A Tale of Commodities and Climate-driven Disasters

Filippo Pellegrino*

Adam Smith Business School

Center for Sustainable Finance at Católica-Lisbon

Click here for the most recent version

This version: January 16, 2026

Abstract

This paper examines the impact of climate-driven disasters on commodity prices. Using extensive archival sources including census data and declassified CIA intelligence reports, I develop a global geospatial dataset to identify the locations of key commodity-producing sites at subnational level since the 1970s. By linking these regions to climate disaster events, I find that, over time, production has become increasingly concentrated in high-risk areas. Leveraging this dataset, I analyze how commodity futures respond to climate-driven shocks and uncover significant cross-sectional differences. Specifically, a long-only portfolio of vulnerable commodities yields a significant monthly α of 0.90%, whereas that of resilient commodities, albeit still significant, is negative at –0.43%, reflecting a premium paid for protection against climate shocks. Furthermore, I find that vulnerable commodities experience slower recoveries from past shocks.

Keywords: Commodities, Climate-driven disasters, Geospatial data.

JEL Classification: C58, G13, Q54.

*Email: mailto:filippo.pellegrino@glasgow.ac.uk. Website: www.filippopellegrino.eu. I am deeply grateful to Andrea Buraschi, Matteo Barigozzi, Lucrezia Reichlin, Domenico Giannone and Pasquale Della Corte for their invaluable guidance. I also wish to thank Enrico Biffis, Matteo Burato, Simone Cenci, Nicolo Ceneda, Veronica De Falco, Massimo Dello Preite, Lorenzo Garlappi, Stefano Giglio, Thomas Hasenzagl, Marcin Kacperczyk, Robert Kosowski, Xuelai Li, Ziang Li, Robert Litterman, Alexander Michaelides, Mirabelle Muuls, Chiara Perricone, Ailsa Roell, Magda Rola-Janicka, Lee Seltzer, Savitar Sundaresan, Fabio Trojani, Ali Uppal and Sumudu W. Watugala, as well as the participants of the 2025 SFS Cavalcade North America, 2025 ICEEE, 2025 IFABS conference, and of the seminar series of City St George's University of London, NYU Abu Dhabi, University College Dublin, University of Geneva, University of Glasgow and the University of Reading for their helpful comments on a earlier draft of this article. Furthermore, I am thankful to the organisers of the 2025 Liechtenstein Workshop of Sustainable Finance for honoring this work with the Best Paper Award. Finally, I would like to express my gratitude to the Leonardo Centre on Business for Society, Livio Scalvini and Maurizio Zollo.

1 Introduction

Over the past 50 years, the frequency and intensity of major climatic and meteorological disasters have more than doubled, posing significant risks to commodity production and distribution. Anecdotal evidence suggests that such disasters impact commodity production beyond the traditional agricultural sphere. For instance, Winter Storms Uri (2021), Elliott (2022), and, more recently, Heather (2024) caused substantial disruptions in natural gas production, leading to supply shortages and price spikes (Troy and Ober, 2024). However, the limited availability of data on the locations of key commodity-producing sites has hindered systematic evaluation of the links between climate-driven disasters and commodity prices. This is a fundamental issue, as commodities are integral to the global economy, and little is known about how investors price similar risks. Existing literature rarely adopts a financial perspective, instead focusing on specific disasters and commodities with an emphasis on scientific and engineering solutions.

Without data identifying these critical geographies, it is challenging to pinpoint disasters likely to disrupt output and, therefore, to study the subsequent impact on prices. From a demand-supply perspective, these events should create scarcity and increase spot prices. However, predicting futures price movements is more complex. The response of futures prices hinges on market participants' expectations regarding the duration, severity, and impact of each disaster. If the market anticipates that the supply shortage will be prolonged, futures prices may rise alongside spot prices. Conversely, if the disruption is expected to be temporary, futures prices might not increase significantly or, in the most extreme cases, they could even decline. Therefore, from a cross-sectional perspective, it is unclear which commodities are most resilient to climate-driven disasters. Moreover, it is unknown whether market participants fully comprehend the extent of these disasters immediately after they occur, as producers might require time to assess and report the

damages caused by major climate events. In this paper, I overcome the data limitations and address each of these research questions.

I start by developing a comprehensive geospatial dataset to identify the locations that have been key for the production of cocoa, coffee, corn, cotton, soybeans, sugar, wheat, crude oil, natural gas, aluminum, copper, nickel, gold, and silver. Due to the lack of data on the geographical distribution of production activities, I leverage a broad range of archival documents including published papers, governmental reports and commentaries, census records, ministerial reports, yearbooks and declassified reports from the Central Intelligence Agency (CIA). By means of manual annotation of thousands of these documents and related maps, I construct a geospatial dataset with subnational level resolution and global coverage since the 1970s. To the best of my knowledge, it is the first of its kind with this level of granularity and harmonized data for such a broad set of commodities.

Linking the resulting dataset with EM-DAT and GDIS (Sapir and Misson, 1992; Rosvold and Buhaug, 2021), well-known databases on disaster events, I uncover that the annual number of major climate-driven catastrophes affecting top commodity-producing regions increased from 286 (1970 to 1980) to 764 (2012 to 2021). These escalating figures have not been driven by a single event type but are rather broad based. Perhaps most importantly, I find that, over time, the production of most commodities has been concentrating in zones particularly at risk. This is especially true since the mid-2000s and extends far beyond traditional agricultural producers, affecting energy and mining.

Having established that climate-driven disasters impact a substantial number of producers and commodities, I investigate their effect on future prices. The vulnerability and resilience of commodities to climate shocks are influenced by two main components: first, whether production is located in regions that frequently experience disasters, and second, how strongly prices respond to these events. For the first component, I further leverage the data to determine whether production occurs in high- or low-risk locations. For the

second component, I use a regression design based on this geospatial dataset to assess how futures prices react to an increase in the number of disasters in the key production sites of their underlying commodities. While it is perhaps expected that typical agricultural products such as corn and wheat are among the most vulnerable commodities, I find that some hard commodities, such as natural gas, also exhibit high levels of vulnerability.

I study the cross-sectional differences between these two classes implementing the long-short strategy VmR (Vulnerable minus Resilient), the CAPM α remains significant and reaches a value of 1.10% per month. Besides, this α remains elevated, significant, and at a level of 1.08% per month even when considering the three factors of Bakshi et al. (2019) and the value factor in Asness et al. (2013). If I were to consider only the first component – the number of disasters – to distinguish vulnerable from resilient commodities, the significance of α would still hold, though its magnitude would be reduced by approximately half. This reduction is expected because the second component uses futures prices to infer which production chains are more exposed to risks, such as vulnerabilities in transportation and storage. Moreover, if I were to run this analysis using solely the first component and country-level data, as available to most agents, the α would not be statistically different from zero. This result further strengthen the value of the proposed subnational dataset.

Furthermore, I find that the futures prices at the nearest expiries tend to slowly recover from past shocks. In a more formal sense, when allowing for a delayed response of the futures prices to past disasters, the monthly α of a long-only portfolio on vulnerable commodities increases from 0.50% to 0.90%, while the one associated to resilient commodities remains rather stable at -0.43%. As a result, the α of the VmR strategy increases from 1.08% to 1.36% per month, and remains statistically significant. Perhaps more interestingly, the negativity of the α associated to resilient commodities suggests that they can be seen as hedging assets for protection against climate-driven shocks.

The paper is organized as follows. I start with a review of the literature (section 1.1). Next, I introduce the data (section 2), beginning with the proposed geospatial dataset on commodity production. I proceed with the disaster data from EM-DAT and GDIS, and conclude this section by describing the commodity futures and relevant technical details. For ease of readability, the historical research underlying the geospatial dataset is reported in the appendix (appendix A). In the results section (section 3), I present my findings. I first explore how climate-driven disasters have affected commodity production over time. I then focus on futures prices response to climate-driven disasters. Finally, I conclude (section 4).

1.1 Literature review

This article contributes to the literature on climate finance and commodities.

Empirical climate finance. This work closely relates to the empirical climate finance literature. In equities, most focus has been on regulatory climate risk (e.g., Bolton and Kacperczyk, 2021). Other asset classes, such as municipal bonds and real estate, have been studied with a stronger focus on physical risk, aligning more closely with the issues discussed in this article. Municipal bonds are tied to specific geographies, and when these areas are exposed to climate risks, they offer higher yields to attract investors (Painter, 2020). Similarly, real estate values are linked to their locations, facing risks from sea-level rise (Bernstein et al., 2019; Giglio et al., 2021b), hurricanes and floods (Ortega and Taşpınar, 2018; Eichholtz et al., 2019), wildfires (McCoy and Walsh, 2018), and impacts on mortgages (Ouazad and Kahn, 2019; Keenan and Bradt, 2020).

Climate risk and commodities. Research on commodities and climate risk is notoriously scarce, most likely due to serious data bottlenecks. This asset class is absent from recent reviews in climate finance (Giglio et al., 2021a). Most studies focus on specific

commodities, mostly agricultural and from a scientific perspective (Battisti and Naylor, 2009; Bonan and Doney, 2018; Chatzopoulos et al., 2020; Ding and Xu, 2023). The closest work to my article is perhaps Jia et al. (2023), as the authors use a broad range of commodities. However, they measure exposure to climate risk using an equity index of low-carbon stocks, thereby overlooking the impact of localized climate-driven disasters on key commodity-producing areas. Papers studying similar phenomena from a forecasting perspective, such as Nam (2021), suffer from the same data resolution issues.

Hedging portfolios. The negative α associated with commodities resilient to climate risk relates to the literature on hedging (Engle et al., 2020). This finding suggests that such commodities may serve as hedging instruments against physical risks. It is compatible with theoretical models featuring positive climate risk premia, where investors demand higher expected returns for bearing additional climate-related risks (see Giglio et al., 2021a, for a review). Investors might accept lower returns from resilient commodities because these assets provide protection following climate-driven disasters.

Gradual information diffusion. This paper is also linked to the equity literature on gradual information diffusion (Hong and Stein, 1999), as it reveals that futures prices tend to recover slowly from past shocks. This gradual recovery is likely due to the time producers need to assess and disclose the actual extent of damages caused by disasters.

Fundamentals. Finally, this paper also connects to the empirical literature on commodities related to production and inventories (Gorton and Rouwenhorst, 2006; Gorton et al., 2013; Rouwenhorst and Tang, 2012) and, albeit loosely, to the Theory of Storage (Kaldor, 1939; Working, 1949; Brennan, 1958). In particular, to the branch focused on inferring information on fundamentals from futures prices (e.g., Fama and French, 1987).

2 Data

2.1 Geographic distribution of commodity producers

To study the effect of climate-driven disasters on commodities, it is essential to have historical data on the subnational geographical distribution of producers. As mentioned in the introduction, such data is extremely scarce. It varies from country-level figures, whose aggregation rules do not allow to infer whether localized climate events pose a risk to main production sites, to relatively high-resolution data covering small subsamples of commodities and time periods. This is a clear bottleneck that must be overcome to link commodities and climate-driven disasters.

I construct a new dataset with the aim of resolving this fundamental problem. Although the readily available data is indeed problematic to use, there is a wide range of reliable archival sources. I gather historical data spanning from 1970 to 2021 on the geographical distribution of cocoa, coffee, corn, cotton, soybeans, sugar, wheat, crude oil, natural gas, aluminum, copper, nickel, gold, and silver. These commodities are actively traded in the derivatives markets of major exchanges such as CME, ICE, and LME. Furthermore, they constitute the core components of the Commodity Research Bureau (now part of the FTSE offering) and Bloomberg Commodity indices, which serve as the principal benchmarks for commodity performance across global markets. Specifically, for each year, I manually annotate the main production areas for the top three producers of each commodity of interest.[1] This process is based on data from published papers (both international and, when unavailable, local peer-reviewed journals), governmental sources (e.g., central bank commentaries, statistical office data, census data, ministerial reports, and yearbooks), and international institutions (e.g., IMF, World Bank, and FAO). For data related to the Soviet Union, I use declassified CIA reports, primarily from the Of-

[1] These correspond to the areas where the relevant futures deliverables are produced.

fice of Economic Research (OER), to ensure the accuracy and comprehensiveness of the dataset. I also employ CIA documents to understand the grain production distribution across China's provinces in the 1980s. For each year and commodity, I determine the top three producers based on comparable country-level production data. Namely, agricultural data is from FAOStat, energy data is from the Energy Institute, and metals data is from the British Geological Survey. These sources provide data spanning from 1970 to 2021, at an annual frequency. The top producers at each given month are determined based on the data available at the time, in a real-time fashion. These sources publish data with a significant delay of approximately one year with respect to the reference period of the production figures. Therefore, I consider the data for year t to become available in December of year $t + 1$, following their official publication calendar.

The manual annotation process, summarized in figure 1, involves reviewing thousands of pages and maps to identify the subnational areas within these top-producing countries where production has historically been concentrated. The subnational areas are first-order GADM 3.6 (the Database of Global Administrative Areas) administrative units, which denote areas the size of regions, states, or provinces. To the best of my knowledge, this is the first harmonized dataset of its kind. The dataset does not show significant changes in core production areas within countries, as such changes rarely occurred when segmenting countries at first-order administrative units. The only exceptions worth mentioning, and thus tracked by the dataset, are the administrative units of the USSR that became independent states after its dissolution in 1991. Consistent with the treatment of other regions, they were considered subnational administrative units prior to the Soviet collapse and treated as independent states afterward.

It is worth stressing that the majority of documents covering production dynamics in the 1970s and 1980s are in the form of scans and, therefore, stored as images. Consequently, the collection process requires studying these archival documents. Working with

these older documents, such as censuses and yearbooks, presents unique challenges. Many of these sources are not easily searchable, and the quality of the scans varies, requiring careful inspection to ensure accuracy. Despite these challenges, these historical records provide invaluable insights that modern data sources cannot match. In particular, they allow me to determine the historical geographical distribution of production centers, including within countries that no longer exist, such as the Soviet Union. Moreover, they permit me to study sectors that, over time, either disappeared or became marginal, such as coffee production in Angola prior to the civil war in 1975. In recent years, more options have become available to determine the geographical distribution of interest. For a subset of agricultural commodities, including grains and cotton, the USDA (United States Department of Agriculture) provides maps showing the share of production at the regional level. However, other soft commodities are mostly absent from USDA crop production maps. Similarly, the Global Energy Monitor provides free access to maps describing oil and gas exploration. In addition, Refinitiv Workspace includes data on the start and end dates of energy projects and mining operations, further refining the identification of key regions of interest. A detailed analysis on the geographical and historical developments of each commodity of interest is reported in appendix A.

2.2 Climatological and meteorological disasters

EM-DAT is the most established disaster-related dataset. It is maintained by the Centre for Research on the Epidemiology of Disasters (CRED), part of the University of Louvain. EM-DAT has a global scope and contains data on more than 23,000 major disasters, spanning from 1900 to the present.[2] However, it does not always provide precise information on the location of each disaster, which is particularly problematic for data prior to 1992,

[2]EM-DAT disaster records disasters that meet at least one of the following criteria: at least ten deaths (including dead and missing); at least 100 affected (people affected, injured, or homeless); a call for international assistance or an emergency declaration.

when the dataset was first published (Sapir and Misson, 1992).[3] More recently, a group of researchers released the Geocoded Disasters (GDIS) dataset, an updated version of EM-DAT that includes major disasters for which geocoded location information was recoverable (Rosvold and Buhaug, 2021). The data, distributed by NASA's Socioeconomic Data and Applications Center (SEDAC), covers the period from 1960 to 2018. Furthermore, it is already mapped with GADM and its highest spatial resolution corresponding to administrative level 3 (usually, districts or villages).

I leverage both datasets to link the core producing regions of the 14 commodities in section 2.1 with all climatological and meteorological natural disasters that occurred in these areas between 1970 and 2021. From 1970 to 2018 I employ the GDIS dataset. For the periods that are not yet cover by GDIS (2019 to 2021), I gather disaster data from EM-DAT. In the latter case, I focus on events with geocoded locations and map them to first-order administrative units. In cases where latitude and longitude data are not available, I rely on EM-DAT's information on the first-order administrative unit where the disaster occurred, whenever available.[4] In doing so, I discard EM-DAT disasters from 2019 to 2021 that cannot be mapped to a GADM administrative unit of order 1.

I refine the resulting dataset to include only sudden onset climate-driven disasters classified as extreme temperature, flood, glacial lake outburst flood, storm, or wildfire, with precise start and end dates.[5] This excludes disasters of similar nature, such as droughts, whose effects are all but sudden and usually prolonged over many months or years. As a result, I obtain a final dataset comprising 6,127 disasters. This covers most of the climate-driven sudden-onset disasters in EM-DAT, which amount to a total of

[3] EM-DAT was initially designed in 1988, featured in UNDRO News in 1991, and had its first reference paper published in a peer-reviewed journal in 1992.

[4] Since EM-DAT uses Global Administrative Unit Layers (GAUL), an outdated alternative to GADM, I convert the GAUL information as necessary. This is done by calculating the centroid coordinates of the GAUL administrative unit and mapping these to the corresponding GADM subregions.

[5] According to the official EM-DAT documentation, sudden onset disasters usually have complete start and end dates. Disasters with incomplete dates are generally not of this type.

7,615 between 1970 and 2021, including those that cannot be geocoded. Of these 6,127 disasters, 33.71% have occurred in one or more of the core commodity-producing areas mentioned in section 2.1.

The data is sampled at monthly frequency to monitor active disasters and construct commodity porfolios as a function of disaster metrics discussed in the following sections.

2.3 Commodity futures returns

I collect daily settlement prices and expiries for the nearest and second nearest futures contracts for each of the 14 target commodities from Bloomberg. This data is used to calculate fully-collateralized commodity futures returns, based on the assumption that investors hold the nearest contract until 35 days before the last trade date (or if the generic curve moves down by one contract before then), at which point they roll over to the second nearest contract. This mechanism, in line with Bakshi et al. (2019) and the official contract specifications of the commodity futures of interest, allows to rollover prior to the first notice date (i.e., the first point in time when investors may be required to physically deliver the underlying commodity). Additionally, the data is used for computing a range of metrics described directly in the main corpus of the results. These figures are derived from the point-in-time futures price curve. When needed, the curve is geometrically interpolated between the first and second nearest expiries. All futures-based data is then aggregated at monthly frequency by taking end-of-period returns to align it with climate-driven disaster data.

Data on cocoa, coffee, cotton, and sugar are sourced from the Intercontinental Exchange (ICE); corn, soybeans, and wheat from the Chicago Board of Trade (CBOT); crude oil and natural gas from the New York Mercantile Exchange (NYMEX); aluminum and nickel from the London Metal Exchange (LME); and copper, gold, and silver from the Commodity Exchange (COMEX). The Chicago Board of Trade (CBOT), New York

Mercantile Exchange (NYMEX), and Commodity Exchange (COMEX) are part of the Chicago Mercantile Exchange (CME). The sample period spans from January 1970 to December 2021, with at least six commodities available at any given time, nine available by 1975, twelve by 1990, and all fourteen included by 1997.

The relevant excess returns are computed relative to the US risk-free rate, as provided by Prof. Kenneth R. French's website. This rate corresponds to the one-month Treasury bill, originally sourced from Ibbotson Associates.

3 Results

3.1 Prevalence of climate-driven disasters

It is challenging to determine, a priori, which commodities are most vulnerable to climate-driven disasters, as they are produced globally in vastly different regions. Specific types of disasters are tied to localized weather patterns and environmental factors. For example, some regions may be more prone to floods or frosts, increasing the threat to local crops or leading to the shutdown of critical facilities like smelters. Coastal areas, meanwhile, are more likely to suffer from hurricanes or rising sea levels, which can disrupt shipping routes and supply chains. Additionally, the resilience of the infrastructure supporting commodity production and transportation plays a key role in determining vulnerability. Commodities backed by robust infrastructure, such as mining operations in stable environments, may be less vulnerable than those dependent on more sensitive, weather-dependent systems like agriculture. However, when severely disrupted, recovery may take longer due to the complexity of rebuilding. Besides, even if the infrastructure is resilient to the effects of climate-driven disasters, workers may still be severely impacted and this might lead to closures. Furthermore, these risks are not static; they evolve due to shifting weather patterns and the increasing frequency of extreme events can heighten

vulnerabilities.

As a result, to start addressing the research questions of this manuscript, it is crucial to study the areas where commodities are being produced. Table 1 provides a detailed overview of the prevalence of major climate-driven disasters. Besides, it offers insight into their evolution over time and across different geographical regions. Column 1 lists the descriptors for key statistics, column 2 shows data for the pre EM-DAT period (1970 to 1991, included), and columns 3 through 5 divide the subsequent sample into three equal-sized subperiods of 10 years. Finally, column 6 summarizes the results for the entire sample period. Panels A through E narrow the focus onto specific geographies.

Over the past 50 years, the data reveals an increase of roughly 113.97% in the number of major climate-driven disasters globally (Panel A). Additionally, over the full sample period, 78.95% of the world's countries was affected by these events. The rate of growth was more modest from the mid-1990s, as the number of disasters increased of about 23.58%. Similarly, the proportion of countries affected by similar disasters increased by 7.46 percentage points in the same window.

Shifting the focus to key commodity-producing countries, the situation appears far more severe (Panel B). In these countries, the prevalence of disasters surged by an astounding 167.13% since the 1970s. More strikingly, even between the mid-1990s and 2021, there was an increase of 71.30%. The number of leading climate-producing countries affected by these disasters is also significant. From 1970 to 2021, 84.38% of these major commodity-producing countries were impacted by climate-driven disasters. These figures suggest that while the global rate of disasters may be slowing, geographies critical to commodity production continue to face escalating risks.

The data further indicates that most disasters have historically affected areas dedicated to agricultural production (Panel C). However, it stands out that there has been a staggering rise in climate-driven disasters affecting regions engaged in energy production

and mining operations (Panels D and E). Specifically, disaster occurrences in energy-producing areas increased by 80.95%, while areas involved in metals production experienced a 304.16% rise overall. Since the mid-1990s alone, disasters increased by 31.03% in energy-related areas and 135.15% in metal-producing regions. This trend underscores a critical shift: while climate-driven disasters once predominantly affected agricultural regions, they are now becoming a growing concern for sectors outside this traditional sphere. At the current rate of increase, climate-driven disasters are expected to affect energy and mining operations combined more than agricultural areas within the next few decades.

Table 2 provides a more granular look at major climate-driven disasters affecting leading commodity producers, breaking down these events by type and reporting their temporal dynamics using the same column format as table 1. This table uncovers that the most prevalent disasters in these countries are meteorological in nature, with floods and storms comprising more than 90% of the total disasters. The third most common category of climate-driven disasters includes extreme temperature events, such as heat waves and cold spells. Perhaps more concerning is the observation that the rising prevalence of climate-driven disasters in energy- and metal-producing regions is broad-based (Panels B and C). All types of disasters have shown heightened growth. The frequency of floods surged by 68.18% and 442.11% in energy- and metal-producing countries, respectively, since the 1970s. Even in the more recent period, from the mid-1990s onward, flood occurrences increased by 42.31% in energy-producing regions and by 249.15% in metal-producing regions. Similarly, the incidence of storms and extreme temperature events has also risen dramatically.

This broad-based increase across different disaster types suggests that the risks faced by key commodity-producing regions are escalating. However, it is unclear from tables 1 to 2, whether this trend is driven by specific production areas impacting a subset of

commodities, or if it reflects a more widespread phenomenon. Figure 2 explores this issue further, showing the global occurrence of climate-driven disasters by subregion. In a structure similar to the previous tables, it breaks down these occurrences across the same post EM-DAT subperiods of interest.

The first takeaway from this figure is that countries more exposed to climate-driven disasters are the United States, India and China. Secondary hotsposts are located in South-East Asia and Australia. These regions share several climatic and meteorological characteristics that contribute to their heightened vulnerability to climate-driven disasters. Many of these areas are situated within similar climate zones, such as tropical and subtropical regions, which are prone to extreme weather events like hurricanes, typhoons, and monsoons. For instance, India and Southeast Asia experience intense monsoon seasons that can lead to severe flooding, while China and the United States are frequently affected by typhoons and hurricanes, respectively. Additionally, these regions often face further commonalities, such as high humidity and temperature extremes, which can exacerbate the impact of climate events.

Second, there has been a widespread increase in the occurrence of climate-driven disasters globally since the mid-1990s. This trend is particularly evident in the previously mentioned hotspots. Indeed, in the United States, the average number of climate-driven disasters by subregion increased from 17.27 (1992 to 2001) to 23.0 (2012 to 2021). Besides, in India and China the increase was more significant and went from 4.75 and 8.06 (1992 to 2001) to 8.86 and 25.19 (2012 to 2021), respectively.

These findings align well with tables 1 to 2. Since the 1970s, traditional producers of many key commodities have been increasingly impacted by climate-driven disasters. In the United States alone, for instance, a broad set of production areas for corn, cotton, soybeans, crude oil and natural gas have been located within some of the most at-risk

areas.[6] China has also been leading production, particularly for corn, cotton and wheat with production operations predominantly in the Yangtze River Valley, Yellow River Valley, and northwest inland areas. Similarly, India's primary sugar-producing regions in the Northeast and Southwest have been among the country's most disaster-prone areas. Moreover, coffee-producing regions in the northwestern departments of Colombia and the southeastern areas of Brazil have also been facing exposure to climate-related risks, though to a somewhat lesser degree.

The heightened occurrence of climate-driven disasters highlighted in tables 1 to 2 is also linked to the emergence of new leading producers in parts of the globe with adverse climatic and meteorological conditions. This is the case, for instance, of India for cotton, wheat, and aluminum. Other examples include China for sugar, aluminum, copper, gold and silver; Indonesia for cocoa, coffee, copper, and nickel; and the Philippines for nickel. As a result, commodities ranging from agricultural products to energy and metals have become more exposed to climate-driven disasters. Figure 3 summarizes these historical shifts in commodity production, thus further stressing how difficult it is to avoid climate risk when investing in commodities.

3.2 The cross-sectional effect of disasters

I then proceed by studying the cross-sectional differences in futures prices between commodities whose main production areas are hit by climate-driven disasters (H) versus those that are spared (S).

To distinguish between these groups, I leverage data from section 2.1, which provides details on the core production subregions within the top three nations for each commodity and point in time in my sample. Specifically, for each month t and commodity i, I quantify the disaster impact on each leading nation c using the disaster count variable $I_{c,i,t}$. This

[6] Key production centers for wheat, aluminum, gold, and silver have been situated in vulnerable parts of the US, although their global significance has diminished over time.

country-level measure is then aggregated to the commodity level by computing a weighted average, where the weights $w_{c,t}$ are based on each nation's market share.[7] Formally, the commodity-level disasters impact

$$D_{i,t} := \sum_{c \in \mathcal{C}_{i,t}} w_{c,t} \, I_{c,i,t}, \tag{1}$$

where $\mathcal{C}_{i,t}$ represents the set of the top three nations producing item i at time t. It is worth noting that the selection of the top producing countries is conducted in a real-time fashion, leveraging the actual publication delays of the data providers, as described in section 2.1.

I then construct a long-short portfolio designed to capitalize on the potential price appreciation of commodities experiencing climate-driven supply shocks while shorting those that are less affected. Namely, this portfolio, denoted as HmS (Hit minus Spared), involves taking a long position on commodities whose producing areas were most affected by climate disasters (top $D_{i,t}$ tercile) and a short position on commodities that were less impacted (bottom $D_{i,t}$ tercile). The analysis employs monthly rebalancing and holding periods, with commodities equally weighted within each basket. The portfolios are carefully constructed to avoid any type of look-ahead bias.

Table 3 reports the annualized Sharpe and Sortino ratios for the HmS strategy, alongside long-only portfolios on hit and spared commodities, and a series of benchmarks.[8] The table is structured as follows. Column 1 includes the portfolio descriptors, Columns 2 and 3 distinguish between the latest two EM-DAT subsamples discussed in section 3.1, whereas Column 4 focuses on the period ranging from 2002 to 2021.[9] The benchmarks are an equal-weighted portfolio of all commodities (serving as a market proxy), as well as

[7] The latter is derived from annual production figures available at time t described in section 2.1.

[8] The Sortino ratio better accounts for highly inflationary periods, such as the one that started post COVID-19, where commodity prices jump and rally over a sustained amount of time.

[9] The first EM-DAT subsample is not reported, as it is used as a presample in analyses reported below.

carry, momentum, and value portfolios. Carry is defined as a portfolio long on the most backwarded commodities and short on the most contangoed (Bakshi et al., 2019). Momentum builds on a tradition originated in equity (Jegadeesh and Titman, 1993; Asness, 1994) and it is defined as portfolio long on winners, commodities with the highest cumulated returns in the recent past, and short on losers, the opposite spectrum of assets. The cumulated returns are computed over a lookback period of 12 months, excluding the most recent (Jegadeesh, 1990; Lo and MacKinlay, 1990; Boudoukh et al., 1994; Asness, 1994; Grinblatt and Moskowitz, 2004; Asness et al., 2013; Buraschi and Pellegrino, 2024).[10] Value is portfolio long on the most undervalued commodities and short on the most overvalued. The relevant sorting variable builds on Asness et al. (2013) who proposed to calculate it as the logarithm of the average spot price from 4.5 to 5.5 years ago divided by the most recent spot price.[11] The long-short benchmarks are constructed using terciles to ensure internal consistency with the HmS strategy.

The long-only portfolio on hit commodities performed well throughout the sample period and outperformed all relevant benchmarks. Indeed, although the Sharpe ratio (Panel A) decreased from 0.75 to 0.29 over 20 years, the market went from 0.5 to -0.08 in the same time period. The long-only portfolio on spared commodities went even more negative to -0.24, starting from 0.47. Similarly, according to the Sortino ratio (Panel B), the long hit portfolio outperformed all comparables across every subsample. Shifting the focus now on the long-short portfolios, it emerges that, albeit the HmS strategy was outperformed by momentum and value in specific subsamples, its performance has been quite reliable throughout. The benchmarks, instead, were quite unstable with performance metrics jumping from positive to negative territory wildly. As a result, the performance of the HmS was higher than all benchmarks (and the market) over the

[10] Bakshi et al. (2019) computes the total returns over a lookback period of 6 months. However, results within this sample do not vary much when choosing this parametrization for the momentum signal.

[11] The only difference between my calculations and Asness et al. (2013) relates to the choice of not taking logarithms of the price, due to the impact of COVID-19 pushing oil prices into negative territory.

longer sample spanning from 2002 to 2021.

Table 4 reports performance attribution analyses for the HmS portfolio, as well as its individual components, over the sample period from 2002 to 2021. Several empirical models have been proposed to explain the cross-section of commodity returns. Each panel differentiates between empirical asset pricing models for commodities. Column 1 presents the relevant estimated coefficients. Columns 2–3 report results for long-only portfolios on hit and spared commodities, while Column 4 focuses on the HmS portfolio.

The CAPM α (Panel A) for the long-only portfolios on hit and spared commodities are 0.54% and -0.05% per month, with only the former being significant at the 1% confidence level. Both portfolios have significant β_{mkt} of 1.05 and 0.96, respectively, both significant at the 1% confidence level. The HmS strategy, instead, has a CAPM α of 0.59% per month significant at the 10% confidence level and mostly driven by the long tail. The HmS is market neutral, and its returns are not well explained by the CAPM, as evident from the adjusted R^2 of 0.10%.

Bakshi et al. (2019) proposed to complement the CAPM with carry and momentum in the case of commodity futures. I implement this approach in incremental steps, first reporting results for a model with excess market returns and carry (Panel B), and then with both carry and momentum (Panel C). Results show an even higher HmS α of 0.65% (Panel B) and 0.67% (Panel C) per month, significant at the 5% confidence level. Again, this is mostly driven by hit commodities, as they exhibit α of 0.57% (Panel B) and 0.58% (Panel C) per month, both significant at the 1% confidence level.

Finally, I consider an extension of the factor model of Bakshi et al. (2019), with the addition of a fourth factor proxying value (Panel D). Also in this case, results are largely unchanged, as value does not significantly explain the return profile of either the long-only portfolios on hit and spared commodities or the HmS strategy. Thus, the return profile of the HmS strategy is still poorly explained by the model, as evident from its

adjusted R^2 of 4.12%.

Table 5 provides further context and shows what the α would have been when disregarding the subnational level granularity of the dataset in section 2.1. In other words, using the data available to most agents. Solely for the purpose of this analysis, I construct the sorting variable by considering all climate-driven disaster occurrences within top-producing countries. This process counts disasters both in locations with and without key producing sites, which is intuitively problematic for producers with large land areas such as China, India, and the U.S. In doing so, I find that the data in section 2.1 provides a substantial informational advantage. Indeed, all the α of HmS computed with low-resolution data are not statistically different from zero, according to the CAPM regression (Panel A) and its extensions with carry (Panel B), carry and momentum (Panel C), as well as carry, momentum, and value (Panel D).

3.3 Implied information on vulnerability and resilience

Although these results are encouraging, there are still a few points to address in order to fully answer the research questions of this article. First, a commodity whose core producing areas suffer more climate-driven disasters is not inherently more vulnerable than one that is less impacted, as the resilience of its production processes might be inherently stronger. Thus, it remains unclear whether the futures prices post-shock will be higher or lower than before. Ultimately, it is an empirical question that primarily hinges on market participants' expectations regarding the long-term effects of the climate-driven disaster. Second, resilient commodities might recover faster from past shocks as their supply becomes less scarce. Third, the extent of disruptions caused by climate-driven events in month t may be fully understood at the beginning of month $t+1$. Indeed, it might take companies some time to assess the post-shock status of their production and storage facilities. Consequently, futures markets may not fully price these disasters until

month $t+1$.

It follows that one would want to account for some measure of production vulnerability, market expectations on disasters and their temporal dynamics to understand the magnitude of such events on commodity futures. While regional data on commodity production and inventories is scarce and hardly harmonized, the futures prices provide an indirect view on the market reaction to climate-driven disasters.[12] To obtain this information, I propose to estimate the regression model

$$J_{i,\,t,\,s} = \alpha_i + controls + \gamma_i \sum_{k=0}^{p} \Delta D_{i,\,t-k} + \varepsilon_{i,\,t,\,s}, \tag{2}$$

where $J_{i,\,t,\,s} = \frac{F_{i,\,t,\,s} - F_{i,\,t-1,\,s}}{|F_{i,\,t-1,\,s}|}$ is the price jump of the futures price for the i-th commodity between time t and $t-1$ with expiry of s months, $D_{i,\,t}$ represents the number of disasters at time t in the regions involved in producing the i-th commodity, p is an integer number of lags and $\varepsilon_{i,\,t,\,s}$ is the error term.[13] The γ_i coefficients measure the effect of an increase in the number of climate-driven disasters onto the futures prices for the i-th commodity at the s expiry. Having fixed an expiry of interest s, it allows to determine which commodities are more likely to have positive shifts at that expiry, following present and past climate-driven disasters. If $p = 1$, for instance, this regression model allows to account for disasters concluded in $t-1$. Given that contracts are held until 35 days before the last trade date and then rolled over to the second nearest expiry, $s = 1$ is the most important expiry for any investing strategy in this manuscript.

Table 6 reports the performance of an improved version of HmS, called Vulnerable minus Resilient (or, VmR). This portfolio is constructed by double sorting over the

[12] Indirect inference based on commodity futures has a long tradition (e.g., Fama and French, 1987). Country-level inventory data can be constructed for a subset of nations following Gorton et al. (2013). Production data can also be reconstructed at country level. However, they are both largely unavailable at subnational level.

[13] In order to avoid problems linked with the time-varying composition of top producers, the change in the number of disasters is computed similarly to equation 1.

number of disasters, as in the case of HmS, and the estimated futures sensitivity to changes in disasters, $\widehat{\gamma}_i$. Following Fama and French (1992), I estimate the model using a rolling five-year window. At each point in time, I standardize the regression coefficients to eliminate the effects of differing units of measurement in futures prices and disasters across commodities. The structure of table 6 is similar to table 3, with the only addition of a column distinguishing between p zero and one, and a few extra panels for relevant controls.[14] Namely, I control, in turn, for the market jump for contracts on all n commodities $J_{mkt,\,t,\,s} := \frac{1}{n}\sum_{i=1}^{n} J_{i,\,t,\,s}$ (Panels A1 and B1), the excess market returns (Panels A2 and B2) and the Baltic Dry Index in month-on-month percentage terms (Panels A3 and B3). The BDI is a shipping-freight cost index often employed as a proxy for global activity (Kilian, 2009).

When controlling for $J_{mkt,\,t,\,s}$, the VmR strategy outperforms HmS both in terms of Sharpe and Sortino ratios (Panels A1 and B1). This is evident throughout the sample and particularly accentuated in the first half. Indeed, the VmR strategy returned an annualized Sharpe of 0.40 ($p = 0$) and 0.82 ($p = 1$) between 2002 and 2011. In that subsample, the HmS had a Sharpe of 0.32. The outperformance is much stronger when measured in terms of Sortino ratios, with the HmS having about half the risk-adjusted performance of the VmR strategy. Narrowing the focus on the long-only portfolios at the base of the VmR, the one on vulnerable commodities outperformed the market in all subsamples and, overall, by a larger margin than the one long on hit commodities. Similarly, the long R portfolio underperformed long S, thus reinforcing that the higher risk-adjusted returns of the VmR are reasonably broad-based. Besides, it implies higher cross-sectional dispersion between vulnerable and resilient commodities, than the one observed between hit and spared. Considering delayed effects of climate-driven disasters

[14] Although one could consider more lags, estimating the model on a standard rolling window of 60 months might lead to high collinearity. Besides, unreported results show that results are largely unchanged when increasing p from 1 to 2.

improves performance. The VmR has indeed a Sharpe of 0.57 and 0.65 over the sample spanning from 2002 to 2021, with $p = 0$ and $p = 1$ respectively. Similarly, the Sortino ratios are 0.93 and 1.13.

Results are robust to changes in the control variables. Indeed, using the excess market returns instead of $J_{mkt,t,s}$ gives a very similar output, with Sharpe now of 0.52 and 0.64 for the VmR with $p = 0$ and $p = 1$ respectively. Results are marginally worse when controlling for the Baltic Dry Index, at 0.43 and 0.62 for equivalent VmR parametrizations. Overall, the big picture does not change much and, thus, going forward I will focus on the VmR versions where I control for $J_{mkt,t,s}$, as this is more in line with the target variable of the regression model in equation 2.

Table 7 reports the performance attribution results for the VmR strategy with $p = 0$ and its components with a structure similar to table 4. The CAPM α (Panel A) are 0.48% and -0.57% per month for V and R, respectively, with the latter being significant at the 5% confidence level. Additionally, their market β are equal to 1.01 and 1.04, and statistically different from zero at the 1% confidence level. Both V and R exhibit relatively low adjusted R^2 values at 50.49% and 62.39%, indicating that excess market returns only partially explain their dynamics. The VmR strategy has an outstanding CAPM α of 1.10% per month, significant at the 5% confidence level. Besides, it is not well explained by the CAPM, with a market β that is not statistically different from zero and an adjusted R^2 of -0.38%. This α is almost twice as large as the one for HmS.

When considering the market excess return and carry (Panel B), the α for V and R are 0.50% and -0.55% per month, both statistically different from zero at the 10% and 5% significance levels, respectively. The β_{mkt} are 1.01 and 1.05, both significant at the 1% level, and the adjusted R^2 are 50.47% and 62.54%. The coefficients associated with carry are not statistically significant. In other words, results from the long-only portfolios are largely unchanged from the CAPM regressions. Similarly, the VmR portfolio maintains a

large and significant α of 1.10% per month. Neither the excess market returns nor carry have associated coefficients that differ from zero in a statistical sense. Therefore, the performance is largely unchanged, as also evident from the adjusted R^2 value of -0.82%. Moving on to the full three-factor model of Bakshi et al. (2019) does not change the overall picture (Panel C). Overall, coefficients remain relatively stable, and neither carry nor momentum help explain the long-only or the long-short returns.

Although value does not have significant associated coefficients for either V, R, or VmR, its joint effect with the other factors changes the estimated α in magnitude, albeit marginally, for R from -0.56% (Panel C) to -0.54% per month (Panel D). As a result, the α of VmR moves from 1.10% to 1.08% per month, but remains significant at the 5% confidence level. From a big-picture perspective, however, results are similar to the simpler CAPM regressions.

Table 8 replicates, in turn, the performance attribution in the case of $p = 1$. Figures contained in the table show two stylized facts. First, the α associated with V increases by about 0.35 percentage points per month, while the one for R decreases (in magnitude) marginally when accounting for delayed effects. This leads to an increase in the VmR α, ranging from 0.1 (Panel A) to 0.26 (Panel D) percentage points compared to the case with $p = 0$. In other words, futures prices of vulnerable commodities tend to slowly recover from past shocks. Second, a long-only portfolio of vulnerable commodities yields a significant monthly α of 0.90%, whereas that of resilient commodities, albeit still significant, is negative at –0.43%, reflecting a premium paid for protection against climate shocks (Panel D).

Tables 9 to 10 further add context by reporting the number of times each commodity was sorted under V and R, as well as their differential, over the full sample ranging between 2002 and 2021 while distinguishing between $p = 0$ and $p = 1$. Interestingly, regardless on the number of lags, there are more than just agricultural commodities in

V. Natural gas, gold, corn and wheat are considered among the most vulnerable. At the opposite end, there are both agricultural commodities such as coffee and cocoa, as well as metals including nickel and copper. Heterogeneity across commodities of the same type and produced in relatively similar geographical areas is also captured by this approach. Indeed, I find, for instance, that natural gas is more vulnerable than crude oil to climate-driven disasters. In this instance, it depends on several factors affecting their vulnerabilities to climate disasters, including that natural gas transportation is more reliant to pipelines, which are susceptible to damage from extreme weather events like hurricanes and floods. Another factor relevant to this example is that natural gas requires more complex storage solutions, such as high-pressure tanks or liquefaction, making it more sensitive to infrastructure disruptions compared to crude oil. This information is embedded in futures prices and extracted via the regression approach described in equation 2.

4 Concluding remarks

This study provides a comprehensive analysis of the impact of climate-driven disasters on commodity markets, addressing pivotal questions that have remained largely unexplored. First, I resolve a fundamental bottleneck in the literature by constructing a novel global geospatial dataset identifying the locations of the main commodity-producing sites since the 1970s. Building on it, I document a significant and sustained increase in the occurrence of climate-driven disasters affecting core commodity-producing areas. This uptick is not confined to traditional agricultural commodities but extends to energy production and mining, highlighting the broadening scope of climate risk.

Second, further leveraging this dataset, I investigate how these disasters impact the futures prices of affected commodities, especially at the nearest expiries. Findings from a

novel regression design proposed in this article reveal that commodities more vulnerable to climate-driven disasters appreciate more than the most resilient following these events.

The performance of the VmR (Vulnerable minus Resilient) portfolio substantiates these observations. The portfolio deliver superior risk-adjusted return. Most importantly, these returns have a CAPM α of 1.10% per month. An additional insight from this analysis is that allowing for delayed adjustment of futures prices increases the CAPM α to 1.20%. Besides, the α of a factor model controlling for excess market returns, carry, momentum and value reaches 1.36% per month. Interestingly, the α of vulnerable commodities is positive and significant, whereas the one for resilient assets, albeit still significant, is negative. This suggests a premium for protection against climate shocks.

While this study advances the understanding of the direct effects of climate-driven disasters on commodity futures, it opens up important avenues for future research on the macroeconomic implications of these supply shocks, particularly through the lens of inflation.

Figures and tables

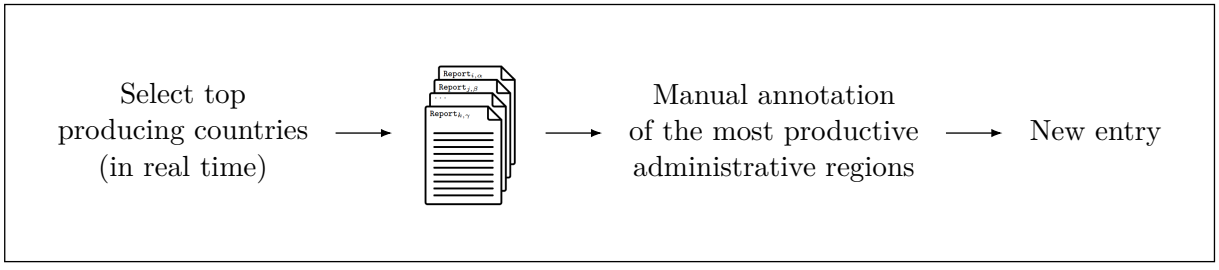


Figure 1: Data processing.

Stylized data processing steps to generate a new entry for every year and commodity.

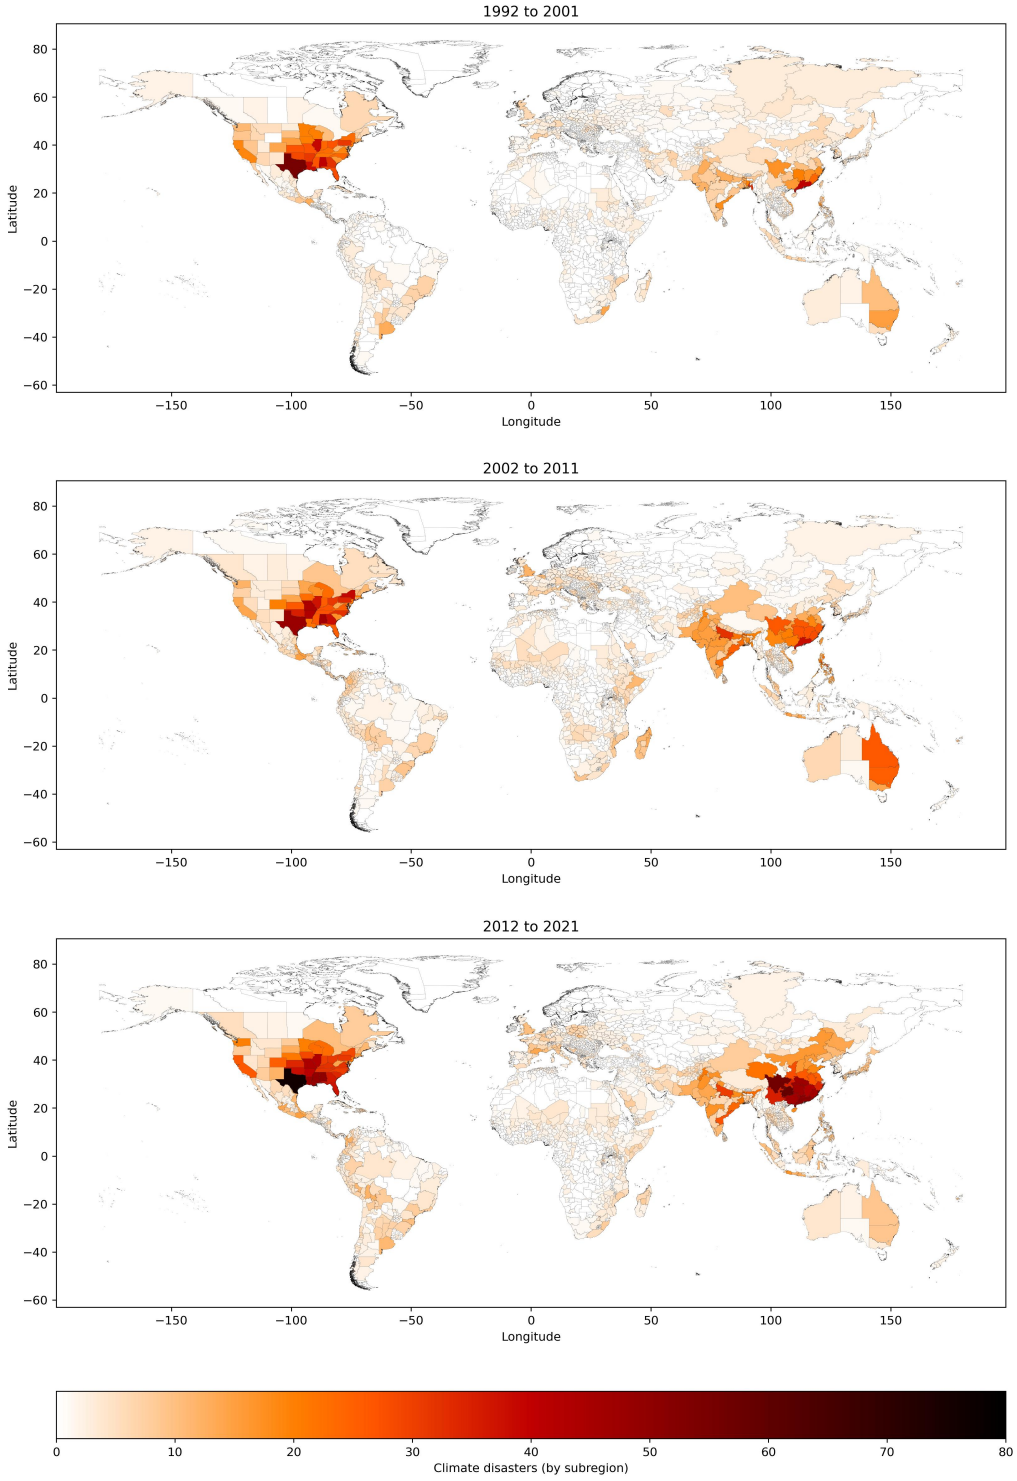


Figure 2: Geographical distribution of climate-driven disasters (over time).

Global occurrence of climate-driven disasters by subregion. The subnational boundaries correspond to first-order administrative units in GADM 3.6.

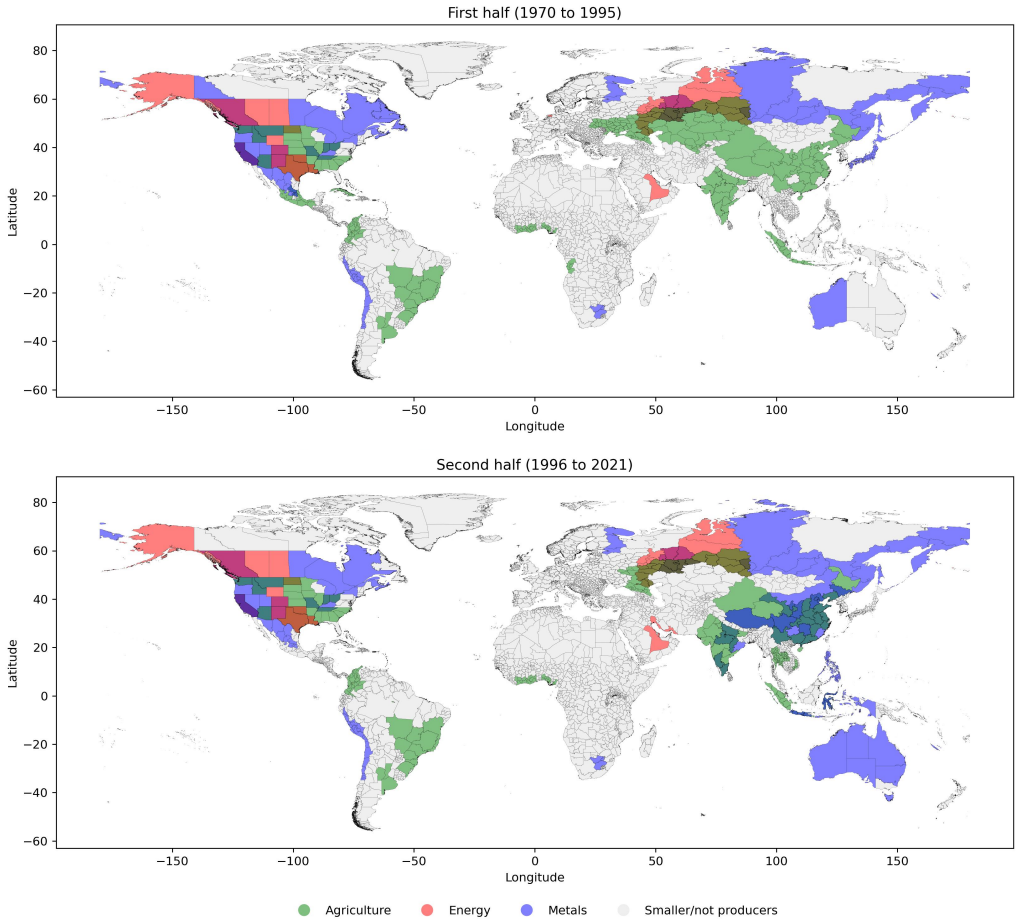


Figure 3: Core producing areas of agricultural commodities, energy resources and metals.

This figure displays the core commodity producing areas of all countries that, in any given year between 1970 and 2021, ranked among the top three global producers for one of the 14 considered in this manuscript. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6. Commodity-level maps are reported in appendix A.

	Pre EM-DAT	1992 to 2001	2002 to 2011	2012 to 2021	Full sample
Panel A: World's view					
Number of major disasters	823	1425	1995	1761	5998
Countries affected by major disasters (%)	46.93	61.84	69.74	69.30	78.95
Panel B: Top commodity-producing countries					
Number of major disasters	286	446	572	764	2066
Countries affected by major disasters (%)	50.00	56.25	62.50	62.50	84.38
Panel C: Top agricultural producers					
Number of major disasters	241	383	491	604	1717
Countries affected by major disasters (%)	50.00	55.00	65.00	60.00	90.00
Panel D: Top energy producers					
Number of major disasters	84	116	110	152	461
Countries affected by major disasters (%)	42.86	42.86	42.86	71.43	71.43
Panel E: Top metal producers					
Number of major disasters	96	165	248	388	897
Countries affected by major disasters (%)	64.29	71.43	78.57	78.57	92.86

Table 1: Occurrence of major climate-driven disasters.

Descriptive statistics on the occurrence of climate-driven disasters. Panel A includes global statistics on the number of major disasters and the proportion of countries affected by them across a range of subsamples indicated in columns 2 to 6. Panel B narrows the focus to major commodity-producing countries. Panels C to E further disaggregate the analysis from Panel B, breaking down the data into key geographies for the production of agricultural and energy commodities, and metals. The pre-EM-DAT period spans from 1970 to 1991 (included), covering early disaster reporting prior to the formalization of the global EM-DAT database. Hence, the full sample ranges from 1970 to 2021.

	Pre EM-DAT	1992 to 2001	2002 to 2011	2012 to 2021	Full sample
Panel A: Top agricultural producers					
Extreme temperature	8	9	14	17	48
Flood	90	162	277	295	822
Glacial lake outburst flood	0	0	0	0	0
Storm	143	212	200	282	837
Wildfire	0	0	0	10	10
Panel B: Top energy producers					
Extreme temperature	1	3	5	3	12
Flood	22	26	39	37	123
Glacial lake outburst flood	0	0	0	0	0
Storm	61	87	66	104	318
Wildfire	0	0	0	8	8
Panel C: Top metal producers					
Extreme temperature	5	7	10	9	31
Flood	38	59	133	206	436
Glacial lake outburst flood	0	0	0	0	0
Storm	53	99	105	169	426
Wildfire	0	0	0	4	4

Table 2: Occurrence of climate-driven disasters (by type) in commodity-producing countries.

Descriptive statistics on the occurrence of climate-driven disasters (by type) in commodity-producing countries. Panels A to C break down the data into key geographies for the production of agricultural and energy commodities, and metals, respectively. The pre-EM-DAT period spans from 1970 to 1991 (included), covering early disaster reporting prior to the formalization of the global EM-DAT database. Hence, the full sample ranges from 1970 to 2021.

	2002 to 2011	2012 to 2021	2002 to 2021
	Panel A: Annualized Sharpe ratios		
Market	0.50	-0.08	0.25
Carry	0.83	-0.39	0.24
Momentum	0.48	-0.31	0.14
Value	-0.32	0.94	0.17
H	0.75	0.29	0.54
S	0.47	-0.24	0.18
HmS	0.32	0.53	0.41
	Panel B: Annualized Sortino ratios		
Market	0.64	-0.13	0.36
Carry	1.78	-0.61	0.44
Momentum	0.92	-0.46	0.24
Value	-0.46	1.87	0.23
H	1.21	0.56	0.95
S	0.63	-0.39	0.26
HmS	0.52	0.91	0.66

Table 3: Performance metrics of the hit minus spared (HmS) strategy and benchmark portfolios.

This table reports the hit minus spared long-short portfolio, which involves taking a long position on commodities from areas most affected by climate disasters (hit) and a short position on commodities from less impacted areas (spared). Long-only portfolios for hit and spared commodities are also reported, alongside a long-only market portfolio that equally weights all commodities considered in this study. The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS).

	Long H	Long S	HmS
	Panel A: CAPM		
α	0.54**	-0.05	0.59*
	(0.01)	(0.76)	(0.08)
β_{mkt}	1.05***	0.96***	0.08
	(0.00)	(0.00)	(0.25)
Adjusted R^2	67.59	73.20	0.10
	Panel B: CAPM plus carry		
α	0.57***	-0.08	0.65**
	(0.01)	(0.62)	(0.05)
β_{mkt}	1.05***	0.96***	0.09
	(0.00)	(0.00)	(0.17)
β_{carry}	-0.10**	0.09**	-0.19***
	(0.01)	(0.04)	(0.01)
Adjusted R^2	68.19	73.92	3.15

(*Continued*)

Table 4: Performance attribution for HmS.

This table evaluates the performance of the HmS strategy and the long-only strategies on H and S commodities. Carry is a long-short portfolio going long on the top tercile of backwardated commodities and short on the bottom tercile. Momentum is a long-short portfolio that goes long on the commodities in the top tercile of cumulative returns over the last 12 months (excluding the most recent month) and short on those in the bottom tercile. Value is a long-short portfolio going long on the top tercile of undervalued commodities and short on the bottom tercile. Similarly to Asness et al. (2013), the value signal is defined as the average spot price from 4.5 to 5.5 years ago, divided by the most recent spot price. Unlike Asness et al. (2013), I do not take logarithms, as the impact of COVID-19 pushed oil prices into negative territory. The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS). Note that the significance level of the coefficients is indicated with the notation $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. The numbers in brackets are the p-values. The alphas and R^2 are reported in percentage points (e.g., 0.5 is equivalent to 0.5%).

	Long H	Long S	HmS
	Panel C: Bakshi et al. (2019)		
α	0.58***	-0.09	0.67**
	(0.01)	(0.58)	(0.04)
β_{mkt}	1.05***	0.96***	0.09
	(0.00)	(0.00)	(0.20)
β_{carry}	-0.08*	0.06	-0.14*
	(0.07)	(0.16)	(0.06)
$\beta_{momentum}$	-0.06	0.08**	-0.15*
	(0.26)	(0.03)	(0.08)
Adjusted R^2	68.34	74.54	4.82
	Panel D: Bakshi et al. (2019) plus value		
α	0.53**	-0.05	0.59*
	(0.01)	(0.75)	(0.07)
β_{mkt}	1.05***	0.96***	0.09
	(0.00)	(0.00)	(0.21)
β_{carry}	-0.07	0.07	-0.14*
	(0.15)	(0.17)	(0.09)
$\beta_{momentum}$	-0.05	0.08*	-0.13
	(0.38)	(0.05)	(0.14)
β_{value}	0.01	0.00	0.01
	(0.79)	(0.95)	(0.88)
Adjusted R^2	68.63	74.71	4.12

Table 4 – *Continued*

	Long H	Long S	HmS
	Panel A: CAPM		
α	0.03	-0.09	0.12
	(0.88)	(0.60)	(0.70)
β_{mkt}	1.00***	1.02***	-0.02
	(0.00)	(0.00)	(0.78)
Adjusted R^2	70.49	74.33	-0.39
	Panel B: CAPM plus carry		
α	0.07	-0.13	0.20
	(0.71)	(0.42)	(0.51)
β_{mkt}	1.00***	1.01***	-0.01
	(0.00)	(0.00)	(0.92)
β_{carry}	-0.12***	0.12***	-0.24***
	(0.00)	(0.00)	(0.00)
Adjusted R^2	71.60	75.61	5.38

(*Continued*)

Table 5: Performance attribution for HmS (country-level benchmark).

This table evaluates the performance of the HmS strategy and the long-only strategies on H and S commodities, using country-level data. For the purpose of this analysis, I construct the sorting variable by considering all climate-driven disaster occurrences within top-producing countries. This process counts disasters both in locations with and without key producing sites. Carry is a long-short portfolio going long on the top tercile of backwardated commodities and short on the bottom tercile. Momentum is a long-short portfolio that goes long on the commodities in the top tercile of cumulative returns over the last 12 months (excluding the most recent month) and short on those in the bottom tercile. Value is a long-short portfolio going long on the top tercile of undervalued commodities and short on the bottom tercile. Similarly to Asness et al. (2013), the value signal is defined as the average spot price from 4.5 to 5.5 years ago, divided by the most recent spot price. Unlike Asness et al. (2013), I do not take logarithms, as the impact of COVID-19 pushed oil prices into negative territory. The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS). Note that the significance level of the coefficients is indicated with the notation $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. The numbers in brackets are the p-values. The alphas and R^2 are reported in percentage points (e.g., 0.5 is equivalent to 0.5%).

	Long H	Long S	HmS
	Panel C: Bakshi et al. (2019)		
α	0.07	-0.14	0.22
	(0.69)	(0.39)	(0.47)
β_{mkt}	1.00***	1.01***	-0.01
	(0.00)	(0.00)	(0.89)
β_{carry}	-0.10**	0.09**	-0.19***
	(0.01)	(0.02)	(0.01)
$\beta_{momentum}$	-0.05	0.09**	-0.14**
	(0.19)	(0.01)	(0.04)
Adjusted R^2	71.74	76.24	7.17
	Panel D: Bakshi et al. (2019) plus value		
α	0.03	-0.15	0.17
	(0.89)	(0.39)	(0.58)
β_{mkt}	1.00***	1.02***	-0.02
	(0.00)	(0.00)	(0.81)
β_{carry}	-0.09**	0.12***	-0.21***
	(0.04)	(0.01)	(0.01)
$\beta_{momentum}$	-0.04	0.11***	-0.14**
	(0.39)	(0.00)	(0.04)
β_{value}	0.03	0.06	-0.03
	(0.54)	(0.16)	(0.67)
Adjusted R^2	72.15	76.72	6.83

Table 5 – *Continued*

	p	2002 to 2011	2012 to 2021	2002 to 2021
Panel A1: Annualized Sharpe Ratio (controlling for $J_{mkt,t,s}$)				
V	0	0.54	0.33	0.43
	1	1.02	0.20	0.64
R	0	0.14	-0.50	-0.13
	1	0.19	-0.29	0.00
VmR	0	0.40	0.73	0.57
	1	0.82	0.45	0.65
Panel B1: Annualized Sortino Ratio (controlling for $J_{mkt,t,s}$)				
V	0	0.78	0.63	0.72
	1	1.49	0.36	1.06
R	0	0.19	-0.83	-0.19
	1	0.25	-0.52	0.00
VmR	0	0.65	1.22	0.93
	1	1.47	0.75	1.13

(*Continued*)

Table 6: Performance metrics of the vulnerable minus resilient (VmR) strategy and benchmark portfolios.

This table reports the vulnerable minus resilient long-short portfolio. Long-only portfolios for vulnerable and resilient commodities are also reported. The Baltic Dry Index is employed as a proxy for global activity following Kilian (2009). The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS).

	p	2002 to 2011	2012 to 2021	2002 to 2021
Panel A2: Annualized Sharpe Ratio (controlling for excess market returns)				
V	0	0.53	0.23	0.38
	1	0.97	0.19	0.60
R	0	0.17	-0.54	-0.13
	1	0.16	-0.35	-0.03
VmR	0	0.37	0.66	0.52
	1	0.80	0.47	0.64
Panel B2: Annualized Sortino Ratio (controlling for excess market returns)				
V	0	0.77	0.42	0.62
	1	1.34	0.36	0.97
R	0	0.23	-0.90	-0.19
	1	0.20	-0.61	-0.05
VmR	0	0.60	1.08	0.84
	1	1.40	0.78	1.11
Panel A3: Annualized Sharpe Ratio (controlling for the Baltic Dry Index, % changes)				
V	0	0.66	0.17	0.40
	1	0.79	0.35	0.59
R	0	0.29	-0.40	-0.02
	1	0.21	-0.38	-0.03
VmR	0	0.44	0.42	0.43
	1	0.57	0.70	0.62
Panel B3: Annualized Sortino Ratio (controlling for the Baltic Dry Index, % changes)				
V	0	1.01	0.21	0.54
	1	1.12	0.70	0.94
R	0	0.40	-0.77	-0.03
	1	0.28	-0.69	-0.04
VmR	0	0.79	0.50	0.57
	1	1.02	1.20	1.06

Table 6 – *Continued*

	Long V	Long R	VmR
	Panel A: CAPM		
α	0.48	-0.57**	1.10**
	(0.11)	(0.02)	(0.01)
β_{mkt}	1.01***	1.04***	-0.03
	(0.00)	(0.00)	(0.70)
Adjusted R^2	50.49	62.39	-0.38
	Panel B: CAPM plus carry		
α	0.50*	-0.55**	1.10**
	(0.10)	(0.02)	(0.01)
β_{mkt}	1.01***	1.05***	-0.03
	(0.00)	(0.00)	(0.70)
β_{carry}	-0.06	-0.07	0.01
	(0.38)	(0.28)	(0.90)
Adjusted R^2	50.47	62.54	-0.82

(*Continued*)

Table 7: Performance attribution with $p = 0$.

This table evaluates the performance of the VmR strategy and the long-only strategies on V and R commodities. Carry is a long-short portfolio going long on the top tercile of backwardated commodities and short on the bottom tercile. Momentum is a long-short portfolio that goes long on the commodities in the top tercile of cumulative returns over the last 12 months (excluding the most recent month) and short on those in the bottom tercile. Value is a long-short portfolio going long on the top tercile of undervalued commodities and short on the bottom tercile. Similarly to Asness et al. (2013), the value signal is defined as the average spot price from 4.5 to 5.5 years ago, divided by the most recent spot price. Unlike Asness et al. (2013), I do not take logarithms, as the impact of COVID-19 pushed oil prices into negative territory. The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS). Note that the significance level of the coefficients is indicated with the notation $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. The numbers in brackets are the p-values. The alphas and R^2 are reported in percentage points (e.g., 0.5 is equivalent to 0.5%).

	Long V	Long R	VmR
	Panel C: Bakshi et al. (2019)		
α	0.50*	-0.56**	1.10**
	(0.10)	(0.02)	(0.01)
β_{mkt}	1.01***	1.05***	-0.04
	(0.00)	(0.00)	(0.66)
β_{carry}	-0.02	-0.09	0.07
	(0.72)	(0.15)	(0.50)
$\beta_{momentum}$	-0.09	0.07	-0.16
	(0.19)	(0.15)	(0.12)
Adjusted R^2	50.78	62.73	0.18
	Panel D: Bakshi et al. (2019) plus value		
α	0.50	-0.54**	1.08**
	(0.11)	(0.02)	(0.02)
β_{mkt}	1.00***	1.04***	-0.04
	(0.00)	(0.00)	(0.65)
β_{carry}	-0.05	-0.09	0.04
	(0.46)	(0.20)	(0.72)
$\beta_{momentum}$	-0.11	0.06	-0.17
	(0.15)	(0.24)	(0.12)
β_{value}	-0.08	-0.01	-0.07
	(0.29)	(0.81)	(0.52)
Adjusted R^2	51.77	62.61	-0.18

Table 7 – *Continued*

	Long V	Long R	VmR
	Panel A: CAPM		
α	0.81***	-0.35	1.20***
	(0.00)	(0.16)	(0.01)
β_{mkt}	1.09***	1.06***	0.02
	(0.00)	(0.00)	(0.83)
Adjusted R^2	61.22	61.87	-0.41
	Panel B: CAPM plus carry		
α	0.84***	-0.40	1.27***
	(0.00)	(0.11)	(0.00)
β_{mkt}	1.09***	1.06***	0.03
	(0.00)	(0.00)	(0.74)
β_{carry}	-0.12**	0.13***	-0.25***
	(0.02)	(0.01)	(0.00)
Adjusted R^2	61.91	62.89	2.97

(*Continued*)

Table 8: Performance attribution with $p = 1$.

This table evaluates the performance of the VmR strategy and the long-only strategies on V and R commodities. Carry is a long-short portfolio going long on the top tercile of backwardated commodities and short on the bottom tercile. Momentum is a long-short portfolio that goes long on the commodities in the top tercile of cumulative returns over the last 12 months (excluding the most recent month) and short on those in the bottom tercile. Value is a long-short portfolio going long on the top tercile of undervalued commodities and short on the bottom tercile. Similarly to Asness et al. (2013), the value signal is defined as the average spot price from 4.5 to 5.5 years ago, divided by the most recent spot price. Unlike Asness et al. (2013), I do not take logarithms, as the impact of COVID-19 pushed oil prices into negative territory. The risk-free rate is sourced from Kenneth R. French's website and accessed via Wharton Research Data Services (WRDS). Note that the significance level of the coefficients is indicated with the notation $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$. The numbers in brackets are the p-values. The alphas and R^2 are reported in percentage points (e.g., 0.5 is equivalent to 0.5%).

	Long V	Long R	VmR
	Panel C: Bakshi et al. (2019)		
α	0.84***	-0.40	1.28***
	(0.00)	(0.10)	(0.00)
β_{mkt}	1.09***	1.06***	0.03
	(0.00)	(0.00)	(0.75)
β_{carry}	-0.11*	0.12**	-0.22**
	(0.06)	(0.02)	(0.01)
$\beta_{momentum}$	-0.03	0.04	-0.07
	(0.64)	(0.41)	(0.46)
Adjusted R^2	61.80	62.87	2.90
	Panel D: Bakshi et al. (2019) plus value		
α	0.90***	-0.43***	1.36***
	(0.00)	(0.09)	(0.00)
β_{mkt}	1.09***	1.06***	0.03
	(0.00)	(0.00)	(0.76)
β_{carry}	-0.14**	0.14**	-0.28***
	(0.01)	(0.02)	(0.00)
$\beta_{momentum}$	-0.07	0.06	-0.13
	(0.33)	(0.34)	(0.26)
β_{value}	-0.11*	0.05	-0.16
	(0.07)	(0.48)	(0.13)
Adjusted R^2	63.09	63.31	4.05

Table 8 – *Continued*

Commodity	Vulnerable	Resilient	Difference
Natural gas	104	25	79
Wheat	99	42	57
Gold	57	20	37
Corn	76	40	36
Cotton	50	15	35
Sugar	71	56	15
Soybeans	56	70	-14
Aluminum	42	59	-17
Crude oil	47	89	-42
Copper	24	72	-48
Silver	30	80	-50
Cocoa	16	75	-59
Nickel	25	85	-60
Coffee	21	136	-115

Table 9: Vulnerable and resilient commodities with $p = 0$.

This table reports the number of times each commodity is considered to be vulnerable or resilient, across the sample ranging from 2002 to 2021. The last column reports the difference.

Commodity	Vulnerable	Resilient	Difference
Soybeans	110	25	85
Natural gas	96	29	67
Corn	101	34	67
Wheat	103	41	62
Gold	61	8	53
Cotton	39	19	20
Silver	60	60	0
Aluminum	44	59	-15
Crude oil	40	79	-39
Sugar	37	83	-46
Copper	26	108	-82
Cocoa	17	112	-95
Nickel	26	125	-99
Coffee	19	143	-124

Table 10: Vulnerable and resilient commodities with $p = 1$.

This table reports the number of times each commodity is considered to be vulnerable or resilient, across the sample ranging from 2002 to 2021. The last column reports the difference.

References

T. Akiyama and A. Nishio. Sulawesi's cocoa boom: Lessons of smallholder dynamism and a hands-off policy. *Bulletin of Indonesian Economic Studies*, 33(2):97–121, 1997.

E. Allen and C. Valdes. Brazil's corn industry and the effect on the seasonal pattern of U.S. corn exports. Technical report, USDA, 2016.

C. S. Asness. *Variables that explain stock returns.* Ph.d. dissertation, University of Chicago, 1994.

C. S. Asness, T. J. Moskowitz, and L. H. Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013.

G. Bakshi, X. Gao, and A. G. Rossi. Understanding the sources of risk underlying the cross section of commodity returns. *Management Science*, 65(2):619–641, 2019.

D. S. Battisti and R. L. Naylor. Historical warnings of future food insecurity with unprecedented seasonal heat. *Science*, 323(5911):240–244, 2009.

A. Bernstein, M. T. Gustafson, and R. Lewis. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272, 2019.

D. Blaise and K. Kranthi. Cotton production in India. *Cotton Production*, pages 193–215, 2019.

P. Bolton and M. Kacperczyk. Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549, 2021.

G. B. Bonan and S. C. Doney. Climate, ecosystems, and planetary futures: The challenge to predict life in earth system models. *Science*, 359(6375):eaam8328, 2018.

J. Boudoukh, M. Richardson, and R. F. Whitelaw. Industry returns and the Fisher effect. *The Journal of Finance*, 49(5):1595–1615, 1994.

M. J. Brennan. The supply of storage. *The American Economic Review*, 48(1):50–72, 1958.

J. Breunese, H. Mijnlieff, and J. Lutgert. The life cycle of the Netherlands' natural gas exploration: 40 years after Groningen, where are we now? In *Geological Society, London, Petroleum Geology Conference Series*, volume 6, pages 69–75. The Geological Society of London, 2005.

A. Buraschi and F. Pellegrino. From extinction to resurrection: Unveiling the role of short-selling frictions in cross-sectional momentum. *Available at SSRN*, 2024.

C. Cacciuttolo and E. Atencio. Past, present, and future of copper mine tailings governance in Chile (1905–2022): A review in one of the leading mining countries in the world. *International Journal of Environmental Research and Public Health*, 19(20): 13060, 2022.

M. Canadian Department of Energy and Resources. *Canadian Minerals Yearbook 1970*. Information Canada, 1972.

K. Chakraborty, D. Hudson, D. E. Ethridge, S. K. Misra, and G. Kar. An overview of the cotton and textile industries in India. Technical Report 53178, Texas Tech University, Department of Agricultural and Applied Economics, 1999.

K. Chao and J. C. Y. Chao. *The Introduction of Cotton into China and Its Spread*, volume 74, pages 4–27. Harvard University Asia Center, 1 edition, 1977.

T. Chatzopoulos, I. P. Domínguez, M. Zampieri, and A. Toreti. Climate extremes and agricultural commodity markets: A global economic analysis of regionally simulated events. *Weather and Climate Extremes*, 27:100193, 2020.

Z. Cheng, X. Xie, W. Yao, J. Feng, Q. Zhang, and J. Fang. Multi-element geochemical mapping in southern China. *Journal of Geochemical Exploration*, 139:183–192, 2014.

CIA. Nickel supply in the Soviet block. Declassified Report DOC_0000496313, Central Intelligence Agency, 1954. URL https://www.cia.gov/readingroom/docs/DOC_0000496313.pdf. Declassified Intelligence Document.

CIA. Soviet cotton production in the postwar period. Declassified Report DOC_0000496067, Central Intelligence Agency, 1955. URL https://www.cia.gov/readingroom/docs/DOC_0000496067.pdf. Declassified Intelligence Document.

CIA. Development of the aluminum industry of the USSR. Declassified Report CIA-RDP81-01043R004000180001-3, Central Intelligence Agency, 1959. URL https://www.cia.gov/readingroom/docs/CIA-RDP81-01043R004000180001-3.pdf. Declassified Intelligence Document.

CIA. The current gold position of the USSR. Declassified Report DOC_0000307719, Central Intelligence Agency, 1962. URL https://www.cia.gov/readingroom/docs/DOC_0000307719.pdf. Declassified Intelligence Document.

CIA. Soviet copper: An overview. Declassified Report DOC_0000484025, Central Intelligence Agency, 1974. URL https://www.cia.gov/readingroom/docs/DOC_0000484025.pdf. Declassified Intelligence Document.

CIA. China: Provincial grain production. Declassified Report CIA-RDP83B00227R000100070005-7, Central Intelligence Agency, 1982. URL https://www.cia.gov/readingroom/docs/CIA-RDP83B00227R000100070005-7.pdf. Declassified Intelligence Document.

D. A. Cranstone. *Canada: A History of Mining and Mineral Exploration and the Outlook for the Future*. Minister of Public Works and Government Services Canada, 2002.

M. de Araújo, T. Goes, R. Marra, and M. O. d. Souza. Sugarcane the evolution of technological thinking in agriculture. *Revista de Política Agrícola*, 19(4):65–77, 2010.

P. Desrochers and J. Szurmak. Long distance trade, locational dynamics and by-product development: Insights from the history of the American cottonseed industry. *Sustainability*, 9(4):579, 2017.

J. H. DeYoung. *International Strategic Minerals Inventory Summary Report: Nickel*. Department of the Interior, US Geological Survey, 1986.

Y. Ding and J. Xu. Global vulnerability of agricultural commodities to climate risk: Evidence from satellite data. *Economic Analysis and Policy*, 80:669–687, 2023.

P. Eichholtz, E. Steiner, and E. Yönder. Where, when and how do sophisticated investor respond to flood risk? 2019.

R. F. Engle, S. Giglio, B. Kelly, H. Lee, and J. Stroebel. Hedging climate change news. *The Review of Financial Studies*, 33(3):1184–1216, 2020.

S. Ermolaev. The formation and evolution of the Soviet Union's oil and gas dependence. Technical report, British Embassy Moscow, 2017. This paper is a part of "Comparative Analysis and Policy Proposals Aimed at Diversifying the Russian Economy and Enhancing Prosperity" project, supported by the UK Foreign and Commonwealth Office.

H. C. Evans. Witches' broom disease (Moniliophthora Perniciosa): History and biology. *Cacao Diseases: A History of Old Enemies and New Encounters*, pages 137–177, 2016.

E. F. Fama and K. R. French. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. *The Journal of Business*, 60(1):55–73, 1987.

E. F. Fama and K. R. French. The cross-section of expected stock returns. *Journal of Finance*, 47(2):427–465, 1992.

Ghana Cocoa Board. The cocoa story, 2024. URL https://cocobod.gh/cocoa-story. Accessed: 2024-08-31.

S. Giglio, B. Kelly, and J. Stroebel. Climate finance. *Annual review of financial economics*, 13(1):15–36, 2021a.

S. Giglio, M. Maggiori, K. Rao, J. Stroebel, and A. Weber. Climate change and long-run discount rates: Evidence from real estate. *The Review of Financial Studies*, 34(8): 3527–3571, 2021b.

D. Giovannucci, J. Leibovich, D. Pizano, G. Paredes, S. Montenegro, H. Arévalo, and P. Varangis. Colombia coffee sector study. *Centro de Estudios Sobre Desarrollo Economico Documento*, 2007.

G. Gorton and K. G. Rouwenhorst. Facts and fantasies about commodity futures. *Financial Analysts Journal*, 62(2):47–68, 2006.

G. B. Gorton, F. Hayashi, and K. G. Rouwenhorst. The fundamentals of commodity futures returns. *Review of Finance*, 17(1):35–105, 2013.

Government of India. Union budget of india, 1991.

M. Grinblatt and T. J. Moskowitz. Predicting stock price movements from past returns: The role of consistency and tax-loss selling. *Journal of Financial Economics*, 71(3): 541–579, 2004.

H. Hong and J. C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *The Journal of Finance*, 54(6):2143–2184, 1999.

H. Hsu and F. Gale. Regional shifts in China's cotton production and use. Technical report, USDA, 2002.

IMF. *A study of the Soviet economy*. International Monetary Fund, 1991.

N. Jegadeesh. Evidence of predictable behavior of security returns. *The Journal of Finance*, 45(3):881–898, 1990.

N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993.

S. Jia, X. Chen, L. Han, and J. Jin. Global climate change and commodity markets: A hedging perspective. *Journal of Futures Markets*, 43(10):1393–1422, 2023.

F. I. Johnson. Sugar in Brazil: Policy and production. *The Journal of Developing Areas*, 17(2):243–256, 1983.

N. Kaldor. Speculation and economic stability. *The Review of Economic Studies*, 7(1): 1–27, 1939.

S. Kansal. Factors determining Indian sugar production and its comparative advantage. Technical report, FAO, 1997.

J. M. Keenan and J. T. Bradt. Underwaterwriting: from theory to empiricism in regional mortgage markets in the us. *Climatic Change*, 162(4):2043–2067, 2020.

L. Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American economic review*, 99(3):1053–1069, 2009.

F. N. Kogan. Grain production in the USSR: Present situation, perspectives for development and methods for prediction. Technical report, USDA, 1981.

S. Kolavalli and M. Vigneri. *Cocoa in Ghana: Shaping the success of an economy*. World Bank, 2011.

M. Krishnamurthy. States of wheat: the changing dynamics of public procurement in Madhya Pradesh. *Economic and Political Weekly*, pages 72–83, 2012.

R. M. Levine and H. G. Van Oss. Russia's silver resources: A post-Soviet assessment. *International Geology Review*, 35(10):974–982, 1993.

Y. Li. China: An emerging sugar super power. *Sugar Tech*, 6(4):213–227, 2004.

Y. Li, W. Zhang, L. Ma, L. Wu, J. Shen, W. J. Davies, O. Oenema, F. Zhang, and Z. Dou. An analysis of China's grain production: Looking back and looking forward. *Food and Energy Security*, 3(1):19–32, 2014.

A. W. Lo and A. C. MacKinlay. When are contrarian profits due to stock market overreaction? *The Review of Financial Studies*, 3(2):175–205, 1990.

C. Lu, A. Liu, Y. Xiao, X. Liu, G. Xie, and S. Cheng. Changes in China's grain production pattern and the effects of urbanization and dietary structure. *Journal of Resources and Ecology*, 11(4):358–365, 2020.

S. MacDonald, F. Gale, and J. Hansen. Cotton policy in China. Technical report, USDA, 2015.

S. J. McCoy and R. P. Walsh. Wildfire risk, salience & housing demand. *Journal of Environmental Economics and Management*, 91:203–228, 2018.

R. McStocker. The Indonesian coffee industry. *Bulletin of Indonesian Economic Studies*, 23(1):40–69, 1987.

C. W. Merrill. *Summarized Data of Silver Production*. Economic Paper 8. United States Government Printing Office, 1930.

G. M. Mudd. Gold mining in Australia: linking historical trends and environmental and resource sustainability. *Environmental Science & Policy*, 10(7-8):629–644, 2007.

K. Nam. Investigating the effect of climate uncertainty on global commodity markets. *Energy Economics*, 96:105123, 2021.

D. Nestel. Coffee in Mexico: international market, agricultural landscape and ecology. *Ecological Economics*, 15(2):165–178, 1995.

M. Oladokun and A. Daramola. Nigerian cocoa development programme: A comparative efficiency of fund utilisation. *Savings and Development*, pages 175–184, 1990.

F. Ortega and S. Taşpınar. Rising sea levels and sinking property values: Hurricane sandy and new york's housing market. *Journal of Urban Economics*, 106:81–100, 2018.

A. Ouazad and M. E. Kahn. *Mortgage finance in the face of rising climate risk*. National Bureau of Economic Research, 2019.

M. Painter. An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2):468–482, 2020.

J. Pereira, L. De Almeida, and S. Santos. Witches' broom disease of cocoa in Bahia: Attempts at eradication and containment. *Crop Protection*, 15(8):743–752, 1996.

M. Renard. The Mexican coffee crisis. *Latin American Perspectives*, 37(2):21–33, 2010.

E. L. Rosvold and H. Buhaug. GDIS, a global dataset of geocoded disaster locations. *Scientific Data*, 8(1):61, 2021.

K. G. Rouwenhorst and K. Tang. Commodity investing. *Annual Review Financial Economics*, 4(1):447–467, 2012.

F. Ruf, P. Ehret, and Yoddang. *Smallholder cocoa in Indonesia: why a cocoa boom in Sulawesi?* Springer, 1996.

B. Y. S. Sabas, K. G. Danmo, K. A. T. Madeleine, and J. Bogaert. Cocoa production and forest dynamics in Ivory Coast from 1985 to 2019. *Land*, 9(12), 2020.

R. J. Samuels. The industrial destructuring of the Japanese aluminum industry. *Pacific Affairs*, 56(3):495–509, 1983.

D. G. Sapir and C. Misson. The development of a database on disasters. *Disasters*, 16 (1):74–80, 1992.

M. J. H. Sly. The Argentine portion of the soybean commodity chain. *Palgrave Communications*, 3(1):1–11, 2017.

V. Smirnov. Alumina production in Russia part I: Historical background. *Journal of The Minerals, Metals & Materials Society*, 1996.

Tropical Forest Aliance. Harnessing the prospect of Indonesia's sustainable cocoa commodity. https://lpem.org/wp-content/uploads/2022/11/Investor_Brief_Cocoa_Designed.pdf, 2022. Accessed: 2024-08-31.

C. Troy and M. Ober. Winter storms have disrupted U.S. natural gas production, 2024. URL https://www.eia.gov/todayinenergy/detail.php?id=61563. Accessed: 2024-11-01.

F. C. Tuan, C. Fang, and Z. Cao. China's soybean imports expected to grow despite short-term disruptions. Technical report, USDA, 2004.

U.S. Congress, Office of Technology Assessment. Nonferrous metals: Industry structure. Technical report, U.S. Government Printing Office, Washington, DC, September 1990.

B. Volsi, T. S. Telles, C. E. Caldarelli, and M. R. G. d. Camara. The dynamics of coffee production in Brazil. *PloS One*, 14(7):e0219742, 2019.

H. Working. The theory of price of storage. *The American Economic Review*, 39(6): 1254–1262, 1949.

World Bank. ICO/CFC study of marketing and trading policies and systems in selected coffee producing countries: Angola country profile, 2000.

C. Wu, G. Bai, and L. Xu. Types and distribution of silver ore deposits in China. *Mineralium Deposita*, 28:223–239, 1993.

X. Xie, X. Wang, Q. Zhang, G. Zhou, H. Cheng, D. Liu, Z. Cheng, and S. Xu. Multi-scale geochemical mapping in China. *Geochemistry: Exploration, Environment, Analysis*, 8 (3-4):333–341, 2008.

T. Zhou, R. J. Goldfarb, and N. G. Phillips. Tectonics and distribution of gold deposits in China–an overview. *Mineralium Deposita*, 37:249–282, 2002.

Appendix

A Geographic distribution of commodity producers

A.1 Soft commodities

Soft commodities are those that are cultivated or grown. This paper focuses on some of the most significant agricultural commodities within this category. Namely, cocoa, coffee, cotton, grains (corn, wheat, and soybeans) and sugar. While the cultivation and production of these products are spread across many nations, Brazil, China, India, and the United States stand out for their contributions to the markets, as they are leading producers of multiple commodities from this list.

Brazil has long been a dominant player in the global production of soft commodities, leveraging its vast land resources. The country is perhaps best known for being a leading coffee producer for over a century. Cultivation in Brazil is concentrated in the southeastern regions, particularly in states like Minas Gerais, Espírito Santo and São Paulo (Volsi et al., 2019). These areas, known for their ideal growing conditions, produce primarily Arabica coffee. In addition to coffee, Brazil is a major producer of sugar and grains. Historically rooted in the country's northeast, sugarcane cultivation shifted to the southeast in the 1970s, particularly in São Paulo, where it now accounts for the vast majority of production (Johnson, 1983; de Araújo et al., 2010). The country's involvement in the production of biofuels, particularly ethanol derived from sugarcane, is another critical element of Brazil's agricultural strategy. In terms of grains, Brazil ranks among the top producers of corn and soybeans, with cultivation spread across the southern and central regions (Allen and Valdes, 2016). Finally, cocoa cultivation was once significant in the northeastern state of Bahia, but the sector suffered greatly in the 1990s due to the outbreak of the Witches' Broom disease (Pereira et al., 1996; Evans, 2016).

China, another major player in the production of soft commodities, has a long history in the cultivation of various crops. In the realm of cotton, China has a history of more than two millennia (Chao and Chao, 1977) and it is one of the world's leading producers. Cotton cultivation is concentrated in the Yangtze River Valley, the Yellow River Valley, and the northwest inland regions (Hsu and Gale, 2002; MacDonald et al., 2015). Beyond cotton, China is a significant producer of grains, particularly corn and wheat. Corn is grown primarily in the Northeast China Plain and the North China Plain, while wheat is cultivated in both the northern plains, as well as in the mid and lower Yangtze River region (CIA, 1982; Li et al., 2014; Lu et al., 2020). Although China still produces soybeans, it has increasingly relied on imports due to rising domestic demand. Additionally, since the 1990s, China has become a leading sugar producer, with its most significant cultivation areas located in the south (Li, 2004).

India is another major player in the global soft commodities market, particularly in cotton, wheat and sugar production. India's rise as a cotton and wheat producer has been significant since the 1990s. Cotton is grown in various regions, including northern, central, and southern states (Chakraborty et al., 1999). Wheat cultivation is primarily concentrated in central and northern states (Government of India, 1991; Krishnamurthy, 2012). India is also a leading producer of sugar since the 1970s, with sugarcane cultivation spread across both tropical and subtropical regions (Kansal, 1997).

The United States is another powerhouse in the soft commodities sector, particularly in the production of grains and cotton. The country's Corn Belt, located in the Midwest, has been a major hub for corn production for decades. States such as Iowa, Illinois, and Nebraska are key contributors to the US corn industry. Corn is also used extensively for biofuel production, with ethanol being a major byproduct of US corn farming. Soybeans are often grown in rotation with corn, with production areas extending further south than the Corn Belt, reaching into Arkansas. The US is also a top producer of wheat,

with major cultivation areas located in the Great Plains and Northern Plains. In addition to grains, the United States has a strong cotton industry, which is concentrated in the south. Historically, cotton was grown primarily in southeastern states, but the Cotton Belt has expanded westward since the 1970s, reaching as far as California (Desrochers and Szurmak, 2017).

In addition to these key players, several other countries have made significant contributions to the global soft commodities market. West African nations such as the Ivory Coast and Ghana have been leading the cocoa market since the 1970s. Indonesia has emerged as a leading producer of both cocoa and coffee. Vietnam, another rising star in the soft commodities market, has become a global leader in coffee production, particularly Robusta coffee, with its main cultivation areas in the Central Highlands. Argentina, too, has become a major player in soybean production, capitalizing on China's growing demand for imports. The Soviet Union and, later, Russia have played a pivotal role in the global production of soft commodities, particularly cotton and wheat. A detailed commodity-level breakdown is reported below.

A.1.1 Cocoa

During the 1970s, the majority of the world's cocoa was produced in the Ivory Coast, Ghana, and Nigeria, as well as in Brazil.

In West Africa, cocoa cultivation was concentrated in the southern regions of Ivory Coast (Sabas et al., 2020) and Ghana (Kolavalli and Vigneri, 2011), and in south-western Nigeria (Oladokun and Daramola, 1990). These areas, characterised by their forested landscapes, were more suitable for cocoa cultivation compared to the northern Savannah regions. Despite a near collapse in Ghana's cocoa production during the early 1980s, the Western African industry managed to recover through the implementation of various reforms. By 2021, both Ghana and Ivory Coast remained among the top three cocoa pro-

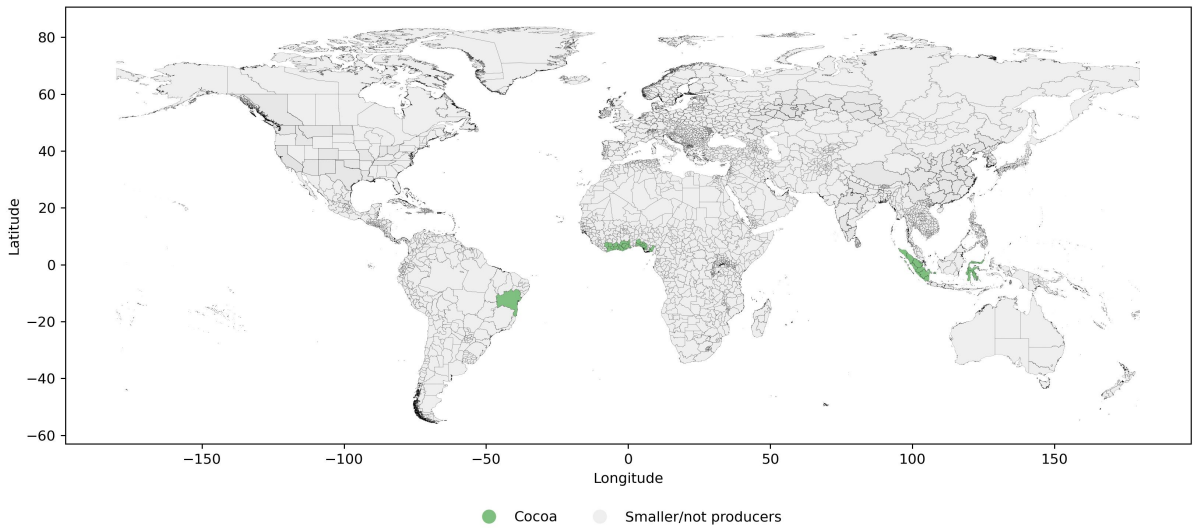


Figure 4: Core producing areas of cocoa.

This figure displays the core producing areas of cocoa for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

ducers globally. Their core production regions survived the test of time, as documented by the official website of the Ghana Cocoa Board (2024) and, again, Sabas et al. (2020). Nigeria remained a top producers, but exited the global podium after 2002.

Brazil, once a leading cocoa producer, saw its prominence decline in 1995 due to the devastating impact of the Witches' Broom disease in Bahia, its main region for this crop (Pereira et al., 1996; Evans, 2016). In the wake of these changes, Indonesia emerged as a major cocoa producer, with primary production concentrated in Sulawesi and, to a lesser extent, in Sumatra. Indonesia's rise as a significant cocoa producer was driven by several factors. The development of this industry was mostly driven by local pioneers, smallholder farmers, who make up the majority of the cocoa producers in the country (Ruf et al., 1996). Low production costs, a highly competitive marketing system driven by the government's policy of limited intervention, and a relatively well-developed transport infrastructure have all contributed to the sector's success (Akiyama and Nishio, 1997). Sulawesi and Sumatra continued to be the leading regions for cocoa production

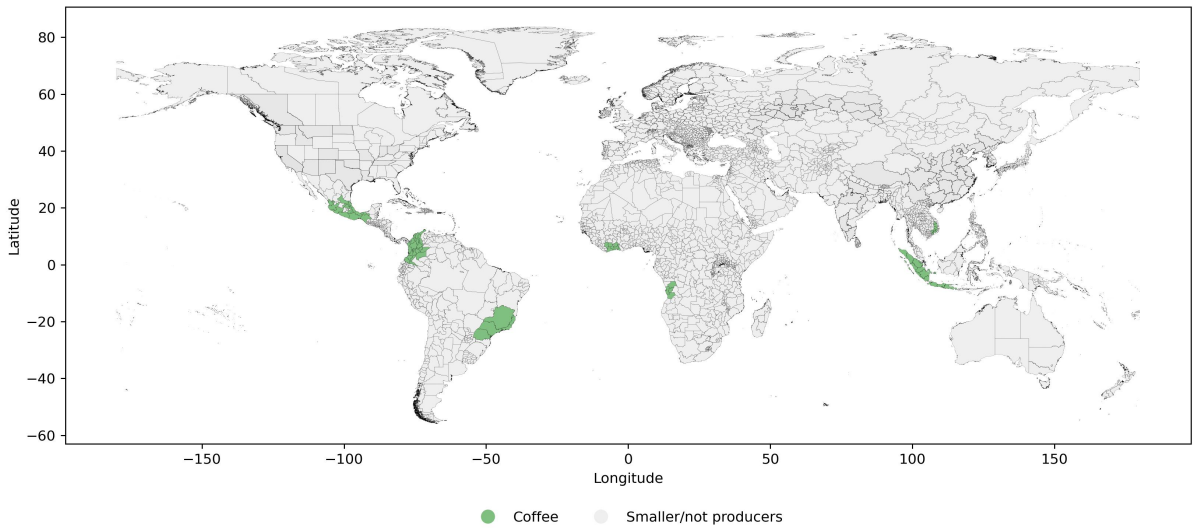


Figure 5: Core producing areas of coffee.

This figure displays the core producing areas of coffee for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

in Indonesia through 2021 (Tropical Forest Aliance, 2022).

A.1.2 Coffee

In the 1970s, Arabica coffee was predominantly produced in Brazil, Colombia, and Mexico, while Robusta coffee was primarily cultivated in Angola and the Ivory Coast.

Brazil cultivated coffee mostly in the southeastern regions, spanning from Minas Gerais to the eastern boarder (west to east) and from Minas Gerais to Paraná (north to south) (Volsi et al., 2019). The relevant USDA crop production maps confirm similar land use for coffee production up to 2021.

At the beginning of the 20th century, Colombia's coffee production was concentrated in the northwestern departments, particularly in North Santander, Antioquia, and Caldas. By the 1970s, coffee cultivation had already expanded further into most western regions, solidifying Colombia's status as one of the leading producers (Giovannucci et al., 2007). Analogous areas constitute present-day core productive centers, as documented by the

relevant USDA crop production maps.

In Mexico, coffee was predominantly cultivated in both the central and southern regions (Nestel, 1995). However, Mexico's dominance was short-lived, ending with the liberalisation of the sector and the onset of the so-called Mexican Coffee Crisis in 1990 (Renard, 2010). Similarly, Angola and the Ivory Coast experienced a brief period of prominence as top coffee producers, but their positions were undermined by both internal and external factors. Angola saw its coffee industry collapse following the outbreak of civil war in 1975. The prolonged conflict devastated the country's agricultural infrastructure and led to the abandonment of coffee plantations (World Bank, 2000). Ivory Coast's shift towards cocoa production curtailed their prominence in coffee.

Consequently, Indonesia and Vietnam gradually ascended to the top ranks in global coffee production, primarily due to their strong Robusta output. Indonesia has a long history with coffee that dates back to the Dutch East Indies period and traditional production centers in Sumatra and Java (McStocker, 1987). Over time, it often alternated positions first with Mexico and then Colombia. In 2021, it once again secured a spot on the podium. In contrast, Vietnam, a relative newcomer, rose to global prominence in the late 1990s, with its primary production areas in the Central Highlands, and consistently maintained its place on the podium since then. USDA GAIN reports indicate that these production areas have remained central to coffee production in both Indonesia and Vietnam well into 2021.

A.1.3 Cotton

In the 1970s, global cotton production was dominated by three major players: China, the United States, and the Soviet Union.

China's involvement in cotton production has deep historical roots, dating back to as early as 200 BC (Chao and Chao, 1977). During the 1970s and 1980s, the most productive

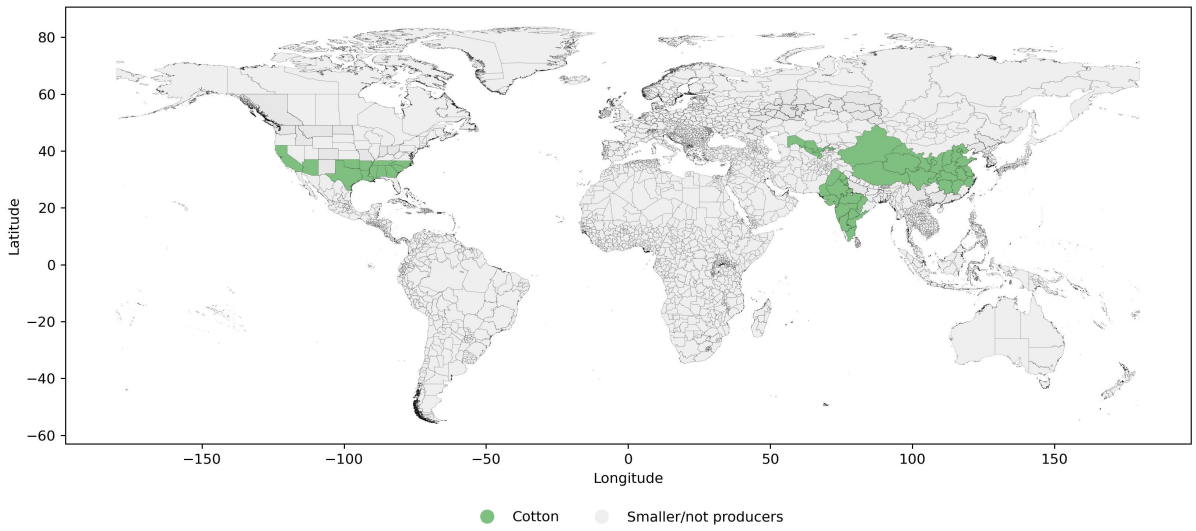


Figure 6: Core producing areas of cotton.

This figure displays the core producing areas of cotton for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

cotton-growing regions, ranked by yield, were the Yangtze River Valley, the Yellow River Valley, and the northwest inland areas (Hsu and Gale, 2002; MacDonald et al., 2015). Although these core regions have remained significant in cotton production, the internal ranking of yields has dramatically shifted, with the northwest inland areas leading in productivity. This shift is further corroborated by USDA crop production maps.

Before the Civil War, cotton production in the United States was confined to the traditional southern states (Desrochers and Szurmak, 2017). However, in the early twentieth century, the Cotton Belt expanded westward, extending as far as California, with the only exception of New Mexico and Florida, marginal producers relative to other states. These areas have remained crucial for cotton production up to the present day. This is further supported by USDA survey data on state-level cotton production over the past five decades, as well as their specialized crop production maps.

The Soviet Union, another major cotton producer, primarily concentrated its cotton cultivation in the Uzbek territory (CIA, 1955). After USSR's collapse in 1991, India

emerged as a global cotton powerhouse. India's rise in cotton production was also bolstered by several technological advancements, including the introduction of hybrid cotton varieties resistant to pests, which were implemented in the 1970s and further improved in the early 2000 (Blaise and Kranthi, 2019). Since the early 1990s, India's major cotton production has been primarily concentrated in three key regions: the northern states of Punjab, Haryana, and Rajasthan; the central states of Madhya Pradesh, Gujarat, Chhattisgarh, and Maharashtra; and the southern states of Telangana, Andhra Pradesh, Karnataka, and Tamil Nadu (Chakraborty et al., 1999).[15] It is notable that India was temporarily surpassed by Pakistan in cotton production from 2000 to 2002, largely due to the latter's successful cultivation in the Punjab and Sind regions.

A.1.4 Grains

Grains are among the most vital agricultural commodities, serving as essential food sources for both humans and livestock. Corn, wheat, and soybeans are often grown in rotation to mitigate crop-specific pests, take advantage of different growing seasons, respond to market demand, and improve soil quality, thus enhancing yields. Hence, top grain-producing countries typically cultivate multiple types of grains.

Since the 1970s, Brazil, China, and the United States have consistently been the top three corn producers. In Brazil, corn production has predominantly been concentrated in the southern and central regions (Allen and Valdes, 2016).[16] China's corn production has been located in the provinces close to the Northeast China Plain and the North China Plain (CIA, 1982; Li et al., 2014; Lu et al., 2020).[17] The United States has been

[15] When Chakraborty et al. (1999) was published, Chhattisgarh and Telangana were part of Madhya Pradesh and Andhra Pradesh, respectively. They gained statehood, in turn, in 2000 and 2014.

[16] This includes centre-west areas once considered unsuitable for agricultural development as they contained the Cerrados, Brazilian lands akin to the Savannah. The main driver of grain cultivation in Brazil's Center-West since the 1960s has been EMBRAPA's development of ad-hoc corn hybrids and soybean varieties. These advances enabled the cultivation of previously unsuitable land, significantly expanding Brazil's agricultural frontier since the 1970s (Allen and Valdes, 2016).

[17] These regions gained prominence in the late 1950s as a result of government policies promoting

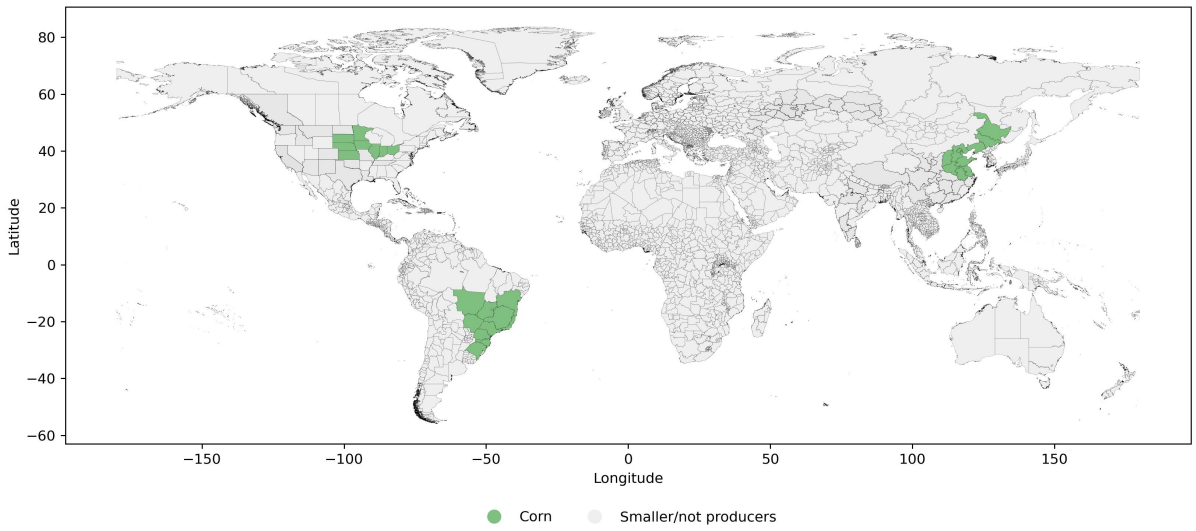


Figure 7: Core producing areas of corn.

This figure displays the core producing areas of corn for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

cultivating corn predominantly in the Midwest, as per USDA Survey Data on state-level corn production. USDA crop production maps confirm these core production areas for both Brazil, China and the United States in 2021.

Brazil, China, and the United States have also long dominated global soybean production. Soybean cultivation typically occurs in areas similar to those used for corn, as the two crops are often grown in rotation. However, in the United States, the traditional soybean-growing region extends slightly further south than the Corn Belt, reaching into Arkansas, as per USDA survey data. Since the early 1980s, China, where soybean production has been primarily concentrated in the northeast, has faced an increasing gap between production and demand. In the 1990s, after significant policy changes, China began to rely heavily on soybean imports, opening opportunities for other countries to become leading producers (Tuan et al., 2004). Argentina capitalized on this shift, developing robust soybean cultivation and surpassing China. From 1998 to 2021, Argentina

agricultural development, enabling the cultivation of all major grains (Li et al., 2014).

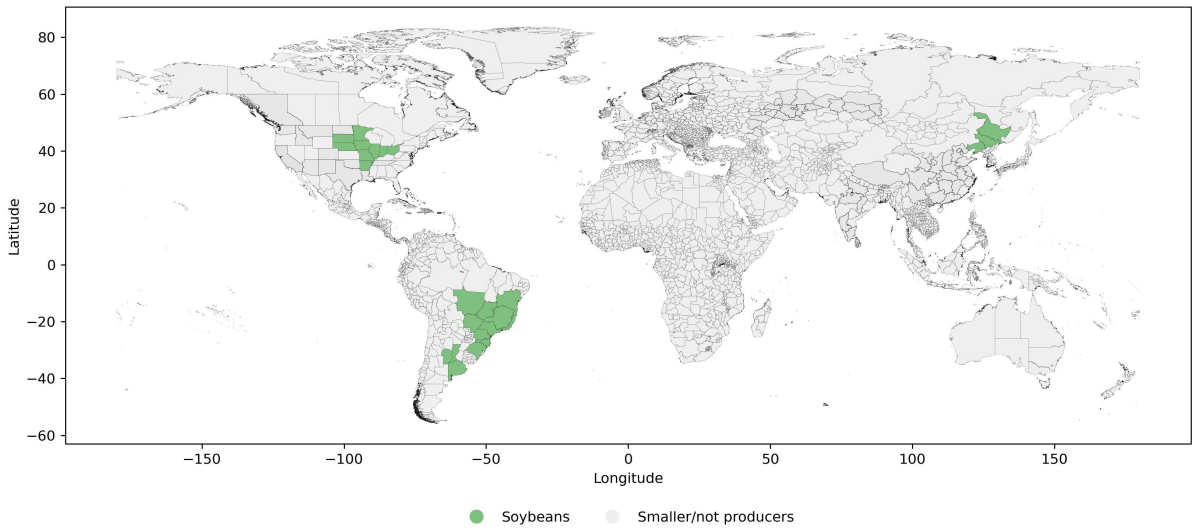


Figure 8: Core producing areas of soybeans.

This figure displays the core producing areas of soybeans for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

consistently maintained its production levels. Soybean farming, which began in the 1970s along the border between Buenos Aires, Córdoba and Santa Fe consolidated in these regions since then (Sly, 2017).

China, the Soviet Union and the United States have traditionally been the largest wheat producers. In China, wheat is primarily grown in the mid and lower Yangtze region and in provinces close the hereinbefore mentioned plains in the north (CIA, 1982; Li et al., 2014; Lu et al., 2020). The Soviet Union, and then Russia, has also been a major wheat producer, with cultivation concentrated in both the European and Asian parts of the country. Winter wheat production has been primarily concentrated in the European territory of the USSR and later the Russian Federation, with key areas including Ukraine (prior to its independence) and the southern parts of Russia's economic regions such as the North Caucasus, Southern, Central, Volga, and Urals. Spring wheat production, on the other hand, has been focused in the eastern territories of both the USSR and the Russian Federation, especially in Kazakhstan (before its independence) and in the

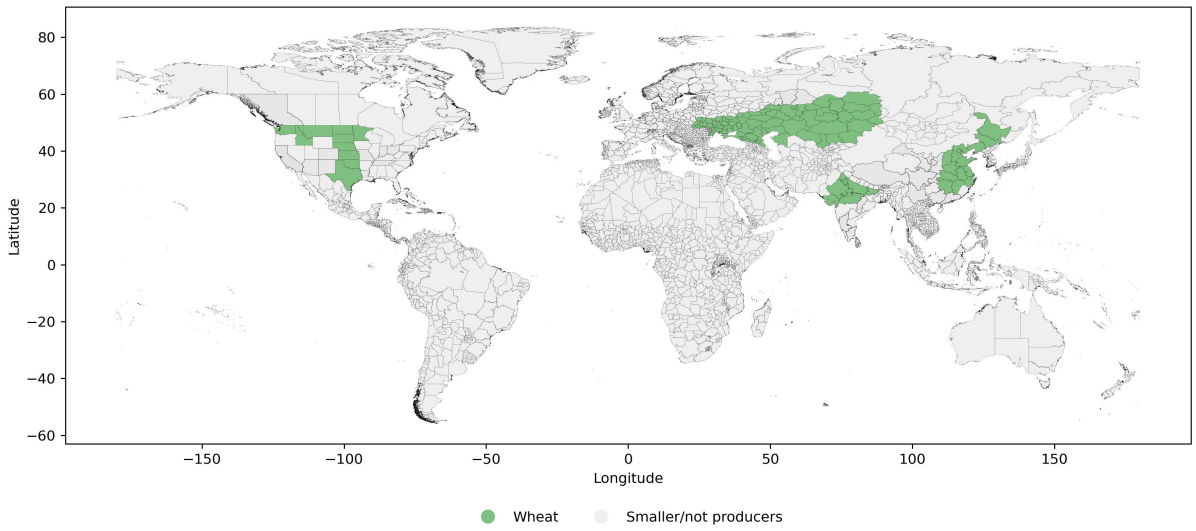


Figure 9: Core producing areas of wheat.

This figure displays the core producing areas of wheat for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

southern parts of Russia's economic regions of the Volga, Urals, and Siberia (Kogan, 1981). USDA crop production maps confirm that the main wheat production regions in Russia were among the most significant during the USSR era. In the United States, wheat cultivation encompasses states including Kansas, Nebraska, Oklahoma, South Dakota, and Texas within the Great Plains, as well as Idaho, Minnesota, Montana, North Dakota, and Washington in the Northern Plains, according to USDA survey data spanning the last five decades. In 1991, India emerged as a leading producer, with cultivation efforts primarily concentrated in the central and northern regions (Government of India, 1991; Krishnamurthy, 2012). These regions primarily have been including Punjab, Haryana, Uttar Pradesh, Rajasthan, Bihar, Madhya Pradesh and Gujarat, as further confirmed by recent USDA production maps. Since India's emergence as a leading producer, Russia and the United States have vied for the position of the third-largest producer, a title Russia ultimately secured in 2014.

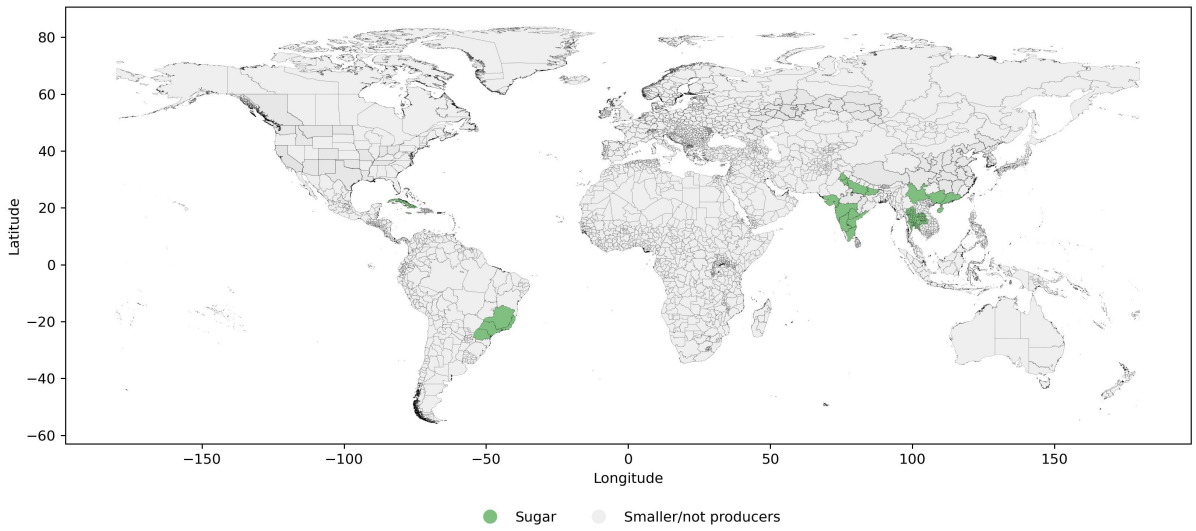


Figure 10: Core producing areas of sugar.

This figure displays the core producing areas of sugar for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

A.1.5 Sugar

During the 1970s, the leading sugar-producing countries were Brazil, Cuba, and India.

Brazil's sugar industry has roots going back to the 1500s when sugarcane was first cultivated in the country's northeastern region. After World War II, the potential for exporting sugar emerged, leading to a rapid expansion of sugarcane cultivation in São Paulo. As a result, São Paulo's sugar industry gained dominance, largely due to its proximity to key consumer markets. In the early 1970s, the southeast region overtook the northeast as Brazil's primary sugar-producing area (Johnson, 1983; de Araújo et al., 2010). Since then, the majority of Brazil's sugar production concentrated there. As a point of reference, the southeast accounted for over 70% of 2021 total production according to USDA GAIN reports.

India's sugar production initially began in the subtropical northwest. However, starting in the 1950s, with the support of the government's Five-Year Plans, sugarcane cultivation was strategically expanded to tropical regions in the southern part of the country.

These regions saw sugarcane thrive, being better suited to tropical climates, resulting in higher yields (Kansal, 1997). As of 2021, sugarcane continues to be cultivated in both subtropical and tropical regions, according to USDA GAIN reports.

Cuba, once a dominant force in the global sugar industry, lost its prominent position following the collapse of the USSR, which had been a key trade partner. In the years since, China has emerged as a major player in the global sugar market, with its sugar production primarily concentrated in the provinces of Yunnan, Guangxi, Guangdong, and Hainan (Li, 2004). Thailand has also ascended as a leading sugar producer, achieving a spot among the top three global producers in 2018 and 2019. Much of Thailand's sugarcane cultivation is centered in the northern and central plains, as documented by USDA GAIN reports.

A.2 Hard commodities

Hard commodities are those that are extracted or mined from the earth. This paper focuses on some of the most traded within this category: crude oil, natural gas, aluminum, copper, nickel, gold, and silver. Similar to appendix A.1, this part of the article is dedicated to reviewing the countries that are leading producers of more than one of these commodities, starting with Australia, China, Russia, and the United States.

Australia has long been a leading producer of industrial and precious metals, particularly nickel, gold, and silver. In terms of nickel, the discovery of significant mineral deposits in the Kambalda region of Western Australia in the mid-1960s marked the beginning of the country's industry (DeYoung, 1986). Since then, Australia has remained one of the world's largest producers, with much of the activity still concentrated in Western Australia. Regarding gold, Australia became a major player as early as the 1850s, with significant goldfields located primarily in Victoria. However, these fields were later overshadowed by the landmark discoveries at Coolgardie and Kalgoorlie in central Western

Australia during the 1890s. This firmly established the country's position as a leading global gold producer, with Western Australia's abundant reserves maintaining Australia's dominance in the gold market to the present day (Mudd, 2007). Silver mining operations, on the other hand, are primarily situated in Australia's central and eastern regions.

By the late 20th century, China has become a dominant player in hard commodities, surpassing traditional producers of aluminum, copper, gold, and silver. This rise has been driven by its vast natural resource base and aggressive industrial policies. Aluminum has been mainly produced in the eastern provinces, such as Shanxi, Shandong, and Henan, as well as in the southern regions of Chongqing, Guizhou, and Guangxi. Gold has been mostly mined in the North China Craton, a geologically rich area spanning central to northeastern provinces (Zhou et al., 2002). Meanwhile, silver has mainly been produced in the south, typically as a by- or co-product from polymetallic and gold mines (Wu et al., 1993; Xie et al., 2008; Cheng et al., 2014).

The Soviet Union was another leading producer of hard commodities, both in the realm of energy and metals. The discovery of massive oil and gas reserves in the Volga-Urals region and West Siberia in 1950s and 1960s laid the foundation for the Soviet energy sector (Ermolaev, 2017). In the metals sector, the development of the Norilsk mining and smelting complex, along with major plants in the Kola Peninsula and the Urals, was pivotal in boosting aluminum, copper, and nickel production (CIA, 1954, 1959, 1974; IMF, 1991). In addition, precious metals, especially gold, were primarily extracted from Eastern Siberia and the southern part of the Far Eastern region (CIA, 1962; IMF, 1991; Levine and Van Oss, 1993). Following the collapse of the Soviet Union, Russia inherited much of this infrastructure and has continued to play a significant role in the production of these commodities.

The United States has long been a leading producer of energy commodities, as well as industrial and precious metals. For the past 50 years, it has been a dominant force

in crude oil and natural gas production, with most of this activity concentrated in the south. From 1970, the US was a top producer of aluminum, primarily in the northwestern and northeastern states (U.S. Congress, Office of Technology Assessment, 1990), while copper production was concentrated in Arizona, Utah, and New Mexico. The US also remained a key player in gold production, with more than 70% of its output since the 1980s coming from southwestern states.

Moreover, various other nations have played significant roles in the hard commodities market. Canada, for example, was once a leading producer of a wide range of industrial and precious metals, as well as natural gas. However, by the mid-1990s, increasing international competition led to a decline in Canada's dominance in the production of many of these commodities. Despite this, Canada remained one of the top three global producers of aluminum, nickel, and natural gas until the mid-2010s, when it was overtaken by other nations. Furthermore, other countries continue to rank among the top three producers but tend to focus on just one or a few specific commodities, such as Chile and Peru for copper, Indonesia and the Philippines for nickel, and Saudi Arabia for crude oil. A detailed commodity-level breakdown is provided below.

A.2.1 Energy

This paper focuses on two key energy commodities: crude oil and natural gas. For simplicity, it does not consider refined products such as gasoline, diesel or jet fuel.

Since the 1970s, Saudi Arabia, the Soviet Union (then, Russia) and the United States have been the largest producers of crude oil. The discovery of oil in Saudi Arabia in 1938 was a significant success, jump-starting the industry in the region, particularly in the Eastern Province, where it remains concentrated as of 2021. In the 1970s, most of the Soviet Union's oil production was concentrated in the Volga-Urals region and in West Siberia. Following the collapse of the Soviet Union, Russian oil production continued

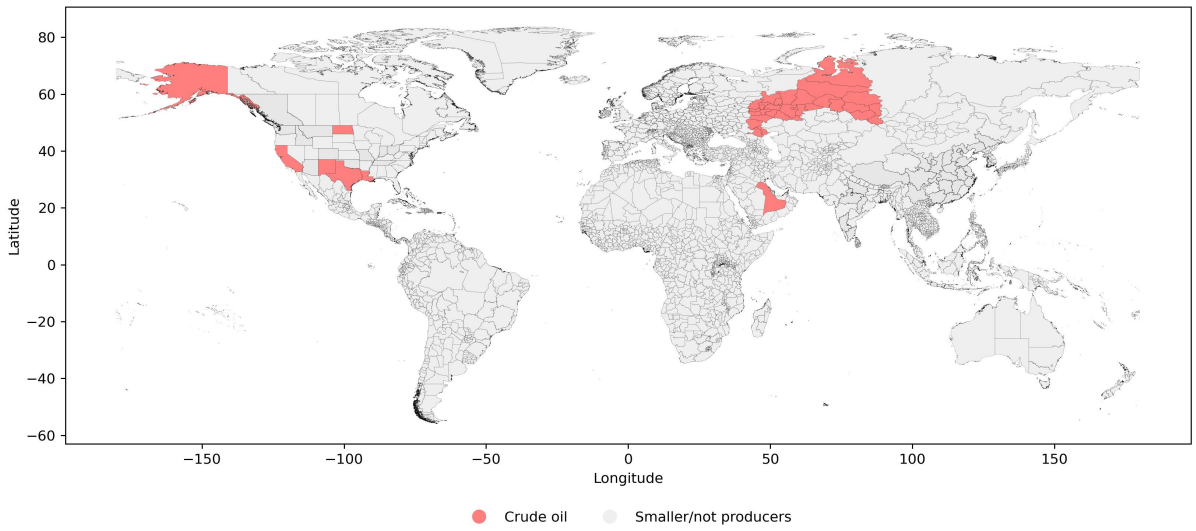


Figure 11: Core producing areas of crude oil.

This figure displays the core producing areas of crude oil for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

to be prominent in these same regions (Ermolaev, 2017). Finally, in the United States, most of the industry concentrated in Alaska, California, Louisiana, New Mexico, North Dakota and Texas according to EIA data on state-level production.

The natural gas market of the 1970s was mostly dominated by the Soviet Union and the United States. The Soviet Union's gas production was also concentrated in the Volga-Urals region and in West Siberia. After the Soviet collapse, Russia maintained a dominant position in the market, largely leveraging these areas (Ermolaev, 2017). In the United States, the main production area since the 1970s has been in the south.[18]. Overall, the US has remained dominant up to the present day. Canada and the Netherlands competed for third place, with the latter eventually securing it in 1974. The Netherlands' core production areas were primarily in the northeast around Groningen (Breunese et al., 2005). However, the Netherlands lost its dominant position in 1982, again, to Canada.

[18]It is worth noting that since the early 2010s, Ohio and Pennsylvania have also emerged as major natural gas producers, thanks to the development of new horizontal drilling techniques used to exploit the Marcellus shale formations. However, their 2021 production was significantly lower than that of the traditional southern states, accounting for less than 30% of total US production.

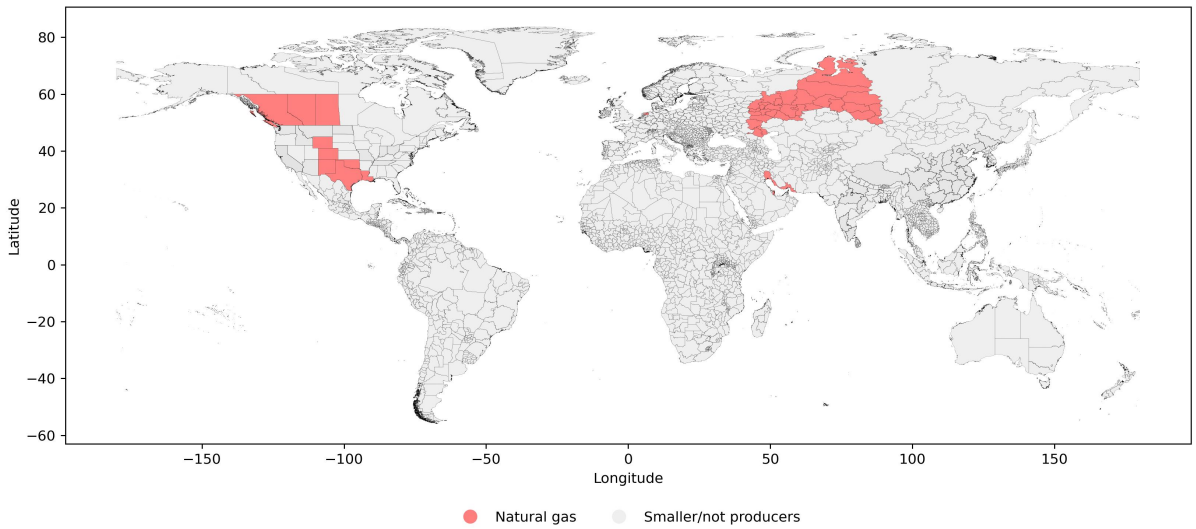


Figure 12: Core producing areas of natural gas.

This figure displays the core producing areas of natural gas for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

The latter maintained this position for nearly two decades until Qatar (2012 to 2013) and then Iran surpassed it.

A.2.2 Industrial metals

Aluminum, copper, and nickel production has been dominated by various countries over the past decades, shaped by the availability of natural resources, technological advancements, and geopolitical shifts.

In the 1970s, the leading aluminum-producing countries were Canada, Japan, the Soviet Union, and the United States. Canada concentrated its production in Québec, where hydroelectric energy was abundant and affordable (Canadian Department of Energy and Resources, 1972; Cranstone, 2002). Japan's aluminum industry grew rapidly, surpassing Canada in 1972, but a severe aluminum crisis in the late 1970s devastated Japan's industry, allowing Canada to regain the top spot in 1981 (Samuels, 1983). At the same time, the United States was also a major aluminum producer, with significant operations

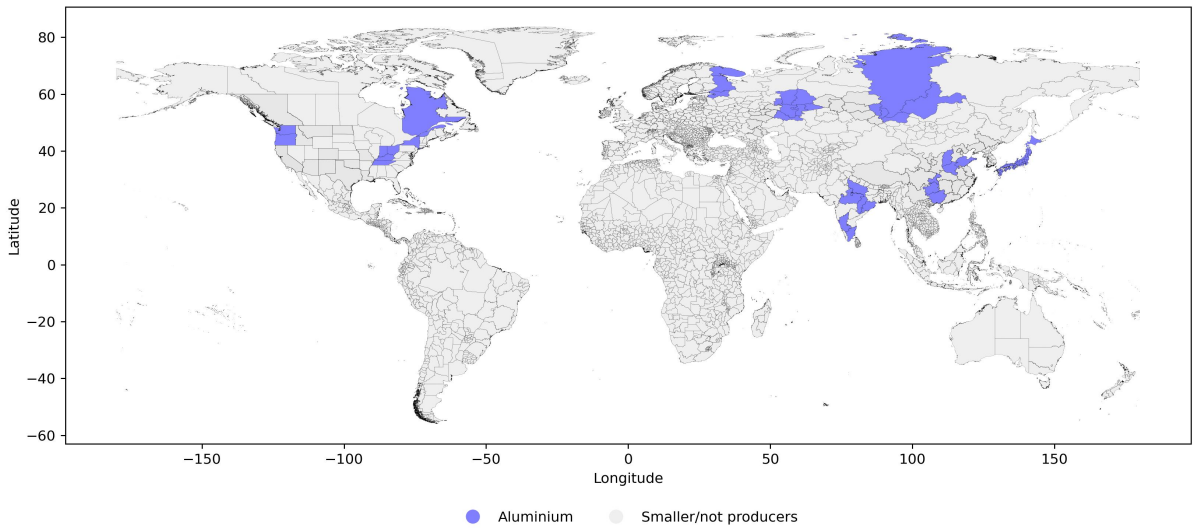


Figure 13: Core producing areas of aluminium.

This figure displays the core producing areas of aluminium for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

spread across the northwestern and northeastern states (U.S. Congress, Office of Technology Assessment, 1990). The Soviet Union, meanwhile, maintained a strong aluminum industry, supported by a network of large production facilities that had been expanded and modernized after World War II (CIA, 1959). After the Soviet Union's dissolution in 1991, Russia inherited its aluminum infrastructure and remained a major producer, with key facilities spread across regions such as Leningrad, Karelia, Murmansk, the Urals, and Eastern Siberia, maintaining the geographical distribution established during the Soviet era (Smirnov, 1996). The other leading producers remained the same until the 2000s. Indeed, the United States, once a global leader, lost its position of dominance and was surpassed by China in 2001. Besides, Canada was surpassed by India in 2016.

In the 1970s, Canada, Chile, the Soviet Union, and the United States were the leading copper-producing nations. Chile capitalised on its extensive copper reserves in the central and northern Andes Mountains (Cacciuttolo and Atencio, 2022). The Soviet Union maintained substantial production in the Kola Peninsula, Urals, and Krasnoyarsk (CIA,

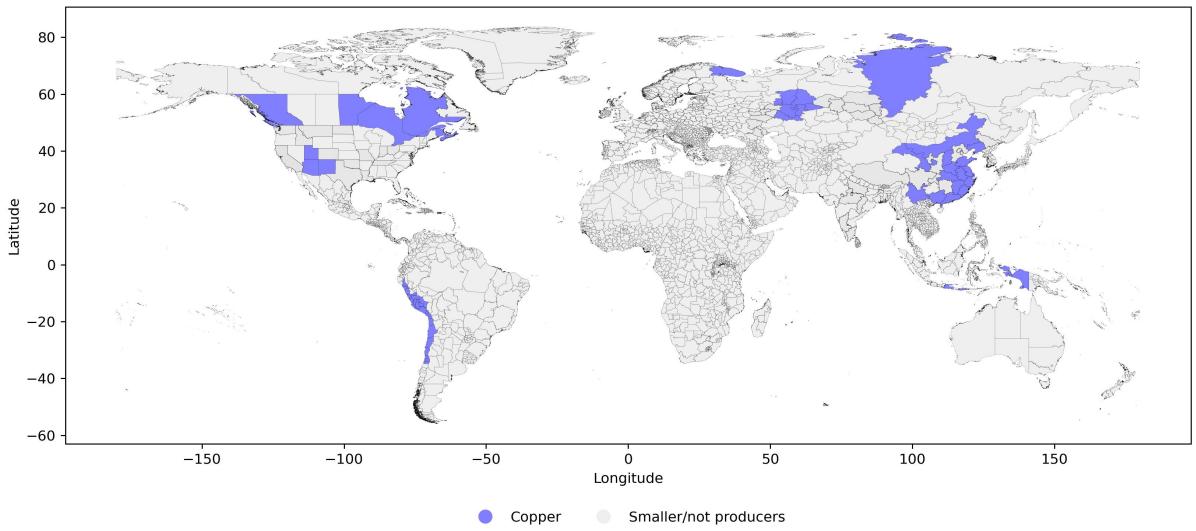


Figure 14: Core producing areas of copper.

This figure displays the core producing areas of copper for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

1954, 1974; IMF, 1991). The United States was also a key player, with major operations in Arizona, Utah, and New Mexico, according to relevant USGS Minerals Yearbooks. In the early 1990s, following the Soviet Union's collapse, Canada rose to prominence, overtaking the former superpower in the copper industry. However, Canada's dominance lasted only a few years. By the early 2000s, several other countries were competing for the podium. Eventually, China and Peru claimed it in the mid-2010s. In 2021, China's position was overtaken by the Democratic Republic of Congo. Chile remained in the podium since the early 1970s.

New Caledonia, the Soviet Union, and Canada dominated nickel production in the 1970s. The Soviet Union's nickel industry, like its copper counterpart, was concentrated in the Kola Peninsula, Urals, and Krasnoyarsk regions (CIA, 1954, 1974; IMF, 1991). Canada's nickel production was primarily centered in Ontario and Manitoba, benefiting from extensive deposits and well-established infrastructure (e.g., Canadian Department of Energy and Resources, 1972). In the late 1970s, Australia emerged as a significant player,

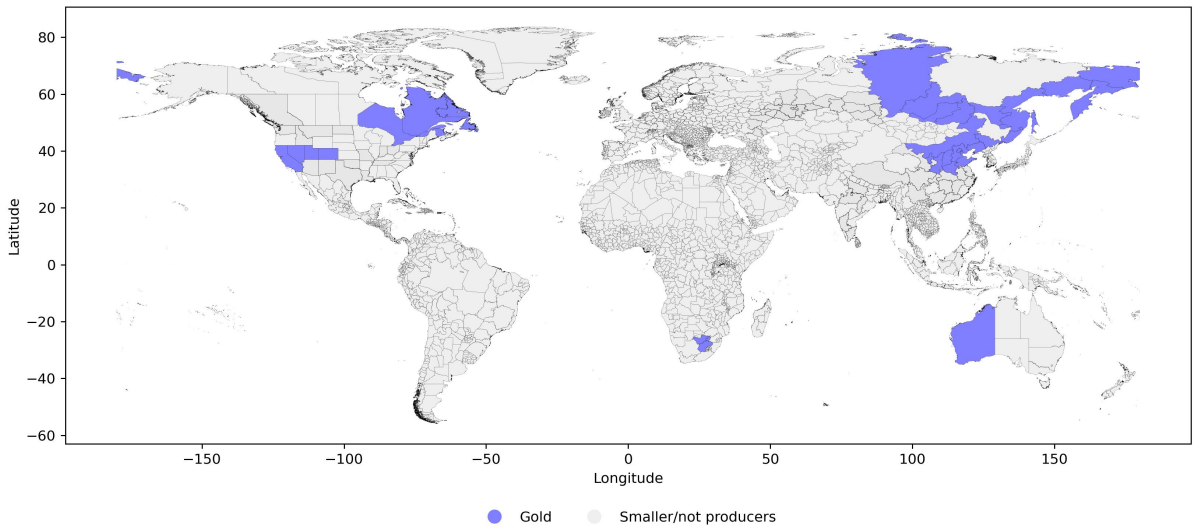


Figure 15: Core producing areas of gold.

This figure displays the core producing areas of gold for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

frequently competing with Canada for a place on the podium across the following decades. However, the landscape shifted in the early 2010s when Indonesia and the Philippines rose to prominence, overtaking both Canada and Australia. Russia, inheriting the Soviet Union's production facilities, continued to be a major producer, securing a spot on the podium alongside Indonesia and the Philippines.

A.2.3 Precious metals

From the 1970s to the mid-1980s, Canada, the Soviet Union, and South Africa were the leading gold producers. In Canada, gold is found across the Canadian Shield, in British Columbia and Nunavut, and on the island of Newfoundland. However, the majority of Canadian production has come from provinces east of Manitoba (Canadian Department of Energy and Resources, 1972, and subsequent yearbooks). In the Soviet Union, gold was predominantly found in the eastern part of the country, beginning in Krasnoyarsk and extending to the southern regions of the Far East (CIA, 1962). South Africa's produc-

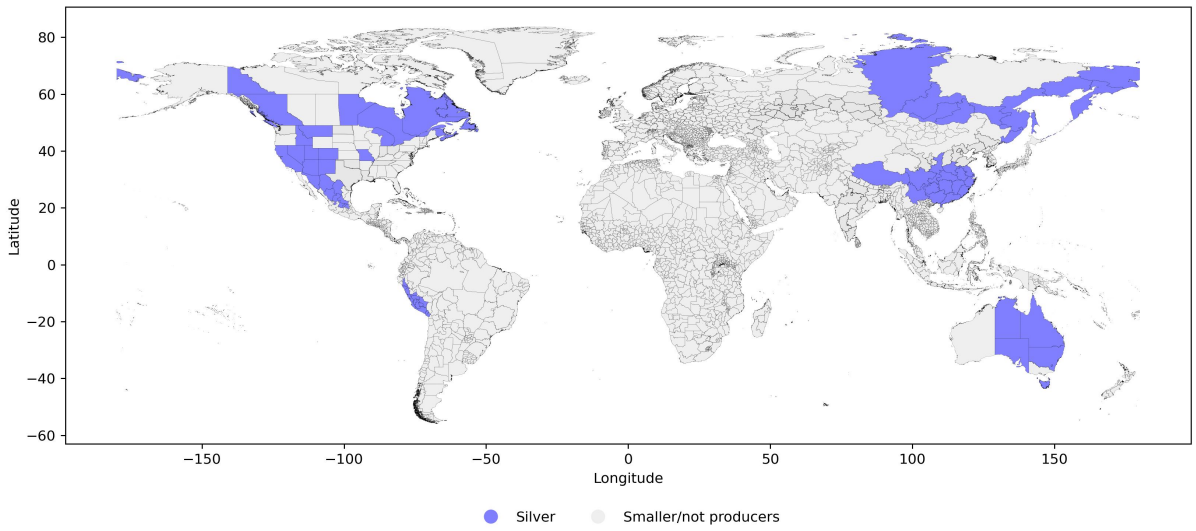


Figure 16: Core producing areas of silver.

This figure displays the core producing areas of silver for all countries that, in any given year between 1970 and 2021, ranked among the top three global producers. These core regions are mapped at the subnational level corresponding to first-order administrative units in GADM 3.6.

tion was centered around the Witwatersrand Basin in the northeast. In the following decade, while South Africa remained a top producer, Canada and the Soviet Union were surpassed by Australia, primarily due to production centers in Western Australia (Mudd, 2007), and the United States. In the United States, several western states produce gold. However, over 70% of gold production has come from southwestern states since the 1980s. Finally, in the early 2010s, South Africa and the United States were overtaken, due to the emergence of China and the resurgence of Russia. Whereas Russia's gold production remains concentrated in the same areas utilized during the Soviet era, China's gold reserves are more unevenly distributed across its vast territory. The most significant deposits are found in the region known as the North China Craton, a geologically rich area that stretches across the central to northeastern provinces, including regions like Shandong, Henan, and Hebei (Zhou et al., 2002).

During the 1970s, Canada, Mexico, and the Soviet Union were the leading producers of silver. Canada's rich silver deposits were located in both eastern and western provinces

(Canadian Department of Energy and Resources, 1972, and subsequent yearbooks). Mexico's silver production was primarily concentrated in the Silver Belt, a historic mining region in the central-northern states, where silver extraction has been significant since colonial times (Merrill, 1930). In the Soviet Union, silver production was concentrated in regions where gold was also mined (Levine and Van Oss, 1993). From the 1980s through the late 1990s, Mexico, Peru, and the United States were dominant players in the global market. In 2000, however, the United States was overtaken by Australia, which itself was surpassed by China in 2001. Since then, China has consistently remained among the top producers, alongside Mexico and Peru. China has been producing silver predominantly as a by- or co-product from polymetallic and gold mines, mostly in its central and southern provinces (Wu et al., 1993; Xie et al., 2008; Cheng et al., 2014).