealth on psychological well-being,” *The Review of Economic Studies*, 87, 2703–2726.
- MALMENDIER, U. AND S. NAGEL (2011): “Depression babies: do macroeconomic experiences affect risk taking?” *The Quarterly Journal of Economics*, 126, 373–416.
- PETER, L. J. AND R. HULL (1969): *The peter principle*, vol. 4, Souvenir Press London.
- REDDING, S. J. (2022): “Trade and geography,” *Handbook of International Economics*, 5, 147–217.

Chapter I



Quantifying Waterway Supply Chain Shocks: Regional Propagation in the Rhine Area*

Bjarne Horst

Yann Müller

October 2025

Abstract

Transportation restrictions on waterways due to high or low water level events lead to disruptions of supply chains, which are exogenous to the current state of the economy. This paper proposes a novel method to exploit and quantify these surprising transportation restrictions which lead to regional supply chain disruptions and applies the method to the river Rhine. A surprising decrease of the Rhine's shipping capacity leads to a short-lived but significant decrease in economic activity, not only in the bordering federal states but entire Germany. This effect is more pronounced in industries and regions that rely more heavily on the Rhine and the goods shipped on it. One channel through which these disruptions propagate are changes in energy prices. We find that energy marketplaces that are dependent on Rhine transportation show a price increase while others do not react. Also, we document a substitution towards suppliers that do not rely on the Rhine to deliver their goods.

Keywords: Regional Supply Chain Disruptions, Waterway Restrictions, Rhine.

JEL Codes: E32, Q41.

*Horst: Department of Economics, University of Mannheim, bjarne.horst@uni-mannheim.de. Bjarne Horst acknowledges financial support from the German Research Foundation (DFG) through CRC TR 224 (Project C02) and from Stiftung Geld & Währung. Müller: Department of Economics, University of Mannheim, yann.mueller@uni-mannheim.de. Special thanks for guidance, support, and useful advice to Miren Azkarate-Askasua, Ursula Berresheim, Antonio Ciccone, Thibault Cézanne, Harald Fadinger, Andreas Gulyas, Lukas Hack, Marina Hoch, Minki Kim, Matthias Meier, Ana Moreno-Maldonado, Max Riegel, Ulrich Roschitsch, Hannes Twieling, and all participants of the ENTER Jamboree 2025 in Stockholm, the Mannheim Doctoral Colloquium, and the Mannheim Macro Lunch.

1.1 Introduction

In recent years, there has been a surge in academic efforts to understand the effects of supply chain disruptions, but a robust identification of these shocks has turned out to be challenging. [Bai et al. \(2024\)](#) present a promising approach to capture the causal effect of global supply chain disruptions. Yet, supply chain disruptions do not necessarily only occur at the country-level, but also between regions within a country. The lack of convincingly identified regional supply chain shocks in the literature has prevented research from exploring the resilience of countries to these shocks.¹ Therefore, we propose a novel microfounded approach to quantify regional supply chain shocks by exploiting an exogenous variation in the transportation capacity of waterways.

To this end, we compute a shock measure that captures transportation capacity restrictions caused by high- and low-level water events and employ our method to the Rhine, the most important river in Germany in terms of transported goods volume. These waterway transportation restrictions, caused by high or low water level events and exogenous to the current state of the economy, occur with notable frequency. This allows us to provide a comprehensive and microfounded analysis of the economic consequences of disruptions on the Rhine, which is also replicable to other waterway systems.

The objective of our research is to assess the macroeconomic impact by estimating the response of industrial production to transport capacity restrictions in a local projection framework. We find a significant temporary contraction in industrial production of 0.14% following a waterway transportation capacity shock of one standard deviation. This negative effect is larger and more persistent in energy-intensive sectors. We also find evidence for a regionally dispersed effect on producer energy prices. Against our supposition, we do not find any transportation substitution from cargo ships to trucks or trains on impact, which might be due to the typically short duration of restrictions in combination with the special nature of the goods transported. However, we find a substitution pattern with respect to the countries of origin for the import of fossil fuels. On impact, imports from Russia increase, which are primarily transported via pipelines. In contrast, imports from the Netherlands, which heavily depend on waterway transportation, decrease. These findings raise questions about future welfare effects of regional supply chain shocks since fossil fuel imports from Russia to the EU have been halted in 2022.

Related Literature. Following the transformation towards international value chains and increased trade in intermediate inputs since the 1980s ([Antràs and Chor, 2022](#)), there has been a growing interest in studying the macroeconomic impact of global supply chain

¹There is no formal definition of global and regional supply chain shocks, and these concepts could be applied both with magnitude or geographic requirements. Within the scope of our paper, we make the distinction based on the effect on global container volumes: Global shocks have a significant effect on global volumes, while regional shocks do not.

disruptions. The empirical literature attempts to identify exogenous variations in global supply chains by utilizing SVAR models with sign and narrative restrictions (e.g., [Finck and Tillmann, 2022](#); [Bai et al., 2024](#)). Their findings show that these disruptions can result in a significant reduction in trade and economic output. Another avenue of the literature analyzes the impact on global supply chains induced by trade cost shocks (e.g., [Feyrer, 2021](#)). Since the onset of the COVID-19 pandemic, policy makers and academic researchers became more keen about consumer goods shortages caused by supply chain disruptions and congestion effects. [Meier and Pinto \(2024\)](#) provide empirical evidence on the economic effects following the anti-COVID-19 measures. [Bonadio et al. \(2021\)](#) find that a one-sided repeal of these measures in large economies could lead to a GDP growth of up to 2.5% for their smaller trading partners.²

Another strand of the literature examines regional events that propagate to the global economy, as outlined in [Elliott and Golub \(2022\)](#). For example, [Carvalho et al. \(2021\)](#) demonstrate how the propagation of the Great East Japan Earthquake was shaped by input–output linkages. Similarly, [Barrot and Sauvagnat \(2016\)](#) analyze major natural disasters in the US and their transmission through production networks.³ However, to the best of our knowledge, there is no approach to capture systematic exogenous supply chain disruptions at a specific location. [Ademmer et al. \(2023\)](#) have previously proposed using water level data at single measuring stations as a proxy for trade flow disruptions. While the method can provide some suitable intuition for the effect of low water events, it does not account for potential sources of bias such as regionally heterogeneous transportation restrictions and anticipation effects. Having this in mind, our approach accounts for several dimensions of heterogeneity by constructing a shock series that utilizes data from multiple measuring stations along the river and incorporates a trade flow weighting system to account for the relevance of specific passage closures. Consequently, our project will complement the literature by providing a methodology to construct a time-varying measure that captures regional supply chain shocks by exploiting waterway transportation restrictions.⁴

The structure of the paper is as follows. First, we provide an overview of the Rhine river and the trade flows it supports. Next, we present our approach for quantifying transport restrictions on the Rhine using administrative data. We incorporate these findings into a local projection framework to estimate the macroeconomic impacts of waterway restrictions. Finally, we conclude the paper with a summary of our findings.

²The pandemic has also induced a surge in the theoretical literature regarding supply chain disruptions (e.g., [Baqae and Farhi, 2022](#); [Bigio and La'o, 2020](#))

³Firm-level linkages have been proposed as major determinant in the transmission of these disruptions (e.g., [Acemoglu and Tahbaz-Salehi, 2025](#); [Elliott et al., 2022](#); [Alfaro-Urena et al., 2022](#)).

⁴Our project also relates to papers from the field of transport economics, which have analyzed the impact of low water levels and the interaction with climate change on waterway transport companies (e.g., [Jonkeren et al., 2007](#); [Beuthe et al., 2014](#)).


1.2 Data

Our approach requires daily data on water levels at multiple measuring stations and information on the trade flows at specific sections of the waterway. We have chosen the Rhine to conduct our analysis, which is a major European river running through Switzerland, France, Germany, and the Netherlands.⁵ It is one of the longest rivers in Europe and has a total length of approximately 1,230 kilometers. The Rhine is an important transportation route and largely used for primary and intermediary goods. The following subsections present information regarding the traded goods on the Rhine together with water level data. This is indicative for which sectors and regions one should expect to be more severely affected by transportation restrictions.

1.2.1 Trade Flow

To construct the capacity shock series and to enhance the interpretation of our results, it is useful to have an impression of the transportation patterns along the Rhine. This subsection presents trade flow data for the Rhine obtained from the *Federal Statistics Office Germany*. The data is available at a yearly frequency and contains information regarding the transport volume of 20 goods categories in seven Rhine subsections.

Figure 1.1: Overview of Goods Transported for each Rhine Section.




Note: Trade flows on the seven administrative Rhine sections - ordered from South to North - in 2021. Goods categories according to NST-2007 classification. (Source: Own illustration)

Figure 1.1 summarizes the data. The horizontal axis reports the different Rhine sections. The corresponding bars indicate the respective trade volume per category, whereby the four most important categories (in terms of volume) are singled out from the remaining

⁵Other rivers frequently encountering waterway transportation restrictions stemming from high or low water levels, and thus suitable for our approach, include the Mississippi in the United States, the Danube in Europe, and the Yangtze River in China.

Figure 1.2: Means of Transportation for certain Goods Categories.



Note: This figure shows the share of goods transported by different means of transportation. The shares are computed relative to the total ton-miles in that goods category. (Source: Own illustration)

categories.

Two observations stand out. First, the largest share of transported goods is processed in the two sections of the Lower Rhine close to the Dutch border. Second, the most important goods categories are primary and intermediary goods, especially energy sources. A detailed breakdown for each goods category at the specific Rhine sections is presented in Appendix 1.B.

The Rhine accounts for 86% of the amount of goods shipped on German inland waterways. This amount constitutes to 5%-8% of the total goods transportation in Germany.⁶ This share, however, varies considerably across goods categories. Figure 1.2 highlights the differences in means of transportation for specific categories. The Rhine's four most important goods categories are disproportionately transported through waterways, with the exception of ores and mining products. For example, more than 40% of coal, petroleum, and natural gas are transported via waterways. Additionally, according to the freight statistics of German inland shipping, the majority of goods transported via the Rhine are shipped to or from overseas ports in the Netherlands.⁷

1.2.2 Water Levels

In the following subsection, we present the water level data employed in our analysis to quantify high and low water events. We obtain raw water levels at a daily frequency from the *Federal Waterways and Shipping Agency (WSV)*, which provides information for all administratively relevant measuring stations along the Rhine.⁸ Additionally, we collect information from legal and historical documents about the reference water level, the depth


⁶Compare press statement #148 from the Federal Statistics Office Germany in April 2022. [[Online Accessible](#) - Last accessed: 27.08.2025]

⁷More information available in the freight statistics reports of German inland shipping. [[Online Accessible](#) - Last accessed: 27.08.2025]

⁸Daily water levels represent the average of three measurements each day at the same time.

at the reference level and the high level threshold at each station. After dropping stations without such specified values, we obtain a sample of 13 stations with data available from November 1971 to October 2022. The locations of the measuring stations are depicted in Figure 1.3.

Figure 1.3: Administratively relevant measuring stations along the Rhine.



Note: Measuring stations (indicated by red dots) for which reference water level, depth at the reference level and the high level threshold are available. From South to North: Maxau, Speyer, Mannheim, Worms, Mainz, Kaub, Koblenz, Andernach, Köln, Düsseldorf, Ruhrort, Wesel, Emmerich. (Source: Own illustration)

However, raw water levels are not comparable across stations.⁹ Therefore, we compute the *fairway depth*, which is comparable across time and locations, as follows:


$$depth_{dj} = level_{dj} - reference_level_{dj} + reference_depth_{dj},$$

where $level_{dj}$ is the raw water level at day d measured in location j . The $reference_level_{dj}$ is a reference value determined every ten years for each station representing the water level at which the $reference_depth_{dj}$ is secured, which corresponds to the fairway depth

⁹Water levels refer to the tide gauge zero which does not contain comparable information regarding the passability for cargo ships.

at the reference water level. Figure 1.4 illustrates these nautical definitions.

Figure 1.4: Nautical Definitions.



Note: Illustration of how the nautical definitions of the fairway depth, the reference depth, the reference level, and the water level are related. (Source: Own illustration)

1.3 Shock Construction

The raw water level and fairway depth data presented in the previous section do not provide direct insight into the potential supply chain disruptions following low or high water events on the Rhine. To this end, this section proposes a two-step approach to quantify surprising shipping capacity restrictions. In the first step, we introduce a measure of shipping capacity for a given day and Rhine section. In the second step, we aggregate and clean this measure to obtain a monthly series for the German Rhine that accounts for seasonal and autoregressive effects.

1.3.1 Capacity Restrictions

Two important challenges arise in quantifying shipping capacity restrictions. First, the measurement has to be consistent for high as well as for low water events. Second, we do not only want to determine the presence of a restriction but also quantify its severity. To resolve these problems, we define the daily shipping capacity at day d and location j as $cap_{dj} \in [0, 100]$. By normalizing the magnitude in this way, we measure the shipping capacity of the respective Rhine section in percent relative to a day where no legal restriction was in place. Three different events are possible:

For high water events there is a clear legal framework. If the water level at a station j exceeds the legal high water threshold $\bar{\ell}_j$, shipping the respective section is prohibited (§ 10.01 RheinSchPV). Consequently, we set

$$cap_{dj} = 0 \quad \text{if } level_{dj} > \bar{\ell}_j$$

For low water events, the restrictions evolve more smoothly. For instance, even if the water level is unusually low, it might still be possible to ship 50 % of the goods that would have been shipped under normal conditions. Crucially, there are no sharp legal thresholds for low water closures; rather, each skipper decides on the draft of their ship based on the current fairway depth. To formalize this relationship, we define a continuous function f capturing (effective) loading restrictions due to low water such that:

$$cap_{dj} = f(depth_{dj}) < 100 \quad \text{if } depth_{dj} < \underline{\ell}$$

We determine the function's properties by exploiting information on low water events from Contargo.¹⁰ $\underline{\ell}$ is the lowest fairway depth at which there is just no restriction in place.¹¹

In all other cases, no restriction is in place, i.e., we define

$$cap_{dj} = 100 \quad \text{otherwise}$$

Figure 1.5 shows the historic restrictions for three selected locations. For better accessibility, the figure only shows the total number of days with restrictions (i.e., days with $cap_{dj} < 100$) within each year. The key observation is a very pronounced regional heterogeneity: While Kaub faced 24 years with more than 50 days with restrictions, there is only a single year where there were more than 50 days with restrictions in Duisburg.¹² However, as discussed in section 2, there is significantly more traffic on the Rhine sections around Duisburg. Consequently, even few days with restrictions there have a large impact on the total shipping volume of the Rhine.

1.3.2 Shock Construction


Aggregation. Until now, the capacity measure we have constructed is location-specific. However, to examine, for example, how industrial production in Germany responds to Rhine capacity restrictions, it is necessary to have an aggregate capacity series for the entire German Rhine. Given the pronounced regional heterogeneity with respect to traffic shares (Figure 1.1) and restrictions (Figure 1.5), selecting a single measuring station only gives an incomplete image of the current shipping restrictions on the German Rhine. Hence, we compute the aggregate capacity measure by taking a weighted average of the location-specific measures. The weights depend on two factors:

¹⁰Contargo is a major Rhine shipping company that published information on low water surcharge. [[Online Accessible](#) - Last accessed: 27.08.2025]

¹¹We fix it at the fairway depth of 260cm at which Contargo starts charging low water fees. We set $f(depth_{dj}) = 0$ once Contargo is not obliged to ship anymore, which corresponds to a fairway depth of ≈ 190 cm. In between we interpolate linearly, again following Contargo.

¹²Kaub is located in the flattest German Rhine section between Mainz and Koblenz.

Figure 1.5: Number of Days with Restrictions at Selected Stations.



Note: Number of days with transportation restrictions per year caused by high or low water level events for three selected measuring stations that are located along the Rhine. (Source: Own illustration)

1. Restrictions in sections with high traffic have a greater impact on the shipping capacity of the entire German Rhine compared to restrictions in sections with low traffic. Therefore, we use the data set on trade flows introduced in section 1.2.1 to compute the traffic share of each Rhine section (tons of goods shipped in that section divided by tons of goods shipped on all German Rhine sections)¹³ to quantify the importance of each station.
2. Since the Rhine sections in the data set at hand are broader than the legal Rhine sections, we frequently observe multiple stations within the same section. To prevent overstating the role of these stations, we correct for the number of stations within each section.


Hence, the weights we use are given by $\frac{\text{traffic share in section}}{\#\text{stations in section}}$. Since all outcome variables of interest are at a monthly frequency, we also compute monthly averages of the daily measure to obtain the monthly aggregate capacity series cap_t .

Figure 1.6 shows the resulting time series from 1990 until 2022. The series still has the interpretation of being the Rhine's shipping capacity relative to a month in which there was no restriction in place at any day or station. Three sources of variation drive the series: (I) Variation in the number of days and (II) the number of stations at which there was a restriction, as well as (III) the severity of these restrictions as outlined in subsection 1.3.1.

The average shipping capacity is approximately 95%. The drought between July and December 2018 stands out as the most extreme event, with a capacity reduction down to 23% at the peak in November. Appendix 1.B.4 contains more disaggregated figures of

¹³Given the data set, we compute these weights for all years between 2010 and 2020 and then average over the years. While the absolute quantities vary a lot over time, the shares are relatively stable.

Figure 1.6: Aggregate Monthly Capacity



Note: The figure depicts the aggregate monthly transportation capacity on the Rhine. The horizontal bar in black indicates the mean capacity. (Source: Own illustration)

the capacity measure, which highlight that the shipping capacity is mainly driven by low water events and not by high water events (which occur less frequently and last only for very few days).

Anticipation Effects. Importantly, some of the fluctuations visible in Figure 1.6 might be expected in advance. For example, agents might expect low water events to be more likely in fall than in spring. Moreover, the capacity in a given month might have predictive power for the capacity in the following month. Conducting an impulse response analysis in the following section requires to control for these effects. We do so by regressing the capacity measure on a full set of monthly dummies (to address seasonal effects) and on its first lag (to address its predictive power):

$$cap_t = \alpha'_0 \mathbf{month}_t + \alpha_1 cap_{t-1} + s_t$$


In the upcoming section, we use the residual of this regression s_t as our shock measure, which is illustrated in Appendix 1.B.8. Of course, other specifications of this step are possible, e.g., including temperature or precipitation data. We have conducted a series of robustness specifications and our findings remain unchanged. The results for these specifications can be found in Appendix 1.A.1.

1.3.3 Shock Diagnostics

Ensuring the validity of our constructed shock measure requires the fulfillment of two conditions. First, it has to capture the decrease in transportation capacity and no

confounding factors. And second, it should be exogenous to the current state of the economy and not predictable by other economic or climate indicators.¹⁴ With respect to the first condition, we use our shock series in a local projection framework and estimate the impulse response of cross-border receiving after a negative Rhine capacity shock of one standard deviation ($\approx 8.3\%$).¹⁵ Figure 1.7 plots the results for two common waterway transportation measures, namely freight ton-miles and transported goods in tons. On impact, we observe a significant contraction by about 6.5% and 4%, respectively. In the subsequent months, the outlined transportation measures recover; being not significantly different from zero two months after the shocks onwards, which highlights the short-lived nature of our identified shocks. These findings are reasonable considering the magnitude of the shock and confirm that we are actually capturing transportation contractions induced by unexpected waterway transportation restrictions.¹⁶

Figure 1.7: IRF of Goods Transportation on the Rhine (Cross-Border Receiving)



Note: IRFs (dark blue) of freight ton-miles (left panel) and transported goods (right panel) after a one standard deviation negative Rhine capacity shock following the empirical framework outlined in equation (2). 90% confidence intervals depicted in light blue. (Source: Own illustration)

The second condition requires us to ensure that our shock series is exogenous to the current state of the economy and not predictable because otherwise firms could adapt

¹⁴This approach is related to [Ramey \(2016\)](#), who states that a robust series of high-frequency surprises must fulfill several properties: (I) No auto-correlation, (II) no predictability, and (III) no correlation with other shocks.

¹⁵We have chosen cross-border receiving, which are majorly primary and intermediate goods, and therefore, should induce cascade effects within their supply chains. German exports shipped on the Rhine, however, are primarily final goods and should not induce cascade effects within the value chain.

¹⁶Another concern could be that Rhine capacity restrictions are no fundamental shocks but rather driven by temperature or precipitation shocks that also impact the economy through other channels. In section 1.4.2, we outline our empirical approach that shows how our results are actually driven by waterway transportation restrictions.

their behavior.¹⁷ Given that waterway transportation restrictions critically depend on hydrological events which occur independently of economic activities, we are reasonably confident that our shock series is exogenous to the current state of the economy.¹⁸ However, the most obvious threat to our identification lies in the possibility that water levels, and therefore, waterway transportation restrictions can be predicted. Clearly, a certain part of this threat is addressed by our shock construction itself. We identify a variation in transportation restrictions that are orthogonal to information obtained from the previous period and the seasonality in which the shock occurs.¹⁹ A remaining threat to our identification arises if it would be possible to construct valid long-term predictions for transportation restrictions. There are state-of-the-art machine learning approaches that attempt to forecast water levels months ahead.²⁰ However, both the precision of water levels predictions and even more so the predictions for extensive and intensive margins of transportation restrictions seem to be quite imprecise for a longer horizon. Therefore, we remain confident with respect to the validity of our constructed shock measure.

1.4 Empirical Framework

Given the shock series constructed in the previous section, it is natural to ask for its propagation effects on the economy. More precisely, how does the economy respond to an unexpected change in the Rhine's shipping capacity? To this end, we introduce the relevant outcome variables (1.4.1) and our empirical framework (1.4.2).

1.4.1 Outcome Variables

The first collection of data sets we use contains the index of industrial production, which is a monthly index of the output produced in four 1-digit sectors (Mining & Quarrying, Manufacturing, Energy Supply, Construction).²¹ On the federal level, we obtain data from January 1991 to December 2019 at varying granularity. Since we also want to examine regional heterogeneity in the responses, we also collect data at the state level. Currently, we have data for two out of the four Rhine neighboring states (Baden-Württemberg and North Rhine-Westphalia) from January 2000 to December 2019. Unfortunately, the data is less granular than that at the federal level.

¹⁷For example, if waterway transportation restrictions occur quite frequently in a given month, companies might be inclined to ship their goods in earlier or later months. In such a case, the impulse response would not capture a reaction to the shock but rather a transportation pattern that is observable irrespective of the occurrence of waterway transportation shocks.

¹⁸At this point, one could also check the autocorrelation with monetary, fiscal, or technology shocks. Since it appears very unlikely to find any correlation with these shocks, we did not conduct this analysis.

¹⁹Our robustness checks also include additional information from the period in which the shock occurs.

²⁰[Wee et al. \(2021\)](#) provide an overview of machine learning algorithms to forecast water levels. [Yueling et al. \(2019\)](#) apply one of these models to make within-month water level forecasts.

²¹See Appendix 1.B.3 for a breakdown of the relative sizes of these sectors.

To examine changing transportation patterns in response to Rhine capacity reductions, we collect data on the truck index at the federal level (January 2009 - December 2019) and train transportation data at the federal level (January 1991 - December 2022).²² In the given context, there are two effects on trucks and trains utilization after a Rhine capacity reduction: On the one hand, there is a positive input effect, i.e., firms demand more of these services to cope with the reduced shipping capacity of the Rhine on the input side. On the other hand, there is a negative output effect, i.e., firms produce less, so they need fewer truck and train services to transport their output to their customers. Since our interest lies in the first effect, we will discuss in section 1.4.2 how it can be isolated empirically. In the same vein, we scrutinize the effect on imports of specific goods that are primarily transported on the Rhine. Since we know that a large share of these goods are fossil fuels, we expect to observe some substitution patterns in order to mitigate the impact on the economy. For this exercise, we collected data on the imports of (un-)refined petroleum and crude oil by the country of origin (January 2008 - October 2022).²³ By comparing changes in the import behavior, we attempt to unveil substitution patterns following a waterway transportation restriction. Finally, we want to analyze how prices react to our constructed shock series. To this end, we collect PPI components of goods that are primarily transported on the Rhine (January 2000 - October 2022).

1.4.2 Empirical Strategy

As we already have a series of shocks capturing unexpected changes in the Rhine's shipping capacity, the most natural empirical strategy to estimate impulse response functions is the local projection framework (e.g. [Jordà, 2005](#); [Plagborg-Møller and Wolf, 2021](#)). Since all outcome variables are highly persistent, a model in first differences seems appropriate.²⁴ The local projection framework offers a convenient way to directly estimate cumulative impulse response functions (IRF from now on) with a model in differences in the background. More concretely, the model for the industrial production index reads as

$$\underbrace{\Delta^{h+1} ip_{t+h}}_{=ip_{t+h}-ip_{t-1}} = \alpha^h + \beta^h s_t + \gamma^h \underbrace{\Delta ip_{t-1}}_{=ip_{t-1}-ip_{t-2}} + v_t, \quad \text{for } h = 0, 1, \dots, 12 \quad (1)$$

where ip_t denotes the logarithm of industrial production in period t . In words, we consider a shock to Rhine shipping capacity occurring in period t and its effect on the industrial production index h periods after the shock has occurred. The regressions are augmented by the first lag of the first difference of (log) industrial production following [Ramey \(2016\)](#) and [Olea et al. \(2025\)](#). We run a sequence of these regressions up to a horizon of 1 year.

²²The truck index is computed based on toll data on kilometers of trucks driven on German highways.

²³Unfortunately, data for natural gas imports was not retrievable from the *Federal Statistics Office Germany* its database at the time we have conducted our analysis.

²⁴An ADF test does not reject the null hypothesis of a unit root for any of the series we consider.

By considering log differences of industrial production, the sequence of IRF coefficients $\{\beta^h\}_{h=0}^{12}$ measures the response of industrial production in percent. For the graphical representation later on, we multiply by 100 to obtain results in percentages. We also compute standard errors robust to serial correlation following [Newey and West \(1987\)](#), which are appropriate for this type of model according to [Jordà \(2005\)](#).²⁵

The models for the other (log) outcome variables, i.e., means of transportation and imported goods (y in the following), are specified in a similar way:

$$\underbrace{\Delta^{h+1}y_{t+h}}_{=y_{t+h}-y_{t-1}} = \alpha^h + \beta^h s_t + \gamma^h \underbrace{\Delta y_{t-1}}_{=y_{t-1}-y_{t-2}} + \delta^h \underbrace{\Delta^{h+1}ip_{t+h}}_{=ip_{t+h}-ip_{t-1}} + v_t, \quad \text{for } h = 0, 1, \dots, 12 \quad (2)$$

Note that the regressions are augmented by the change in (log) industrial production between one period before and h periods after the shock has occurred. As discussed in section 1.4.1 and similar to the previous equation, this additional term is added to isolate the input effect from the output effect. Put differently, by controlling for the effect of changing output on the demand for transportation services or imported goods, β^h only captures the effects related to changing transportation or imported goods patterns caused by supply chain disruptions.

Causality. Before turning to our empirical results, we want to address the causality of our estimates. The largest Rhine capacity shocks in our sample (e.g., that of fall 2018) are caused by unexpectedly hot and dry summers. Hence, it is not entirely clear whether the estimated coefficients actually capture the supply chain disruptions related to the Rhine capacity reduction or just weather effects on output (e.g., declining labor productivity due to heat). One possible defense of interpreting our shocks as supply chain disruptions is the lagging behavior of water levels. Referring back to figure 1.6, nearly all of the most extreme capacity reduction (which will govern the estimators' behavior) reach their peak in fall (September - November), when the most severe heat is already over. Since agricultural sectors are not included in the industrial production index, additional effects of low precipitation seem unlikely. Nevertheless, we also conduct a series of robustness checks and include controls for temperature and precipitation in the baseline regressions. The main findings remain unchanged and are presented in Appendix 1.A.2. Therefore, we are reasonably confident that our approach isolates the causal channel of regional supply chain disruptions via waterway transportation restrictions.

²⁵Our results are qualitatively robust to other standard error specifications (e.g., [Herbst and Johannsen, 2024](#)).


1.5 Empirical Results

The following section presents the effect of waterway transportation restrictions on industrial production, truck transportation, train transportation, energy imports, and energy prices.

1.5.1 Industrial Production

This subsection discusses the response of the industrial production index following a Rhine capacity shock. The presentation is organized from broad to specific: We first discuss aggregate industrial production in Germany. In a second step, we explore regional and sectoral heterogeneity in the response patterns.

Figure 1.8: IRFs of German Industrial Production




Note: IRFs (dark blue) of industrial production (left panel) and industrial production of energy-intensive sectors only (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

The left panel of Figure 1.8 shows the impulse response function of the industrial production index in Germany after a negative Rhine capacity shock of one standard deviation ($\approx 8.3\%$). On impact, industrial production decreases by -0.14% . This point estimate is significantly different from zero at the 90% level. However, from one month after the shock onwards the response is no longer significantly different from zero. As shipping on the Rhine mainly involves energy sources, an obvious extension is to focus on sectors that are more heavily reliant on these goods. Therefore, the right panel in Figure 1.8 only includes those subsectors, which are energy-intensive in their production processes. As expected, the impact response of -0.19% is more pronounced than that of the aggregate industrial production index and decreases even further down to -0.29% in the following month. Both

of these coefficients are significantly different from zero, but none at horizons larger than one are.

Figure 1.9: IRF of Industrial Production for Different States and Sectors



Note: The figure shows the IRFs of the industrial production index after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany, NRW: North Rhine-Westphalia, BW: Baden-Württemberg). (Source: Own illustration)

To explore how these results differ across regions and sectors, Figure 1.9 plots the IRF coefficients together with the 90% confidence intervals (again normalized for a negative one standard deviation Rhine capacity shock) for different state-sector pairs. The first three panels compare the results on the federal level with state-level results in North Rhine-Westphalia (NRW in the following; the largest Rhine neighboring state in terms of economic output). The fourth panel compares Germany with Baden-Württemberg (BW in the following). Since none of the coefficients is significantly different from zero from horizon two onwards, the figure only shows the responses at horizon zero and one for better accessibility. Given the shorter data sample at the state level (cf. section 1.4.1), we conduct all regressions on a sample from January 2000 until December 2019.

The first panel compares the response of the manufacturing sector of NRW to that in Germany. Since manufacturing makes up the largest part of the industrial production index (cf. Appendix 1.B.3), the results for Germany are similar to those presented in Figure 1.8: The impact response is at around -0.11% and then decreases further to -0.17%, however, none of them is statistically significant at the 90% level. The impact response of -0.47% in NRW is notably larger than that for entire Germany (and also significantly different from zero). Such a result was to be expected given that NRW appears to be

most dependent on the Rhine for transportation, as discussed in section 1.2.1. However, this difference vanishes one month after the shock.

The second panel again compares the response in NRW and Germany but for the mining sector. While the response in Germany is not significantly different from zero, the impact response in NRW (-0.68%) is the largest point estimate (in absolute terms) among all estimates presented in this section. One possible narrative is again related to the composition of goods shipped on the Rhine (cf. section 1.2.1), since mining activities require energy, chemicals, and gases.


The third panel considers the energy supply sector in NRW and Germany. Here, none of the four coefficients is significantly different from zero at the 90% level. This might be a surprise given that primarily energy sources are shipped on the Rhine. Two potential narratives might rationalize this result: On the one hand, confounding weather effects (cf. section 1.4.2) might be important here, e.g., surprisingly hot and dry periods lead to rising electricity demand. On the other hand, energy suppliers are systemically relevant and consequently better insured against these events, e.g., they might hoard energy sources to cope with shipping restrictions.

The last panel focuses on the manufacturing and mining sector jointly (due to data limitations), comparing BW and Germany. Since the mining sector is close to negligible compared to manufacturing (cf. Appendix 1.B.3), the results for Germany are very similar to those in the first panel. Surprisingly, the coefficients for BW are extremely close to zero (0.04 and -0.02) for both horizons. Referring back to Figure 1.1, there are way fewer goods transported on the Rhine in BW than in NRW (by a factor between 4 and 7), i.e. the supply chains in BW might be less dependent on the Rhine as those in NRW. Moreover, our aggregation scheme introduced in section 1.3.2 gives relatively large weights to stations in NRW. This could imply that the shock series is 'tailor-made' for NRW and thus does not capture the effective restrictions firms face in BW very well. Hence, a potential extension of the current analysis is to construct state-specific shock series using more granular trade flow data.

1.5.2 Truck Index

The results discussed above already suggest that firms are not able to perfectly switch to other means of transportation after a Rhine capacity shock (at least on impact and at equal cost). To further explore the response in transportation patterns, Figure 1.10 shows the IRF of the truck index for Germany. As discussed in section 1.4.2, these regressions control for the change in industrial production to isolate the substitution effect on the input side of production. Hence, one would expect a positive response after a negative Rhine capacity shock due to higher demand for truck services for transporting input goods.

Figure 1.10: IRF of Truck Index



Note: The figure shows the IRF of the truck index for Germany after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. (Source: Own illustration)


One apparent feature of the IRF shown in Figure 1.10 is its sluggish behavior compared to the responses of industrial production. The impact response is close to zero and then builds up over time until it reaches its peak between four and five months after the shock at around 0.15%. Afterwards, the IRF declines back to zero rapidly. Note that the point estimates are significantly different from zero only at horizons four and six (at the 90% level).

Two conclusions emerge from this exercise: First, the periods in which the truck index exhibits a strong and positive response (horizons two to six) correspond to the periods when industrial production starts to recover again (cf. Figure 1.8). Referring to the discussion in section 1.4.2, this provides suggestive evidence that the estimates do indeed capture the supply chain disruptions caused by Rhine capacity shocks and not just other channels of extreme weather events. Second, the results suggest that firms cannot flexibly switch means of transportation right after a Rhine capacity shock occurs. On the one hand, a single Rhine ship can transport as many goods as 150 trucks, which might not be available on the spot. On the other hand, coal, gas, and chemicals are not shipped in containers but on specialized ships. This makes it even more difficult to switch to trucks since special loading terminals and truck types are needed. These arguments could rationalize why the estimated response is close to zero on impact and only builds up over time.

1.5.3 Train Transportation

Another way of substituting the mode of transportation could be to transfer input goods with trains during the waterway transportation restrictions. Therefore, we present IRFs of train transportation in Figure 1.11. Similar to the analysis of the truck index response, we do not find evidence for any substitution patterns on impact. The point estimates of freight ton-miles and transported goods are positive but at no horizon significantly different from zero at a 90% confidence level. As outlined before, one possible explanation for this finding could be the special transportation requirements of the input goods which impede the short-term substitution with trains. Further, not all inland waterway ports are connected to the railway network which makes the substitution of goods to trains impossible.

Figure 1.11: IRF of Train Transportation (Cross-Border Receiving)




Note: IRFs (dark blue) of cross-border receiving via train transportation for Germany measured in freight ton-miles (left panel) and transported goods in tons (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

1.5.4 Energy Imports

Given the importance of the Rhine for the transportation of energy imports, we present evidence on the total imports and the substitution patterns with respect to the countries of origin of these goods. To this end, we compute the IRFs for total German imports of petroleum and crude oil, which are important primary goods for the German economy. Figure 1.12 illustrates the results. For unrefined imports, the impulse response in the left panel indicates a contraction on impact (-1.1%) and one month after (-2.1%) while recovering subsequently. Besides horizon 1, however, the estimates are not statistically

Figure 1.12: IRF of Petroleum and Crude Oil Imports



Note: IRFs (dark blue) of total German imports of unrefined (left panel) and refined (right panel) petroleum and crude oil after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)


significant. The pattern for refined imports in the right panel looks quite different. Here, we observe a prolonged increase of imports reaching its peak at around 2.1% at horizon three and remaining at this level until horizon six before normalizing afterwards. The point estimates are statistically significant in the periods three, five, and six. These findings are surprising. On impact, we would have expected a decrease for both kinds of goods given the relevance of the Rhine for the transportation of fossil fuels as outlined in section 1.2.1.

To rationalize the increase in refined crude oils, we compute the IRFs of German imports from each exporting country separately. Figure 1.13 presents the IRFs for the Netherlands and Russia, which had been the major countries of origin for German imports of refined petroleum and crude oil. For imports from the Netherlands, we find a statistically significant decrease of 2.4%. From horizon 1 onwards, the point estimates are approximately zero and not statistically significant.²⁶ For Russian imports, however, the results indicate a different pattern. On impact, we observe a statistically significant increase of 4.4%. Thereafter, Russian imports continued to increase up until horizon 5, reaching a peak of 8.1% and normalizing afterwards. These findings indicate that there is a substitution pattern in German imports with quite an overshooting following the negative transportation capacity shock. There are two lines of arguments that could rationalize these outcomes. First, the short-term increase in imports from Russia is only

²⁶One exception is the estimate for horizon 10. However, due to the relatively short sample size, we assume that this finding is more of a statistical nature rather than being associated with any economical meaning.

possible with certain contractual agreements that require a prolonged increase in imports of those goods. Another possible explanation is that economic agents become aware of these shocks and start stockpiling due to a precautionary motive.

Figure 1.13: IRF of Refined Petroleum and Crude Oil Imports by Country




Note: IRFs (dark blue) of total German imports of refined petroleum and crude oil from the Netherlands (left panel) and Russia (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

1.5.5 Energy Prices

Following our findings, it is reasonable to assume that the goods shortages and short-term substitution patterns are associated with increasing energy costs. Therefore, we examine whether firms are affected through price changes following the regional supply chain disruptions. To this end, we compute the impulse response for the PPI components of brown coal, gasoil, and heating oil. The results are presented in Figure 1.14. On impact we observe a significant increase of the brown coal PPI component by 0.21%. The impulse response estimates for the gasoil and the heating oil PPI components are also positive but not significantly different from zero at a 90% level. However, one month after the shock, the results for the PPI components of gasoil and heating oil indicate an increase of 0.7% and 0.67%, respectively, and are statistically significant at a 90% level. Also the point estimate of the brown coal PPI component increases further one month after the shock but is no longer statistically significant.

Given the importance of the Rhine for the transportation of these goods, the non-significance of the reaction at some horizons might seem surprising. There are two possible explanations that can lead to these results. Firstly, there is a substitution pattern and

Figure 1.14: IRF of PPI components



Note: The figure shows the IRFs of various PPI components after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany). (Source: Own illustration)


recipients are able to find alternative providers of these goods. And secondly, there is a regionally heterogeneous response to our Rhine transportation capacity shocks which leads to non-significant results in the aggregate. Regarding the former, we presented evidence that shows a substitution towards Russian fossil fuels after the occurrence of a regional supply chain shock. This observation may partly explain the aggregate PPI response. Nevertheless, we also want to provide some evidence regarding the second possible explanation. To this end, we compute the impulse response of the PPI component heating oil at various regional marketplaces.²⁷ A summary of the regional heterogeneity is illustrated in Figure 1.15. There are five marketplaces that depend on transportation via the Rhine (Düsseldorf, Frankfurt, Mannheim, Stuttgart, and Munich) and five marketplaces that are independent of heating oil transportation via the Rhine.²⁸ The split circles indicate the sign of the impulse response for every marketplace at horizon 0 and horizon 1, respectively. We observe that all Rhine-dependent marketplaces experience an increase in the PPI heating oil component on impact. Besides Düsseldorf, all of these markets also show an increase in prices at horizon 1. In contrast, none of the Rhine-independent markets shows a significant reaction on impact or at horizon 1. These results indicate that firms that are linked to the Rhine face higher energy costs following a regional supply chain disruption. This is another explanation regarding the non-significant results of the aggregate outcomes of the PPI components and showcases how companies are differently

²⁷This includes supplies of 500 tonnes or more to wholesalers or supplies in tanker trucks to consumers. Other components of the PPI which are transported are not reported at such a granular level.

²⁸The Federal Statistics Office Germany also provides the Rhine as area of reference in which the averages of Düsseldorf, Frankfurt, and Mannheim are computed.

affected by our regional supply chain shocks.

Figure 1.15: PPI Heating Oil by Marketplace




Note: The figure indicates the signs of the IRFs for the PPI component heating oil at various marketplaces after a one standard deviation negative Rhine capacity shock together. The marketplaces labeled with (1) - (5) represent markets that depend on the delivery through the Rhine directly (Düsseldorf, Frankfurt, and Mannheim) or indirectly through tributaries (Stuttgart and Munich). The markets labeled with (6) - (10) are independent of the Rhine delivery. (Source: Own illustration)

We also present the results for all Rhine-dependent markets separately in Figure 1.16 to showcase the magnitudes of the PPI increase. We can observe that the increase of the PPI component heating oil at all Rhine-dependent market places is on average about 1.2% on impact. At horizon 1, the response indicates an increase of about 1%. The detailed results for the Rhine-independent markets can be found in Appendix 1.A.19.

1.6 Conclusion

We present a new approach to measure regional supply chain disruptions by exploiting waterway transportation restrictions and apply it to the river Rhine. By using administrative trade flow data, we assign weights to the impact of certain passage closures and calculate an overall shock measure that can be used to assess the macroeconomic effects of disruptions in waterway supply chains.

Figure 1.16: IRF of PPI Heating Oil by Marketplace



Note: The figure shows the IRFs of the PPI component heating oil at various Rhine-dependent marketplaces after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred. (Source: Own illustration)

By incorporating this shock series into a local projection framework, we find that transportation restrictions have a temporary contractionary effect on industrial production in bordering federal states and Germany as a whole. This effect is more pronounced and more permanent when considering only energy-intensive sectors, which rely heavily on fossil fuels, and therefore, disproportionately affected by transport restrictions on waterways. Our results also suggest significant differences in the impact of disruptions across states and sectors, which may be due to variations in industry structures or reliance on waterway transportation in the supply chains of different states.

Controlling for the effect of reduced output, we do not find evidence that cargo is redirected through a ship-truck or ship-train substitution pattern. This observation might be due to the specialized nature of the goods being transported, such as fossil fuels that require specific terminals and cannot be easily transferred from ships to trucks.

However, we observe a substitution pattern with respect to the country of origin of imported goods. During supply chain disruptions caused by waterway transportation restrictions, fossil fuel imports from the Netherlands, often shipped through the Rhine, decrease. In contrast, imports from Russia increase, which are primarily transported via pipelines. The observed increase in Russian imports is also quite persistent, for which there are two possible explanations. First, the short-term demand increase of these goods is only

possible with a contractual obligation to increase imports for a prolonged horizon. The second explanation of this finding could be a precautionary motive, such that economic agents insure themselves against possible regional supply chain shocks in the near future. These findings raise the question of the extent to which these regional supply chain shocks will affect the economy going forward, since the possibility to substitute with Russian imports will likely no longer be available.

Finally, we also present evidence on how firms that do not directly receive input goods via the Rhine can be affected by our regional supply chain disruptions. We find that producer energy prices on various Rhine-dependent marketplaces increase in the presence of a shock while the prices at Rhine-independent marketplaces do not react. This pattern highlights the regional response heterogeneity.

Our method to construct a shock series paves the way for more research on the effects of regional supply chain disruptions on waterways other than the Rhine. In particular, research that employs firm-level data to analyze the propagation through firm linkages seems promising.

References

- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2025): “The Macroeconomics of Supply Chain Disruptions,” *The Review of Economic Studies*, 92, 656–695.
- ADEMMER, M., N. JANNSEN, AND S. MEUCHELBOCK (2023): “Extreme Weather Events and Economic Activity: The Case of Low Water Levels on the Rhine River,” *German Economic Review*, 24, 121–144.
- ALFARO-URENA, A., I. MANELICI, AND J. P. VASQUEZ (2022): “The effects of joining multinational supply chains: New evidence from firm-to-firm linkages,” *The Quarterly Journal of Economics*, 137, 1495–1552.
- ANTRÀS, P. AND D. CHOR (2022): “Global value chains,” *Handbook of International Economics*, 5, 297–376.
- BAI, X., J. FERNÁNDEZ-VILLAYERDE, Y. LI, AND F. ZANETTI (2024): “The causal effects of global supply chain disruptions on macroeconomic outcomes: evidence and theory,” *NBER Working Paper 32098*.
- BAQAEE, D. AND E. FARHI (2022): “Supply and demand in disaggregated Keynesian economies with an application to the Covid-19 crisis,” *American Economic Review*, 112, 1397–1436.
- BARROT, J.-N. AND J. SAUVAGNAT (2016): “Input specificity and the propagation of idiosyncratic shocks in production networks,” *The Quarterly Journal of Economics*, 131, 1543–1592.
- BEUTHE, M., B. JOURQUIN, N. URBAIN, I. LINGEMANN, AND B. UBBELS (2014): “Climate change impacts on transport on the Rhine and Danube: A multimodal approach,” *Transportation Research Part D: Transport and Environment*, 27, 6–11.
- BIGIO, S. AND J. LA’O (2020): “Distortions in production networks,” *The Quarterly Journal of Economics*, 135, 2187–2253.
- BONADIO, B., Z. HUO, A. A. LEVCHENKO, AND N. PANDALAI-NAYAR (2021): “Global supply chains in the pandemic,” *Journal of International Economics*, 133, 103534.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2021): “Supply chain disruptions: Evidence from the great east japan earthquake,” *The Quarterly Journal of Economics*, 136, 1255–1321.
- ELLIOTT, M. AND B. GOLUB (2022): “Networks and economic fragility,” *Annual Review of Economics*, 14, 665–696.
- ELLIOTT, M., B. GOLUB, AND M. V. LEDUC (2022): “Supply network formation and fragility,” *American Economic Review*, 112, 2701–2747.
- FEYERER, J. (2021): “Distance, trade, and income—The 1967 to 1975 closing of the Suez Canal as a natural experiment,” *Journal of Development Economics*, 153, 102708.
- FINCK, D. AND P. TILLMANN (2022): “The macroeconomic effects of global supply chain disruptions,” *BOFIT Discussion Paper 14/2022*.

- HERBST, E. P. AND B. K. JOHANNSEN (2024): “Bias in local projections,” *Journal of Econometrics*, 240, 105655.
- JONKEREN, O., P. RIETVELD, AND J. VAN OMMEREN (2007): “Climate Change and Inland Waterway Transport: Welfare Effects of Low Water Levels on the river Rhine,” *Journal of Transport Economics and Policy (JTEP)*, 41, 387–411.
- JORDÀ, (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- LANGSCHIED, K. (1990): “Der Rhein - Vom Urstrom zur Wasserstraße. Wasserbauliche Maßnahmen am Mittelrhein,” *Jahrbücher des Nassauischen Vereins für Naturkunde*, 85–101.
- MEIER, M. AND E. PINTO (2024): “Covid-19 supply chain disruptions,” *European Economic Review*, 162, 104674.
- MEURER, R. (2000): “Der Rheinausbau seit der Jahrhundertwende,” in *Wasserbau und Wasserwirtschaft in Deutschland: Vergangenheit und Gegenwart*, Wiesbaden: Vieweg+Teubner Verlag, 151–154.
- NEWBY, W. K. AND K. D. WEST (1987): “A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.
- OLEA, J. L. M., M. PLAGBORG-MØLLER, E. QIAN, AND C. K. WOLF (2025): “Local projections or VARs? a primer for macroeconomists,” Tech. rep., National Bureau of Economic Research.
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica*, 89, 955–980.
- RAMEY, V. A. (2016): “Macroeconomic Shocks and Their Propagation,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 71–162.
- WEE, W. J., N. B. ZAINI, A. N. AHMED, AND A. EL-SHAFIE (2021): “A review of models for water level forecasting based on machine learning,” *Earth Science Informatics*, 14, 1707–1728.
- YUELING, M., E. MATTA, D. MEISSNER, H. SCHELLENBERG, R. HINKELMANN, ET AL. (2019): “Can machine learning improve the accuracy of water level forecasts for inland navigation? Case study: Rhine River Basin, Germany,” in *38th IAHR World Congress Panama City 2019, Water-Connecting the world*, International Association for Hydro-Environment Engineering and Research, 1979–1989.

Appendix

1.A Robustness Checks

1.A.1 Shock Regression


As outlined in subsection 1.3.2, there are other reasonable ways to formulate the regression from which we derive our shock series. To mitigate concerns that our results hinge on the chosen regression specification of our shock construction, we present robustness specifications and replicate the empirical analysis from the main text.

Precipitation. A major concern is that companies may be able to predict waterway transportation restrictions due to precipitation in the same month. Therefore, we adjust our shock regression as follows:

$$cap_t = \alpha'_0 \mathbf{month}_t + \alpha_1 cap_{t-1} + \alpha_2 prec_t + s_t^{prec},$$


where $prec_t$ is the average level of precipitation in Germany in period t . As before, we use s_t^{prec} as the shock measure. When computing the IRF of waterway cross-border receiving as shock diagnostics, we can again observe a statistically significant contraction on impact of freight ton-miles and transported goods in tons. Both the magnitude and subsequent recovery are similar to the results of our baseline shock specification. In the following we repeat the empirical analysis from the main text and present the outcomes without further comments. The main takeaway from these figures is that our conclusions derived previously do not hinge on the design of our shock regression.

Figure 1.A.1: IRF of Goods Transportation on the Rhine with Adjusted Shock Series (Cross-Border Receiving)




Note: IRFs (dark blue) of freight ton-miles (left panel) and transported goods (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.3: IRF of Industrial Production for States and Sectors with Adjusted Shock Series




Note: The figure shows the IRFs of the industrial production index after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany, NRW: North Rhine-Westphalia, BW: Baden-Württemberg). (Source: Own illustration)

Figure 1.A.2: IRF of German Industrial Production with Adjusted Shock Series




Note: IRFs (dark blue) of industrial production (left panel) and industrial production of energy-intensive sectors only (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.4: IRF of Truck Index with Adjusted Shock Series




Note: The figure shows the IRF of the truck index for Germany after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. (Source: Own illustration)

Figure 1.A.5: IRF of Train Transportation with Adjusted Shock Series (Cross-Border Receiving)




Note: IRFs (dark blue) of cross-border receiving via train transportation for Germany measured in freight ton-miles (left panel) and transported goods in tons (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.6: IRF of Petroleum and Crude Oil Imports with Adjusted Shock Series




Note: IRFs (dark blue) of total German imports of unrefined (left panel) and refined (right panel) petroleum and crude oil after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.7: IRF of Refined Petroleum and Crude Oil Imports by Country with Adjusted Shock Series



Note: IRFs (dark blue) of total German imports of refined petroleum and crude oil from the Netherlands (left panel) and Russia (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.8: IRF of German PPI components with Adjusted Shock Series




Note: The figure shows the IRFs of various PPI components after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany). (Source: Own illustration)

Precipitation x Temperature Interaction. The effect of precipitation could also be mediated through temperature in the respective month. For this reason, we include an interaction term to account for this effect.²⁹

$$cap_t = \alpha'_0 \mathbf{month}_t + \alpha_1 cap_{t-1} + \alpha_2 prec_t + \alpha_3 prec_t \times temp_t + s_t^{inter}$$


Figure 1.A.9: IRF of Goods Transportation on the Rhine with Adjusted Shock Series (Cross-Border Receiving)



Note: IRFs (dark blue) of freight ton-miles (left panel) and transported goods (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)


²⁹Temperature is to a larger degree captured by the monthly dummy variable compared to precipitation. Therefore, we do not include temperature as single variable.

Figure 1.A.10: IRF of German Industrial Production with Adjusted Shock Series




Note: IRFs (dark blue) of industrial production (left panel) and industrial production of energy-intensive sectors only (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.11: IRF of Industrial Production for States and Sectors with Adjusted Shock Series



Note: The figure shows the IRFs of the industrial production index after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany, NRW: North Rhine-Westphalia, BW: Baden-Württemberg). (Source: Own illustration)

Figure 1.A.12: IRF of Truck Index with Adjusted Shock Series



Note: The figure shows the IRF of the truck index for Germany after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. (Source: Own illustration)

Figure 1.A.13: IRF of Train Transportation with Adjusted Shock Series (Cross-Border Receiving)




Note: IRFs (dark blue) of cross-border receiving via train transportation for Germany measured in freight ton-miles (left panel) and transported goods in tons (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.14: IRF of Petroleum and Crude Oil Imports with Adjusted Shock Series




Note: IRFs (dark blue) of total German imports of unrefined (left panel) and refined (right panel) petroleum and crude oil after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.15: IRF of Refined Petroleum and Crude Oil Imports by Country with Adjusted Shock Series



Note: IRFs (dark blue) of total German imports of refined petroleum and crude oil from the Netherlands (left panel) and Russia (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

Figure 1.A.16: IRF of German PPI components with Adjusted Shock Series



Note: The figure shows the IRFs of various PPI components after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred (GER: Germany). (Source: Own illustration)


1.A.2 Baseline Specification

To identify the causal effect of our series of regional supply chain disruptions, we need to ensure that the effect on economic output is actually driven by our identified series of supply chain shocks and not some other confounding variables. A valid concern has been that instead of the waterway transportation restrictions, it is actually the temperatures that lead to a drop, since high temperatures and waterway restrictions might be correlated. To mitigate this concern we enhance our baseline regression as follows:

$$\Delta^{h+1}ip_{t+h} = \alpha^h + \beta_1^h s_t + \gamma^h \Delta ip_{t-1} + \beta_2^h temp_t + v_t, \quad \text{for } h = 0, 1, \dots, 12$$


By accounting for temperature $temp_t$, we ensure that output contractions are not driven by temperature changes. The results are presented in Figure 1.A.17 and Figure 1.A.18 without further comments.

Figure 1.A.17: IRF of German Industrial Production



Note: IRFs (dark blue) of industrial production with the baseline specification (left panel) and industrial production controlling for temperature (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)


Figure 1.A.18: IRF of German Industrial Production (Energy-Intensive Sectors Only)



Note: IRFs (dark blue) of industrial production in energy-intensive sectors only with the baseline specification (left panel) and industrial production controlling for temperature (right panel) after a one standard deviation negative Rhine capacity shock. 90% confidence intervals depicted in light blue. (Source: Own illustration)

1.A.3 Additional Results

Figure 1.A.19: IRF of PPI Heating Oil by Marketplace




Note: The figure shows the IRFs of the PPI component heating oil at various Rhine-independent marketplaces after a one standard deviation negative Rhine capacity shock together with 90% (Newey-West) confidence bands. The left panels show the response on impact; the right panels show the response one month after the shock occurred. (Source: Own illustration)

1.B Data Details

1.B.1 Trade Flow Data

The trade flow data is published by the *Federal Statistical Office Germany* and we retrieved the data from *Destatis*.³⁰ The data set includes yearly trade volume for seven defined Rhine sections: (1) Rheinfelden-Straßburg, (2) Straßburg-Neuburgweiler, (3) Neuburgweiler-Mannheim, (4) Mannheim-Bingen, (5) Bingen-Lülsdorf, (6) Lülsdorf-Orsoy, (7) Orsoy-Dutch Border. The trade data subdivides the trade volume into 20 goods categories following the NST-2007 classification. Further, the data differentiates between hazardous and non-hazardous goods. For the purpose of our project, we work with the total transport volume. The data is available for the years 2011-2021. The absolute volume of the transported goods changes over the years; however, the relative shares remain relatively stable. Figures 1.B.1, 1.B.2, and 1.B.3 summarize the trade flow data.


Figure 1.B.1: Trade Flows by Category and Rhine section in 2021 (1/3)



Note: Each figure contains the trade flows for a specific goods category on every of the seven Rhine sections. The Rhine sections are ordered from South (Swiss border) to North (Dutch border). (Source: Own illustration)


³⁰Table 46321-0013 [[Online Accessible](#) - Last accessed: 28.09.2025]

Figure 1.B.2: Trade Flows by Category and Rhine section in 2021 (2/3)



Note: Each figure contains the trade flows for a specific goods category on every of the seven Rhine sections. The Rhine sections are ordered from South (Swiss border) to North (Dutch border). (Source: Own illustration)

Figure 1.B.3: Trade Flows by Category and Rhine section in 2021 (3/3)



Note: Each figure contains the trade flows for a specific goods category on every of the seven Rhine sections. The Rhine sections are ordered from South (Swiss border) to North (Dutch border). (Source: Own illustration)

1.B.2 Water Level and Related Data

Water Levels. We obtain the raw water level data from the *German Federal Waterways and Shipping Agency*. On request, they provided us with daily observations for administratively relevant measuring stations along the Rhine. The daily water level represents an average of three observations taken each day at every station at the same time. The period of observations varies across the measuring stations, going back as far as to the beginning of World War II.

Reference Level and High Level Threshold. The legal high water level thresholds did not change over time and are available at the *German Federal Waterways and Shipping Agency*.³¹ The reference water levels ('Gleichwertiger Wasserstand'/'GIW' in German) are determined every ten years by the *Central Commission for the Navigation of the Rhine*. However, some of these changes were irregular or delayed. To precisely determine the dates of the changes, we rely on resolutions of the *Central Commission for the Navigation of the Rhine*, which we received upon request.³²

Reference Depth. For the reference depth ('Solltiefe unter GIW' in German) we rely on two types of sources. The current reference depth (minimum guaranteed fairway depth at the reference level) is publicly available at the *German Federal Waterways and Shipping Agency*.³³ For the historic changes in the reference depth, we exploit the information contained in two historical studies by [Langschied \(1990\)](#) and [Meurer \(2000\)](#).

As outlined in the paper, we include only those measuring stations for which the previous variables have been defined. As a result, we obtain water level data for 13 relevant measuring stations along the Rhine (sorted South to North): Maxau, Speyer, Mannheim, Worms, Mainz, Kaub, Koblenz, Andernach, Köln, Düsseldorf, Ruhrort, Wesel, Emmerich.

1.B.3 Outcome Variables

Industrial Production. The industrial production index is published by the *Federal Statistical Office Germany* and by some of the statistical offices of the states. It includes sectors B (Mining and Quarrying), C (Manufacturing), D (Energy Supply) and F (Construction) according to the WZ2008 classification, which is closely related to the international NACE Rev. 2 classification.³⁴ All series used in our analysis are normalized to 100 in 2015.

³¹Overview of stations and high water thresholds are available on ELWIS. [[Online Accessible](#) - Last accessed 28.09.2025]

³²The relevant resolutions are 1932-II-12, 1952-II-18, 1962-IV-49, 1973-I-28, 1976-II-40, 1984-II-40, 1992-I-32, 1996-I-34, 1998-I-27, 2002-II-26, 2012-II-18, 2014-II-17.

³³Overview reference levels are available on ELWIS. [[Online Accessible](#) - Last accessed 28.09.2025]

³⁴More information available in the quality reports of the Federal Statistics Office Germany. [[Online Accessible](#) - Last accessed: 28.09.2025]

Germany. We retrieve the data from *Destatis*.³⁵ The index is available for all four sectors as an aggregate and also for several subsets of sectors, including the ones discussed in section 1.5.1. We use the series that is cleaned for seasonal and calendar effects according to the X-13 JDemetra+ procedure. The series spans a time interval from January 1991 to September 2022.

NRW. We retrieve the data from the database of the *Statistical Office in North Rhine-Westphalia*.³⁶ The index is available for sectors B, C and D individually (or the corresponding sectors in the preceding sector classification schemes) but not as an aggregate index for all four sectors. The raw data contains overlapping subsamples corresponding to different base years. We unify the base year by rescaling the individual series using a simple splicing method to obtain a time series spanning January 1995 - September 2022. The raw data is already cleaned for calendar effects. We use the X-13 procedure in JDemetra+ to clean for seasonal effects using the same specification as the one described in the quality reports of the Federal Statistics Office Germany.

BW. We obtain data upon request from the *Statistical Office in Baden-Württemberg*. The index is only available for sectors B and C jointly, but not as an aggregate index for all four sectors, nor for individual sectors. The raw data contains overlapping subsamples corresponding to different base years. We unify the base year by rescaling the individual series using a simple splicing method to obtain a time series spanning January 2000 - September 2022. We use the X-13 procedure in JDemetra+ to clean for seasonal and calendar effects using the same specification as the one described in the quality reports of the Federal Statistics Office Germany.

Value Added. Figure 1.B.4 shows the relative sizes of the four sectors contained in the index of industrial production. More precisely, the figure shows the value added generated in each of the four sectors relative to the value added generated in the entire market economy.³⁷ The four sectors included in the industrial production index jointly account for roughly 40% of value added in the market economy. The manufacturing sector accounts for the largest share among them (around 30%). The shares are stable over time, except for 2009 where the manufacturing sector was affected more severely by the great recession in relative terms.


Truck Index. We use the monthly truck index (*LKW-Maut-Fahrleistungsindex*) published by the *Federal Statistical Office*. The index uses toll data on kilometers of trucks

³⁵Table 42153-01 [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁶Tables 42153-03i, 42153-06i, 42153-09i, 42153-16i [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁷The data is taken from EUKLEMS, Release 2021, Growth Accounts Basic, Germany. [[Online Accessible](#) - Last accessed 28.09.2025]

Figure 1.B.4: Sector Shares of Value Added



Note: Each bar represents the relative size of the value added of the four industrial production sectors (B, C, D, and F) from 2005 to 2019. (Source: LUISS Lab of European Economies)

driven on German highways. We retrieve the data from *Destatis*.³⁸ We use the series that is cleaned for seasonal and calendar effects according to the X-13 JDemetra+ procedure. The series spans a time interval from January 2005 to October 2022.

Train Transportation. We obtain information on train transportation in Germany by the *Federal Statistical Office*. The series spans a time interval from January 2008 to December 2022. We retrieve data from *Destatis*.³⁹

Imported Goods. To illustrate the effects of supply chain disruptions on the Rhine on imported goods, we use data of refined and unrefined crude oil and petroleum imports published by the *Federal Statistical Office*. We retrieve data from *Destatis*.⁴⁰ The series spans a time interval from January 2008 to December 2022. The data provides information of the total amount of imports together with a partitioning of imports by countries of origin.

PPI. We use the monthly PPI publications from the *Federal Statistical Office Germany*. The data provides information for each sub-category of the PPI basket. Within our analysis, we consider the components *Brown Coal*, *Gasoil*, and *Heating Oil*. The data is retrieved from *Destatis*.⁴¹

³⁸Table 42191-01 [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁹Table 46131-0004 [[Online Accessible](#) - Last accessed: 28.09.2025]


⁴⁰Table 51000-0011 [[Online Accessible](#) - Last accessed: 28.09.2025]

⁴¹Table 61241-0004 [[Online Accessible](#) - Last accessed: 28.09.2025]

1.B.4 Capacity and Shock Series

Figure 1.B.5 presents an extension of Figure 1.5, which distinguishes between restrictions due to high vs. low water. The figure shows the total number of days at which a restriction (due to low water in the upper panel, due to high water in the lower panel) was in place within the respective year at three selected stations. Two key observations stand out. First, the regional heterogeneity is very pronounced for both types of events. Second, low water events occur way more frequently than high water events. In addition, high water restrictions typically only last for few days, while low water restrictions can last for several months. Even though high water restrictions lead to more severe capacity restrictions (complete closure of the passage as opposed to low water events that usually just limit the loading capacity), this implies our results are mostly driven by low water events.

Figure 1.B.5: Number of days with restrictions at selected stations




Note: Number of days with waterway transportation restrictions caused by low level water events (top panel) and high water level events (bottom panel) at the measuring stations Maxau, Kaub, and Duisburg-Ruhrort from 1990 until 2022. (Source: Own illustration)

Figures 1.B.6 and 1.B.7 replicate figures 1.5 and 1.B.5 but average over locations using the weighting scheme presented in section 3. The conclusions drawn above and in section 3 carry over to these figures.


Figure 1.B.8 shows the final shock series used for the impulse response analysis. The series is obtained by regressing the aggregate capacity series (cf. figure 1.6) on a full set of monthly dummies and on its first lag, as discussed in section 1.3.2. The qualitative behavior is similar to that of the aggregate capacity series, with the notable difference that the final shock series also contains positive shocks (mostly occurring at the end of low water episodes).

Figure 1.B.6: Average Number of Days with Restrictions




Note: Average number of days with restrictions at all measuring stations for every year between 1990 and 2022. (Source: Own illustration)

Figure 1.B.7: Average Number of Days with Restrictions (High vs. Low Water Levels)



Note: Average number of days with restrictions caused by low water level events (top panel) and high water level events (bottom panel) at all measuring stations for every year between 1990 and 2022. (Source: Own illustration)

Figure 1.B.8: Final Rhine Capacity Shock Series



Note: Resulting shock series from our constructed series. The red line indicates the monthly shock realizations between 1990 and 2022 and the black line indicates the mean over the entire observation period. (Source: Own illustration)

Chapter II



The Amateur's Curse of Positive Income Shocks*

Yann Müller

October 2025

Abstract

How do soccer clubs react to positive income shocks? Exploiting quasi-random variation induced by penalty shootouts in *Verbandspokal* finals, a German amateur tournament, we find that clubs receiving a positive income shock are unable to improve their final league position in the following years. Even more surprising, these teams start to significantly underperform three years and onwards after the income shock. We observe an associated increase in insolvency proceedings, indicating financial struggles after the *Verbandspokal* final. We identify a causal link to longer managerial tenure, which suggests the emergence of a bias in the decision-making process regarding whether or not to appoint a new team manager. The long-run decline in club performance also carries political externalities, as the majority party in the municipal council of the club's municipality significantly underperforms in subsequent elections.

Keywords: Soccer, Amateur Clubs, Income Shocks.

JEL Codes: E21, Z2.

*Müller: Department of Economics, University of Mannheim, yann.mueller@uni-mannheim.de. Special thanks for guidance, support, and useful advice to Antonio Ciccone, Felix Köhler, Maximilian Mähr, Max Riegel, and participants of the Mannheim Doctoral Colloquium.

2.1 Introduction

The reaction of economic agents to income shocks is a well-studied subject in the Economics literature. For individuals (e.g., [Lindqvist et al., 2020](#)), households (e.g., [Hall, 1978](#); [Cesarini et al., 2017](#)), families ([Cesarini et al., 2023](#)), firms (e.g., [Davis et al., 2025](#); [Lee et al., 2011](#)), and governments (e.g., [Brollo et al., 2013](#); [Litschig and Morrison, 2013](#)), a wide range of studies examine responses to income shocks and investigate potential deviations from theoretical predictions. However, there is limited research on how sport clubs, particularly soccer clubs, respond to such shocks as economic agents.

We address this gap by asking a straightforward question: How do soccer clubs perform after a positive income shock? We exploit the institutional setting of German amateur soccer to investigate this question.¹ To begin, it is necessary to establish how soccer clubs can be conceptualized as economic agents. [Szymanski \(2003\)](#) discusses the premise whether club owners act as profit or win maximizers. While this is quite debatable in an international context (e.g., [Dabscheck, 1975](#)), we assume that German amateur soccer clubs behave as constrained win maximizers.² Standard theoretical models would predict that relaxing the budget constraint should, on average, improve performance (compare [Fort, 2019](#)).

However, exploiting a quasi-natural experiment in the context of *Verbandspokal* finals, yearly regional tournaments in Germany, we find the opposite. Clubs that receive a positive income shock do not initially improve their league rankings. Even more surprising, three years after the shock, these clubs significantly underperform compared to control clubs that did not receive such a shock. Also, the treated clubs exhibit a higher likelihood of insolvency proceedings following the income shock. Exploring underlying mechanisms, we identify a link to longer managerial tenure. For example, two years after the *Verbandspokal* final, the average tenure duration of team managers from treated clubs increases by almost one year compared to control clubs. This finding indicates a systematic bias in the process by which clubs assess the need to appoint a new team manager. The income shocks also generate political externalities. While local politicians cannot benefit electorally from their club's success in the *Verbandspokal* final, they appear to be punished when club performance subsequently declines. Specifically, the majority party in the municipal council at the time of the *Verbandspokal* final underperforms by four percentage points in the long run, relative to unaffected municipalities.

¹Soccer is the most popular sport in Europe, and in Germany alone, the national soccer association, the *Deutscher Fußball-Bund* (DFB), counts approximately 8 million registered members in 2025. In context, this corresponds to about 10% of the German population, making the DFB the largest single-sport federation in the country.

²German amateur soccer clubs are usually constituted as *e.V.* (eingetragener Verein - registered association), a legal form that requires the pursuit of non-profit objectives and prohibits profit distribution.

Related Literature. Our paper contributes to the broader literature on how economic agents respond to income shocks. In macroeconomics, this literature is rooted in the Permanent Income Hypothesis (PIH), which states that individuals smooth consumption based on permanent income expectations, responding less to transitory shocks. Seminal contributions by [Friedman \(1958\)](#) and later empirical investigations (e.g., [Hall, 1978](#); [Campbell and Deaton, 1989](#); [Carroll, 1997](#)) have analyzed the extent to which households behave in accordance with the PIH. More recent work has shown that liquidity constraints ([Deaton, 1991](#)), myopia ([Laibson, 1997](#)), or misperception of income persistence ([Malmendier and Nagel, 2011](#)) can cause agents to overreact to temporary shocks, leading to suboptimal consumption or investment decisions.

Our setting provides a novel example of such behavior in the context of amateur soccer clubs, where the positive income shock appears to affect managerial decisions and ultimately increases the likelihood of insolvency proceedings. This complements a growing literature on unintended consequences in organizational settings, such as a decline in managerial quality, following a positive income shock.³ For example, [Brollo et al. \(2013\)](#) show how income shocks increase corruption and decrease average education of mayors in Brazil.⁴ In a similar vein, [Blanchard et al. \(1994\)](#) analyze how corporate managers react to such income shocks. A possible mechanism in our situation is the existence of a current performance bias, where management decisions are based on recent accolades instead of considering other observable characteristics, and is also known as the Peter principle.⁵ Such a bias is documented in the literature (e.g., [Benson et al., 2019](#)) and could explain the increase in team manager tenure duration.

Our paper also relates to the literature on sports economics. In this context, the theoretical foundations of soccer clubs are particularly relevant to our research question.⁶ The properties of soccer clubs as economic agents are quite different from those of a standard firm. [Sloane \(1971\)](#) has been the first to make the argument about soccer clubs being win rather than profit maximizers. Further, these clubs have a shared interest in preserving some degree of competition in their league to maximize viewership ([Sloane, 1976](#)) and to maintain access to a sufficient pool of talent ([Hoehn and Szymanski, 1999](#)).⁷ [Fort \(2019\)](#) provides an overview of the commonly used objective functions for soccer clubs. Given

³Although the terminology differs slightly from the PIH literature (e.g., *cash windfalls*), it refers to the same concept of income shocks.

⁴A more aggregated consideration is presented by [Caselli and Tesei \(2016\)](#), who show that the political system itself is changing after income shocks.

⁵The Peter principle is a management concept introduced by [Peter and Hull \(1969\)](#) stating that there is a tendency to incompetence at management roles since promotions are primarily based on current job performance instead of observable characteristics that could indicate suitability for the future job.

⁶For a broader perspective we refer to [Dobson et al. \(2001\)](#) and [Szymanski \(2003\)](#).

⁷The second argument applies specifically to clubs that also participate in tournaments at a European level and require a sufficient pool of potential players to be competitive. However, [Hoehn and Szymanski \(1999\)](#) also argue that this is responsible for a decline in the national league's competitiveness. With the internationalization of professional soccer, one could also make the argument that the relevance of a national talent pool becomes less relevant.

these properties, our findings of declining performance after a positive income shock reveal a puzzle that cannot be explained by standard models of soccer clubs.

A possible rationalization of our findings can be drawn from the literature on positive assortative matching, particularly as formulated in [Becker \(1973\)](#) in the context of the marriage market. According to this framework, optimal matches occur when individuals with complementary or similar traits pair together.⁸ In our context, the equivalent would be a productive alignment of team manager and player squad with matching skill sets in every season.⁹ Such an optimal match would maximize the subsequent wins given the financial constraints. Our finding of extended team manager tenure after a random event could indicate a breakdown in optimal matching, possibly due to the above-mentioned performance bias. The concept of positive assortative matching has also been documented in other contexts. For example, [Mendes et al. \(2010\)](#) find that long-lived firms exhibit strong sorting between worker skills and firm productivity, consistent with the logic of optimal matches. We do not attempt to provide a definitive test of this mechanism in our setting but leave this avenue for future research. Instead, we focus on empirically documenting the causal effect of income shocks on amateur soccer clubs.

Finally, our paper connects to the literature on the externalities of soccer. For example, [Depetris-Chauvin et al. \(2020\)](#) show how collective experiences induced by soccer tournaments can help in building a national identity in sub-Saharan African countries.¹⁰ Other research focuses more on the regional impact. [Köhler \(2025\)](#) analyzes the effect on public finances of municipalities, showing how the relegation of soccer clubs affects budgets and council election outcomes. Additional evidence suggests that the characteristics of cities themselves affect the soccer industry. For example, [Kleven et al. \(2013\)](#) explore how local tax laws influence player migration, showing that taxation affects team composition by altering player inflows and outflows.

The remainder of the paper is organized as follows. First, we introduce the institutional background of the *Verbandspokal* tournament. Next, we outline our data sources. Then, we introduce our empirical framework and discuss the results. The final section concludes.

⁸[Becker \(1973\)](#) states IQ, education, height, attractiveness, skin color, and ethnic origin as such traits.

⁹Examples of such skills could be the understanding of tactics or the ability to read a match situation.


¹⁰In the same vein, [Allison \(2000\)](#) analyzes the relationship between professional sports and nationalism.

2.2 Background: Verbandspokal

German soccer is organized under the national association *Deutscher Fußball-Bund* (DFB) and 21 regional associations. Every year, the DFB hosts the German Cup for the professional clubs in the first three tiers of the league system, which is associated with sizeable revenue streams. Furthermore, the regional associations organize the *Verbandspokal*, regional tournaments that allow lower-tier clubs to qualify for the following year's German Cup. Generally, winning the *Verbandspokal* and subsequently participating in the German Cup is perceived as a great opportunity for amateur clubs to build a stronger financial foundation and transition into higher leagues in the future.

Clubs that qualify through the *Verbandspokal* are usually knocked out in the first stage of the German Cup. Nevertheless, the first-round participation alone guarantees more than 200.000€ in revenues in 2025.¹¹ In addition, the participating clubs receive income from ticketing, merchandise, and potential new sponsors. For many amateur clubs, the participation revenues in the German Cup can amount to a sizeable fraction of the annual budget.¹² Given this income stream, it is natural to ask how the performance of clubs evolve after winning the *Verbandspokal* and subsequently participating in the German Cup.

Figure 2.1: League Ranking Before/After Verbandspokal Win



Note: This figure shows average improvements in league rankings before and after a team wins the *Verbandspokal* final. Positive values indicate improvements in final league ranking (i.e., lower numerical ranks). Detailed information on tracking league rankings across leagues and seasons is provided in section 2.3.1. The event year (0) marks the season of the *Verbandspokal* final, with earlier and later seasons shown relative to this point. Dots represent average changes in rankings, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

¹¹Detailed information regarding the German Cup tournament is available at the website of the DFB. [[Online Accessible](#) - Last accessed: 27.09.2025]

¹²As mentioned before, there is no disclosure requirement for financial statements of German amateur soccer clubs. Therefore, we cannot exactly define the magnitude of the positive income shock. However, we argue that the lower bound additional income of 200.000€ can already amount to a sizeable fraction of the budget. According to Sportschau reports, the annual budget of amateur clubs in the 4th tier ranges from 300.000€ to 2.400.000€. [[Online Accessible](#) - Last accessed: 27.09.2025]

Figure 2.1 illustrates the average improvement in league rankings before and after winning the *Verbandspokal*. Two observations stand out. First, following the victory, teams improve their final league ranking by nearly two places on average in the subsequent seasons, with the improvement persisting in the following years. This upward trend may reflect increased motivation, enhanced financial support, or better player recruitment following the participation in the German Cup. While the estimates do not establish a causal link, they suggest that winning the *Verbandspokal* can coincide with a sustained positive momentum in league performance for lower-tier clubs.

The second observation, the visible pre-trend, raises questions regarding the direction of causation. Teams that win the *Verbandspokal* appear to be on an upward trajectory already, and therefore, we do not know whether the subsequent improvement in league rankings is caused by or irrespective of the additional income following the *Verbandspokal* win. It might be the case that the observed pattern is driven by clubs that have already undergone a structural transformation and would have improved their league rankings in subsequent seasons, likely even without the additional income stream. An example of such a case is TSG 1899 Hoffenheim, a club that was acquired by a German billionaire and transitioned from a low-level amateur league into the *Bundesliga*, the highest German soccer league, meanwhile winning the *Verbandspokal* four times in a row.¹³ Few would argue that the club's transition into the *Bundesliga* was a direct consequence of the income associated with the *Verbandspokal* victories rather than of strong financial backing and long-term strategic investments. Our research design attempts to disentangle these channels and to answer the question of how a positive income shock associated with winning the *Verbandspokal* affects future league performance of these amateur clubs.

2.3 Data

The following section outlines how we construct a consecutive ranking that allows us to track and compare the performance of teams across leagues and seasons. Also, we describe the data sources that we use in our empirical framework.

2.3.1 Consecutive Ranking

The main objective of the paper is to track the performance of teams after participating in a regional *Verbandspokal* final. To this end, we construct a consecutive ranking that determines a team's position within the entire German soccer league system.¹⁴ Figure


¹³TSG 1899 Hoffenheim started to receive financial support from Dietmar Hopp, the founder of SAP, in 1990. At this time, the club was playing in the 7th-tier *Kreisliga A*. Within ten years, the club experienced four promotions and won the *Verbandspokal* in 2002, 2003, 2004, and 2005. In the year 2008, the club was promoted to the *Bundesliga*.

¹⁴The data for rankings and league structures are retrieved from www.fussball.de and www.f-archiv.de.

2.2 illustrates the approach for the year 2000. A team's position therefore reflects the number of teams that are ahead in the team's direct pathway up to the first place in the *Bundesliga*. Since the German football league system follows a hierarchical tree structure, multiple teams can occupy the same nominal league position if they compete in parallel divisions at the same tier. In our illustration, this becomes apparent starting with the *Regionalliga Nord* and *Regionalliga Süd*, which functioned as third-tier leagues in the early 2000s. Our analysis covers the period from the 1990/1991 to the 2021/2022 season, during which two major structural reforms to the league system occurred.

First, in 1994, the *Regionalliga* was introduced as the new third tier. Later, in 2008, the *3. Liga* was established, which subsequently replaced the *Regionalliga* as the third tier, relegating the *Regionalliga* to the fourth tier. These changes mechanically decrease the league rankings of lower-tier clubs. We discuss this structural discontinuity later in our empirical design.

Figure 2.2: Illustration Consecutive Ranking in the Year 2000



Note: This figure illustrates the hierarchical structure of the German football league system in the year 2000 and the construction of consecutive league rankings across tiers. The figure shows that multiple leagues exist at the same tier (e.g., *Regionalliga Nord* and *Regionalliga Süd* at the third tier, or multiple *Oberligas* at the fourth tier), meaning that teams playing in different leagues can share the same ordinal league position. The consecutive ranking is constructed by ordering all clubs within their respective tiers, while accounting for the number of clubs above them in higher tiers. This ensures comparability across divisions and enables the use of a continuous league position measure in the empirical analysis. (Source: Own illustration)

2.3.2 Board of Directors

As remarked previously, most German amateur soccer clubs are constituted as registered associations that are not required to release financial statements for the public. However, they have to release information about organizational changes with respect to the board of directors and ongoing insolvency proceedings. We outline the detailed process to retrieve this information in Appendix 2.A.2.

2.3.3 Newspaper Coverage

In our search for newspaper articles related to amateur clubs, we use www.wiso-net.de, which is a comprehensive online research platform providing access to old regional newspaper articles. Specifically, we search for articles related to financial problems of the soccer clubs, and therefore, use the team name together with the word *Schulden*, which is the German translation of debt, in our search query. Since the number of articles varies noticeably with the size of the club, we use the in-built filter function for football articles once our search yields more than 500 articles to avoid unrelated press coverage.¹⁵

2.3.4 Team Managers and Players

We collect information about the team managers and players for each season in our observation period from www.fupa.net and www.transfermarkt.de, two websites containing information on German amateur and semi-professional soccer clubs. A shortcoming of amateur clubs is that the available data is often incomplete, especially for earlier seasons. We append some of the missing information with hand-collected data, however, results based on this data, presented later in section 2.5.3, should be interpreted with caution.

2.3.5 Elections

We also collect election results for the municipal council in the local elections from the statistical offices in the federal states. Missing information from older elections is retrieved from election reports documented on www.statistischebibliothek.de or hand-collected from the websites of the municipalities.

2.4 Empirical Strategy

Our empirical approach exploits variation induced by penalty shootouts in *Verbandspokal* finals. Specifically, we restrict the comparison to those teams that reached the final stage

¹⁵For example, search queries for the amateur team of Werder Bremen result in a number of boulevard articles about the former player Ailton and the financial struggles he has encountered in his private life after his professional career. These kinds of results are excluded with the in-built football filter.

of the competition and in which the outcome was determined by a penalty shootout. This design allows us to isolate plausibly exogenous variation in the treatment assignment, i.e., winning the final and qualifying for the German Cup.

This strategy relies on three core identifying assumptions. First, we assume that, on average, there are no structural differences between teams that reach the shootout stage of a *Verbandspokal* final. Second, the outcome of a penalty shootout can reasonably be treated as a quasi-random event. While not purely random, shootouts introduce an element of chance that is unrelated to underlying team strength or characteristics.¹⁶ Third, in the absence of the German Cup income shock associated with winning the final, we assume that the subsequent trajectories of winners and losers would have followed parallel trends. This identification approach is conceptually similar to strategies used in the quasi-experimental literature, such as close election research designs (e.g., Dell, 2015; Marshall, 2024; Akhtari et al., 2022) or strategies that compare winners to runner-ups in investment bidding processes (e.g., Greenstone et al., 2010).¹⁷

Using our empirical strategy, we identify 64 eligible *Verbandspokal* finals, which are listed in Appendix 2.A.1. Team characteristics for treated and control clubs during the season of the *Verbandspokal* final are presented in Table 2.1. We observe that the characteristics of treated and control teams are reasonably similar across all reported dimensions, supporting our identification strategy and suggesting that we are not systematically comparing structurally different clubs. We also present a comparison of the characteristics of the corresponding board of directors in Appendix 2.A.2.

In our baseline specification, we compare never-treated clubs with first-time-treated clubs.¹⁸ To estimate the dynamic effects of winning a penalty shootout in the *Verbandspokal* final, we employ an event study framework. This allows us to observe the evolution of club outcomes in the years before and after the event. Our baseline empirical specification is given by:

$$y_{i,t} = \mu_i + \theta_t + \sum_{j=-J}^{\bar{J}} \beta_j e_{i,t}^j + \delta \mathbf{X}_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where $y_{i,t}$ denotes the outcome variable, which is the final league ranking of club i in

¹⁶We argue that this assumption is particularly plausible in the context of amateur soccer. If any meaningful advantage has existed, structurally superior teams are unlikely to remain eligible for the *Verbandspokal*, as they typically would have advanced to higher leagues that exclude them from participation. As a result, the outcome of a penalty shootout among amateur teams can reasonably be considered as a quasi-random event.

¹⁷Similar to these designs, there is a potential risk of spillovers in our empirical strategy. Specifically, the positive income shock may allow treated clubs to buy players from control clubs. In that case, the subsequent performance of control clubs would no longer reflect a trajectory independent of the shock. However, since this would reduce the performance of control clubs, such an externality, if present, would align with our finding of declining performance among treated clubs. Therefore, our main results should qualitatively remain unaffected.

¹⁸We have five clubs that won the *Verbandspokal* twice via penalty shootout and also five clubs that lost a final via penalty shootout before winning one.

Table 2.1: Team Characteristics of Treatment and Control Clubs

Variable	Treatment Clubs	Control Clubs
Number of Players	32.1 (5.55)	30.8 (4.15)
New players (%)	44.67 (14.13)	49.87 (18.48)
Player tenure (years)	1.65 (0.577)	1.36 (0.566)
Team manager tenure (years)	2.40 (2.23)	2.62 (1.81)
Player age (years)	25.8 (1.35)	25.2 (1.70)
German players (%)	79.92 (10.98)	83.29 (7.76)

Note: This table reports summary statistics describing the composition of treatment and control teams in the season of the *Verbandspokal* final. Treatment teams are those that won the *Verbandspokal* final in a penalty shootout and qualified for the German Cup. Reported values are means, with standard deviations in parentheses. Information is retrieved following the procedures presented in section 2.3.4.

season t in our main specification. The term μ_i captures club-level fixed effects, accounting for time-invariant differences between clubs. The term θ_t represents time fixed effects, absorbing shocks that affect all clubs in a given season, such as changes in the league structure.¹⁹

The key variables of interest are the event-time indicators $e_{i,t}^j$, which take the value of 1 if the observation corresponds to j years before or after the event, defined as winning a *Verbandspokal* final via penalty shootout, which results in participation in the German Cup, and 0 otherwise. The associated coefficients β_j capture the dynamic treatment effects in the event window. We include $\underline{J} = 5$ pre-treatment periods and $\overline{J} = 9$ post-treatment periods.²⁰

We normalize the coefficient on the first lead β_{-1} to zero, and therefore, the resulting

¹⁹While this statement is true, it only partially addresses our concern outlined in section 2.3.1 with respect to the introduction of new league tiers. For example, let's compare the two possible scenarios of a treated and control team after the introduction of the *3. Liga* in the year 2008 which relegated the *Regionalliga* to the 4th tier: 1) Both clubs experience a mechanical decrease in rankings of the same magnitude. This effect would be captured by time fixed effects. 2) Both clubs are originally in the *Regionalliga*, for example with the final league rankings 9 and 10. Hypothetically, if only the first nine teams were to continue in the *3. Liga* and the others are placed in the new 4th-tier *Regionalliga*, this introduces a mechanical ranking difference stemming from factors other than actual league performance and would not be captured by time fixed effects. This would introduce a downward bias in our estimates. However, to the best of our knowledge, there are only two *Verbandspokal* finals for which a change in league structure affected the teams differently. Repeating our analysis excluding these finals does not qualitatively change our results.

²⁰With a slight abuse of notation: to account for periods beyond the end of the event window, we apply a binning procedure by combining all subsequent periods into a single bin, which captures the long-run effect of the treatment, following a standard practice in the event study literature (compare [Borusyak et al., 2024](#)).

coefficients can be interpreted as differences between treatment and control clubs relative to the year in which the clubs played in the *Verbandspokal* final. We further include a vector of control variables $\mathbf{X}_{i,t}$ to account for observable differences across clubs. This vector comprises indicators for prior *Verbandspokal* victories as well as variables capturing our collected team characteristics, such as average squad age and player turnover.

This event study approach enables us to compute both pre-treatment trends and post-treatment effects, ensuring that our identifying assumptions hold and that the estimated treatment effects are not driven by differential pre-trends. Following [Borusyak et al. \(2024\)](#), we discuss the validity of our event study design and present a series of robustness specifications in Appendix 2.B.2.

2.5 Empirical Results

In the following section, we present the results of our event study approach and discuss the effect on league rankings after the income shock. Also, we present evidence for potential mechanisms that explain the observed results.

2.5.1 Effect on League Ranking


Figure 2.3 illustrates the core result of our paper, showing the average change in final league rankings for clubs that won the *Verbandspokal*, relative to control clubs, in the years before and after participating in the German Cup and receiving the temporary income shock. The outcome variable is defined such that positive values represent improvements in league rankings, while negative values indicate a decline in performance. Error bands indicate a 95% confidence interval.

In the first year after the *Verbandspokal*, treated and control clubs perform almost identically. Afterwards, the treated clubs perform progressively worse in the league relative to the control group. The coefficients become increasingly negative and statistically significant over time, suggesting a persistent decline in league performance. For example, six or more years after the German cup participation, treated clubs persistently underperform control clubs by more than ten places in the consecutive ranking, relative to their initial difference at the time of the *Verbandspokal*.

This pattern indicates that, rather than improving club outcomes, the temporary income shock associated with winning the *Verbandspokal* and the subsequent German Cup qualification may have unintended negative consequences for future performance. The absence of significant pre-trends supports the validity of the parallel trends assumption and strengthens the interpretation of these results as a causal effect of the income shock.²¹

²¹For transparency, we want to note the following: There is a possible confounding exposure effect that could arise in addition to the income shock that results from participation in the German Cup.

Figure 2.3: League Ranking (Event Study)



Note: This figure shows average improvements in league rankings before and after a team wins the *Verbandspokal* final. Positive values indicate improvements in final league positions (i.e., lower numerical ranks). The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in rankings, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

Given the sizeable point estimates and the average number of teams in a league, it is conclusive to ask whether the treated clubs actually get relegated more often.

Figure 2.4 presents the results from an event study where the outcome variable is the league tier of the clubs. The graph shows that treated clubs experience a relative decline in their league tier in the seasons following their *Verbandspokal* win. Taken together, these results imply that the *Verbandspokal*-related income shock does not shield clubs from sporting decline and may, in fact, precede a relegation in the league system.


For completeness, we also estimate corresponding event studies using relegation and promotion as outcome variables. These results, shown in Figure 2.B.6 and 2.B.7 of Appendix 2.B, reveal an increase in the likelihood of relegation. To further contextualize these findings, we explore changes in the full distribution of league rankings over time. Figures 2.B.1 and 2.B.2 in Appendix 2.B present kernel density plots of league rankings for treated and control clubs in pre- and post-treatment periods. These plots reveal that treated clubs not only decline relative to the control group but also exhibit an absolute drop in performance as their average league ranking deteriorates over time, whereas the average ranking of control clubs improves.

2.5.2 Effect on Financial Stability

Given that winning the *Verbandspokal* constitutes a positive income shock, we analyze how this affects the financial stability of clubs in subsequent years. Due to the disclosure

Specifically, participating in the German Cup exposes amateur players to professional clubs. It could be the case that this results in the best players leaving their amateur clubs due to offers from higher-league clubs. While we cannot rule this out, there is a solid argument against the actual impact of this exposure effect. If there is such a player buyout, it would further increase the financial capabilities of treated clubs. Therefore, it cannot explain entirely why these clubs subsequently underperform.

Figure 2.4: League Tier (Event Study)



Note: This figure shows average improvements in league tiers before and after a team wins the *Verbandspokal* final. Positive values indicate improvements, meaning movement to higher tiers (i.e., numerically lower tier values). The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in league tiers, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

requirements for soccer clubs, we cannot access financial statements directly.²² Instead, we investigate whether the income shock influences the likelihood of clubs entering ongoing insolvency proceedings.²³ Figure 2.5 presents the results of the corresponding event study. In the first one to three years following German Cup participation, the estimated coefficients suggest an increase in the likelihood of insolvency proceedings among treated clubs by approximately 3 percentage points, although these estimates are not statistically significant.²⁴ However, between four and seven years after participation, the estimated effects rise to a peak of 8.3 percentage points and become statistically significant at the 95% confidence level.

This increase in financial distress coincides with the observed decline in league performance. However, we cannot identify whether the deteriorating financial situation is a cause or a consequence of that sporting decline. Nevertheless, it is noteworthy, and somewhat unexpected, that clubs begin to face significant financial difficulties several years after receiving a positive income shock.

It is important to note that insolvency proceedings represent a severe regulatory intervention, and clubs may already experience financial distress prior to such formal action. As a result, our event study based on insolvency filings may capture only part of the financial reality faced by clubs.


To complement this analysis, we examine debt-related coverage in local and regional

²²As mentioned before, these registered associations are not required to release financial statements for the public.

²³The outcome variable is now a binary indicator equal to 1 if a club has ongoing insolvency proceedings in a given season, and 0 otherwise. This design allows us to interpret our regression estimates as the change in likelihood of an ongoing insolvency proceeding.

²⁴During the season of the *Verbandspokal* final, 2.5% and 5.5% of treated and control clubs, respectively, experienced an ongoing insolvency proceeding.

Figure 2.5: Insolvency Proceedings (Event Study)



Note: This figure shows average changes in the likelihood of insolvency proceedings before and after a team wins the *Verbandspokal* final. The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in percentage points, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

newspapers. Specifically, we identify articles that mention both the club name and the keyword *debt*. These articles often appear prior to the initiation of formal insolvency proceedings and can serve as an early indicator of financial difficulties. We then use the log-transformed count of debt-related newspaper articles as the outcome variable in an event study framework, analogous to the previous analysis.²⁵


Figure 2.6 illustrates the results. In the years following the *Verbandspokal* final, treated clubs experience a sustained increase in debt-related newspaper coverage relative to control clubs. Although none of the post-treatment coefficients are statistically significant, the estimates remain positive across most years, suggesting a pattern of elevated financial visibility following the income shock.

One possible explanation for the limited statistical significance across the post-treatment years is that financial difficulties and their visibility in newspaper coverage tend to be temporally concentrated. In practice, once financial problems surface, e.g., during licensing processes, they are likely to trigger a brief window of elevated media scrutiny, typically lasting 1–3 years. After this period, coverage may decline again as the club either recovers or media attention shifts elsewhere.

This pattern may suggest that the income shock from winning the *Verbandspokal* is followed by increased financial strain, which results in higher media attention related to debt. Therefore, even though no post-treatment coefficient is statistically significant, the overall pattern of positive coefficients following the *Verbandspokal* victory supports our previous results that financial challenges may follow after the temporary income shock.

²⁵Since the outcome variable is available at daily frequency, we define the year prior to the *Verbandspokal* final as the reference period. Each period spans from April to March of the following year, aligning with the typical timing of the *Verbandspokal* finals, which usually occur in April or May. To account for variation in baseline coverage across clubs, we normalize the variable by computing a log-transformation for each observation period. The results are qualitatively similar when computing z-scores instead.

Figure 2.6: Debt-Related Newspaper Coverage (Event Study)




Note: This figure shows the average difference in debt-related newspaper coverage between teams that won the *Verbandspokal* final (event year 0) and control teams, across seasons before and after the final. The outcome variable is the log-transformed number of relevant newspaper articles. Dots represent average effects, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

This supports the notion that the observed underperformance of treated clubs in later years may be partly driven by financial mismanagement or overspending following the temporary income shock.

2.5.3 Effect on Team Manager Tenure and Team Composition

Mechanisms that could explain the decline in performance include changes in team composition and team manager tenure. Figure 2.7 shows the average change in team manager tenure duration before and after a team wins the *Verbandspokal* final.²⁶

Figure 2.7: Team Manager Tenure Duration (Event Study)



Note: This figure shows the average difference in manager tenure between teams that participated in the German Cup (event year 0) after winning the *Verbandspokal* final and control teams. The outcome variable is measured in years and captures how long the manager has already been at the club in a given season. Coefficient values therefore reflect differences in tenure duration in years. For example, a coefficient of 0.5 indicates that managers at treated teams had, on average, half a year longer tenure than those at control teams in that season. Dots represent average effects, and vertical bars show 95% confidence intervals. (Source: Own illustration)

²⁶We define tenure duration as the number of seasons a player or team manager has already been with the club at a given point in time.

We observe a clear increase in team manager tenure duration in the seasons immediately following the income shock, with managers staying significantly longer at treated clubs compared to control clubs.²⁷ This increase is particularly concentrated in the first years after the final. Our estimates can be the result of two different developments. First, it could be the case that control clubs increase their likelihood of hiring a new coach. Second, treated clubs could decrease their likelihood of hiring a new coach in the seasons after the *Verbandspokal* final. Both scenarios would lead to the presented results. Computing the unconditional team manager tenure duration of treated and control clubs, we find that it is actually the latter, i.e., treated clubs become less likely to hire a new team manager.²⁸ A possible explanation for this pattern is that the team manager who wins the *Verbandspokal* is more likely to be retained by the club's management. This is a somewhat surprising result, considering that the decisive moment in our setup, the penalty shootout, is plausibly random and not necessarily indicative of superior managerial quality. Nevertheless, it appears to anchor the club's perception of the team manager's performance and suitability for future seasons. Such a behavior is well documented in the literature (e.g., [Benson et al., 2019](#)), and is known as the Peter principle, i.e., promotions, or continued appointments, are based on current performance rather than on observable characteristics indicating suitability for future roles. However, besides our observations hinting at the Peter principle, they cannot entirely be explained by it, because our estimates remain elevated, even though not statistically significant anymore, beyond the term of most team managers at the time of the *Verbandspokal* final. There are several possible reasons one could argue for why this pattern of longer tenure duration extends to team managers that start in their role only after the *Verbandspokal* final. However, due to missing data for a more in-depth analysis, we remain agnostic about it and just document the pattern at this point.

Yet, given that the club management must reevaluate the optimal match between team manager and squad each season, this tendency to stick with the existing team manager may hinder necessary adjustments. Over time, this could contribute to a non-optimal composition of team and team manager, which might help explain the underperformance observed in league outcomes.

Another channel could be a change, in various possible forms, to the player structure of the treated clubs. As outlined in section 2.3.4, the data for amateur clubs is far from complete for every season and the results must be viewed with caution. Therefore, we conduct an event study using the number of new players as the outcome variable and include all observations where the team size is at least 25.²⁹ Figure 2.8 presents the


²⁷As outlined in Table 2.1, the average team manager tenure duration at the time of the *Verbandspokal* final for treated and control clubs is 2.4 and 2.6 years, respectively.

²⁸In addition, we do not observe an increase in outliers which we define as clubs who have the same team manager for five or more years.

²⁹Even though 25 might seem like an already large team size, it is important to note that the scraped

average change in the number of new players per season following the *Verbandspokal* final. We observe no apparent change in the number of incoming players in response to the income shock compared to control clubs. However, this does not necessarily imply that the composition of spending remains unchanged. It is possible that clubs choose to invest the additional income in more expensive or higher-profile players rather than increasing the number of transfers. Since we do not observe transfer fees or player wages, we cannot directly test this hypothesis.

Figure 2.8: New Players per Season (Event Study)



Note: This figure shows the average change in new players per season between teams that participated in the German Cup (event year 0) after winning the *Verbandspokal* final and control teams. The outcome variable is measured in absolute numbers and captures the amount of new players for every club in a given season. Coefficient values therefore reflect differences in new players. Dots represent average effects, and vertical bars show 95% confidence intervals. (Source: Own illustration)


In the next step, we want to analyze whether there is a shift in the team composition besides the amount of new players. For example, the club might be inclined to release players that have been part of the club for longer but do not seem to match with the increased ambitions anymore following the income shock after the *Verbandspokal* final. To this end, we conduct our event study with the average player tenure as outcome variable. Figure 2.9 illustrates the results. We observe coefficients close to zero across all post-treatment periods, indicating no significant differences in player tenure between treated and control clubs. Given these findings, we conclude that the basic structure of team composition is not the reason for the observed decline in league performance. However, as noted above, we cannot test for player quality or player wages with the available data and therefore place limited emphasis on results based on player data.

2.5.4 Effect on Board of Directors

It is not only the team managers and players that affect club performance, but the club management itself could also play a pivotal role. We analyze possible structural

data includes all players that have been at least once in the squad at a match day during the entire season. Therefore, smaller team sizes usually indicate that some players are missing from the data.


Figure 2.9: Average Player Tenure (Event Study)



Note: This figure shows the average change in player tenure per season between teams that participated in the German Cup (event year 0) after winning the *Verbandspokal* final and control teams. The outcome variable is measured in years and reflects the average number of seasons that players have been part of the club. A decline in player tenure indicates a higher turnover and more new players joining the team. The coefficient values therefore reflect differences in the average stability or continuity of team composition. Dots represent average effects, and vertical bars show 95% confidence intervals. (Source: Own illustration)


changes in the club management by conducting event studies on several publicly available characteristics of the board of directors. Specifically, we are interested in proxies that could indicate a change in managerial quality. Therefore, we conduct event studies with the average age and the number of board members with doctoral degrees as outcome variables. Figures 2.10 and 2.11 illustrate the results. In both cases, we cannot identify any significant changes. Other characteristics of the board of directors also show no clear pattern, as presented in Appendix 2.B.4.

Figure 2.10: Age of Board Members (Event Study)



Note: This figure shows average changes in the age of board members before and after a team wins the *Verbandspokal* final. Positive values indicate an increase in the average age of board members. Event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final, with earlier and later seasons shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals. (Source: Own illustration)

Figure 2.11: Board Members with Doctoral Degrees (Event Study)




Note: This figure shows average changes in the number of board members with a doctoral degree before and after a team wins the *Verbandspokal* final. Positive values indicate an increase in the number of board members with a doctoral degree. Event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final, with earlier and later seasons shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals. (Source: Own illustration)

2.5.5 Effect on Political Elections

Finally, we examine whether the income shock associated with a *Verbandspokal* victory generates political spillovers. Prior research, such as Köhler (2025), finds that soccer-related events like relegations and stadium developments can influence local election outcomes. Figure 2.12 presents the estimated effect of winning the *Verbandspokal* on the electoral performance of the local council's majority party, specifically, changes in its seat share over time. While estimates remain close to zero in the early years after the German Cup participation, a clear downward trend emerges from year 4 onward. From this time onward, the estimates become increasingly negative, reaching a maximum decline of approximately 4 percentage points in the long run, which is statistically significant. This pattern suggests that voters in treated municipalities initially do not value the *Verbandspokal* victory more than voters in control municipalities value the participation in the *Verbandspokal* final — if they value it at all. Given that the distinction between treatment and control clubs was determined by a random event, unrelated to performance, this reaction appears reasonable. However, since the decline in electoral performance seems to coincide with the decline in sporting performance, it appears that voters attribute some responsibility to the majority party. This could reflect disappointment over unfulfilled expectations following the German Cup participation or a reaction to the actual deterioration in the club's league performance. This interpretation is supported by an additional empirical test which we describe in detail in Appendix 2.B.5. To summarize, we augment our event study with the change in league ranking as a control variable. Once we account for this factor, the difference in election outcomes between treated and control clubs disappears. This result aligns with our argument that it is the deterioration in league performance, rather than the cup victory itself, that voters partly attribute to local politicians, thereby

reducing their electoral support.

Figure 2.12: Seat Share of Majority Party (Event Study)



Note: This figure shows average changes in the seat share of the majority party before and after a club wins the *Verbandspokal* final. The outcome variable is the change in the seat share (in percentage points) of the party that held the majority at the time of the *Verbandspokal* final. Event year (0) marks the season of the German Cup participation, with earlier and later years shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals. (Source: Own illustration)

2.6 Conclusion

This paper analyzes how amateur soccer clubs in Germany respond to a temporary income shock resulting from winning the *Verbandspokal* final and qualifying for the German Cup. Exploiting the quasi-random nature of penalty shootouts in *Verbandspokal* finals to identify treatment and control groups, we observe club outcomes in the years before and after the shock. While we do not find evidence for persistent improvements in performance, we instead document a gradual decline in league rankings among treated clubs. This downward trend suggests that the income shock is not only ineffective in improving sporting success but is actually followed by a decline in league rankings.

A potential explanation is behavioral adjustments that suggest a departure from forward-looking decision-making. Treated clubs are more likely to retain their team manager in the years following the income shock. This pattern is consistent with a performance-based bias in managerial evaluation (e.g., [Benson et al., 2019](#)), where recent success, such as winning the *Verbandspokal* after a penalty shootout, is overemphasized relative to more informative indicators of managerial quality. This misjudgment aligns with the logic of the Peter principle and may explain why clubs fail to recalibrate their team manager and player squad composition when subsequent performance declines.

These behavioral responses and the subsequent decline in performance may relate to the concept of positive assortative matching ([Becker, 1973](#)). If long-term success depends on the compatibility between managerial strategy and team characteristics, then extending a match that was initially successful under specific circumstances may hinder future

performance. In this context, the temporary nature of the income shock may lead clubs to misinterpret short-term success as a sign of optimal matching—resulting in inefficient persistence with a suboptimal composition. Taken together, our findings complement earlier work challenging the predictions of the Permanent Income Hypothesis (PIH) by showing that organizations, like individuals, may overreact to transitory shocks and fail to adjust their expectations or strategies accordingly.

Overall, the evidence suggests that temporary income shocks at the amateur level do not generate long-term improvements in league rankings. Instead, they may trigger organizational decisions and behavioral biases that induce a decline in league rankings.

References

- AKHTARI, M., D. MOREIRA, AND L. TRUCCO (2022): “Political turnover, bureaucratic turnover, and the quality of public services,” *American Economic Review*, 112, 442–493.
- ALLISON, L. (2000): “Sport and nationalism,” *Handbook of sports studies*, 344–355.
- BECKER, G. S. (1973): “A theory of marriage: Part I,” *Journal of Political economy*, 81, 813–846.
- BENSON, A., D. LI, AND K. SHUE (2019): “Promotions and the peter principle,” *The Quarterly Journal of Economics*, 134, 2085–2134.
- BLANCHARD, O. J., F. LOPEZ-DE SILANES, AND A. SHLEIFER (1994): “What do firms do with cash windfalls?” *Journal of financial economics*, 36, 337–360.
- BORUSYAK, K., X. JARAVEL, AND J. SPIESS (2024): “Revisiting event-study designs: robust and efficient estimation,” *Review of Economic Studies*, 91, 3253–3285.
- BROLLO, F., T. NANNICINI, R. PEROTTI, AND G. TABELLINI (2013): “The political resource curse,” *American Economic Review*, 103, 1759–1796.
- CAMPBELL, J. AND A. DEATON (1989): “Why is consumption so smooth?” *The Review of Economic Studies*, 56, 357–373.
- CARROLL, C. D. (1997): “Buffer-stock saving and the life cycle/permanent income hypothesis,” *The Quarterly Journal of Economics*, 112, 1–55.
- CASELLI, F. AND A. TESEI (2016): “Resource windfalls, political regimes, and political stability,” *Review of Economics and Statistics*, 98, 573–590.
- CESARINI, D., E. LINDQVIST, M. J. NOTOWIDIGDO, AND R. ÖSTLING (2017): “The effect of wealth on individual and household labor supply: evidence from Swedish lotteries,” *American Economic Review*, 107, 3917–3946.
- CESARINI, D., E. LINDQVIST, R. ÖSTLING, AND A. TERSKAYA (2023): “Fortunate families? The effects of wealth on marriage and fertility,” Tech. rep., National Bureau of Economic Research.
- DABSCHECK, B. (1975): “Sporting equality: Labour market vs. product market control,” *Journal of Industrial Relations*, 17, 174–190.
- DAVIS, S. J., S. HANSEN, AND C. SEMINARIO-AMEZ (2025): “Macro Shocks and Firm-Level Response Heterogeneity,” Tech. rep., National Bureau of Economic Research.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–48.
- DELL, M. (2015): “Trafficking networks and the Mexican drug war,” *American Economic Review*, 105, 1738–1779.
- DEPETRIS-CHAUVIN, E., R. DURANTE, AND F. CAMPANTE (2020): “Building nations through shared experiences: Evidence from African football,” *American Economic Review*, 110, 1572–1602.
- DOBSON, S., J. A. GODDARD, AND S. DOBSON (2001): *The economics of football*, vol. 10, Cambridge University Press Cambridge.

- FORT, R. (2019): “Economic objective functions in team sports: A retrospective,” *The SAGE Handbook of Sports Economics*, 125–134.
- FRIEDMAN, M. (1958): “A Theory of the Consumption Function,” .
- GREENSTONE, M., R. HORNBECK, AND E. MORETTI (2010): “Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings,” *Journal of Political Economy*, 118, 536–598.
- HALL, R. E. (1978): “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence,” *Journal of Political Economy*, 86, 971–987.
- HOEHN, T. AND S. SZYMANSKI (1999): “The americanization of European football,” *Economic Policy*, 14, 204–240.
- KLEVEN, H. J., C. LANDAIS, AND E. SAEZ (2013): “Taxation and international migration of superstars: Evidence from the European football market,” *American Economic Review*, 103, 1892–1924.
- KÖHLER, F. (2025): “Money Pit(ch): The Impact of Soccer-Induced Municipal Debt,” *Working Paper*.
- LAIBSON, D. (1997): “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 112, 443–478.
- LEE, K., W. KANG, AND R. A. RATTI (2011): “Oil price shocks, firm uncertainty, and investment,” *Macroeconomic Dynamics*, 15, 416–436.
- LINDQVIST, E., R. ÖSTLING, AND D. CESARINI (2020): “Long-run effects of lottery wealth on psychological well-being,” *The Review of Economic Studies*, 87, 2703–2726.
- LITSCHIG, S. AND K. M. MORRISON (2013): “The impact of intergovernmental transfers on education outcomes and poverty reduction,” *American Economic Journal: Applied Economics*, 5, 206–240.
- MALMENDIER, U. AND S. NAGEL (2011): “Depression babies: do macroeconomic experiences affect risk taking?” *The Quarterly Journal of Economics*, 126, 373–416.
- MARSHALL, J. (2024): “Can close election regression discontinuity designs identify effects of winning politician characteristics?” *American Journal of Political Science*, 68, 494–510.
- MENDES, R., G. J. VAN DEN BERG, AND M. LINDEBOOM (2010): “An empirical assessment of assortative matching in the labor market,” *Labour Economics*, 17, 919–929.
- PETER, L. J. AND R. HULL (1969): *The peter principle*, vol. 4, Souvenir Press London.
- SLOANE, P. J. (1971): “Scottish journal of political economy: the economics of professional football: the football club as a utility maximiser,” *Scottish Journal of Political Economy*, 18, 121–146.
- (1976): “RESTRICTION OF COMPETITION IN PROFESSIONAL TEAM SPORTS.” *Bulletin of Economic Research*, 28.
- SZYMANSKI, S. (2003): “The economic design of sporting contests,” *Journal of Economic Literature*, 41, 1137–1187.

Appendix

2.A Data

2.A.1 *Verbandspokal* Finals

In Germany, there are 21 regional associations that host a *Verbandspokal* tournament every year.³⁰ The winner qualifies for participation in the German Cup in the following season. In table 2.A.1, we list all finals that were decided in a penalty shootout between the 1990/1991 and 2021/2022 seasons. We do not include the matches of the Lower Saxony Cup since both the winner and the runners-up are automatically qualified to participate in the German Cup in the following season.

Table 2.A.1: Regional *Verbandspokal* Finals Decided by Penalty Shootout

Date	Winner	Runners-up	Result
<i>Baden</i>			
11 May 2010	SV Sandhausen II	FC Nöttingen	7-6 (p)
23 May 2012	FC Nöttingen	SpVgg Neckarelz	4-3 (p)
25 May 2017	Karlsruher SC	1. CfR Pforzheim	5-4 (p)
<i>Bavaria</i>			
1998 ^a	SV Schalding-Heining	Post/Süd Regensburg	6-4 (p)
27 July 2004	SSV Jahn Regensburg II	TSV Aindling	6-5 (p)
18 July 2008	SpVgg Unterhaching	SpVgg Ansbach	6-5 (p)
9 May 2013	TSV 1860 Rosenheim	Wacker Burghausen	6-5 (p)
14 May 2014	Würzburger Kickers	SV Schalding-Heining	4-2 (p)
20 May 2015	SpVgg Unterhaching	SpVgg SV Weiden	6-5 (p)
5 Sept 2020	TSV 1860 München	Würzburger Kickers	4-1 (p)
27 June 2021	Türkgücü München	FV Illertissen	8-7 (p)
21 May 2022	FV Illertissen	TSV Aubstadt	4-3 (p)
<i>Berlin</i>			
4 May 2005	Tennis Borussia Berlin	BFC Alemannia 90 Wacker	5-4 (p)
<i>Brandenburg</i>			
2 June 2010	SV Babelsberg 03 II	Brandenburger SC Süd 05	4-3 (p)
5 June 2013	FSV Optik Rathenow	SV Altlüdersdorf	4-3 (p)
<i>Bremen</i>			
1991 ^a	Bremer SV	Werder Bremen II	9-8 (p)
25 May 2007	Werder Bremen II	FC Oberneuland	4-3 (p)
25 May 2017	Leher Turnerschaft	Bremer SV	9-8 (p)

³⁰Some regional associations canceled the tournament in the first year of COVID-19.

Date	Winner	Runners-up	Result
5 Sept 2020	FC Oberneuland	Blumenthaler SV	5-4 (p)
<i>Hamburg</i>			
22 June 1995	1. SC Norderstedt	FC St. Pauli II	5-4 (p)
16 May 1998	Hamburger SV II	Vfl 93	4-2 (p)
2002 ^a	USC Paloma	SC Victoria Hamburg	4-3 (p)
2014 ^a	USC Paloma	SC Condor	3-2 (p)
<i>Hesse</i>			
31 May 1998	SG Hoechst	SC Neukirchen	4-1 (p)
19 June 2014	Kickers Offenbach	SV Darmstadt 98	4-2 (p)
25 May 2017	SV Wehen Wiesbaden	SV Rot-Weiß Hadamar	4-3 (p)
<i>Lower Rhine</i>			
24 May 1997	Borussia Mönchengladbach II	Rot-Weiß Oberhausen	5-4 (p)
15 May 2015	Rot-Weiss Essen	Rot-Weiß Oberhausen	6-5 (p)
<i>Mecklenburg-Vorpommern</i>			
2009 ^a	Torgelower SV Greif	TSG Neustrelitz	3-2 (p)
28 May 2016	Hansa Rostock	FC Schönberg 95	4-3 (p)
21 May 2022	TSG Neustrelitz	Greifswalder FC	6-5 (p)
<i>Middle Rhine</i>			
10 June 2009	FC Germania Dattenfeld	Bonner SC	3-1 (p)
28 May 2016	FC Viktoria Köln	SC Fortuna Köln	6-5 (p)
<i>Rhineland</i>			
1994 ^a	FSV Salmrohr	SV Wittlich	6-3 (p)
30 May 2001	Eintracht Trier	TuS Koblenz	6-4 (p)
19 May 2004	TuS Mayen	SG Betzdorf	4-1 (p)
3 June 2015	FSV Salmrohr	Spvgg Burgbrohl	5-4 (p)
25 May 2019	FSV Salmrohr	TuS Koblenz	4-3 (p)
<i>Saarland</i>			
28 April 2005	FC Kutzhof	1. FC Saarbrücken II	7-6 (p)
2007 ^a	SV Hasborn	FC Homburg	4-1 (p)
5 Sept 2020	SV Elversberg	FC Homburg	5-4 (p)
<i>Saxony</i>			
02 May 1998	Chemnitzer FC	Erzgebirge Aue	5-4 (p)
2000 ^a	Erzgebirge Aue	VfB Leipzig	5-3 (p)
31 May 2002	Erzgebirge Aue	FSV Zwickau	5-4 (p)
<i>Saxony-Anhalt</i>			
4 May 2005	Magdeburger SV 90 Preußen	VfB Sangerhausen	4-2 (p)

Date	Winner	Runners-up	Result
14 May 2008	Hallescher FC	1. FC Magdeburg	4-3 (p)
<i>Schleswig-Holstein</i>			
16 May 2014	Holstein Kiel	ETSV Weiche	13-12 (p)
2022 ^a	VfB Lübeck	TSB Flensburg	8-7 (p)
<i>South-Baden</i>			
1992 ^a	Freiburger FC	SV Linx	4-3 (p)
1993 ^a	VfB Gaggenau	SC Freiburg II	4-3 (p)
25 May 2005	FC 08 Villingen	SC Freiburg II	3-1 (p)
4 June 2008	SC Pfullendorf	FC 08 Villingen	9-8 (p)
<i>South-West</i>			
2000 ^a	TSG Pfeddersheim	1. FC Kaiserslautern II	4-2 (p)
2012 ^a	TSG Pfeddersheim	Arminia Ludwigshafen	7-6 (p)
5 Sept 2020	1. FC Kaiserslautern	SV Alemannia Waldalgesheim	5-3 (p)
<i>Thuringia</i>			
28 May 1993	FC Carl Zeiss Jena II	SV Jenaer Glas	5-3 (p)
22 May 2001	FC Rot-Weiß Erfurt	FC Carl Zeiss Jena	7-5 (p)
28 April 2004	FC Carl Zeiss Jena	FC Rot-Weiß Erfurt II	5-3 (p)
4 May 2005	FC Rot-Weiß Erfurt II	FC Carl Zeiss Jena	7-6 (p)
17 May 2011	ZFC Meuselwitz	1. SC 1911 Heiligenstadt	4-3 (p)
<i>Westphalia</i>			
26 July 2005	Sportfreunde Siegen	VfL Bochum II	5-3 (p)
2015 ^a	Sportfreunde Lotte	SC Verl	4-3 (p)
21 May 2022	SV Rödinghausen	Preußen Münster	3-2 (p)
<i>Württemberg</i>			
24 May 2006	Stuttgarter Kickers	SSV Ulm 1846	7-6 (p)
29 May 2021	Stuttgarter Kickers	SSV Ulm 1846	5-4 (p)

Notes: This table lists all regional *Verbandspokal* finals decided by a penalty shootout.

^a Exact day of the match is not known.

^b We exclude VfB Leipzig from the analysis because the club was dissolved shortly after the event.

2.A.2 Board of Directors

Most amateur clubs are governed under an *e. V.* structure (*eingetragener Verein*, registered association). These associations are usually not required to publish financial statements but must report changes in the board of directors and indicate insolvency proceedings.³¹ We retrieved these documents to build our dataset.³² In Figure 2.A.1, we present an illustrative example for the soccer club FSV Zwickau. The provided information allows us to construct a dataset consisting of board size, average age, tenure duration, share of members with doctoral degrees, and insolvency proceedings. Missing information was hand-collected through online obituaries, newspaper articles, and career networks (LinkedIn).

Figure 2.A.1: Illustrative German Commercial Register Document

The document is a page from the German Commercial Register (Handelsregister) for the club FSV Zwickau. It is divided into several sections, with three specific areas highlighted by red circles and numbered 1, 2, and 3.

Section 1: This section lists the board of directors (Vorstand) of the club. It includes the names of the board members and their dates of appointment or resignation. The entries are as follows:

- Präsident:** [Name], 10.07.1950, Zwickau
- Schatmeister:** [Name], 21.07.1960, Langenweißbach
- Präsidentenmitglied:** [Name], 21.07.1960, Langenweißbach
- Präsidentenmitglied:** [Name], 26.11.1949, Zwickau
- Präsidentenmitglied:** [Name], 05.10.1964, Reichenbach
- Präsidentenmitglied:** [Name], 01.03.1966, Zwickau

Section 2: This section contains additional notes, including information on insolvency proceedings (Insolvenzverfahren). It mentions that the club has been declared insolvent and that the insolvency proceedings have been opened. The text states: "Über das Öffnen des Vereins wurde vom 28.12.1999 (Hr. Mz. 1N 1999/99) wurde eine vorläufige Insolvenzverwaltung bestellt. Von Amts wegen gem. § 75 Nr. 2 BGB eingetragene." It also mentions that the club has been declared insolvent and that the insolvency proceedings have been opened.

Section 3: This section reports the disclosure date (Eintragungsdatum). It states that the club was entered into the register on 12.11.2004.

The document is signed by the club's representative (Neudeck) and includes a stamp from the court (Amtsgericht Zwickau).

Note: This image illustrates the structure of documents from the German commercial register that are publicly available. Section (1) indicates changes in the board of directors. Section (2) contains additional notes, such as information on insolvency proceedings. Section (3) reports the disclosure date. (Source: Excerpt from the commercial register entry for "Fußball-Sport-Verein Zwickau e.V.")


³¹See the German Commercial Register: www.handelsregister.de.

³²No commercial register documents were available at the time of the *Verbandspokal* final for the following clubs: Erzgebirge Aue, Sportfreunde Siegen, FC Schönberg 95.

2.B Additional Results


2.B.1 Ranking Distributions

Figure 2.B.1: Ranking Distribution in (Pre-)Treatment Years



Note: This figure presents kernel density estimates of league ranking improvements for treated and control teams across pre-treatment years and the treatment year. Positive values indicate ranking improvements. Treated teams (blue) are those that won the *Verbandspokal* in the pre-treatment year 1, which is normalized to zero, while control teams (red) are those that lost in the final via penalty shootout. Horizontal bars represent 95% confidence intervals for the group means, with dots indicating the mean values. Density estimates are truncated at ± 40 to focus on the central distribution. (Source: Own illustration)

Figure 2.B.2: Ranking Distribution in Post-Treatment Years




Note: This figure presents kernel density estimates of league ranking improvements for treated and control teams across post-treatment years. Positive values indicate ranking improvements. Treated teams (blue) are those that won the *Verbandspokal* in the post-treatment year 1, which is normalized to zero, while control teams (red) are those that lost in the final via penalty shootout. Horizontal bars represent 95% confidence intervals for the group means, with dots indicating the mean values. Density estimates are truncated at ± 40 to focus on the central distribution. (Source: Own illustration)

2.B.2 Event Study Robustness Specifications

In the following, we conduct a series of robustness checks for our baseline specification of the event study following the guidelines discussed by [Borusyak et al. \(2024\)](#). We abstain from a detailed discussion since the results of all robustness specifications are fairly similar to our baseline specification.

Binning. In our baseline specification, we apply a binning procedure by combining all post-event periods beyond the event horizon into a single coefficient. In this robustness specification, we also combine all pre-event periods before the horizon into a single coefficient, represented at event time -5 . Results are illustrated in Figure 2.B.3.

Figure 2.B.3: Robustness Specification - Pre-Treatment Binning




Note: This figure shows average changes in league rankings before and after a team wins the *Verbandspokal* final. Positive values indicate improvements in final league positions (i.e., lower numerical ranks). The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in rankings, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

Sample Restriction. Because of the number of post-treatment periods, not all observations are available for estimating all event coefficients, especially in the case of later *Verbandspokal* finals. To ensure that this does not affect our results, we conduct our baseline event study but only consider *Verbandspokal* finals that occurred before 2015. Results are illustrated in Figure 2.B.4.


First Time Treated/Not Yet Treated. In this specification, we compare clubs that are treated for the first time with clubs that have not yet been treated, using the latter as the control group. Results are illustrated in Figure 2.B.5.

Figure 2.B.4: Robustness Specification - Sample Restriction Pre-2015 Finals



Note: This figure shows average improvements in league rankings before and after a team wins the *Verbandspokal* final. The sample is restricted to *Verbandspokal* finals before 2015. Positive values indicate improvements in final league positions (i.e., lower numerical ranks). The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in rankings, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

Figure 2.B.5: Robustness Specification - Not Yet Treated Control Group




Note: This figure shows average improvements in league rankings before and after a team wins the *Verbandspokal* final. The sample is adjusted to first time treated and not yet treated clubs as treatment and control clubs, respectively. Positive values indicate improvements in final league positions (i.e., lower numerical ranks). The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in rankings, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

2.B.3 Event Study Relegation/Promotion


Figure 2.B.6 presents the results from an event study, where the outcome variable is a binary indicator for relegation. The graph shows that treated clubs face an increased risk of relegation in the seasons following their *Verbandspokal* win. Several post-treatment coefficients are positive and statistically significant, suggesting that the likelihood of relegation rises after the positive income shock by up to 10 percentage points. Some pre-treatment periods also feature relatively large point estimates, though they are not statistically different from zero. This may partly reflect the limited number of actual relegation events in our sample. Taken together, the results suggest that the positive temporary income shock does not offer protection against sporting decline, and may even be followed by higher relegation risk.

Figure 2.B.6: Likelihood of Relegation Before/After *Verbandspokal* Win (Event Study)



Note: This figure shows average changes in the likelihood of relegation before and after a team wins the *Verbandspokal* final. The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in percentage points, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

For completeness, we also compute the same event study, where the outcome variable is a binary indicator for promotion. The results are illustrated in Figure 2.B.7. The pattern is less clear in this case. We do not observe systematic changes following the treatment, as most post-treatment point estimates are negative but statistically insignificant. We interpret this as suggestive evidence that in addition to the increased likelihood of relegation, the treatment clubs also experience a decreased likelihood of promotion after the *Verbandspokal* victory.

Figure 2.B.7: Likelihood of Promotion Before/After *Verbandspokal* Win (Event Study)


Note: This figure shows average changes in the likelihood of promotion before and after a team wins the *Verbandspokal* final. The event year (0) marks the season of participating in the German Cup after winning the *Verbandspokal* final in year (-1), with earlier and later seasons shown relative to this point. Dots represent average changes in percentage points, and vertical bars indicate 95% confidence intervals. (Source: Own illustration)

2.B.4 Event Study Board of Directors


Table 2.B.1: Board of Directors Characteristics of Treatment and Control Clubs

Variable	Treatment Clubs	Control Clubs
Number of Members	3.58 (1.03)	2.86 (1.06)
Average tenure (years)	5.47 (2.61)	4.40 (1.98)
Average age (years)	51.1 (7.3)	53.4 (6.9)
Share Doctoral Degrees (%)	4.3 (10.2)	3.3 (15.3)

Note: This table reports summary statistics describing the composition of the board of directors for treatment and control teams in the season of the *Verbandspokal* final. Treatment teams are those that won the *Verbandspokal* final in a penalty shootout and qualified for the German Cup. Reported values are means, with standard deviations in parentheses. Data is retrieved following the instructions presented in Appendix 2.A.2.

Figure 2.B.8: Board of Directors Size Before/After *Verbandspokal* Win (Event Study)

Note: This figure shows average changes in the size of the board of directors before and after a team wins the *Verbandspokal* final. Positive values indicate an increase in members of the board of directors. Event year (0) marks the season of the German Cup participation, with earlier and later seasons shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals. (Source: Own illustration)

Figure 2.B.9: Tenure Duration of Board Members Before/After *Verbandspokal* Win (Event Study)

Note: This figure shows average changes in the tenure duration of board members before and after a team wins the *Verbandspokal* final. Positive values indicate an increase in the average tenure duration of board members. Event year (0) marks the season of the German Cup participation, with earlier and later seasons shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals. (Source: Own illustration)


2.B.5 Event Study Elections Results

We show in Section 2.5.5 that the majority parties in the municipal council of treated clubs experience a decline in their voting share compared to the control clubs. Our supposition is that the onset of the treated club's performance deterioration is at least partly attributed to the local politicians and we want to identify whether clubs that performed similarly after the income shock also experienced a comparable outcome. To this end, we extend our baseline event study regression and control for the linear and quadratic effect of a club's change in league ranking. The regression equation is now as follows:

$$y_{i,t} = \mu_i + \theta_t + \sum_{j=-J}^J \beta_j e_{i,t}^j + X_{i,t} + \gamma_1 \Delta rank_{i,t} + \gamma_2 \Delta rank_sq_{i,t} + \varepsilon_{i,t}, \quad (2)$$


where $\Delta rank_{i,t}$ is the change in rankings relative to the season of the *Verbandspokal* final, and $\Delta rank_sq_{i,t}$ is its squared term. The results are illustrated in Figure 2.B.10. Even though the resulting coefficients for the treatment effect on election outcomes are still negative, they are not statistically different from zero anymore. Therefore, we interpret this finding as evidence for the hypothesis that the decline in the seat share of the majority party is actually in part related to the decline in sporting performance years after the *Verbandspokal* final. We also repeated this analysis including an indicator for ongoing insolvency proceedings, and the results remained almost unchanged.

Figure 2.B.10: Change in Seat Share of Majority Party After Controlling for League Ranking (Event Study)



Note: This figure shows average changes in the seat share of the majority party before and after a club wins the *Verbandspokal* final. The outcome variable is the change in the seat share (in percentage points) of the party that held the majority at the time of the *Verbandspokal* final. Additionally, we control for both the linear and quadratic effects of changes in league rankings. Event year (0) marks the season of the German Cup participation, with earlier and later years shown relative to this point. Dots represent average effects, and vertical bars display 95% confidence intervals based on profile likelihood estimates. (Source: Own illustration)

Chapter III



The Causal Effect of Domestic Market Potential*

Yann Müller

October 2025

Abstract

Historically, the access to potential markets was an important determinant for local development, and therefore, shaped the spatial distribution of economic activity. However, the evidence for the contemporaneous effect of market potential is less conclusive, especially for emerging economies. Therefore, our paper attempts to identify the causal effect of market potential in the emerging economy of Morocco. To this end, we construct a novel instrument that exploits a location's exposure to geographic conditions favoring the cultivation of an oasis. Utilizing Moroccan census data, we find that until 2014, municipalities with higher market potential experienced higher annual employment growth. After 2014, however, the effect loses significance in our IV estimates—while naïve OLS still shows a positive correlation—suggesting that any post-2014 link is driven by confounding factors rather than the actual impact of market potential.

Keywords: Market Potential, Morocco, Oases.

JEL Codes: O1, R11, R12.

*Müller: Department of Economics, University of Mannheim, yann.mueller@uni-mannheim.de. Special thanks for guidance, support, and useful advice to Miren Azkarate-Askasua, Antonio Ciccone, Andreas Gulyas, Bjarne Horst, Minki Kim, Maximilian Mähr, Ana Moreno-Maldonado, David Müller, Max Riegel, Pooja Singh, Ayman Tibi, and participants of the Mannheim Macro PhD Seminar and other doctoral colloquiums. I also want to highlight the work of the team around Violante Pallavicino at the [LabOasis Foundation](#) that has done an incredible job for the awareness and preservation of oases in Northern Africa and the Middle East.

3.1 Introduction

Why does economic activity concentrate in certain locations while remaining absent in others? And has the relevance of the determining factors changed over time? These questions have been central to spatial economics. While the existing literature has identified multiple factors that can explain these spatial patterns in developed economies, much less is known about the current impact of these factors on today's emerging economies.


The factors determining the spatial distribution of economic activity can be classified as locational fundamentals, i.e., first-nature geography, and the endogenous location choices of economic agents, i.e., second-nature geography (compare [Redding and Rossi-Hansberg, 2017](#)). Among the latter, market potential has been prominently featured in the literature. Market potential is a measure for a location's access to nearby markets and is usually computed as a distance or trade cost weighted average of nearby markets.¹ Empirical studies have established a link between locations' differential market potential and several economic outcome variables. For example, [Hanson \(2005\)](#) finds a strong correlation between market potential and wages in the US. [Redding and Sturm \(2008\)](#) exploit the post World War II division of Germany to show the negative effect of an induced change in market potential on population growth. Further, [Hanson \(1996\)](#) shows how changes in relative market potential within Mexico have led to a reallocation of economic activity towards the US border. Also, there has been a link identified between market potential and agricultural land values ([Donaldson and Hornbeck, 2016](#)). These papers illustrate the broad range of effects that are attributed to differences in market potential. However, establishing a causal relationship between market potential and economic outcomes has always been challenging due to the difficulty of disentangling its effects from those of other locational fundamentals or institutional factors ([Redding, 2022](#)).

Our paper contributes to the literature by providing new causal estimates for the effect of market potential today on the distribution of economic activity in the emerging economy of Morocco. To this end, we exploit the locations of micro-climates in the Maghreb region that historically enabled the cultivation of water and oases in arid regions. These locations have subsequently evolved into modern cities. However, we argue that a technological shock, the introduction of motorized mechanical wells, has made the ability to cultivate water within these micro-climates dispensable. In support of our argument, we document changes in population and agricultural output which indicate a catch-up of oasis provinces, that were previously constrained by their traditional water cultivation systems. We guide

¹The literature features multiple terms that refer to a similar concept (e.g., market access and market size). There is no clear distinction between these terms but usually they refer to different theoretical underpinnings. For example, the term market potential is mostly used in papers with the inverse distance or trade costs as weighting element in the estimation equation (e.g., [Brühlhart et al., 2020](#)). Meanwhile, the term market access has been primarily used in papers that derive an estimation equation from a Gravity model (e.g., [Redding and Venables, 2004](#); [Jacks and Novy, 2018](#)).

the findings with a simple model of agricultural adaptation after a technological shock. Also, we provide evidence for an increasing relevance of access to potential markets after the technological shock. We then use these insights to inform our instrumental variable construction.

Our approach relies on three propositions which we will elaborate more in detail in chapter 3.4.1. I) Historically, oases were cultivated at locations that featured a specific favorable combination of first-nature conditions, II) the combination of these conditions has no effect on today's economic outcomes other than through the locations of the oases, and III) the economic performance of cities today does not depend on whether or not they are located at a historical oasis location. Given these premises, our theory can be summarized as follows:



This observation enables us to exploit an exogenous variation in market potential arising from a location's exposure to micro-climates. Using a two-stage least squares approach, we find that market potential has a positive effect on employment growth. However, this effect is statistically significant only during the 2004–2014 period and not in the subsequent decade, 2014–2024. These results are consistent with prior literature suggesting a declining effect of market potential over time. Our findings are robust across a range of model specifications and robustness checks.

Related Literature. [Harris \(1954\)](#) has introduced the concept of market potential as a measure of the accessibility of a location to economic markets. Harris defines market potential as a weighted sum of the purchasing power of surrounding regions, where the weights reflect the inverse of the distance between the location and these regions. This approach to compute the market potential has over time been enhanced by using infrastructure-informed distances or trade costs rather than the mere distance between locations.² For example, [Brühlhart et al. \(2020\)](#) employ a lowest cost algorithm that computes the trade cost between different grid cells using information on roads, railroads, and waterway infrastructure. In a similar vein, [Donaldson and Hornbeck \(2016\)](#) put their focus on the cost of transported tonnes by utilizing information on railroad and waterway infrastructure.³

Another approach to derive an estimation equation for market potential was introduced by

²It is important to note, however, that some papers specifically outline how the results appear to be robust to different specifications of the employed distance measuring method (e.g., [Daniele et al., 2018](#)).

³[Gambuli \(2023\)](#) provides a summary of different modeling approaches in addition to the ones mentioned in our paper.

[Redding and Venables \(2004\)](#). They use an international trade Gravity model and provide a clear theoretical underpinning for the market potential equation. Others have built on this approach. [Head and Mayer \(2011\)](#) combine this kind of model with insights from the New Economic Geography literature (compare [Krugman, 1991](#)). [Liu and Meissner \(2019\)](#) derive an estimation equation for market potential from a Gravity model that allows for non-homothetic preferences.

As outlined before, the range of studied effects stemming from market potential is broad. One strand of the literature has analyzed the historical importance of cross-country differences in market potential; specifically, how lower trade barriers have increased market potential and GDP per capita in many developed countries. [Jacks and Novy \(2018\)](#) collect data on 51 countries and find a significant role of market potential on global income growth over the 20th century. [Liu and Meissner \(2019\)](#) also analyze the role of market potential in the early 20th century and find a correlation with GDP per capita. They find that specifically smaller European countries profited from a higher market potential. Also, there appears to be an effect on regions within a country with differences in market potential. [Gambuli \(2023\)](#) show that a major part of the North-South divide in Italy during the industrialization stems from differences in market potential. These papers draw a quite convincing picture of market potential and its role for today's developed economies.

Besides the historical evidence, the situation today is less clear. [Brühlhart et al. \(2020\)](#) utilize a global employment dataset and nightlight-identified market potential to show that the correlation with local growth is weakening especially in developed economies. However, according to their findings, the correlation between market potential and local growth is still sizeable in developing economies. We want to provide new causal evidence whether regions in emerging economies profit in the same way today from a higher market potential.

The main difficulty in the context of market potential is to establish a causal link that plausibly rules out confounders like locational fundamentals or institutional factors. In the historical context, many papers have utilized exogenous variations in trade costs to establish that link (e.g., [Jacks and Novy, 2018](#)). A common approach in these papers was to exploit an exogenous variation in transportation technologies (also compare [Pascali, 2017](#); [Feyrer, 2009](#)). Identifying such a variation today seems quite challenging. For this reason, our approach does not exploit changes in the trade costs between markets but rather an exogenous variation in the location of these markets itself.

The underlying idea of our instrument is related to insights from the persistence literature, i.e., we exploit historically relevant circumstances that do not have an effect on economic outcomes today other than through their historical impact. For example, [Bleakley and Lin \(2012\)](#) show that historical portage hubs in the US remain centers of economic activity, even though the historical advantage of their natural features has become dispensable.

[Davis and Weinstein \(2002\)](#) analyze how the atomic bombings of Japan affected the population distribution. They observe a rapid recovery of the affected regions, indicating the persistence in distribution of economic activity. In our case, the traditional water cultivation methods at oasis locations have become dispensable. With the introduction of motorized mechanical wells in the historical oasis locations since the 1970s, the unique combination of geographic conditions favoring the oasis cultivation have become dispensable while historical oasis locations still determine the locations of modern cities.⁴

The structure of the paper is as follows. In section 2, we provide an overview of the functioning and geographic requirements of an oasis together with the impact resulting from the introduction of motorized mechanical wells. A simple model of agricultural adaptation is used to guide the observations. In section 3, we present our approach for constructing an instrument for market potential out of these observations. In section 4, we employ a two-stage least squares regression approach to identify the causal effect of market potential on employment growth. After, we conduct a series of robustness checks. Finally, we conclude the paper with a summary of our findings.

3.2 Background

North Africa was historically a region with favorable living conditions for humans. However, the desertification that occurred between 6200 and 2200 BC triggered a shift in the allocation of human activity (compare [Bini et al., 2019](#); [Clarke et al., 2016](#)).⁵ As water became increasingly scarce, innovative cultivation methods were developed, most notably those utilized in mountain and wadi oases. These oases became the foundation for the emergence of oasis civilizations (e.g., [Kaczmarek et al., 2024](#)). Over time, many of these settlements grew into cities, with modern economic activity developing around their agricultural cores. The following chapter explores the history and functioning of oases, with a particular focus on their role after the introduction of motorized mechanical wells.

3.2.1 Oases


Oases in North Africa have long captivated the human imagination, often romanticized as verdant havens amidst the harsh desert. As early as 400 BC, the Greek historian Herodotus chronicled his fascination with the oases of Egypt and Libya in his travelogues,

⁴The determining factors for motorized mechanical wells are the groundwater level or the access to aquifers. We will discuss this matter in more detail in chapter 3.4.

⁵In archaeology, one way to determine the location of human activity in the far past is to utilize charcoal findings together with radiocarbon dating. Data on the Global Paleofire Database indicates such a shift of the spatial allocation of human activity in the mentioned time frame. [[Online Accessible](#) - Last accessed: 19.05.2025]

referring to one as “the Fountain of the Sun.”⁶ However, beyond the myths and allure, the existence of oases is not a random mystery in the vast desert. Instead, it represents a human exploitation of specific geographic conditions that allow for the cultivation of water in (semi-)arid environments.

Figure 3.1: Overview of Oases in Morocco and Illustration of a Wadi Oasis.



Note: Panel (a) shows every oasis in Morocco with the corresponding oasis classification. Panel (b) shows a wadi oasis and the dependent agricultural system. (Sources: Own illustration; [LabOasis Foundation, 2021](#))


In our dataset, an oasis is defined as a date palm grove supported by a traditional irrigation system, accompanied by a human settlement of at least 30 inhabitants with a history spanning over 100 years ([LabOasis Foundation, 2021](#)). Panel (a) in Figure 3.1 illustrates the locations of oases in Morocco, along with their corresponding classifications, which are based on the irrigation systems employed at each oasis.⁷ Morocco is home to a total of 162 oases, categorized into three main types: 142 are classified as wadi oases, 16 as mountain oases, and 4 as erg oases. The majority of these oases are located south of the Atlas Mountains.

We explain the functioning of an oasis based on the illustration of a wadi oasis shown in panel (b) of Figure 3.1. This type of oasis utilizes wadis, prehistoric riverbeds that remain dry for most of the year but still replenish groundwater in the subsurface layers and provide better access to groundwater than other places. Aeration shafts are strategically placed to facilitate maintenance and access to a drainage gallery, which captures and directs the groundwater for irrigation. This water is distributed to surrounding fields and

⁶From *The Histories*, around 430 BC. [[Online Accessible](#) - Last accessed: 19.05.2025]

⁷Our oases dataset has been documented and categorized by the [LabOasis Foundation](#). The geolocation was conducted by the [Institute of Cartography and Geoinformation](#) at ETH Zurich.

Figure 3.2: Population Development in Morocco.



Note: The figure shows the development of the urban and rural populations according to the annual reports of the National Statistics Office in Morocco for oasis provinces in panel (a) and non-oasis provinces in panel (b). The dashed vertical line marks the approximate onset of the spread of mechanical wells around 1970 according to the [LabOasis Foundation \(2021\)](#). (Source: Own illustration)

settlements via a water distribution system.


The cultivated water is then used to irrigate an ecosystem based on autogenic succession, which refers to multiple layers of vegetation, where each layer affects the ecological conditions of the successive layer ([Cowles, 1899](#)). In the case of the oasis, it is an ecosystem with three layers. The first layer consists of the date palm (*Phoenix dactylifera*), which provides shade and protection from storms for crops in the middle layer. The crops in the middle layer could then provide the basis to engage in livestock farming in the final layer.

3.2.2 Introduction of Motorized Mechanical Wells

After World War II, agriculture benefited from new technological innovations, specifically mechanical wells with motorized pumps. [Hornbeck and Keskin \(2014\)](#) summarize the technological progress of motorized mechanical wells and the effect on agriculture in the Great Plains region in the US after these wells made the Ogallala aquifer accessible in the 1950s. The widespread introduction in Morocco occurred a few years later but started to affect agriculture around 1970 ([LabOasis Foundation, 2021](#)). Before, most of the oasis locations conducted subsistence farming and the vast majority of the population in oasis provinces was considered rural. With the introduction of motorized mechanical wells in oasis regions, water became less scarce, enabling structural demographic and economic changes, which we illustrate with data from annual reports of the National Statistics Office in Morocco.⁸ The former development is illustrated in Figure 3.2. The figure shows

⁸However, it is important to mention that there was quite some heterogeneity in the development of oases after the introduction of motorized mechanical wells. [De Haas \(1998\)](#) presents different cases on how oasis locations developed after the introduction of the new technology.

Figure 3.3: Agricultural Development: Livestock



Note: The figure shows the development of livestock before and after the introduction of motorized mechanical wells. Goats are displayed in panel (a), sheep in panel (b), and cattle in panel (c). Numbers are based on the annual reports from the National Statistics Office in Morocco. Numbers are aggregated for oasis and non-oasis provinces and represented in dashed and solid lines, respectively. We normalize the numbers in the year 1970 to 100. The y-scale is log-transformed. The dashed vertical line marks the approximate onset of the spread of mechanical wells around 1970 according to the [LabOasis Foundation \(2021\)](#). (Source: Own illustration)

the development of the rural and urban populations of oasis and non-oasis provinces from 1960 to 2024 in panels (a) and (b), respectively.⁹ Two observations stand out. First, it is noticeable that the urban population in oasis provinces started to increase around the time these new motorized mechanical wells were introduced in the 1970s. Second, while the rural population in oasis provinces kept increasing, it appears to stabilize at the 1970 level in non-oasis provinces.

Further, we observe a noticeable increase in agricultural output. Figure 3.3 presents the development of livestock before and after the introduction of motorized mechanical wells. The numbers are normalized to the year 1970. For goats, sheep, and cattle we observe an increase after the introduction of mechanical wells around 1970, as shown in panel (a), (b), and (c), respectively. The figure also highlights the difference between oasis and non-oasis provinces. It becomes apparent that the relative increase in livestock is more pronounced in oasis regions.¹⁰ This development could be a result of the reduced water constraint that previously limited the water-intensive livestock farming. Due to inconsistencies in the publication of other aggregate variables in the annual reports, we cannot provide further descriptive evidence beyond the development of tourism. However, there is no visible trend and we refer to Figure 3.A.10 in Appendix 3.A.3 for more detailed information.

⁹The classification of oasis and non-oasis province is based on whether there is a documented oasis within a given province. We provide more detailed information on the workflow of constructing the aggregated time series in Appendix 3.A.3.

¹⁰Sheep and cattle already show an upward trend before 1970. For sheep, however, the growth accelerates further after 1970. In any case, our observation of a more pronounced relative increase after 1970 for oasis provinces is unaffected by this fact.

3.2.3 Agricultural Adaptation after Technological Shock

We guide the observations of agricultural output by a simple two-region model of agricultural adaptation after a technological shock that relaxes the constraint of water endowments.¹¹ We assume that the agricultural output in region $c \in \{1, 2\}$ is given by $Y_c = A_c L_c^\alpha$ with $\alpha \in (0, 1)$, where labor L_c and natural endowments A_c are used in production.¹² Total labor supply is given by $\bar{L} = L_1 + L_2$. The natural endowment index is a CES aggregate consisting of water $W_c > 0$ and other natural endowments $Z_c > 0$:¹³

$$A_c = \left[\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho \right]^{1/\rho} \quad \text{with } \rho \leq 1 \text{ and } \gamma_c \in (0, 1). \quad (1)$$

The optimal production decisions are given by $L_c^*(A_1, A_2)$. The introduction of a new technology leads to an increase in available water endowments with $W'_c = W_c + \Delta$ and $\Delta > 0$, which increases agricultural output:

$$\frac{\partial Y_1^*}{\partial W_1} > 0, \quad \frac{\partial Y_2^*}{\partial W_2} > 0. \quad (2)$$

Specifically, we consider a situation where a new technology lifts the water constraint to the same extent in both regions. For $\gamma_1 = \gamma_2 \equiv \gamma$ with $\kappa \equiv \frac{1-\gamma}{\gamma}$, this affects agricultural output as follows:

$$d \ln Y_1^* - d \ln Y_2^* > 0 \quad \Longleftrightarrow \quad W_2 \left(1 + \kappa \left(\frac{Z_2}{W_2} \right)^\rho \right) > W_1 \left(1 + \kappa \left(\frac{Z_1}{W_1} \right)^\rho \right), \quad (3)$$

i.e., the relative increase depends on the initial water level and the ratio of water to other endowments. This simple model attempts to illustrate the larger relative increase in agricultural output of oasis provinces after the introduction of motorized mechanical wells that increased water endowments in all regions. In the theoretical Appendix 3.D, we provide a proof for the comparative statics after a technological shock as shown in equation (2) and (3).

3.2.4 Change in Population Structure

Given the descriptive evidence that motorized mechanical wells enabled the urban population development in oasis provinces, we are interested in the spatial patterns of the population allocation. Specifically, we want to know whether access to markets—in this

¹¹Our model is similar to the model of agricultural adaptation by [Hornbeck and Keskin \(2014\)](#). However, we emphasize differences in regional responses instead of the response of crop usage.

¹²Agricultural output modeled as Cobb-Douglas function is common in the literature (e.g., [Lin, 1992](#); [Adamopoulos and Restuccia, 2014](#)). For simplicity, we abstract from modeling capital. We refer to [Mundlak et al. \(2012\)](#) for a detailed discussion on how to identify the agricultural production function.

¹³The complementarity of water and other endowments is discussed in the agricultural and ecological literature (e.g., [Cai et al., 2008](#); [Mundlak et al., 2012](#); [Calzadilla et al., 2010](#)).

case, specifically, access to other oases—has been a significant driver. The underlying idea is that agricultural specialization was enabled with this new technology. As a result, farmers required places to trade their goods and intermediary exporting industries, like packaging, to sell their products to domestic and international buyers.

From an economic perspective, these places should arise in central locations. There is anecdotal evidence that it was indeed centrally located municipalities and those which already served as marketplaces before, so-called *souks* (Arabic: market), that have grown most after the introduction of motorized mechanical wells. For example, the region and province capital Errachidia (formerly *Ksar es-Souk*) has been growing by a factor larger than ten since 1960. Today, the city is home to many packaging companies and hosts the International Dates Fair in its province.

To provide some formal evidence regarding the effect of access to nearby oases on population, we conduct simple regressions, similar to [Bleakley and Lin \(2012\)](#), around the time of the technological shock, i.e., the introduction of motorized mechanical wells:

$$population_m = \alpha + \beta_1 \times proximity_m + \beta_2 \times exposure_m + \epsilon_m, \quad (4)$$

where $population_m$ is the population in municipality m , $proximity_m$ measures the distance in kilometers to the closest oasis and is intended to account for the possibility that the effect of oases exposure on population is more likely to be confounded the further a municipality is away from an actual oasis. Our coefficient of interest is β_2 , which captures the effect of oases exposure in municipality m on the population level. We define $exposure_m$ as weighted average that accounts for the amount and proximity of nearby oases:

$$exposure_m = \sum_{\substack{n \in N, \\ n \neq m}} \mathbb{1}[Oasis \text{ in } n] \times \tau_{mn}^{-\gamma}, \quad (5)$$

where the indicator function equals 1 if municipality n is an oasis location, and 0 otherwise. τ_{mn} is the Euclidean distance between municipalities m and n with the trade elasticity $\gamma = 1$. We discuss these market potential style equations in detail later in section 3.3.2.

Table 3.1: Regression of Population on Oases Exposure

Population (000)	1959	1971	1994	2004	2014	2024
Oasis Exposure	−2.04	−2.60	10.75*	13.70**	19.43**	24.02***
(Percentile)	(2.15)	(3.07)	(6.17)	(5.80)	(7.70)	(8.57)
Proximity closest Oasis	0.01**	0.02***	0.13***	0.14***	0.19***	0.22***
(km)	(0.00)	(0.00)	(0.01)	(0.01)	(0.02)	(0.02)

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the estimation of equation (4). Standard errors in parentheses. All Moroccan municipalities are included for which census data is available in the respective census years. Detailed information regarding the census data is presented in Appendix 3.A.3.

We conduct this regression for every year in which a general census was conducted and present the results in Table 3.1.¹⁴ We observe that in the years before the adoption of motorized mechanical wells there was no significant association between oases exposure and the population of a municipality. However, in the years after, our estimates show an increasing and statistically significant relationship. This observation suggests that once water became less of a constraint with the introduction of motorized mechanical wells, the allocation of human activity was driven by the access to nearby oasis locations. We will exploit this insight later in our identification strategy.

3.2.5 Moroccan Economy

Since the empirical analysis in the following sections focuses on the effect of market potential today, we highlight some key developments in the Moroccan economy over the past three decades. As is the case for many emerging economies, tracking and assessing economic activity poses challenges due to limited data availability. A widely adopted approach in the literature to approximate economic development at the subnational level is the use of satellite-based nightlight luminosity, which serves as a proxy for the spatial allocation of economic activity (compare [Henderson et al., 2012](#)).


Panels (a) and (b) of Figure 3.4 display nightlight luminosity in Morocco for the years 1992 and 2021, respectively. Two key patterns emerge. First, economic activity is predominantly concentrated along the coastal corridor. In particular, notable industrial hubs include the Casablanca–Rabat axis along the Atlantic Ocean, the Tangier–Tétouan region near the Strait of Gibraltar, and Nador in the northeastern Mediterranean region. Second, economic activity in the interior is relatively more dispersed, largely shaped by the geography of the Atlas Mountains. Moreover, a substantial number of areas that exhibited no nightlight activity in 1992 show luminosity in 2021. While this shift partly reflects demographic expansion, it is also a direct consequence of Morocco’s rural electrification program. According to World Bank estimates, fewer than 15% of rural households had access to electricity in the early 1990s.¹⁵ This figure increased to nearly 100% by 2023. For this reason, we restrict the use of nightlight data to the construction of a cross-sectional market potential measure and do not employ it as a time-varying outcome variable. As emphasized in the development literature, changes in nightlight may confound the effect of market potential with the mechanical effects of an electrification roll-out. For more information on the use of nightlight, we refer to [Chen and Nordhaus \(2011\)](#) and [Michalopoulos and Papaioannou \(2013\)](#).

To evaluate economic development over time, we instead utilize microcensus data from

¹⁴We do not possess the municipality-level data for the RGPH census in 1984.

¹⁵From the World Bank Online Database: The Energy Progress Report. [[Online Accessible](#) - Last accessed: 28.09.2025]

Figure 3.4: Overview of Economic Development in Morocco.



Note: The figure shows nightlight luminosity in Morocco in 1992 and 2021 based on satellite records from [Li et al. \(2020\)](#) in panels (a) and (b), respectively. Panel (c) shows the employment growth between 2004-2014, and panel (d) the employment growth between 2014-2024, both based on the RGPH, which is described in detail in Appendix 3.A.3. Gray areas indicate that no data is available. (Source: Own illustration)

Morocco, which provides municipality-level employment statistics at decennial intervals.¹⁶ Panels (c) and (d) of Figure 3.4 illustrate employment growth between 2004–2014 and 2014–2024, respectively. The earlier period presents a highly heterogeneous pattern: employment increased in some provinces but declined in others, without a clear spatial structure. In contrast, the later period reveals a more systematic decline in employment across most provinces, with meaningful growth observed only in metropolitan coastal areas. While multiple factors likely contribute to these dynamics, our primary interest lies in the role of market potential. Specifically, we investigate whether municipalities with higher market potential experienced stronger employment growth, and whether this relationship holds across both observation periods. Additional information on socio-economic characteristics are provided in Table 3.2.

Table 3.2: Overview of Population Development in Morocco, 2004–2024.

Indicator	2004	2014	2024
Population (millions)	31.0	35.9	38.9
Active Population (millions)	11.2	12.4	11.9
Unemployment Rate (%)	12.4	16.2	21.1
Share Males (%)	49.3	49.8	49.8
Share Age 0-14 (%)	30.5	27.5	25.8
Age First Marriage (yrs)	28.7	28.7	28.7
Share Prim./Sec. Educ. (%)	49.7	52.8	56.4
Fertility Rate (#)	2.5	2.2	2.0

Note: Summary of the Moroccan population over time, including employment and socio-economic information. The data is based on the RGPH in the years 2004, 2014, and 2024, which is described in detail in Appendix 3.A.3.

3.3 Data

The following section presents the data that is employed in our paper. First, we present information about the identification and construction of micro-climates. Second, we outline our approach to measure market potential. The geographic data is available at a 30 arcsecond grid cell resolution.¹⁷

3.3.1 Micro-Climates

Our identification strategy utilizes the existence of micro-climates, which we define as a specific combination of geographic conditions within (semi-)arid regions that favor the cultivation of an oasis. In theory, an oasis could be cultivated at many locations; however,

¹⁶Detailed information on the microcensus data are presented in Appendix 3.A.3.

¹⁷At the latitude of Morocco, this corresponds to grid cells of roughly 0.9×0.9 km, i.e., about 1 km².

the returns, i.e., agricultural output, vary substantially with the prevailing conditions. Our premise is that oases, and therefore, modern cities are more likely to be located in these micro-climates since these sites had a historical advantage due to their suitability for agricultural production.¹⁸

These optimal geographic conditions are chosen based on two factors. First, conditions that are a necessity to operate the water cultivation systems at an oasis. Second, conditions that can be identified in the agricultural literature and affect the yield and resilience of a date palm.¹⁹ In the following, we briefly outline the optimal geographic conditions that favor the cultivation of an oasis and the reasoning behind some of our choices. Table 3.3 summarizes the identified requirements and in Appendix 3.A.1, we present maps indicating the locations that fulfill every outlined geographic requirement together with the data sources.

Table 3.3: Summary of Optimal Geographic Conditions for an Oasis Ecosystem

Geographic Conditions	Requirements	Source	Data Available
<i>Environmental factors</i>			
Temperature	21°C - 24°C	Alotaibi et al. (2023)	✓
Irrigation Water	1500 mm/y - 2400 mm/y	Alazba (2004)	-
Wind	3 m s ⁻¹ - 6 m s ⁻¹	Alotaibi et al. (2023)	-
Salinity	< 4.0 dS m ⁻¹ (or even 3)	Tripler et al. (2011)	-
<i>Soil factors</i>			
Texture	Sand, Sandy Loam, or Clay	Allbed et al. (2017)	✓
<i>Topology</i>			
Slope	10th percentile area slope is below 10%	Own definition	✓
Water Access	Wadi within 5km	Own definition	✓

Note: Summary of optimal geographic conditions for the cultivation of an oasis ecosystem identified from information in the agricultural and hydrological literature. The environmental factors and soil factors are related to the optimal growth and yield of date palms while the topology conditions are related to the employed traditional water irrigation system.

Water Access. In Figure 3.1 we have illustrated that the most common oasis type in Morocco is the wadi oasis, and therefore, we have chosen the distance to a wadi as a

¹⁸We want to highlight again that we mean a comparative advantage within (semi-)arid regions, not a comparison between such regions and, for example, coastal areas.

¹⁹It is without question that many of our conditions continue to affect economic outcomes today. However, we seek to emphasize that our analysis does not center on the intrinsic geographic conditions of a specific location but rather on its exposure to other locations characterized by this specific combination of geographic conditions. We will address this matter in more detail in section 3.4.1.

critical requirement.²⁰ Since we operate on a grid cell level and an oasis can encompass many square kilometers, we decided on a radius of 5km to align well with the scale of a typical oasis ecosystem. This choice is reasonable because it ensures that the proximity to a wadi is broad enough to encompass the vast majority of oases without excluding areas that could still reasonably be influenced by their presence. Notably, this choice does not lead to the exclusion of many locations, as most grid cells in the dataset still meet the criterion due to the widespread presence of wadis.

Temperature. The date palm is a very resilient plant and native to hot and dry environments. It can endure temperatures up to 50°C and tolerate cold periods down to -5°C, but the optimal average daily temperatures for fruit yield are between 21°C and 24°C (Alotaibi et al., 2023).²¹ While it would likely be possible to further refine optimal temperature conditions, for example with respect to average daily temperature variance, this is the most precise measure available from our data source.

Irrigation Water. The effective amount of irrigation water available for date palms in a given location results from the sum of precipitation, capillary rise, and soil water storage. While Alazba (2004) present optimal ranges, we do not possess data for all three sources of irrigation water, and therefore, do not include it in our identification strategy.

Wind. The date palm is more resilient to desert winds compared to other palm species (Dowson et al., 1982). However, strong winds could damage the fruit bunches. Yet, Brunel et al. (2006) find evidence for a wind-induced stomatal closure mechanism which reduces evapotranspiration (ET). This leads to a situation in which the hourly ET averages are effectively identical for wind speed magnitudes of 1.5 m s⁻¹ and 5.5 m s⁻¹. However, to the best of our knowledge there is no wind data available at a 30 arcsec resolution, and therefore, we do not include it in our identification strategy.

Salinity. Date palms were assumed to be salt tolerable (compare Maas and Hoffman, 1977). However, recent studies suggest that both date palm growth and fruit yield are severely affected by increased salinity levels, which cause a restriction in the root's water uptake stemming from decreased osmotic potential. Carr (2013) provides an overview of the literature. However, to the best of our knowledge there is no salinity data available at a 30 arcsec resolution, and therefore, we do not include it in our identification strategy.

²⁰Since we still include other oasis types in our empirical approach, this will add some noise to our estimates but also highlights the robustness of the empirical strategy.

²¹One reasonable critique regarding the use of average daily temperatures is that annual temperature averages have changed over recent decades. Unfortunately, there is no granular climatic data available for the time the oases emerged.

Texture. Soil texture is a crucial determinant of physical properties that affect crop yield (Alotaibi et al., 2023). The date palm can be cultivated in sand, sandy loam, and clays (Allbed et al., 2017). Since there is no further scientific evidence on whether any of these soil textures have advantages over the others, we include all three of them as optimal texture conditions in our analysis.

Slope. Conducting agriculture, and especially livestock farming, becomes increasingly more expensive with a location's area slope. To exclude locations with unreasonable topology, we require that at least 10% of a grid cell's area has a slope below 10%.²² For a more detailed discussion on land slope and terrain ruggedness we refer to Nunn and Puga (2012).

3.3.2 Market Potential

We compute market potential in the spirit of Harris (1954):

$$MP_{m,t} = \sum_{n \in N} Y_{n,t} \tau_{mn}^{-\gamma}, \quad (6)$$

where $MP_{m,t}$ is the market potential in municipality m at time t , $Y_{n,t}$ is a measure of the market size in municipality n , and τ_{mn} is a measure of the distance between municipality m and municipality n . γ is the trade elasticity which we, following Brühlhart et al. (2020), set equal to 1. As in many emerging economies, there is no granular data available for national income, which requires us to use a proxy for the market size $Y_{n,t}$. To this end, we use harmonized yearly satellite data by Li et al. (2020) and spatially intersect the gridded nightlight observations with the official municipality boundaries.²³ For each municipality, we assign grid cells whose centroids fall within the respective polygon. In cases where a grid cell intersects multiple municipalities, we weight its contribution by the share of its area lying within the polygon. Aggregation is then performed by summing the nightlight values of all assigned grid cells to approximate total light emissions.

Distance Measure. We construct two specifications for the distance measure τ_{mn} , namely a baseline specification using Euclidean distance and a robustness check based on the lowest-cost route. The lowest-cost route is computed using a 2D Fast Marching

²²This condition primarily excludes regions in the Atlas mountains that are unsuitable for human settlements in any case.

²³This data is available from 1992-2021 and combines records of nightlight luminosity from the DMSP (1992–2013) and the VIIRS (2012–2021) missions. Especially in development economics, there is an intense discussion about the appropriate cases for the use of nightlight data. Since we are using it as a proxy for national income in a developing economy and it has been used in a similar context (e.g., Brühlhart et al., 2020), we do not provide a more detailed discussion here and instead refer to Henderson et al. (2012) for this matter. The data is available at a 30 arcsecond resolution, allowing for a granular computation of market potential.

algorithm and the technical details are outlined in Appendix 3.C.1.

3.4 Empirical Strategy

As outlined before, estimating the causal effect of market potential on any outcome variable is challenging due to potential endogeneity. In this chapter, we present our instrumental variable approach to address this concern and discuss the validity of our instrument. After, we outline the empirical framework to estimate the causal effect of market potential on several outcome variables.

3.4.1 Exposure to Optimal Geographic Conditions

In chapter 3.3.1 we present the existence of micro-climates, i.e., a unique combination of specific geographic conditions, that favor the cultivation of an oasis.²⁴ We also outline how modern cities evolved around the location of these historical oasis sites. However, with the introduction and spread of mechanical wells since the 1970s, the comparative geographic advantage of locations within these micro-climates disappeared.²⁵ Following our line of argument that this unique combination of geographic conditions determined the location of a modern city through the outlined channel, but it no longer provides advantages today, we are able to employ an identification strategy following a similar reasoning utilized in the persistence literature (e.g., [Bleakley and Lin, 2012](#)).

To this end, we construct a measure capturing the exposure of a location to other locations that feature the optimal geographic conditions for oasis cultivation. This measure is of a similar nature to the market potential measure presented in equation (6) but utilizes an exogenous variation in the economic activity of other locations, i.e., a unique non-linear combination of geographic features in those locations.

Instrument Construction. We use the market potential formula in equation (6) as reference and construct our instrument as follows:


$$EXPO_i = \sum_{\substack{j \in J, \\ j \neq i}} GC_j d_{ij}^{-\gamma}, \quad (7)$$

where GC_j is a binary variable indicating whether grid cell j is within a micro-climate

²⁴There are also other factors which we cannot measure that ultimately define the location of an oasis, e.g., protection from threats of wildlife or other humans, location relative to other oases, and more. This does not invalidate our approach as long as sufficient variation in oasis locations is explained by geography.

²⁵To be precise, the individual conditions, e.g., temperature at a location still affects outcomes today. However, since we control for these factors individually in our regression equation later on, it is only important that the non-linear combination of these factors does not affect outcomes today except through its historical effect.

Figure 3.5: Instrument Construction: Exposure to Optimal Geographic Conditions.



Note: Panel (a) depicts oasis locations in Morocco together with grid cells that inherit optimal conditions for oasis cultivation as outlined in chapter 3.3.1. Panel (b) depicts the exposure to other grid cells that inherit the same combination of optimal geographic conditions based on equation (7). (Source: Own illustration)

that has optimal geographic conditions for oasis cultivation, i.e., all identified geographic requirements presented in chapter 3.3.1 are fulfilled.²⁶ d_{ij} is a weighting element using the distance between grid cells i and j together with the trade elasticity γ .²⁷ A subtle difference in the structure compared to equation (6) is the exclusion of grid cell i in the summation to further mitigate endogeneity concerns.²⁸ Figure 3.5 illustrates the results. In panel (a) we depict the location of oases in Morocco and indicate whether a grid cell inherits optimal geographic conditions for the cultivation of an oasis, i.e., whether the binary variable GC_j takes on the value 1 in that specific grid cell. Panel (b) shows the resulting $EXPO_i$ measure from equation (7). To work with our exposure measure on a municipality-level in our empirical approach later on, we compute the municipality exposure by averaging the exposure values of all grid cells that intersect with the municipality's boundaries.

Exclusion Restriction. The issue in the literature when identifying an exogenous cross-sectional variation in market potential is often due to a violation of the exclusion restriction. Redding (2022) outlines the difficulties of disentangling the causal effect of market access from other prominent determinants of economic performance such as institutions and locational fundamentals. In the following, we discuss potential threats to the validity of our instrument.

²⁶To avoid confusion: Since we now operate at the grid cell rather than municipality level, we use i and j as subscripts instead of m and n .

²⁷Consistent with the market potential measure, we set $\gamma = 1$.

²⁸Since we control for geographic conditions in location i , including GC_i in the summation only has a minor impact on our results but does not change them qualitatively.

Using the exposure to optimal geographic conditions for the cultivation of an oasis as an instrument rather than the location of an oasis itself allows us to mitigate several concerns. First, the decision to cultivate an oasis is driven by both first- and second-nature geography. While the former are measurable and not affected by human choice, the latter depend on the interaction of humans with the environment, which is not measurable. Therefore, focusing only on the optimal geographic conditions mitigates endogeneity concerns. Second, using the location of oases would induce measurement error because there is the possibility of a systematic vanishing or emergence of oases over time that would add noise to the oasis locations captured in our dataset.²⁹

Another potential threat to our identification stems from climatic changes in our region of interest since the desertification, i.e., the observed geography of a location today may not be the same as it was during the time the oases started to be cultivated. For three out of our four identified geographic conditions, this concern can be ruled out with certainty. Slope of landmass, location of wadis, and the texture of soil have not changed over time. Our last condition, the mean daily temperature of a location, however, is more likely to have changed. We cannot rule this out and want to highlight this for transparency.

Further, we address potential confounding variables that could affect both our instrument and our outcome variable. We argue that this is not a large concern since our instrument is a non-linear combination of very specific geographic conditions such that after controlling for these individual components in a specific location the threat of a confounding variable should be minimal.

Finally, one might argue that locations with optimal geographic conditions for oasis cultivation are more likely to undergo structural transformation since they are more likely to have higher groundwater levels. We account for this concern by controlling for the groundwater level. In addition, we conduct a robustness check by analyzing the interaction of oasis locations with our outcome variable. The detailed approach is outlined in chapter 3.6.1.

Alternative Instrument Construction. Our instrument approach classifies a location as having optimal geographic conditions based on fulfilling specific requirements for oasis cultivation. Without question, there are many other ways one could conduct this binary classification, e.g., fulfilling at least half of the requirements. To demonstrate that the results do not qualitatively depend on the structure of our classification, we employ a machine learning algorithm that computes for every location the probability of having an oasis based on the geographic conditions. Then, we use this continuous

²⁹The most likely reason for such a systematic vanishing would be wars and raids that rendered these oases uninhabitable. However, since the oases are located in less accessible regions that have not been central to any known large-scale conflicts, we assume that this is not a major concern. The Correlates of War datasets allow for an analysis of conflicts in the past 200 years. From this data we cannot observe any clear pattern indicating a systematic vanishing of oases. [[Online Accessible](#) - Last accessed: 19.05.2025]

measurement instead of the binary classification to compute the exposure to optimal geographic conditions. We present the details of the method and results of this robustness check in chapter 3.6.4.

3.4.2 Empirical Framework

We estimate the first-stage equation in two repeated cross-sections corresponding to the outcome periods 2004–2014 and 2014–2024. With a slight abuse of notation, the market potential $MP_{m,r,t}$ refers to values computed two years earlier, i.e., from 2002 and 2012, to avoid potential confounding effects. In the second stage, these estimates will be linked to employment growth rates over the periods 2004–2014 and 2014–2024. The first-stage equation is as follows:

$$MP_{m,r,t} = \beta EXPO_{m,r} + \gamma f(E_{m,r,t}) + \delta \mathbf{X}_{m,r,t} + \epsilon_{m,r,t}, \quad (8)$$

where $MP_{m,r,t}$ is the market potential percentile of municipality m in region r at period t , computed from equation (6). $EXPO_{m,r}$ is the exposure to optimal geographic conditions percentile of municipality m computed from equation (7). $f(E_{m,r,t})$ is a quadratic function of employment density $E_{m,r,t}$ expressed in percentiles.³⁰ $\mathbf{X}_{m,r,t}$ is a matrix of geographic and socio-economic control variables and regional fixed effects.³¹ Unexplained variation is captured by $\epsilon_{m,r,t}$.

In the second step, we estimate equation (9) separately for the two repeated cross-sections (i.e., growth over 2004–2014 and 2014–2024), using $\widehat{MP}_{m,r,t}$ predicted from equation (8):

$$\frac{E_{m,r,t+1} - E_{m,r,t}}{E_{m,r,t}} = \beta \widehat{MP}_{m,r,t} + \gamma f(E_{m,r,t}) + \delta \mathbf{X}_{m,r,t} + \epsilon_{m,r,t+1}, \quad (9)$$

where the dependent variable is the employment growth rate of municipality m in region r between t and $t + 1$ for the periods 2004–2014 and 2014–2024. $\widehat{MP}_{m,r,t}$ is the predicted market potential percentile from equation (8). Our coefficient of interest is β , capturing the causal effect of market potential on employment growth. $f(E_{m,r,t})$ is again a quadratic function of employment density $E_{m,r,t}$ expressed in percentiles. We also include the same set of controls $\mathbf{X}_{m,r,t}$ as in the first stage.

³⁰We include employment density in our regression design to follow Brühlhart et al. (2020) and make the results comparable. Re-computing our analysis without employment density, however, does not affect our results qualitatively.

³¹A detailed overview of the control variables is provided in Appendix 3.A.1.

3.5 Empirical Results

In the following section, we present the results of our empirical strategy.

3.5.1 Effects on Employment Growth

First Stage. Our identification strategy, outlined in Chapter 3.4.2, relies on a strong and significant relationship between the instrumental variable and the measure of market potential. In Appendix 3.B, Table 3.B.1 reports the first-stage regression results from equation (8). Panel A reports results with the market potential percentile in 2002 as the dependent variable, while Panel B does so for 2012. Across all specifications, we find a statistically significant and positive association between the exposure to optimal geographic conditions and market potential. In the most comprehensive specification in column (4), the municipality with the highest exposure to optimal geographic conditions has, on average, a 0.357 (2002) and 0.344 (2012) higher market potential percentile compared to the municipality with the lowest exposure. The corresponding F-statistics are sufficiently large in both years, supporting the robustness of our instrumental variable.

2SLS/IV. Table 3.4 presents the OLS and 2SLS/IV regression results. Columns (1) and (2) report estimates with employment growth from 2004–2014 as the dependent variable, while Columns (3) and (4) use employment growth from 2014–2024. The OLS estimates show a consistently significant and positive relationship between market potential and employment growth across all specifications and both time periods, although the magnitude of the effect is slightly smaller in the later period. However, the 2SLS (IV) results tell a different story. For 2004–2014, the causal estimates are slightly larger than their OLS counterparts. Conditional on the full set of controls, a municipality at the highest market potential percentile experienced, on average, 0.429 percentage points higher annual employment growth than one at the lowest percentile. This effect is therefore approximately the same size as the estimate reported by [Brühlhart et al. \(2020\)](#) for a similar period in emerging economies.

In contrast, for the period 2014–2024, while the OLS estimates remain positive and significant, the 2SLS estimates lose statistical significance, suggesting a diminishing causal impact of market potential on local employment growth in more recent years. [Brühlhart et al. \(2020\)](#) do not have post-2010 estimates so we cannot make an adequate comparison. Nevertheless, and similar to developed countries, it appears that the role of market potential appears to be diminishing for developing countries with a temporal delay.

Our results are robust with respect to our least-cost market potential computation presented in section 3.3.2. In Appendix 3.B, we present the estimates for the first stage and 2SLS/IV regressions in Figure 3.B.2 and 3.B.3, respectively. The results do not change qualitatively compared to the baseline specification.

Table 3.4: Regression Estimates Distance-Measured Market Potential


	Employment Growth 2004-2014		Employment Growth 2014-2024	
	(1)	(2)	(3)	(4)
Panel A: 2SLS				
Market Potential	0.394** (0.154)	0.429*** (0.146)	-0.174 (0.164)	-0.220 (0.162)
Employment Density	-0.537*** (0.152)	-0.711*** (0.146)	-0.333** (0.163)	-0.575*** (0.163)
Employment Density (sq)	0.666*** (0.126)	0.650*** (0.122)	0.663*** (0.135)	0.837*** (0.135)
Observations	1 499	1 499	1 499	1 499
R²	0.195	0.258	0.166	0.201
Panel B: OLS				
Market Potential	0.287*** (0.046)	0.308*** (0.045)	0.148*** (0.053)	0.163*** (0.052)
Employment Density	-0.480*** (0.130)	-0.649*** (0.127)	-0.509*** (0.138)	-0.771*** (0.163)
Employment Density (sq)	0.642*** (0.121)	0.632*** (0.120)	0.739*** (0.128)	0.891*** (0.131)
Observations	1 499	1 499	1 499	1 499
R²	0.198	0.261	0.187	0.229
Fixed Effects				
Region	✓	✓	✓	✓
Geographic Controls	✓	✓	✓	✓
Socio-Economic Controls	-	✓	-	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the 2SLS/IV estimation of equation (9). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the Euclidean distance.

3.5.2 Predictions by Market Potential and Employment Density

To visualize the relationship between market potential, employment density, and employment growth, we construct heat maps based on our regression estimates. Panel (a) of Figure 3.6 displays predicted annual employment growth for 2004–2014 across the full distribution of market potential and employment density, using a flexible functional form with higher-order terms of employment density. Panel (b) shows the pointwise difference in predicted growth relative to Brühlhart et al. (2020), allowing a direct comparison of the spatial growth patterns.

Figure 3.6: Employment Growth Predictions (2004–2014)




Note: Panel (a) shows the prediction for the annual employment growth based on the regression estimates in Table 3.4 for the period from 2004–2014. Panel (b) shows the percentage points difference compared to the predictions in Brühlhart et al. (2020) for emerging economies in the period 1990–2010. (Source: Own illustration)

Our results indicate that employment growth is highest in municipalities with both high market potential and either very low or very high employment density. This suggests that both emerging agglomerations and established urban areas benefit most from market access. The largest divergence from Brühlhart et al. (2020) occurs in low-density, high-market potential municipalities, where our model predicts substantially stronger growth. Moderate differences also appear in municipalities with medium employment density and medium–low market potential. These deviations should, however, be interpreted with caution, since the underlying country samples differ and the periods of observation do not fully overlap. Moreover, our specification includes only the linear effect of market potential, which may also account for part of the discrepancy.

We extend the analysis to 2014–2024 in Figure 3.7. Panel (a) displays predicted average annual employment growth across the joint distribution of market potential and employment density. Relative to 2004–2014, the growth pattern shifts noticeably. Predicted growth is now concentrated in municipalities with low to medium market potential and high employment density, while areas with high market potential but low density show weaker expected growth. Panel (b) highlights these temporal changes by plotting the

pointwise difference in predictions between the two periods.

Figure 3.7: Employment Growth Predictions (2014-2024)



Note: Panel (a) shows the prediction for the annual employment growth based on the regression estimates in Table 3.4 for the period from 2004-2014. Panel (b) shows the percentage points difference compared to the 2004-2014 predictions. (Source: Own illustration)

3.6 Robustness Checks

3.6.1 Interaction with Oasis Locations

In Section 3.4.1, we outline that our identification strategy relies on the assumption that the combination of optimal geographic conditions favoring the oasis cultivation influence economic outcomes today exclusively through their effect on the location of oases and their relative position to one another. A key threat to this identification arises if the oasis itself directly affects a municipality's economic performance, rather than merely serving as a channel.

As mentioned before, anecdotal evidence from the [LabOasis Foundation \(2021\)](#) indicates that motorized mechanical wells are used in oases today. Therefore, the historical advantage of these specific geographic conditions should not be relevant anymore. To empirically address this concern, we augment our baseline regression in three steps: (i) by controlling for the presence of an oasis within the municipality, (ii) by accounting for the presence of an oasis in adjacent municipalities, and (iii) by including an interaction term between the two.

The corresponding results are presented in Tables 3.B.4 and 3.B.5 in Appendix 3.B. We find that during the first observation period, the presence of an oasis has a statistically significant, albeit modest, effect on employment growth. More importantly, however, the coefficients on market potential remain remarkably stable across specifications, reinforcing the interpretation that our estimates reflect the causal impact of market potential rather than being driven by the direct influence of oases.

3.6.2 Limiting Potential Markets

Computing market potential as in equation (6) mechanically advantages centrally located municipalities. To address this concern and enhance the robustness of our results, we compute market potential by restricting access to nearby markets using a set of pre-defined distance thresholds:

$$MP_{m,t} = \sum_{n \in N} \mathbb{1}\{d_{mn} < \phi\} \cdot Y_{n,t} \tau_{mn}^{-\gamma}, \quad (10)$$


where the indicator function includes municipality n only if its distance from municipality m is less than $\phi \in \{25, 50, 75, 100, 250, 500\}$ kilometers. This approach ensures that our domestic market potential estimates from equation (9) are not biased by proximity to international borders—a concern mitigated especially by the use of smaller distance thresholds.

Figure 3.8 illustrates the resulting regression estimates. Panels (a) and (c) show that the estimates for both market potential measures in 2002 — total Euclidean distance and least-cost route — are robust across threshold choices. For instance, using the total Euclidean distance measure, all coefficient estimates are positive and statistically significant at the 90% confidence level, regardless of the threshold. The point estimates slightly decline from 0.75 (25 km radius) to 0.43 (500 km radius), closely aligning with the unrestricted estimate of 0.42.

Similarly, the 2012 results shown in Panels (b) and (d) exhibit consistent patterns. In both cases, the estimates are negative but statistically insignificant, and converge toward the unrestricted estimate of the market potential effect.

Another way to encounter this issue with the market potential construction would be a re-centering approach as presented by [Borusyak and Hull \(2023\)](#). They propose to bootstrap an average of market potential realizations and then only use the difference between that average and the originally constructed instrument. However, defining the distribution from which every market potential realization is drawn requires relatively strong assumptions due to which we decided to use our outlined approach.

Figure 3.8: Regression Coefficients of Distance-Restricted Market Potential Computations.



Note: Panels (a) and (b) depict the regression coefficients from equation (9) based on the distance-restricted computation of market potential using total Euclidean distance as a measure of τ_{mn} for the years 2002 and 2012, respectively. Panels (c) and (d) present the corresponding results when using a least-cost route algorithm. Each panel displays estimates across varying distance thresholds as specified in equation (10), as well as the unrestricted baseline. Error bars represent 90% confidence intervals. (Source: Own illustration)

3.6.3 Changes in Elasticities

Following up on the previous robustness check, one might ask how our choice for the trade elasticity in equation (6) affects the regression estimates.³² Therefore, we repeat our analysis for different levels of $\gamma \in \{0.25, 0.5, 0.6, 0.7, 0.8, 0.9\}$.³³ We only use lower values of γ compared to our baseline scenario since the previous exercise in section 3.6.2 with restricting access to nearby markets serves a similar purpose already, i.e., putting a lower weight on markets that are far away. The results for lower elasticities are presented in Figure 3.9. We observe, that for the entire range of elasticities, our estimated regression coefficients are almost invariant. Therefore, we are confident in the robustness of our results with respect to the chosen elasticity level.

3.6.4 Alternative Instrument Construction

In chapter 3.4.1, we address concerns that our results could depend on the binary classification of optimal geographic conditions for our instrument construction. To support the robustness of our results, we present a non-binary instrument classification using a machine learning algorithm. Specifically, we employ a gradient boosting algorithm to compute the probability that there is an oasis in a specific location conditional on geographic features. Then, we adjust our instrument construction from equation (7) and compute the exposure to optimal geographic conditions as follows:

$$EXPO_i^{ML} = \sum_{\substack{j \in J_i \\ j \neq i}} Prob_j d_{ij}^{-\gamma}, \quad (11)$$


where $Prob_j$ is the estimated probability to find an oasis in location j . Figure 3.10 illustrates the alternative construction of our exposure measure. Panel (a) shows the resulting percentiles for the exposure the optimal geographic conditions based on our machine learning approach. Panel (b) compares the baseline market potential computation with the machine learning computation for every municipality. Most of the observations are located around the 45° line which is reassuring with respect to our chosen exposure computation approach.

After, we repeat our two-stage least square regression approach outlined in equations (8) and (9) with the new exposure measure. Table 3.B.6 presents the results in Appendix 3.B. We observe that our most involved specification still yields significant positive results for the period 2004-2014, even though they are slightly larger in magnitude.

³²As mentioned before, our baseline specification is computed with $\gamma = 1$ and following [Brühlhart et al. \(2020\)](#). The authors also mention the work by [Disdier and Head \(2008\)](#), who survey the literature for common values of γ .


³³We only change the elasticity in the computation of market potential, not in the computation of our instrumental variable.

Figure 3.9: Regression Coefficients of Market Potential Computations with Varying Elasticities.



Note: Panels (a) and (b) depict the regression coefficients from equation (9) based on the computation of market potential with varying elasticities using total Euclidean distance as a measure of τ_{mn} for the years 2002 and 2012, respectively. Panels (c) and (d) present the corresponding results when using a least-cost route algorithm. Each panel displays estimates across varying elasticity levels as specified in section 3.6.3, as well as the unrestricted baseline. Error bars represent 90% confidence intervals. (Source: Own illustration)

Figure 3.10: Machine Learning Computed Exposure to Optimal Geographic Conditions.



Note: Panel (a) depicts the percentiles of machine learning computed exposure to optimal geographic conditions from equation (11). Panel (b) compares the resulting percentiles of every municipality with the baseline computation method outline in equation (7). (Source: Own illustration)

3.7 Conclusion

This paper investigates the causal relationship between domestic market potential and employment growth in Morocco. Building on insights from the history of oases in North Africa and the agricultural literature, we develop a novel instrumental variable approach to identify an exogenous variation in market potential across Moroccan municipalities. Our instrument exploits exposure to geographic micro-climates that historically favored oasis cultivation and shaped the spatial distribution of economic activity, but which no longer directly affect outcomes since the introduction of motorized mechanical wells around 1970. By isolating variation in market potential that is orthogonal to contemporary locational fundamentals, we provide new evidence on its role in shaping local development.

Our results show that higher market potential significantly increased employment growth between 2004 and 2014. The effect is economically meaningful: municipalities at the highest market potential percentile experienced more than 0.4 percentage points higher annual employment growth than those at the lowest percentile. These effects are stronger than those implied by OLS estimates, and our IV estimates for this earlier period exceed comparable findings for other emerging economies (Brühlhart et al., 2020).

In contrast, the causal effect of market potential weakens in the subsequent decade. For 2014–2024, our IV estimates lose statistical significance, even though the OLS association remains positive and significant. This suggests that the role of market potential in driving local employment growth has diminished over time, with Morocco appearing to follow the trajectory already observed in developed economies.

We remain agnostic regarding the mechanisms behind our findings. Granular firm or migration data at the municipality-level could shed light on possible determinants, as

anecdotal evidence suggests that both have played a central role in recent shifts in economic activity and population structure. Moreover, it is important to note that the exogenous variation identified here captures only the component of market potential stemming from oases. It is plausible that, following the technological shock and the resulting relative catch-up of these oasis locations, the process of convergence toward a new steady state is gradually slowing down.

References

- ADAMOPOULOS, T. AND D. RESTUCCIA (2014): “The size distribution of farms and international productivity differences,” *American Economic Review*, 104, 1667–1697.
- ALAZBA, A. (2004): “Estimating palm water requirements using Penman-Monteith mathematical model,” *Journal of King Saud University*, 16, 137–152.
- ALLBED, A., L. KUMAR, AND F. SHABANI (2017): “Climate change impacts on date palm cultivation in Saudi Arabia,” *The Journal of Agricultural Science*, 155, 1203–1218.
- ALLEN, T. AND C. ARKOLAKIS (2014): “Trade and the Topography of the Spatial Economy,” *The Quarterly Journal of Economics*, 129, 1085–1140.
- ALOTAIBI, K. D., H. A. ALHARBI, M. W. YAISH, I. AHMED, S. A. ALHARBI, F. ALOTAIBI, AND Y. KUZUYAKOV (2023): “Date palm cultivation: A review of soil and environmental conditions and future challenges,” *Land Degradation & Development*, 34, 2431–2444.
- BINI, M., G. ZANCHETTA, A. PERȘOIU, R. CARTIER, A. CATALÀ, I. CACHO, J. R. DEAN, F. DI RITA, R. N. DRYSDALE, M. FINNÈ, ET AL. (2019): “The 4.2 ka BP Event in the Mediterranean region: an overview,” *Climate of the Past*, 15, 555–577.
- BLEAKLEY, H. AND J. LIN (2012): “Portage and path dependence,” *The Quarterly Journal of Economics*, 127, 587–644.
- BORUSYAK, K. AND P. HULL (2023): “Nonrandom exposure to exogenous shocks,” *Econometrica*, 91, 2155–2185.
- BRÜLHART, M., K. DESMET, AND G.-P. KLINKE (2020): “The shrinking advantage of market potential,” *Journal of Development Economics*, 147, 102529.
- BRUNEL, J.-P., J. IHAB, A. M. DROUBI, AND S. SAMAAAN (2006): “Energy budget and actual evapotranspiration of an arid oasis ecosystem: Palmyra (Syria),” *Agricultural water management*, 84, 213–220.
- CAI, X., C. RINGLER, AND J.-Y. YOU (2008): “Substitution between water and other agricultural inputs: Implications for water conservation in a River Basin context,” *Ecological economics*, 66, 38–50.
- CALZADILLA, A., K. REHDANZ, AND R. S. TOL (2010): “The economic impact of more sustainable water use in agriculture: A computable general equilibrium analysis,” *Journal of Hydrology*, 384, 292–305.
- CARR, M. (2013): “The water relations and irrigation requirements of the date palm (*Phoenix dactylifera* L.): A review,” *Experimental Agriculture*, 49, 91–113.
- CHEN, X. AND W. D. NORDHAUS (2011): “Using luminosity data as a proxy for economic statistics,” *Proceedings of the National Academy of Sciences*, 108, 8589–8594.
- CLARKE, J., N. BROOKS, E. B. BANNING, M. BAR-MATTHEWS, S. CAMPBELL, L. CLARE, M. CREMASCHI, S. DI LERNIA, N. DRAKE, M. GALLINARO, ET AL. (2016): “Climatic changes and social transformations in the Near East and North

- Africa during the 'long' 4th millennium BC: A comparative study of environmental and archaeological evidence," *Quaternary Science Reviews*, 136, 96–121.
- COWLES, H. C. (1899): "The Ecological Relations of the Vegetation on the Sand Dunes of Lake Michigan. Part I.-Geographical Relations of the Dune Floras." *Botanical Gazette*, 27, 95–117.
- DANIELE, V., P. MALANIMA, AND N. OSTUNI (2018): "Geography, market potential and industrialization in Italy 1871–2001," *Papers in Regional Science*, 97, 639–662.
- DAVIS, D. R. AND D. E. WEINSTEIN (2002): "Bones, bombs, and break points: the geography of economic activity," *American economic review*, 92, 1269–1289.
- DE HAAS, H. (1998): "Socio-economic transformations and oasis agriculture in southern Morocco," *Looking at maps in the dark: Directions for geographical research in land management and sustainable development in rural and urban environments of the third world*, 65–78.
- DESMET, K., D. K. NAGY, AND E. ROSSI-HANSBERG (2018): "The geography of development," *Journal of Political Economy*, 126, 903–983.
- DISDIER, A.-C. AND K. HEAD (2008): "The puzzling persistence of the distance effect on bilateral trade," *The Review of Economics and statistics*, 90, 37–48.
- DONALDSON, D. AND R. HORNBECK (2016): "Railroads and American economic growth: A "market access" approach," *The Quarterly Journal of Economics*, 131, 799–858.
- DOWSON, V. ET AL. (1982): "Date production and protection with special reference to North Africa and the Near East," .
- FEYRER, J. (2009): "Distance, trade, and income-the 1967 to 1975 closing of the suez canal as a natural experiment," Tech. rep., National Bureau of Economic Research.
- GAMBULI, G. (2023): "Navigating the Geography of Regional Disparities: Market Access and the Core-Periphery Divide," Tech. rep., THEMA (THéorie Economique, Modélisation et Applications), Université de
- HANSON, G. H. (1996): "Localization Economies, Vertical Organization, and Trade," *The American Economic Review*, 86, 1266–1278.
- (2005): "Market potential, increasing returns and geographic concentration," *Journal of international economics*, 67, 1–24.
- HARRIS, C. D. (1954): "The, Market as a Factor in the Localization of Industry in the United States," *Annals of the association of American geographers*, 44, 315–348.
- HEAD, K. AND T. MAYER (2011): "Gravity, market potential and economic development," *Journal of Economic Geography*, 11, 281–294.
- HENDERSON, J. V., A. STOREYGARD, AND D. N. WEIL (2012): "Measuring economic growth from outer space," *American economic review*, 102, 994–1028.
- HORNBECK, R. AND P. KESKIN (2014): "The historically evolving impact of the ogallala aquifer: Agricultural adaptation to groundwater and drought," *American Economic Journal: Applied Economics*, 6, 190–219.

- JACKS, D. S. AND D. NOVY (2018): “Market potential and global growth over the long twentieth century,” *Journal of International Economics*, 114, 221–237.
- KACZMAREK, T., M. VAN DER VEEN, S. IVORRA, D. MATTINGLY, J.-F. TERRAL, AND M. GROS-BALTHAZARD (2024): “Origins and evolution of oasis agriculture in the Sahara: Evidence from morphometric analyses of archaeological date palm seeds,” *The Holocene*, 34, 353–365.
- KRUGMAN, P. (1991): “Increasing returns and economic geography,” *Journal of political economy*, 99, 483–499.
- LABOASIS FOUNDATION (2021): “Atlas of the Saharan and Arabian Oases, LabOasis Foundation.” <https://www.laboasis.org>.
- LI, X., Y. ZHOU, M. ZHAO, AND X. ZHAO (2020): “A harmonized global nighttime light dataset 1992–2018,” *Scientific data*, 7, 168.
- LIN, J. Y. (1992): “Rural reforms and agricultural growth in China,” *The American economic review*, 34–51.
- LIU, D. AND C. M. MEISSNER (2019): “Market potential and economic development with non-homotheticity,” *Journal of Development Economics*, 139, 217–228.
- MAAS, E. V. AND G. J. HOFFMAN (1977): “Crop salt tolerance—current assessment,” *Journal of the irrigation and drainage division*, 103, 115–134.
- MESSAGE, M. L., B. LEHNER, C. COCKBURN, N. LAMOUREUX, H. PELLA, T. SNELDER, K. TOCKNER, T. TRAUTMANN, C. WATT, AND T. DATRY (2021): “Global prevalence of non-perennial rivers and streams,” *Nature*, 594, 391–397.
- MICHALOPOULOS, S. AND E. PAPAIOANNOU (2013): “Pre-colonial ethnic institutions and contemporary African development,” *Econometrica*, 81, 113–152.
- MUNDLAK, Y., R. BUTZER, AND D. F. LARSON (2012): “Heterogeneous technology and panel data: The case of the agricultural production function,” *Journal of Development Economics*, 99, 139–149.
- NUNN, N. AND D. PUGA (2012): “Ruggedness: The blessing of bad geography in Africa,” *Review of Economics and Statistics*, 94, 20–36.
- PASCALI, L. (2017): “The wind of change: Maritime technology, trade, and economic development,” *American Economic Review*, 107, 2821–2854.
- REDDING, S. AND A. J. VENABLES (2004): “Economic geography and international inequality,” *Journal of international Economics*, 62, 53–82.
- REDDING, S. J. (2022): “Trade and geography,” *Handbook of International Economics*, 5, 147–217.
- REDDING, S. J. AND E. ROSSI-HANSBERG (2017): “Quantitative spatial economics,” *Annual Review of Economics*, 9, 21–58.
- REDDING, S. J. AND D. M. STURM (2008): “The costs of remoteness: Evidence from German division and reunification,” *American Economic Review*, 98, 1766–1797.
- TRIPLER, E., U. SHANI, Y. MUALEM, AND A. BEN-GAL (2011): “Long-term growth,

water consumption and yield of date palm as a function of salinity,” *Agricultural Water Management*, 99, 128–134.

VERKAIK, J., E. H. SUTANUDJAJA, G. H. OUDE ESSINK, H. X. LIN, AND M. F. BIERKENS (2024): “GLOBGM v1. 0: a parallel implementation of a 30 arcsec PCR-GLOBWB-MODFLOW global-scale groundwater model,” *Geoscientific Model Development*, 17, 275–300.


Appendix

3.A Data Appendix

In the following chapter we summarize and describe the data that we have utilized in the construction of our instrument or used as outcome and control variables. Further, we explain any transformations that have been conducted and not been discussed in detail during the main part of this paper.


3.A.1 Optimal Geographic Conditions for Oases Emergence

Figure 3.A.1: Geographic Conditions in Morocco



Note: The figure shows the geographic conditions in Morocco that have been utilized to construct our instrument. The strategy for choosing the conditions is explained in chapter 3.3.1. Panel (a) shows the mean annual temperature, panel (b) an approximated average slope, panel (c) the dominant soil texture, and panel (d) the distance to a wadi. The illustration uses a 30 arcsec resolution of grid cells. NAs in panel (c) indicate that the soil texture is not available which is usually the case in cities. (Source: Own illustration)

Figure 3.A.2: Instrument Construction: Fulfilled Geographic Requirements



Note: The figure shows the where the optimal geographic conditions in Morocco are fulfilled according to our requirements defined in section 3.3.1. The figure indicates in green whether the defined requirement is fulfilled in a specific grid cell. The illustration uses a 30 arcsec resolution of grid cells. (Source: Own illustration)

The reasoning for choosing specific geographic variables in our instrument construction is outlined in chapter 3.3.1. In the following we describe the data sources. The geographic conditions are illustrated in figure 3.A.1.

Temperature. We use the mean annual air temperature from the *CHELSEA v2.1* database at a 30 arcsec resolution. The temperature is defined as mean daily temperature averaged over the year.³⁴

Slope. We retrieve information on the slope of a grid cell from the *Harmonized World Soil Database v1.2* (HWSD) at a 30 arcsec resolution. Panel (b) of figure 3.A.1 shows the average slope for every grid cell. The HWSD data documents the share of landmass for a given cell that falls into eight pre-defined slope ranges, which are 0%-0.5%, 0.5%-2%, 2%-5%, 5%-10%, 10%-15%, 15%-30%, 30%-45%, 45%-100%. To compute the average slope in a grid cell, we estimate a weighted average that weighs the mean of each slope range with the corresponding share of landmass.³⁵

Soil Texture. The information on the soil texture is downloaded from the *Harmonized World Soil Database v2.0* (HWSDv2).³⁶ The HWSDv2 data contains the soil texture class at a 30 arcsec resolution. The soil texture classification follows USDA conventions which list the following texture classes: Sand, loamy sand, sandy loam, sandy clay loam, loam, silt loam, silt, silty clay loam, clay, clay loam, sandy clay and silty clay.³⁷

Distance to Wadi. Information on wadis are retrieved from a constructed dataset on non-perennial rivers and streams from [Messenger et al. \(2021\)](#).³⁸ Among these rivers, we choose those that have a predicted probability of more than 90% to be non-perennial for at least 30 days in a year. The threshold here is chosen conservatively since we do not have any references from which point a waterway system is considered a wadi. Using other reasonable thresholds, however, had no qualitative impact on our results. Then, we compute the distance for every grid cell to the nearest wadi, using the *distance()* function from the **Terra** package in R.

³⁴Data from the CHELSA database. [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁵Data from the Harmonized World Soil Database v1.2. [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁶Data from the Harmonized World Soil Database v2.0. [[Online Accessible](#) - Last accessed: 28.09.2025]


³⁷For more information regarding the soil texture classification, we refer to the USDA Field Book for Describing and Sampling Soils. [[Online Accessible](#) - Last accessed: 28.09.2025]

³⁸Data from the global prevalence of non-perennial rivers and streams database. [[Online Accessible](#) - Last accessed: 28.09.2025]

3.A.2 Geographic Control Variables

The following chapter describes the geographic control variables that we use in our regression design. The results are illustrated in figure 3.A.3.

Figure 3.A.3: Geographic Control Variables



Note: The figure shows the geographic conditions in Morocco that have been used as control variables in our empirical framework. Panel (a) shows the distance to the coastline, panel (b) the Terrain Ruggedness Index, and panel (c) the groundwater level. The illustration uses a 30 arcsec resolution of grid cells. (Source: Own illustration)

Distance to Coastline. We compute the distance to the coastline for every grid cell in Morocco by using the `st_union()`, `st_nearest_points()`, and `st_cast()` functions from the `sf` package in R.

Terrain Ruggedness. We use the Terrain Ruggedness Index (TRI) from [Nunn and Puga \(2012\)](#). Using a global elevation model, the authors estimate the terrain ruggedness for every grid cell by employing the following formula: $e_{r,c} = \sqrt{\sum_{i=r-1}^{r+1} \sum_{j=c-1}^{c+1} (e_{i,j} - e_{r,c})^2}$, which is a weighted average of elevation differences to neighboring cells.³⁹

Groundwater Level. We use the *Global-Scale Groundwater Model* (GLOBGM v1.0) published by [Verkaik et al. \(2024\)](#). As outlined in chapter 3.4.1, the groundwater level is subject to seasonal fluctuations. This model computes a steady-state for every grid cell which allows us mitigate the concern that the unique combination of these geographic conditions that favored the cultivation of an oasis still have an effect today since we can control for an average groundwater level, i.e., control for the cost of extracting water with mechanical wells.⁴⁰

³⁹Data from the website of Diego Puga. [[Online Accessible](#) - Last accessed: 28.09.2025]

⁴⁰Data from the Copernicus database. [[Online Accessible](#) - Last accessed: 28.09.2025]

3.A.3 Census and Annual Reports Data

In our paper, we employ two main sources for economic and demographic data in Morocco. First, a variety of municipality level data from the Moroccan *Recensement général* (RGPH), which is a microcensus conducted every ten years in recent times.⁴¹ Second, we include information from the annual reports of *Le Haut Commissariat au Plan* (HCP), the national statistics office in Morocco. Table 3.1 summarizes the available data.

Table 3.1: Overview Census and Annual Reports Data

Variables	1959	1960	1971	1975	1980	1984	1985	1990	1994	1995	2000	2004	2005	2010	2014	2015	2020	2024
<i>Municipality</i>																		
Population	✓	-	✓	-	-	-	-	-	✓	-	-	✓	-	-	✓	-	-	✓
Employment	-	-	-	-	-	-	-	-	-	-	-	✓	-	-	✓	-	-	✓
Demographics	-	-	-	-	-	-	-	-	-	-	-	✓	-	-	✓	-	-	✓
<i>Province</i>																		
Population	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Agriculture	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Tourism	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note: Summary of available data from Morocco on municipality- and province-level.

Data Collection. While newer reports from the HCP are available as *.xls* or readable *.pdf* files, older data is only available as scanned *.pdf* files as shown in Figure 3.A.4. We have used OpenAI’s ChatGPT to extract relevant data fields. One difficulty in merging data over time is the change in administrative boundaries. Since we use shapefiles of the current municipal boundaries, such a restructuring could induce a measurement error in our results. However, the last major communal reform was in the 1990s, so that the main results of our paper should be unaffected by it.

Additional remarks. In the background section, we analyze aggregated province-level data categorized as oasis/non-oasis province. This aggregation is based on whether there is an oasis located within that province. Since the number of provinces increased over time, there is a possibility that our observations are driven for accounting reasons rather than the introduction of motorized mechanical wells. While we cannot rule out this possibility entirely, we can mitigate the concern by stating that, to the best of our knowledge, up until the large municipal reform in 1992 a split of a former non-oasis province resulted always in multiple new non-oasis provinces, and a split of a former oasis province always resulted in multiple new oasis provinces.

⁴¹More information on the RGPH can be found on the website of the Federal Statistics Office Morocco in French and Arabic. [[Online Accessible](#) - Last accessed: 07.04.2025]

Figure 3.A.4: Census Data Illustration


PREFECTURES						العمالات	
N° de code géographique	Communes	Total المجموع	Marocains المغاربة	Etrangers الاجانب	الجماعات	الارقام الجغرافية	
17	Casablanca	965.277	850.806	114.471	الدار البيضاء	17.0.01	
18	Rabat	227.445	193.303	34.142	الرباط	18.0.01	
0.1 PROVINCE D'AGADIR						0.1 إقليم أكادير	
N° de code géographique	Communes	Total المجموع	Marocains المغاربة	Etrangers الاجانب	الجماعات	الارقام الجغرافية	
01.0.01	Ville d'Agadir	16.695	15.828	867	مدينة أكادير	01.0.01	
01.1	Cercle de Goulmine	122.030	122.014	16	دائرة كولمين	01.1	
01.1.01	Akka	10.727	10.727	—	أكا	01.1.01	
01.1.02	Asrir	7.936	7.936	—	أزرير	01.1.02	
01.1.03	Assa	7.082	7.082	—	آسا	01.1.03	
01.1.04	Bou Izakarn	5.953	5.946	7	بوزكارن	01.1.04	
01.1.05	Fask	4.522	4.522	—	فاسك	01.1.05	
01.1.06	Foum-El-Hassane	9.422	9.422	—	فم الحسن	01.1.06	
01.1.07	Goulmine	10.317	10.309	8	كولمين	01.1.07	
01.1.08	Ifrane de l'Anti-Atlas	7.964	7.964	—	يفران الأطلس الصغير	01.1.08	
01.1.09	Ksabi	10.325	10.325	—	القصابي	01.1.09	
01.1.10	Souk Jemâ-n-Tirhirte	18.570	18.570	—	سوق الجمعة تنيغرت	01.1.10	
01.1.11	Souk Tleta des Akhasass	11.719	11.718	1	سوق ثلاثاء الاخصاص	01.1.11	
01.1.12	Souk Tnine d'Adaï	3.029	3.029	—	سوق اثنين الداعي	01.1.12	
01.1.13	Souk Tnine des Aït Erkha	6.322	6.322	—	سوق اثنين آيت رخا	01.1.13	
01.1.14	Tarhijjt	8.142	8.142	—	تفجيحت	01.1.14	
01.2	Cercle d'Inezgane	283.523	282.255	1.268	دائرة انزكان	01.2	
01.2.01	Ahmar	7.222	7.210	12	أحمار	01.2.01	
01.2.02	Aït Baha	4.904	4.902	2	آيت باها	01.2.02	
03	Aït Melloul	24.371	24.207	164	آيت ملول	03	
04	Aksri	8.161	8.161	—	أكصري	04	
05	Biougra	21.750	21.750	—	بيوكرة	05	
06	El Koudia	7.951	7.928	23	الكديّة	06	
07	Imi-El-Had de Tasguedelt	10.338	10.338	—	ايمي الحد تزكدلت	07	
08	Immouzer des Ida-Outamane	8.879	8.873	6	ايموزار ادا وتنان	08	
09	Inounsiss	18.500	18.500	—	ايمونسيس	09	
10	Inchaden	11.671	11.671	—	انشادن	10	
11	Inezgane	6.917	6.333	584	انزكان	11	
12	Isk	13.803	13.803	—	اسك	12	
13	Oulad-Teïma	16.454	16.190	264	أولاد تايمّة	13	

Note: The figure illustrates the structure of the data from the RGPH. It contains information on the municipality and the corresponding population, as well as the number of Moroccans and foreigners. (Source: RGPH 1959, p.11)

In our main analysis, we use the data from the RGPH in 2004, 2014, and 2024. In every census, about 10% of the Moroccan population are surveyed. The data provides information on the demographics, employment, education, and literacy of households. Due to changes in variable availability and categorization across survey years—such as increasingly granular age groupings—we restrict our analysis to variables that are consistently available across all three waves, or to those that can be harmonized through aggregation or recombination of existing categories. In the following, we provide an overview of the socio-economic variables included in our analysis together with some additional figures of province-level data that are not included in the main text.

Share Males. Figure 3.A.5 illustrates the province-level average share of males across municipalities in Morocco. Since 2004, a slight increase in the male share is observed in most provinces. Two potential explanations for this trend are an increase in Sub-Saharan immigration or a decline in Moroccan emigration—both of which would be expected to raise the proportion of males in the population.


Figure 3.A.5: Share of Males in Morocco



Note: The figure shows the average share of males for every province in Morocco. Panel (a) shows the share of males in 2004, panel (b) the 2014 percentage points change compared to 2004, and panel (c) the 2024 percentage points change compared to 2004. Gray areas indicate that no data is available. (Source: Own illustration)

Share Age 0-14. Figure 3.A.6 illustrates the province-level average share of children aged 0-14 among the total population in 2004. We observe that the share of this age group has decreased since then in all provinces, indicating a shift in the population composition with a deceleration in population growth.


Figure 3.A.6: Share of Children Age 0-14



Note: Average share of children age 0-14 among the entire population for every province in Morocco. Panel (a) shows the share of this age group in 2004, panel (b) the 2014 percentage points change compared to 2004, and panel (c) the 2024 percentage points change compared to 2004. Gray areas indicate that no data is available. (Source: Own illustration)

Age First Marriage. Figure 3.A.7 illustrates the province-level average age of first marriage among the total population in 2004. We observe that this average increases in most provinces in 2014 and 2024, especially in the more rural regions south of the Atlas mountains.


Figure 3.A.7: Age First Marriage



Note: Average age of first marriage among the entire population for every province in Morocco. Panel (a) shows the average age in 2004, panel (b) the 2014 change in years compared to 2004, and panel (c) the 2024 change in years compared to 2004. Gray areas indicate that no data is available. (Source: Own illustration)

Share Primary and Secondary Education. Figure 3.A.8 illustrates the province-level share of the population with primary and secondary education in 2004. We observe that this share increases in almost all provinces in 2014 and 2024, especially in the coastal areas.


Figure 3.A.8: Share Primary/Secondary Education



Note: Average share of the population with at least primary or secondary education for every province in Morocco. Panel (a) shows the share in 2004, panel (b) the 2014 percentage points change compared to 2004, and panel (c) the 2024 percentage points change compared to 2004. Gray areas indicate that no data is available. (Source: Own illustration)

Fertility Rate. Figure 3.A.9 illustrates the province-level fertility rate of the population in 2004. We observe that the fertility rate decreases in all provinces in 2014 and 2024.


Figure 3.A.9: Fertility Rate



Note: Fertility rate for every province in Morocco. Panel (a) shows the rate in 2004, panel (b) the 2014 change compared to 2004, and panel (c) the 2024 change compared to 2004. Gray areas indicate that no data is available. (Source: Own illustration)

Tourism. Related to section 3.2.2, we also present descriptive evidence for the evolution of tourism in figure 3.A.10. To this end, we document the number of hotel rooms in oasis and non-oasis provinces. We can observe no visible break in trends around 1970 for both oasis and non-oasis provinces. The figure is also interesting with respect to possible effects of tourism on our main results outlined in section 3.5.1. Since, the magnitude of tourism, here measured in the number of hotel rooms, is kind of stable since the year 2000, we argue that it is unlikely to be an explanation for our observed results.

Figure 3.A.10: Tourism Development: Hotel Rooms



Note: The figure shows the development of available hotel rooms according to the annual reports of the HCP in Morocco for oasis and non-oasis provinces. Administrative changes of province boundaries over time might induce a mechanical measurement error. The dashed vertical line in the year 1971 marks the approximate onset of the spread of mechanical wells according to the [LabOasis Foundation \(2021\)](#). (Source: Own illustration)

3.B Additional Results

Table 3.B.1: Regression Estimates Distance-Measured Market Potential - First Stage

	(1)	(2)	(3)	(4)
Panel A: Market Potential 2002				
Exposure Geographic Conditions	0.310*** (0.028)	0.226*** (0.027)	0.385*** (0.032)	0.394*** (0.032)
Employment Density	-	-0.638*** (0.077)	-0.319*** (0.071)	-0.291*** (0.071)
Employment Density (sq)	-	0.338*** (0.072)	0.068 (0.066)	0.001 (0.067)
Observations	1 499	1 499	1 499	1 499
R²	0.588	0.642	0.742	0.748
F-statistic	121.17	70.68	147.76	153.80
Panel B: Market Potential 2012				
Exposure Geographic Conditions	0.303*** (0.027)	0.213*** (0.026)	0.391*** (0.029)	0.394*** (0.029)
Employment Density	-	-0.729*** (0.074)	-0.322*** (0.065)	-0.283*** (0.067)
Employment Density (sq)	-	0.431*** (0.070)	0.075 (0.060)	-0.016 (0.063)
Observations	1 499	1 499	1 499	1 499
R²	0.608	0.664	0.780	0.787
F-statistic	121.15	66.61	177.15	178.99
Fixed Effects				
Region	-	-	✓	✓
Geographic Controls	✓	✓	✓	✓
Socio-Economic Controls	-	-	-	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the first stage estimation of equation (8). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the Euclidean distance.

Table 3.B.2: Regression Estimates Least-Cost Market Potential - First Stage

	(1)	(2)	(3)	(4)
Panel A: Market Potential 2002				
Exposure Geographic Conditions	0.215*** (0.029)	0.143*** (0.028)	0.345*** (0.032)	0.357*** (0.032)
Employment Density	-	-0.743*** (0.081)	-0.341*** (0.071)	-0.315*** (0.072)
Employment Density (sq)	-	0.493*** (0.076)	0.152** (0.067)	0.087 (0.068)
Observations	1 499	1 499	1 499	1 499
R²	0.561	0.605	0.737	0.744
F-statistic	54.50	25.47	116.26	124.18
Panel B: Market Potential 2012				
Exposure Geographic Conditions	0.204*** (0.028)	0.119*** (0.027)	0.337*** (0.030)	0.344*** (0.030)
Employment Density	-	-0.826*** (0.077)	-0.355*** (0.066)	-0.331*** (0.069)
Employment Density (sq)	-	0.554*** (0.073)	0.143** (0.061)	0.073 (0.064)
Observations	1 499	1 499	1 499	1 499
R²	0.580	0.632	0.773	0.778
F-statistic	51.22	19.00	127.59	131.35
Fixed Effects				
Region	-	-	✓	✓
Geographic Controls	✓	✓	✓	✓
Socio-Economic Controls	-	-	-	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the first stage estimation of equation (8). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the least-cost route.

Table 3.B.3: Regression Estimates Least-Cost Market Potential

	Employment Growth 2004-2014		Employment Growth 2014-2024	
	(1)	(2)	(3)	(4)
Panel A: 2SLS				
Market Potential	0.439** (0.173)	0.474*** (0.163)	-0.202 (0.190)	-0.252 (0.185)
Employment Density	-0.561*** (0.159)	-0.735*** (0.152)	-0.317* (0.172)	-0.554*** (0.172)
Employment Density (sq)	0.706*** (0.132)	0.691*** (0.127)	0.647*** (0.140)	0.815*** (0.139)
Observations	1 499	1 499	1 499	1 499
R²	0.179	0.241	0.166	0.199
Panel B: OLS				
Market Potential	0.230*** (0.047)	0.251*** (0.045)	0.108** (0.053)	0.127** (0.052)
Employment Density	-0.449*** (0.131)	-0.622*** (0.128)	-0.487*** (0.138)	-0.755*** (0.172)
Employment Density (sq)	0.645*** (0.122)	0.641*** (0.121)	0.734*** (0.128)	0.895*** (0.131)
Observations	1 499	1 499	1 499	1 499
R²	0.190	0.253	0.185	0.227
Fixed Effects				
Region	✓	✓	✓	✓
Geographic Controls	✓	✓	✓	✓
Socio-Economic Controls	-	✓	-	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the 2SLS/IV estimation of equation (9). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the least-cost route.

Table 3.B.4: Regression Estimates Distance-Measured Market Potential and Oasis Locations (1/2)

	Employment Growth 2004–2014			Employment Growth 2014–2024		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 2SLS						
Market Potential	0.430*** (0.146)	0.416*** (0.151)	0.415*** (0.151)	−0.219 (0.162)	−0.223 (0.167)	−0.224 (0.168)
Oasis	0.098** (0.041)	0.108*** (0.040)	0.021 (0.089)	0.033 (0.045)	0.036 (0.045)	−0.008 (0.098)
Neighboring Oasis	-	−0.043 (0.038)	−0.054 (0.040)	-	−0.011 (0.041)	−0.017 (0.043)
Oasis × Neighboring	-	-	0.107 (0.096)	-	-	0.054 (0.105)
Employment Density	−0.684*** (0.144)	−0.711*** (0.141)	−0.706*** (0.141)	−0.569*** (0.162)	−0.575*** (0.160)	−0.574*** (0.160)
Employment Density (sq)	0.628*** (0.122)	0.650*** (0.122)	0.645*** (0.122)	0.833*** (0.135)	0.838*** (0.135)	0.837*** (0.135)
Observations	1 499	1 499	1 499	1 499	1 499	1 499
R²	0.261	0.262	0.263	0.201	0.201	0.201
Fixed Effects						
Region	✓	✓	✓	✓	✓	✓
Geographic Controls	✓	✓	✓	✓	✓	✓
Socio-Economic Controls	✓	✓	✓	✓	✓	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the 2SLS/IV estimation of equation (9). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the least-cost route.

Table 3.B.5: Regression Estimates Distance-Measured Market Potential and Oasis Locations (2/2)

	Employment Growth 2004–2014			Employment Growth 2014–2024		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel B: OLS						
Market Potential	0.321*** (0.045)	0.310*** (0.046)	0.307*** (0.046)	0.177*** (0.053)	0.183*** (0.053)	0.182*** (0.053)
Oasis	0.086** (0.038)	0.100** (0.039)	0.005 (0.087)	0.076* (0.041)	0.069 (0.042)	0.051 (0.094)
Neighboring Oasis	-	-0.055 (0.034)	-0.068* (0.036)	-	0.028 (0.037)	0.026 (0.039)
Oasis × Neighboring	-	-	0.116 (0.095)	-	-	0.022 (0.103)
Employment Density	-0.632*** (0.127)	-0.670*** (0.129)	-0.664*** (0.129)	-0.763*** (0.141)	-0.744*** (0.143)	-0.743*** (0.143)
Employment Density (sq)	0.615*** (0.120)	0.643*** (0.121)	0.638*** (0.121)	0.882*** (0.131)	0.867*** (0.132)	0.867*** (0.132)
Observations	1 499	1 499	1 499	1 499	1 499	1 499
R²	0.264	0.265	0.266	0.231	0.231	0.231
Fixed Effects						
Region	✓	✓	✓	✓	✓	✓
Geographic Controls	✓	✓	✓	✓	✓	✓
Socio-Economic Controls	✓	✓	✓	✓	✓	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the 2SLS/IV estimation of equation (9). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the least-cost route.

Table 3.B.6: Regression Estimates Machine Learning Computed Least-Cost Market Potential

	Employment Growth 2004-2014		Employment Growth 2014-2024	
	(1)	(2)	(3)	(4)
Panel A: 2SLS				
Market Potential	0.221 (0.224)	0.581** (0.229)	-0.162 (0.242)	-0.026 (0.231)
Employment Density	-0.225 (0.215)	-0.464** (0.202)	-0.619*** (0.216)	-1.118*** (0.223)
Employment Density (sq)	0.521*** (0.164)	0.573*** (0.148)	0.837*** (0.165)	1.106*** (0.162)
Observations	1 499	1 499	1 499	1 499
R²	0.150	0.233	0.179	0.181
Panel B: OLS				
Market Potential	0.215*** (0.045)	0.244*** (0.044)	0.092* (0.050)	0.127** (0.049)
Employment Density	-0.495*** (0.128)	-0.653*** (0.125)	-0.520*** (0.136)	-0.768*** (0.216)
Employment Density (sq)	0.695*** (0.119)	0.679*** (0.117)	0.774*** (0.126)	0.915*** (0.128)
Observations	1 499	1 499	1 499	1 499
R²	0.188	0.252	0.183	0.225
Fixed Effects				
Region	✓	✓	✓	✓
Geographic Controls	✓	✓	✓	✓
Socio-Economic Controls	-	✓	-	✓

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Results for the 2SLS/IV estimation of equation (9). All Moroccan municipalities are included for which census data is available in the years 2004, 2014, and 2024. Detailed information regarding the exposure to optimal geographic conditions as well as geographic and socio-economic control variables are presented in Appendix 3.A.1, 3.A.2, and 3.A.3, respectively. Market potential is computed as outlined in equation (6) using the least-cost route.

3.C Technical Appendix

3.C.1 Transportation Network: Least-Cost Route

In equation (6), we present our approach to compute market potential. As distance measure we use total Euclidean distance and transportation costs. In the following, we outline our procedure to compute municipality-to-municipality least-cost transport routes that account for the road infrastructure network in Morocco.⁴² To this end, we employ a 2D Fast Marching algorithm which is a common approach in the spatial economics literature (e.g., [Allen and Arkolakis, 2014](#); [Desmet et al., 2018](#)) and proceed as follows.

A regular grid of 5 km × 5 km cells is overlaid on the country's territory, and each grid cell is assigned a transport cost based on whether it intersects with the existing road network. Grid cells that intersect with roads are assigned a base cost of 1, while those without road coverage receive a higher cost of 3.5, reflecting more difficult or time-consuming traversal. Next, the cost surface is transformed into a raster object, with each pixel representing the assigned transport cost. This raster is used to construct a transition matrix using the `gdistance` R package. The transition matrix reflects the ease of movement between adjacent cells and serves as the basis for routing algorithms.

For each origin-destination pair of municipalities, we compute the least-cost path based on the centroids of municipal boundaries. Specifically, we calculate the cumulative cost-distance between each pair and extract the corresponding least-cost path geometry. These paths represent the lowest accumulated transport cost across the grid, rather than the shortest Euclidean distance.⁴³

We illustrate the 2D Fast Marching algorithm by finding the optimal route from Azilal to Casablanca. In Figure 3.C.1, we plot the road transportation network together with the optimal path by overlaying them on the cost surface raster, highlighting both the origin and destination municipalities as well as the least-cost route that connects them.


3.C.2 Predicting Oasis Location with Geodata using Gradient Boosting

In Chapter 3.6.4, we present an alternative instrument construction that does not employ a binary classification scheme but predicts a grid cell's probability of having an oasis cultivated based on geographic factors. Then, this probability is utilized to compute a measure of market potential. In the following, we present our gradient boosting approach employing an XGBoost classifier to determine the probabilities necessary to construct our adjusted exposure measure outlined in equation (11). We use a spatial K-Fold Cross-

⁴²We use the Global Roads Open Access dataset. [[Online Accessible](#) - Last accessed: 28.09.2025]

⁴³The transportation network dataset was published in 2013. Therefore, it could be the case that in our earlier observation period market potential is overestimated because some roads were not built at the time. This would induce a downward bias in our estimates for the effect of market potential on employment growth. Therefore, we are not too much concerned with respect to the validity of our empirical results since they represent a lower bound in that case.

Figure 3.C.1: Least-Cost Path from Azilal to Casablanca



Note: The figure shows the cost of transportation in Morocco induced by the road transportation network on a $5\text{km} \times 5\text{km}$ raster. The red line represents the least-cost transportation route from Azilal to Casablanca following a 2D Fast Marching algorithm. (Source: Own illustration)

Validation approach, where each fold corresponds to a different province in Morocco. Let the full dataset be denoted as:

$$D = \{(x_i, z_i, y_i, p_i)\}_{i=1}^N$$

where:

- $x_i \in \mathbb{R}^2$ are the spatial coordinates (longitude and latitude),
- $z_i \in \mathbb{R}^d$ is a vector of d covariates (all available geodata together with information from neighboring cells),
- $y_i \in \{0, 1\}$ is the binary outcome variable,
- $p_i \in \mathcal{P}$ indicates the province to which observation i belongs.

Let $\mathcal{P} = \{P_1, P_2, \dots, P_K\}$ denote the set of K provinces (folds). For each fold $k \in \{1, \dots, K\}$, we define:

- **The test set:**

$$D_{\text{test}}^{(k)} = \{(x_i, z_i, y_i) \mid p_i = P_k\}$$

- The **training set**:

$$D_{\text{train}}^{(k)} = D \setminus D_{\text{test}}^{(k)}$$

We train a gradient-boosted decision tree model using XGBoost on the training data for every fold. The technical implementation was conducted using the R package `xgboost`.⁴⁴ We use the following hyperparameters: maximum depth = 100, learning rate $\eta = 1$, number of boosting rounds = 100, and objective = `binary:logistic`.

After generating predictions for each grid cell using the gradient boosting model, we construct a spatially smoothed probability surface to facilitate interpretation of regional patterns. While the raw prediction vector is informative at the grid cell level, it often results in a pixelated representation when visualized across the entire province. To highlight broader spatial trends and reduce local noise, we apply a weighted kernel density estimation procedure. In practice, we use the `ks` package in R to estimate this weighted kernel density surface.⁴⁵

Predictions. Figure 3.C.2 presents the computed density maps for every province. The predictive performance of the XGBoost model is overall satisfactory. In provinces without any observed oases, the model consistently assigns near-zero probabilities to the vast majority of grid cells, indicating a high degree of precision in distinguishing true negatives. In provinces where oases are present, the algorithm correctly identifies a substantial share of oasis clusters, assigning elevated probabilities to the appropriate spatial regions. The incidence of false positives, i.e., areas predicted to host oases where none are observed, is low which highlights the robustness of the model's predictions.


Nevertheless, certain oasis clusters are not accurately detected by the model. This likely reflects the fact that while geographic covariates account for a portion of the spatial variation in oasis locations, they do not fully capture the underlying determinants. In particular, second-nature geographic factors, such as patterns of human settlement, are not incorporated in the feature set and may play a non-negligible role in shaping oasis distribution.

For the purpose of constructing an alternative instrument, however, the attained level of predictive accuracy is sufficient. The objective is not to achieve perfect out-of-sample classification performance, but rather to isolate the component of variation in oasis presence that is attributable to exogenous geographic features, which is subsequently used to construct a plausibly exogenous instrumental variable.

⁴⁴We refer to the technical report of the `xgboost` package for more details. [[Online Accessible](#) - Last accessed: 28.09.2025]

⁴⁵We refer to the technical report of the `ks` package for more details. [[Online Accessible](#) - Last accessed: 28.09.2025]

Figure 3.C.2: Province Density Maps of XGBoost-Predicted Oasis Probabilities



Note: The figure shows the density map of XGBoost-predicted oasis probabilities in Morocco by using a K-fold cross validation algorithm in which every province represents a fold. The computation follows the outlined approach in Chapter 3.C.2. (Source: Own illustration)

3.D Theory Appendix

In the following we provide proofs for the results presented in Section 3.2.3. Agricultural output is produced with the following production function:

$$Y_c = A_c L_c^\alpha, \quad \alpha \in (0, 1),$$

where A_c is an effective endowment index and L_c is employed labor in region $c \in \{1, 2\}$. Total labor is fixed, i.e. $L_1 + L_2 = \bar{L}$, and total output is $Y = Y_1 + Y_2$.

The endowment index aggregates water W_c and other endowments Z_c via CES:

$$A_c = \left[\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho \right]^{1/\rho}, \quad \gamma_c \in (0, 1), \quad \rho \leq 1, \quad Z_c > 0,$$

and the regions differ in their water endowments, i.e. $W_2 > W_1 > 0$.

3.D.1 Optimal Allocation

The representative farmer maximizes total output:

$$\max_{L_1, L_2} Y = A_1 L_1^\alpha + A_2 L_2^\alpha \quad \text{s.t.} \quad L_1 + L_2 = \bar{L}.$$

The Lagrangian is

$$\mathcal{L} = A_1 L_1^\alpha + A_2 L_2^\alpha + \lambda(\bar{L} - L_1 - L_2).$$

First-order conditions are

$$\alpha A_1 (L_1^*)^{\alpha-1} = \lambda^*, \tag{1}$$

$$\alpha A_2 (L_2^*)^{\alpha-1} = \lambda^*, \tag{2}$$

$$L_1^* + L_2^* = \bar{L}. \tag{3}$$

Therefore, at the optimum, marginal products of labor are equalized across regions:

$$\alpha A_1 (L_1^*)^{\alpha-1} = \alpha A_2 (L_2^*)^{\alpha-1}.$$

This implies the allocation rule

$$\frac{L_1^*}{L_2^*} = \left(\frac{A_1}{A_2} \right)^{\frac{1}{1-\alpha}}.$$

Using the resource constraint,

$$L_1^* = \frac{\bar{L} A_1^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}, \quad L_2^* = \frac{\bar{L} A_2^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}.$$

3.D.2 Comparative Statics

Proposition 1. At the optimum, total equilibrium output Y^* is increasing in each region's endowment A_c :

$$\frac{\partial Y^*}{\partial A_c} > 0, \quad (c = 1, 2).$$

Proof. Equilibrium output is

$$Y^* = A_1(L_1^*)^\alpha + A_2(L_2^*)^\alpha.$$

Differentiating with respect to A_1 :

$$\frac{\partial Y^*}{\partial A_1} = (L_1^*)^\alpha + \alpha A_1(L_1^*)^{\alpha-1} \frac{\partial L_1^*}{\partial A_1} + \alpha A_2(L_2^*)^{\alpha-1} \frac{\partial L_2^*}{\partial A_1}.$$

From the FOCs (1)–(2),

$$\alpha A_1(L_1^*)^{\alpha-1} = \alpha A_2(L_2^*)^{\alpha-1} = \lambda^*.$$

From the resource constraint (3),

$$\frac{\partial L_1^*}{\partial A_1} + \frac{\partial L_2^*}{\partial A_1} = 0,$$

and therefore, the last two terms cancel:

$$\alpha A_1(L_1^*)^{\alpha-1} \frac{\partial L_1^*}{\partial A_1} + \alpha A_2(L_2^*)^{\alpha-1} \frac{\partial L_2^*}{\partial A_1} = \lambda^* \left(\frac{\partial L_1^*}{\partial A_1} + \frac{\partial L_2^*}{\partial A_1} \right) = 0.$$

This results in:

$$\frac{\partial Y^*}{\partial A_1} = (L_1^*)^\alpha > 0,$$

since $L_1^* > 0$. The case $c = 2$ is identical. □

Proposition 2. An increase in A_1 raises L_1^* and lowers L_2^* , and conversely for A_2 .

Proof. From the closed-form solution,

$$L_1^* = \frac{\bar{L} A_1^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}, \quad L_2^* = \frac{\bar{L} A_2^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}.$$

Differentiating L_1^* w.r.t. A_1

$$\begin{aligned} \frac{\partial L_1^*}{\partial A_1} &= \bar{L} \cdot \frac{\frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}-1} (A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}) - A_1^{\frac{1}{1-\alpha}} \cdot \frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}-1}}{(A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}})^2} \\ &= \frac{\bar{L} \frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}-1} A_2^{\frac{1}{1-\alpha}}}{(A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}})^2} > 0, \end{aligned}$$

since $\bar{L} > 0$, $\frac{1}{1-\alpha} > 0$, and $A_1, A_2 > 0$. By $L_1^* + L_2^* = \bar{L}$,

$$\frac{\partial L_2^*}{\partial A_1} = -\frac{\partial L_1^*}{\partial A_1} < 0.$$

The case for A_2 is symmetric. □

3.D.3 Technological Shock

Now, we want to analyze how a technological shock similar to the introduction of motorized mechanical wells, i.e., $W'_c = g(W_c)$ and $W'_c > W_c$, affects production and labor allocation.

$$\begin{aligned} \frac{\partial A_c}{\partial W_c} &= \gamma_c W_c^{\rho-1} \left[\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho \right]^{\frac{1}{\rho}-1} \\ &= \gamma_c W_c^{\rho-1} A_c^{1-\rho} > 0, \end{aligned}$$

since $\gamma_c > 0$, $W_c^{\rho-1} > 0$ for $W_c > 0$, and $A_c^{1-\rho} > 0$. Clearly, using Proposition 1 and Proposition 2, we know that $\frac{\partial Y_c^*}{\partial W_c} > 0$ and $\frac{\partial L_c^*}{\partial W_c} > 0$. However, we are interested in a shock that affects both regions at the same time.

Additive Shock. We consider an additive water shock in both regions:

$$W'_c = W_c + \Delta, \quad \Delta > 0, \quad c \in \{1, 2\},$$

Proposition 3. In a first-order differential comparison, initial water endowments and the ratio of water to other endowments determine which region experiences the larger relative increase in output with $\gamma_1 = \gamma_2 \equiv \gamma$ and $\kappa = (1 - \gamma)/\gamma$:

$$\begin{aligned} d \ln Y_1^* - d \ln Y_2^* &> 0 \\ \iff W_2 \left(1 + \kappa \left(\frac{Z_2}{W_2} \right)^\rho \right) &> W_1 \left(1 + \kappa \left(\frac{Z_1}{W_1} \right)^\rho \right), \end{aligned}$$

Proof. A first-order approximation of the log-change in A_c from the additive shock is

$$d \ln A_c = \frac{\partial A_c / \partial W_c}{A_c} \Delta = \frac{\gamma_c W_c^{\rho-1} A_c^{1-\rho}}{A_c} \Delta = \underbrace{\frac{\gamma_c W_c^\rho}{\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho}}_{:= s_{W,c} \text{ ("CES Water Share")}} \cdot \frac{\Delta}{W_c} = s_{W,c} \frac{\Delta}{W_c}.$$

For each region, the log-change in output decomposes as

$$d \ln Y_c^* = d \ln A_c + \alpha d \ln L_c^*.$$

We start from the log of the optimal allocation:

$$L_1^* = \frac{\bar{L} A_1^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}} \implies \ln L_1^* = \ln \bar{L} + \frac{1}{1-\alpha} \ln A_1 - \ln \left(A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}} \right).$$

Differentiating both sides, using the following identities:

$$(i) \quad d \ln u = \frac{du}{u}, \quad (ii) \quad d(x^p) = p x^{p-1} dx = p x^p d \ln x,$$

then we set $h(A_1, A_2) = A_1^{1/(1-\alpha)} + A_2^{1/(1-\alpha)}$:

$$d \ln L_1^* = d(\ln \bar{L}) + \frac{1}{1-\alpha} d \ln A_1 - d \ln h,$$

Since $\bar{L}$ is constant, $d(\ln \bar{L}) = 0$. Now, we expand $d \ln h$ via (i):

$$d \ln h = \frac{dh}{h},$$

and compute dh using (ii) term by term:

$$dh = d \left(A_1^{\frac{1}{1-\alpha}} \right) + d \left(A_2^{\frac{1}{1-\alpha}} \right) = \frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}} d \ln A_1 + \frac{1}{1-\alpha} A_2^{\frac{1}{1-\alpha}} d \ln A_2.$$

Substituting dh/h back:

$$d \ln h = \frac{1}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}} \left[\frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}} d \ln A_1 + \frac{1}{1-\alpha} A_2^{\frac{1}{1-\alpha}} d \ln A_2 \right].$$

After, we put this into $d \ln L_1^*$:

$$d \ln L_1^* = \frac{1}{1-\alpha} d \ln A_1 - \frac{1}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}} \left[\frac{1}{1-\alpha} A_1^{\frac{1}{1-\alpha}} d \ln A_1 + \frac{1}{1-\alpha} A_2^{\frac{1}{1-\alpha}} d \ln A_2 \right].$$

Defining labor-share weights induced by A_c :

$$\omega_1 \equiv \frac{A_1^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}, \quad \omega_2 \equiv 1 - \omega_1 = \frac{A_2^{\frac{1}{1-\alpha}}}{A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}}}.$$

Re-writing change in labor allocation using these weights:

$$d \ln L_1^* = \frac{1}{1-\alpha} \left[d \ln A_1 - \omega_1 d \ln A_1 - \omega_2 d \ln A_2 \right] = \frac{1}{1-\alpha} \omega_2 (d \ln A_1 - d \ln A_2).$$

Similarly, from $\ln L_2^* = \ln \bar{L} + \frac{1}{1-\alpha} \ln A_2 - \ln(A_1^{\frac{1}{1-\alpha}} + A_2^{\frac{1}{1-\alpha}})$ we get

$$d \ln L_2^* = \frac{1}{1-\alpha} \omega_1 (d \ln A_2 - d \ln A_1).$$

Plugging these into output changes

$$d \ln Y_1^* = d \ln A_1 + \alpha d \ln L_1^* = d \ln A_1 + \frac{\alpha}{1-\alpha} \omega_2 (d \ln A_1 - d \ln A_2),$$

$$d \ln Y_2^* = d \ln A_2 + \alpha d \ln L_2^* = d \ln A_2 + \frac{\alpha}{1-\alpha} \omega_1 (d \ln A_2 - d \ln A_1).$$

Now, we compare the relative output changes across regions by subtracting the two:

$$\begin{aligned} d \ln Y_1^* - d \ln Y_2^* &= (d \ln A_1 - d \ln A_2) + \frac{\alpha}{1-\alpha} \left[\omega_2 (d \ln A_1 - d \ln A_2) + \omega_1 (d \ln A_1 - d \ln A_2) \right] \\ &= \left(1 + \frac{\alpha}{1-\alpha} \right) (d \ln A_1 - d \ln A_2) = \frac{1}{1-\alpha} (d \ln A_1 - d \ln A_2). \end{aligned}$$

Using $d \ln A_c = s_{W,c} \frac{\Delta}{W_c}$, we obtain

$$d \ln Y_1^* - d \ln Y_2^* = \frac{\Delta}{1-\alpha} \left(\frac{s_{W,1}}{W_1} - \frac{s_{W,2}}{W_2} \right).$$

Condition determining which region experiences the larger relative increase in output:

$$\boxed{d \ln Y_1^* > d \ln Y_2^* \iff \frac{s_{W,1}}{W_1} > \frac{s_{W,2}}{W_2}}$$

Now, we recall that

$$s_{W,c} = \frac{\gamma_c W_c^\rho}{\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho},$$

hence

$$\frac{s_{W,c}}{W_c} = \frac{\gamma_c W_c^{\rho-1}}{\gamma_c W_c^\rho + (1 - \gamma_c) Z_c^\rho} = \frac{1}{W_c} \cdot \frac{\gamma_c}{\gamma_c + (1 - \gamma_c) \left(\frac{Z_c}{W_c}\right)^\rho}.$$

Therefore, a necessary and sufficient condition for the initially water-scarce region ($W_1 < W_2$) to have the larger relative increase in output is

$$\frac{1}{W_1} \cdot \frac{\gamma_1}{\gamma_1 + (1 - \gamma_1) \left(\frac{Z_1}{W_1}\right)^\rho} > \frac{1}{W_2} \cdot \frac{\gamma_2}{\gamma_2 + (1 - \gamma_2) \left(\frac{Z_2}{W_2}\right)^\rho}$$

Assumption. Identical CES weights $\gamma_1 = \gamma_2 \equiv \gamma$. With $\kappa \equiv \frac{1-\gamma}{\gamma} > 0$ and $q_c \equiv \frac{W_c}{Z_c}$, our condition becomes:

$$\begin{aligned} & \frac{1}{W_1} \cdot \frac{\gamma}{\gamma + (1 - \gamma) \left(\frac{Z_1}{W_1}\right)^\rho} > \frac{1}{W_2} \cdot \frac{\gamma}{\gamma + (1 - \gamma) \left(\frac{Z_2}{W_2}\right)^\rho} \\ \iff & \frac{1}{Z_1} \cdot \frac{1}{q_1(1 + \kappa q_1^{-\rho})} > \frac{1}{Z_2} \cdot \frac{1}{q_2(1 + \kappa q_2^{-\rho})} \\ \iff & Z_2 q_2(1 + \kappa q_2^{-\rho}) > Z_1 q_1(1 + \kappa q_1^{-\rho}) \\ \iff & W_2 \left(1 + \kappa \left(\frac{Z_2}{W_2}\right)^\rho\right) > W_1 \left(1 + \kappa \left(\frac{Z_1}{W_1}\right)^\rho\right), \end{aligned}$$

which is the condition for the initially water-scarce region 1 ($W_1 < W_2$) to experience a larger relative increase in agricultural output. $\square$

Resources

Footnote	Link
<i>Chapter I</i>	
6	https://www.destatis.de/DE/Presse/Pressemitteilungen/2022/04/PD22_148_463.html
7	https://www.destatis.de/DE/Themen/Branchen-Unternehmen/Transport-Verkehr/Gueterverkehr/Publikationen/_publikationen-innen-schiffahrt.html
10	https://www.contargo.net/en/business/auxiliary-conditions/low-water
30	https://www-genesis.destatis.de/datenbank/online/statistic/46321/details
31	https://www.elwis.de/DE/Service/Wasserstaende/Wasserstaende-node.html
33	https://www.elwis.de/DE/Service/Telematikanwendungen/Inland-ECDIS/Allgemeines/ARGO/Fahrrinnensituation/Fahrrinnensituation-node.html?__forceHttps=1
34	https://www.destatis.de/DE/Methoden/Qualitaet/Qualitaetsberichte/Industrie-Verarbeitendes-Gewerbe/produktionsindex.pdf?__blob=publicationFile
35	https://www.destatis.de/DE/Methoden/Qualitaet/Qualitaetsberichte/Industrie-Verarbeitendes-Gewerbe/produktionsindex.pdf?__blob=publicationFile
36	https://www.landesdatenbank.nrw.de/ldbnrw/online?operation=find&suchanweisung_language=de&query=42153
37	https://euklems-intanprod-llee.luiss.it/download/
38	https://www-genesis.destatis.de/datenbank/online/table/42191-0001/table-toolbar
39	https://www-genesis.destatis.de/datenbank/online/table/46131-0004/table-toolbar
40	https://www-genesis.destatis.de/datenbank/online/table/51000-0011/table-toolbar
41	https://www-genesis.destatis.de/datenbank/online/statistic/61241/table/61241-0004
42	https://develop.unepgrid.ch/en/datasetproxy/22d523ec-d606-44bb-a0f7-52a2ab324db2

Footnote	Link
<i>Chapter II</i>	
11	https://www.dfb.de/maenner/wettbewerbe/dfb-pokal/wettbewerbsinformationen
12	https://www.sportschau.de/fussball/amateurfussball/regionalliga-deutschlands-unlogischste-fussball-spielklasse,regionalliga-analyse-100.html
<i>Chapter III</i>	
5	https://www.paleofire.org/home
6	https://sourcebooks.fordham.edu/ancient/herod-libya1.asp
7	https://www.laboasis.org/
7	https://ikg.ethz.ch/en/
15	https://data.worldbank.org/indicator/EG.ELC.ACCS.RU.ZS?locations=MA
29	https://correlatesofwar.org/data-sets/
34	https://chelsa-climate.org/
35	https://www.fao.org/soils-portal/data-hub/soil-maps-and-databases/harmonized-world-soil-database-v12/zh/
36	https://www.fao.org/soils-portal/data-hub/soil-maps-and-databases/harmonized-world-soil-database-v20/en/
37	https://www.nrcs.usda.gov/sites/default/files/2025-01/Field-Book-for-Describing-and-Sampling-Soils-v4.pdf
38	https://figshare.com/articles/dataset/Global_prevalence_of_non-perennial_rivers_and_streams/14633022
39	https://diegopuga.org/data/rugged/
40	https://gmd.copernicus.org/articles/17/275/2024/
41	https://www.hcp.ma/Recensement-general-RGPH_r518.html
42	https://cran.r-project.org/web/packages/xgboost/index.html
43	https://cran.r-project.org/web/packages/ks/index.html

Declaration

This dissertation is the result of my own work, and no other sources or means, except the ones listed, have been employed.

Mannheim, April 17, 2025

Yann Müller

Curriculum Vitae

- | | |
|-----------|---|
| 2021–2025 | University of Mannheim (Germany)
<i>Ph.D. in Economics</i> |
| 2018–2021 | Humboldt University of Berlin (Germany)
<i>M.Sc. in Economics</i> |
| 2019 | Universidad de la Habana (Cuba)
<i>Summer School</i> |
| 2013–2016 | Ulm University (Germany)
<i>B.Sc. in Management and Economics</i> |
| 2015–2016 | University of Toronto (Canada)
<i>Undergraduate Exchange Student</i> |
| 2013 | Gymnasium Gernsheim (Germany)
<i>Abitur</i> |